The Chemomentum Data Services – A Flexible Solution for Data Handling in UNICORE

Katharina Rasch[1], Robert Schöne[1], Vitaliy Ostropytskyy[2], Hartmut Mix[1], and Mathilde Romberg[3]

[1] Technische Universität Dresden, Dresden, Germany
[2] University of Ulster, Coleraine, Northern Ireland
[3] Forschungszentrum Jülich GmbH, Jülich, Germany

Abstract. This paper introduces the Chemomentum data services, a UNICORE-based flexible solution for managing large amounts of data and metadata produced in a Grid. In order to store and manage the increasing amounts of data produced in Grid environments, a highly scalable and distributed Grid storage system is needed. However, the simple storage of data is not enough. To allow a comfortable browsing and retrieving of the data, it is crucial that files are indexed and augmented with metadata. This paper analyses integrated solutions that already provide these functionalities for their features and shortcomings. Incorporating the conclusions drawn from the examination, an architecture for a revised data management solution is presented. This system provides the means to store files with augmenting extensible metadata. It allows also to browse through data using metadata, handle ontologies and transparently access external data sources. In the current stage, most of these functionalities are implemented and running in a distributed environment.

1 Introduction

Grid technologies are starting to realise their large potential to provide innovative infrastructures for complex scientific and industrial applications. The UNICORE 6 Grid computing solution [1] already provides a seamless, secure, and intuitive access to distributed Grid resources. However, to make Grids more useful for knowledge-oriented applications such as decision support and risk assessment, more effort in the fields of semantics, metadata and knowledge management in distributed, heterogeneous environments is needed.

The Chemomentum project [2] aims to fill these gaps by taking up and enhancing state-of-the-art Grid technologies and applying them to real-world challenges in computational chemistry and related application areas. The Chemomentum project specifically supports the European REACH (Registration and evaluation of chemicals) [3] initiative aimed at optimising risk assessment strategies.

To be in line with the REACH initiative it is crucial that all results of calculations can be subsequently evaluated and reproduced. This means that not only final experiment results may need to be stored permanently, but also intermediate results and metadata that describe the provenance of the result data,

E. César et al. (Eds.): Euro-Par 2008 Workshops, LNCS 5415, pp. 84–93, 2009.

e.g. the workflow and applications that produced the data and information on the user who ran the workflow.

It is foreseeable that a considerable amount of data and metadata will be produced. Depending on the type of executed workflow and the level of detail, the size of metadata can reach hundreds of kilobytes, the data even up to gigabytes. Therefore a highly scalable, distributed and decentralised Grid storage solution is needed. The storage of data and augmenting metadata is the fundamental requirement for such a storage system. An important feature is the possibility to define sets of metadata, which describe data used or generated within specific workflows. Clearly defined and well structured metadata sets are crucial to provide adapted user interfaces for diverse purposes.

A transparent access to external data sources (e.g. web-based chemical databases) is also desirable. Data from such sources could then be included as input data of a workflow just like data in the internal storage. Ontologies describing synonyms in a specific context can be used to ease the retrieving of metadata. If a user searches for data connected with the chemical name 'H_2O' additional results for 'water' or 'dihydrogen monoxide' could be returned as well. By providing an integrated conversion of values between scientific units, the advantages of using ontologies can be even further extended.

As one solution, the data management system developed in the Chemomentum project is presented in this paper. The system is not limited to the scientific domains in focus of the Chemomentum project, but is a general approach to the challenges of a Grid storage solution.

2 Related Work

Handling of data and metadata in a Grid environment has been described before. Most of these approaches, however, provide only data access but no or insufficient handling of metadata. Nevertheless, there are existing solutions for the Globus[4] and gLite[5] Grid middlewares.

The Storage Resource Broker SRB [6] handles data and metadata in a data Grid, a digital library, a persistent archive, or a distributed file system. It has been developed by the San Diego Supercomputer Center and is commercialised by Nirvana. The SRB provides a uniform data access to different storage types over a network as well as the replication of files. It is often used in Globus Grid projects. The SRB could not be used because the support for metadata extension and schema handling is not sufficient. The need for a commercial license to deploy SRB within Chemomentum was an important disadvantage.

Another solution that offers both, data and metadata handling, is the Arda Metadata Catalogue Project AMGA [7] which is used with the gLite Grid middleware. It supports user-defined metadata to describe the data stored in the system. These schemas are, however, not shared between users, a crucial point against the adoption for the Chemomentum project. The AMGA server is a monolithic C++ implementation, making it platform dependent and inhibiting the installation of just a subset of the functionalities.

A data and metadata management system that also integrates ontologies, the conversion of scientific units and the transparent access to external data sources could not be found.

3 Overview of the Chemomentum Data Management System

The data management system within Chemomentum provides data storage and retrieval functionality and a global data view independent of the actual data location. A crucial point in designing the system was to build lightweight, specialised services dealing with the different types and sources of data and metadata. The services have well-defined interfaces that allow the installation and running of single services or of subsets of the complete system. The interfaces also support the easy plugging in of extensions, e.g. to access other storage systems.

The heart of the data management system is the Documented Data Space DDS (see figure 1). It is composed of metadata databases, data storages and a location database. The data storages contain data in flat files, typically input and output data produced by workflows. The location database acts as a global file location directory by indexing those files and assigning them globally unique logical names. The metadata databases contain metadata that describe the files in the data storages, referencing them by their logical names. Access to the components of the DDS is provided by a set of three specialised services.

The central interface to the data management system is the Data Management System Access Service (DMSAS). It forwards service requests to the appropriate service(s), collects the results and returns them to the requester. The Database

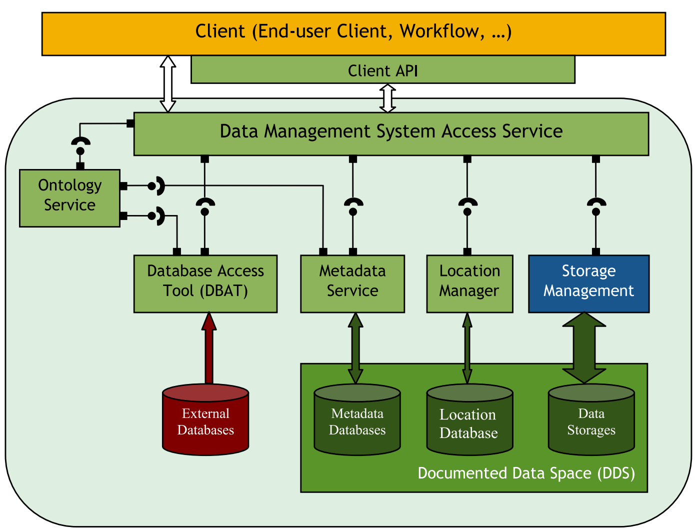

Fig. 1. System architecture

Access Tool (DBAT) serves as a uniform interface to external data sources. It transforms a query to the data management system into the native query language used by the external data source, queries the source and transforms the result back into a format the client understands. The ontology service provides the access to external ontology services and knowledge about types, units and vocabulary to interpret the data.

For the other side of the communication, which includes the workflow system, GUIs and other clients, a special API is provided. The ClientAPI procures objects and methods which enable programmers to access the data management system without any knowledge of the underlying web service schemas and security issues. Built on top of it, special Eclipse RCP [8] views are implemented, which enable the users to create new metadata schemas and to create, query and download data and metadata.

The data management system uses the wsrflite web service framework of the UNICORE 6 Grid middleware. UNICORE 6 implements the Grid standard WSRF, compliant with the Open Grid Forum's Open Grid Services Architecture (OGSA)[9]. It offers strong security based on industry standards such as the X.509 public key infrastructure. The communication with UNICORE 6 services is protected by mutual authentication. All services that form the data management system are implemented as web services that can be run distributed across the Grid.

4 Detailed Architecture

Following the short introduction of the Chemomentum data management solution, this section will now focus in detail on the key characteristics of the system.

4.1 Metadata Modelling

While Chemomentum itself is specifically aimed at the computational chemistry domain, the ambition is to develop Grid technologies applicable in any scientific, economic or other domain. Hence, the ability to cater for arbitrary, extensible metadata schemas instead of limiting the available metadata items to a fixed set was a crucial design criteria for the data management system. The naive approach of using a general metadata model like RDF (Resource Description Format)[10] for this was deemed unfeasible. Having a completely unrestricted modelling of metadata can be more limiting than a fixed set of items. The user needs some knowledge about the data's semantics to interpret the results. Without a restricting metadata schema describing the semantics, a set of metadata triples is just a bunch of information, difficult to interpret and process.

Instead, a metadata model similar to a relational database was chosen for Chemomentum. The metadata schema for a scientific domain is declared (and can also be extended) by a scientific administrator. The schema includes the basic information about names and data types of metadata items, but also further

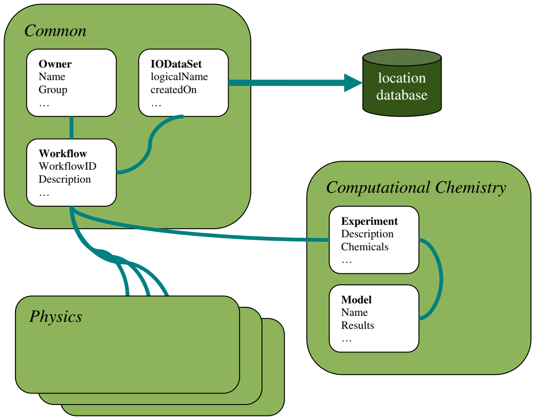

Fig. 2. Metadata modelling

information like a detailed description and units. Attributes can be linked to each other and to logical names of files.

Figure 2 shows an exemplary metadata schema setup. The domain 'common' is a fixed domain that contains basic provenance information, e.g. the workflow that produced the data, the owner of the data and the applications used. The other domains are defined additionally. To gain knowledge about a domain, a user client can request the domain's metadata schema from the data management system and provide the user with a dynamically created graphical interface adapted to the domain.

Aside from manual uploads, most of the metadata is automatically filled by the UNICORE 6 Workflow System that runs the applications and uploads the resulting files into the DDS. While this is a straightforward process for the fixed 'common' domain, the handling of additional domains needs further input from the scientific administrator. The Workflow System has to be provided with the knowledge of how to extract metadata from the input and output of applications and how to map it to the metadata items required by the schema.

Users can manually annotate metadata with additional information like rating and comments for files. Also, metadata which could not be retrieved automatically, can be supplemented later.

4.2 The Documented Data Space

The Documented Data Space (DDS) is split into three types of data storages: the metadata databases, the location database and the file storages. The location database acts as a global file catalogue that maps globally unique logical file names onto the actual physical locations of files and directories. The files are only referenced by their logical names in the system. This mechanism allows for an easy migration of files from one file storage to another. The location database also supports the replication of files, increasing the reliability as well as the performance

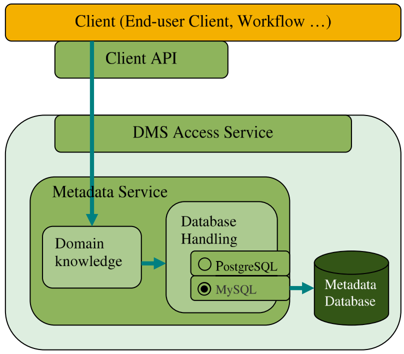

Fig. 3. Database handlers

of the system. Access to the location database is given through the Location Manager. This service provides the functionalities to store, look-up, update and delete locations.

The data files themselves are stored on file storages. The file storages are implemented using the UNICORE 6 Storage Management Services that offer access to storages in a Grid. The current transfer is based on OGSA ByteIO standard, but can be extended to other protocols as well.

Metadata that augment the files in the file storages is kept in the metadata databases of the DDS, as described in section 4.1. The metadata databases are accessed by the Metadata Service, which offers functionalities to store, query, update and delete metadata. The metadata modelling acts as an additional layer above the underlying database systems (see figure 3). All metadata exchanged between the client and the Metadata Service operates on one or more domain schemas and not directly on the database system. Pluggable database handlers provide the mapping between the client data and the specific underlying database technology. Special care is taken for write access to metadata items marked as provenance. The DDS implements a row-based security on the data in the metadata databases. Access control lists mark whether a user, group or virtual organisation has the right to read the data, to write it or to change the permissions. Users with high-level security and privacy needs can also provide their own private DDS on their own site.

4.3 The Database Access Tool

The Database Access Tool (DBAT) provides a comfortable access to external data sources. External data sources can be various database systems, web services or simply flat files. Like the Metadata Service, the DBAT is called only by the Access Service. The DBAT extracts the name of the external data source to be accessed from the query and transforms the query accordingly to the native query standard of the data source. The native query standard is, for example, SQL for data in relational database systems and XML-based service calls for data provided by web services. The transformation is done by using the aforementioned database handlers. The result set of the query is returned to the Access Service for further processing.

Just like the Metadata Service, the DBAT uses metadata schemas to know the correct set up for the scientific domains it supports. In the current status of the project the Ecotoxicology[11] databases Aquire and Terretox as well as the Protein Data Bank (PDB) [12] are implemented.

4.4 The Ontology Service

The Ontology Service supports queries by providing additional, domain specific information about data - for example, synonyms to broaden queries to data services or knowledge about the conversion of values between different units.

To broaden a query, the Ontology Service contacts external ontology services to examine the query for data items that can be represented in a different fashion. At the present stage, for example, the external ontology service ChEBI[13] is used to look up the molecular names of small molecules. The synonyms of data items are then aggregated into the query. When the broadened query is executed, it will return more exhaustive results then the original query.

The Ontology Service also provides the means to automatically convert values in requests between user-provided units and the units used in the databases. The set of units supported by the Ontology Service is not fixed, scientific administrators can set up own unit groups and conversions.

4.5 The Data Management System Access Service

The Data Management System Access Service (or more conveniently Access Service) is the sole entry point into the data management system. Therefore, only one open port is needed to operate the system. The Access Service is a lightweight service that can easily be run in several parallel instances to avoid bottlenecks. It implements the logic to bundle the functionalities of the underlying services to a higher-level interface, that offers comfortable methods to query, store, update and delete data and metadata.

The user is presented with a uniform interface for queries to the DDS and to external databases. Queries made are directed to the appropriate services. The Access Service also manages distributed queries to multiple Metadata Services. With the help of the Ontology Service, queries are automatically broadened to improve the results and any unit conversions necessary between the user's units and the units used in the system are performed. Files requested by the users are looked up in the DDS and made available for an easy download.

In particular the process of storing data and metadata demands the interplay of multiple services. The Access Service coordinates the following tasks:

1. Check the metadata for validity.
2. Perform necessary unit conversions.
3. Upload the file(s) to the storage(s).

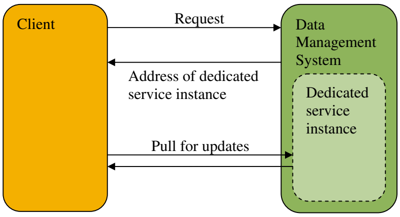

Fig. 4. Asynchronous processing

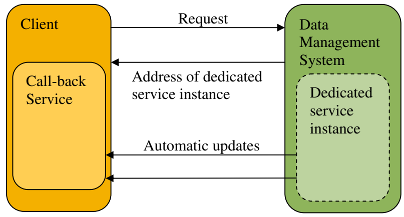

Fig. 5. Automatic call-backs

4. Register the files with the Location Manager.
5. Store the metadata in the metadata database.

A transaction mechanism automatically rolls back all changes already made if one of those steps fails.

Considering the huge amount of data that can be processed in one request, the data management system uses asynchronous execution (see figure 4). The Access Service spawns a dedicated service that handles the request and returns its address to the user. The user can then poll this service for the current status of the processing. Alternatively, the user can provide the address of a call-back service that is updated automatically by the system (figure 5).

4.6 The Client API and the Eclipse Plugin

The Client API provides users and developers with a convenient Java based interface to the data management system. Its key features are the comfortable support for handling large sets of metadata items and for the asynchronous execution mode.

The Client API supports the access to the system in various levels of complexity. It consists of low-level implementations that resemble the web service

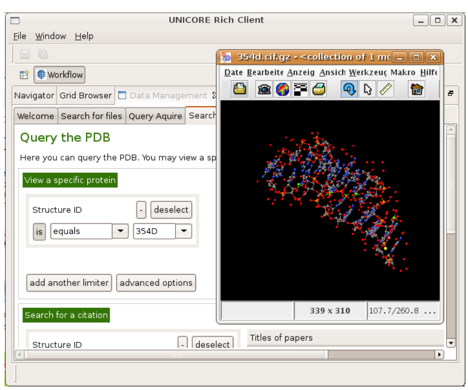

Fig. 6. Query external data sources (e.g. PDB)

interface, but it includes also high-level implementations that allow the client developer to start a request with a single line of code.

On top of the Client API an Eclipse Plugin has been developed, which smoothly integrates with the UNICORE RCP client. It can be used to access the data management system and, for example, upload and download files, browse through metadata or access the external databases (figure 6).

5 Current Status and Future Work

In the current stage of development, all services described here are implemented and can be installed on any server that is running a Java Virtual Machine in version 1.5 or higher. The components have already been installed in a distributed fashion and are running as part of a testbed for the Chemomentum project. GUI components that allow browsing and storing data and metadata are available and are in use for testing the system.

Within the year 2008, the final version of the Chemomentum software – including the data management system – will be available. In addition to the functionality presented in this paper, this version will support more external databases and contain a sophisticated administration interface. Additional GUI components for a simple and flexible access to the system will also be provided.

6 Conclusions

The Chemomentum data management system is a flexible, distributed and user-friendly approach to data management in the Grid. Because of its extensible metadata system, it is not limited to specific use cases, but can be used in arbitrary scientific, economic or other fields.

The data management system can be deployed on any server running the Java Virtual Machine. Its modular design allows also for a distributed deployment on multiple servers and sites. The data in the data management system is protected by sophisticated security and safety solutions based on the UNICORE 6 framework. The data management system supports web service standards like WSRF and OGSA ByteIO. The Eclipse Client and the ClientAPI allow an easy integration into other projects without knowledge of the internal handling of data and meta data. Despite of its UNICORE background, the system is deployable also in non UNICORE based Grid environments. This makes it an excellent solution for Grid data management.

Acknowledgement

This work has been funded by the European Commission under contract no. IST-5-033437.

References

1. Riedel, M., et al.: Web Services Interfaces and Open Standards Integration into the European UNICORE 6 Grid Middleware. In: Proceedings of 2007 Middleware for Web Services (MWS 2007) Workshop at 11th International IEEE EDOC Conference The Enterprise Computing Conference (2007)
2. Schuller, B., Demuth, B., Mix, H., Rasch, K., Romberg, M., Sild, S., Maran, U., Bała, P., del Grosso, E., Casalegno, M., Piclin, N., Pintore, M., Sudholt, W., Baldridge, K.K.: Chemomentum - UNICORE 6 based infrastructure for complex applications in science and technology. In: Bougé, L., Forsell, M., Träff, J.L., Streit, A., Ziegler, W., Alexander, M., Childs, S. (eds.) Euro-Par Workshops 2007. LNCS, vol. 4854, pp. 82–93. Springer, Heidelberg (2008)
3. REACH (Registration and evaluation of chemicals), Regulation (EC) No 1907/2006 of the European Parliament and the Council, http://echa.europa.eu/reach_en.html
4. Foster, I.: Globus toolkit version 4: Software for service-oriented systems. In: Jin, H., Reed, D., Jiang, W. (eds.) NPC 2005. LNCS, vol. 3779, pp. 2–13. Springer, Heidelberg (2005)
5. Laure, E., et al.: Programming the Grid with gLite. Computational Methods in Science and Technology 12(1), 33–45 (2006)
6. Rajasekar, A., et al.: Storage Resource Broker - Managing Distributed Data in a Grid. Computer Society of India Journal, Special Issue on SAN 33(4), 42–54 (2003)
7. Santos, N., et al.: Distributed Metadata with the AMGA Metadata Catalog. In: Workshop on Next-Generation Distributed Data Management HPDC-15, Paris, France (2006)
8. Eclipse Rich Client Platform homepage, http://wiki.eclipse.org/index.php/Rich_Client_Platform
9. Foster, I., et al.: The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration. In: Open Grid Service Infrastructure WG, Global Grid Forum (2002)
10. McBride, B., et al.: The Resource Description Framework (RDF) and its Vocabulary Description Language RDFS. In: International Handbooks on Information Systems, Handbook on Ontologies (2004)
11. Ecotoxicology databases homepage, http://cfpub.epa.gov/ecotox/
12. Berman, H.M., et al.: The Protein Data Bank. Nucleic Acids Research 28, 235–242 (2000)
13. Degtyarenko, K., et al.: ChEBI: a database and ontology for chemical entities of biological interest. Nucleic Acids Research 36, 344–350 (2008)