Chapter 3 Data grids

3.1 Introduction

In the first chapter the reader was introduced to data grids. In this chapter more information about grid computing and data grids is presented with the ultimate goal of showing that the data grid approach as enabler for SDI data sharing is both innovative and new, and also extremely relevant at the current point in time. In reference to Figure 1, this chapter relates mostly to the Computer Science discipline and provides an interpretation of work on data grids in relation to both the Compartimos reference model and, as well as a data grid approach to address databases in national SDI.

In the first section, 3.2, of this chapter the origins, vision and current reality of grid computing are related and concluded with the author's interpretation of an outlook to what the future may hold. Next, in section 3.3, the layered and service-oriented aspects of the common Grid architecture are discussed and related to Compartimos. Subsequently, in section 3.4, data grids, a special case of grid computing, are introduced. A number of existing data grid implementations are analyzed and compared to the requirements for an address data grid in an SDI environment for which the Compartimos reference model, presented in Chapter 4, has been designed. The chapter is concluded in section 3.6 with a discussion of research that is related to the work in this dissertation, confirming that the data grid approach in Compartimos, a reference model for a data grid as enabling platform for address data in an SDI environment, is innovative and new, but also extremely relevant at the current point in time.

3.2 Grid computing

3.2.1 Origins

Grid computing started in the late 1990s as a distributed infrastructure with the main aim of solving specific Grand Challenge applications through high performance computing. The term "Grid" was coined as an analogy to an electrical power grid, envisioning users tapping into a computational grid, similar to consumers currently tapping into an electrical power grid. Since those initial days, the concept of a grid has evolved to address the general need for flexible, secure, coordinated resource sharing among collections of individuals, institutions and resources (Foster and Kesselman 1999).

There is an abundance of definitions for a grid, but one commonly cited definition, and the one that is used in this dissertation, is Foster's (2002) three point check list, stating that a grid is a system that

1. coordinates resources that are not subject to centralized control;
2. delivers non-trivial qualities of service, and
3. uses standard, open, general-purpose protocols and interfaces.

In other words, a grid integrates and coordinates resources and users that live within different control domains such as different administrative units of the same company or different companies altogether, and addresses the issues of security, policy, payment, membership, and so forth that arise in these settings. A grid is built from multi-purpose, standard, open protocols and interfaces that address such fundamental issues as authentication, authorization, resource discovery, and resource access. A grid allows its constituent resources to be used in a coordinated fashion to deliver various qualities of service, resulting in a combined system that is significantly greater than that of the sum of its parts.

Grid systems are used by virtual organizations (VOs) comprising a set of individuals and/or institutions, having direct access to computers, software, data, and other resources for collaborative problem solving or other purposes. The real and specific problem that underlies the Grid concept is this *coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations*, originating from an emerging need for collaborative problem-solving and resource brokering strategies in industry, science and engineering (Foster *et al.* 2001).

3.2.2 The vision

Foster and his co-authors' vision of how virtual organizations can be enabled to collaborate and share resources in order to achieve a common goal is described in terms of a Grid architecture in the two papers, *The Anatomy of the Grid* (Foster *et al.* 2001) and *The Physiology of the Grid* (Foster *et al.* 2002). This Grid architecture has subsequently evolved into the Open Grid Services Architecture (OGSA) published by the Open Grid Forum (2006), a vision of a broadly applicable and adopted framework for integration, virtualization, and management of resources and services within distributed, heterogeneous, dynamic virtual organizations, which is reflected in the definition of a Grid in the OGSA Glossary of Terms (OGF 2007c):

> A ***system***
> that is concerned with the
> ***integration, virtualization, and management of services and resources***
> in a ***distributed, heterogeneous environment***
> that ***supports virtual organizations*** (collections of users and resources) across
> traditional administrative and organizational domains (real organizations).

Standardization is a key requirement for realizing this vision of the Grid so that resources that are provided by different vendors and operated by different organizations can be discovered, coordinated and managed in a grid. The Open Grid Forum (OGF) is an open community committed to driving the rapid evolution and adoption of applied distributed computing and the work of OGF is carried out through community-initiated working groups, which develop standards and specifications in cooperation with other leading standards organizations, software vendors, and users (www.ogf.org).

OGSA is a service-oriented architecture that addresses the needs for standardization by describing the requirements and scope of core capabilities that are required to support Grid applications in industry, engineering and science. Communication in the OGSA architecture happens through Grid services, i.e. Web services that provide a set of well-defined interfaces and follow specific conventions. Grid services are specialized Web services and form the Grid-enabling layer between applications and the resources on a lower level, as illustrated in Figure 7. OGSA thus extends the power of the Web services framework, and integrates the Grid and Web technologies to the extent that the distinction between the two is blurring. OGSA is the 'blueprint' for standards-based grid computing (OGF 2008a) and although these Web services are still in the process of being standardized, the Globus Toolkit, an open source software toolkit with implementations of these Web services is fast becoming the *de facto* standard. Compartimos, the reference model that is presented in Chapter 4, is based on the OGSA architecture.

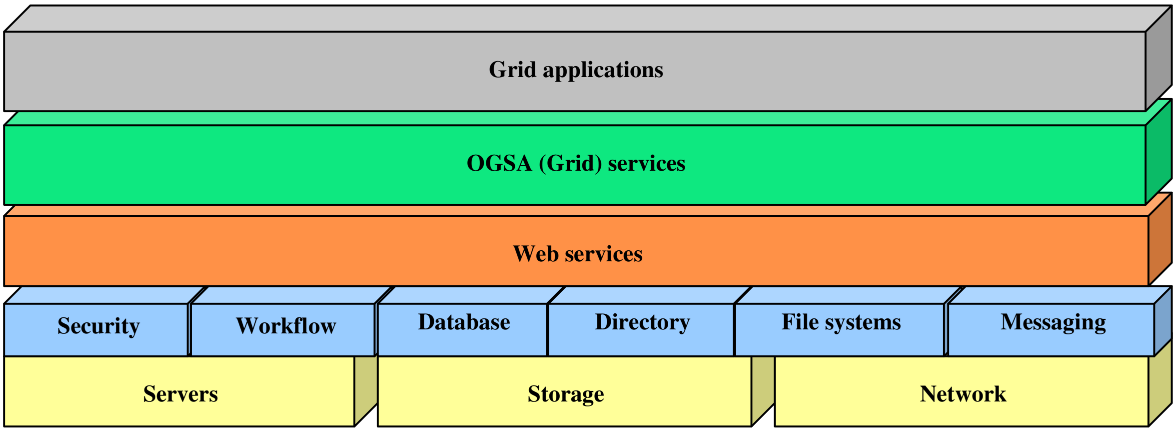

Figure 7. OGSA, adapted from Baker *et al.* (2005)

3.2.3 The current reality

The evolution of a new technology can be divided into an initial developmental phase and a later phase of mass adoption (although some technologies never reach this phase). In the developmental phase the primary concern is the technology itself – how it is built and how it works – and during this stage the users are the experts themselves, the technology is highly specialized and costly to implement. Over time as the technology grows in popularity it is standardized and costs decline until it is adopted by the masses. At this stage the application of the technology, along with the ease of use, reliability, availability and cost become the primary concerns (Wladawsky-Berger 1998).

Grid technology, although popular and fast spreading, still has to become a commodity technology: setting up and maintaining a Grid environment is still quite a complex task. Consequently, skills and resources are highly specialized and therefore limited and expensive. Easy-to-use commodity software is only now starting to emerge with developments such as XtreemOS, an operating system that supports grid applications with native support for setting up and managing virtual organizations, thereby shielding a user from the low-level details of grid middleware and overcoming many of the barriers of entry for establishing virtual organizations. Cloud computing or Web operating systems where users work with Web-based, rather than local, storage and software eliminate the need for skilled resources to set-up a grid in an organization, rather the skilled resources are provided by the 'cloud service provider'. The author suspects this contributes to the current appeal of clouds. Cloud computing is at this point in time becoming more interesting, relevant, and, vendors hope, commercially viable (Coppola *et al.* 2008, Lawton 2008). Commercial viability is a precursor to becoming commodity software.

Without going into a detailed comparison between grids and clouds, with all the recent hype surrounding cloud computing, it is worth noting that clouds have different shapes (Weiss 2007), some of them showing strong resemblance to grids. According to Weiss (2007) the cloud can be seen as a *data center*: cheap, commodity hardware in large numbers controlled by an operating system that is designed to manage resources—hard drive space, memory–to replicate the kinds of intra-server channels that now coordinate events within a single physical machine. This coordination of resources relates to the definition of a grid by both Foster (2002) and later the OGF (2007c). The cloud could be a data center at a single physical location or dozens, hundreds, or thousands of data centers spread around the world, its speed and efficiency is limited by how intelligently it delegates responsibility. Thus alluding to the 'distributed, heterogeneous environment' in the OGF (2007c) definition of a grid and the lack of centralized control in Foster's (2002) definition. Another shape of the cloud is that of the *utility grid*: assuming that a Web application that is hosted in a cloud has been designed intelligently, additional machine instances can be launched on demand so that the application dynamically, and gracefully, scales up. This is exactly the early vision of Foster and Kesselman (1999) that sees users tapping into a computational grid, similar to consumers tapping into an electrical power grid. Finally, the cloud can provide *software as a service* where processing power is centralized in the cloud and liberates users to choose efficient, uncomplicated access machines that run ultra-thin clients. However, this shape of the cloud centralizes computing power (even it draws on distributed computing resources), whereas a grid aims to coordinate distributed computing resources that are not subject to centralized control.

Data grids based on standard, open, general-purpose protocols and interfaces are still in their infancy because standards and easy-to-use tools are still being developed. Standards development is a slow process, and general adoption and implementation of the standards will take another few years. In the mean time, it is worthwhile to prepare applications and application domains for the world of the Grid, because once Grid standards are in general use and easy-to-use tools are available, Grid technology holds the promise of revolutionizing the world in a fashion similar to the Internet and the Web.

3.2.4 The future

Figure 8 shows the author's interpretation of how the development of software has evolved from tightly coupled software with no software re-use to today's grids for virtual organizations that can be built by dynamically integrating loosely coupled objects on different platforms, in distributed geographic locations and over different administrative domains that deliver services that are beyond the capabilities of an individual organization. Standard, open Grid protocols hold the promise of allowing virtualization of resources until the equivalent of a virtual operating system emerges so that the entire aggregation of heterogeneous architectures can be managed in an automated fashion (Wladawsky-Berger 1998). In fact, Web operating systems, such as Fearsome Engine's ZimDesk (www.zimdesk.com) and Sun Microsystem's Secure Global Desktop (SGD) (www.sun.com/software/products/sgd/index.jsp) are already proof that such virtual operating systems, while not completely open in the sense that any distributed resource can be added and virtualized, are technically and commercially viable (Lawton 2008). Other examples are the Amazon Elastic Compute Cloud (EC2) (http://aws.amazon.com/ec2/) and the Windows Azure operating system Microsoft, announced in October 2008 (http://www.microsoft.com/azure/whatisazure.mspx).

Virtual organizations (VOs) have the potential to change dramatically the way we use computers, much as the Web has changed how we exchange information. Standardized Internet protocols made the Web possible, and standard open Grid protocols hold the promise of fostering unprecedented integration of technologies, applications, files, data, and just about any other IT resource. One cannot predict the future but the author regards it as entirely possible that Grid technology and Web services will become fully compatible, and that the distinction between the two will eventually fade. Evidence of this fading can be found in the OGSA Glossary of Terms (OGF 2007c): the term ‘Grid service’ has been deprecated and it is recommended that one refers to it as a ‘Web service that is designed to operate in a Grid environment, and meets the requirements of the Grid(s) in which it participates’.

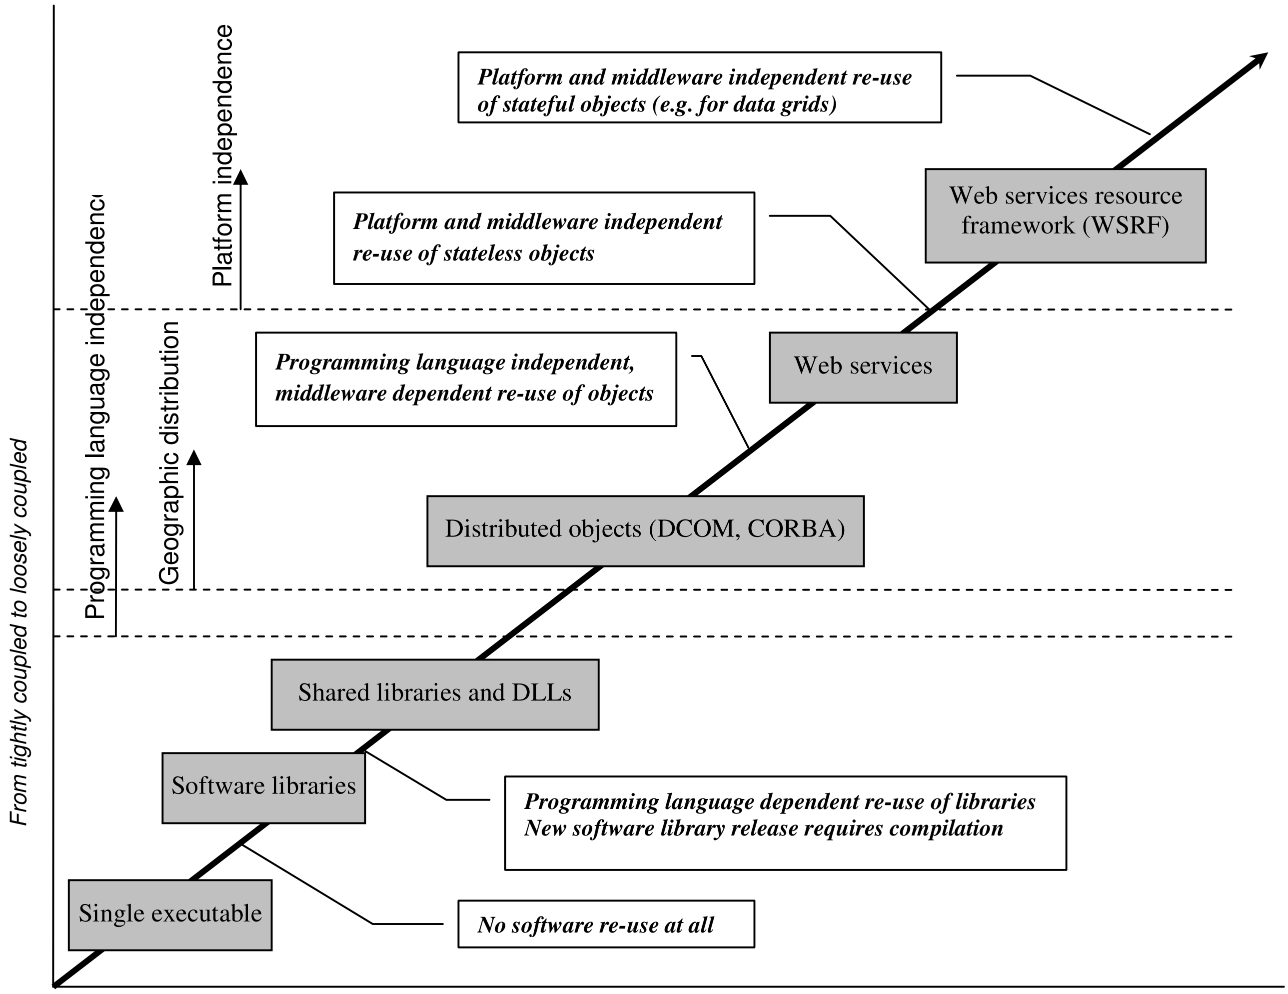

Figure 8. The software evolution

The introduction of the term ‘cloud computing’ in 2007 for the subset of grid computing that includes utility computing already displays this trend of using Grid technology without being aware of it, a sure sign of advancing from the development phase to mass adoption of a new technology.

The future of the 'cloud' is still unclear and it could pan out to include any or all of the following: word processing and similar applications available as services on the Web; enterprise computing in the cloud where, for example, customer relationship management is supplied as a software service; or the cloudy infrastructures rented out by the likes of Amazon Web Services or Google (Hayes 2008). There is still some debate on what the cloud is and what it's not. To some, the cloud looks like Web-based applications, a revival of the thin-client. To others, the cloud looks like utility computing, a grid that charges metered rates for processing time. Then again, the cloud could be distributed or parallel computing, designed to scale complex processes for improved efficiency (Weiss 2007).

Whatever the relationship between the grid and the cloud, and whatever becomes of the cloud in future, the cloud has certainly gained from past grid research, and with everyone from IBM to Google to Amazon to Microsoft to Oracle to OGC claiming their stake in the cloud through press releases and announcements (IBM 2008, Google 2008, Amazon 2008, Microsoft 2008, Oracle 2008, OGC 2008), it is clear that the cloud, and thus also grid computing, is moving into the everyday realm.

Cloud computing also holds advantages for SDIs in future: the cloud as *data center* could be the platform on which data is shared; the cloud as *utility grid* could provide the computing power that is needed for complex spatial analysis; and the cloud that provides *software as a service* would ensure even wider access to spatial data on any number of thin client devices. While the Compartimos reference model has been designed as a data grid, some Compartimos components, such as the AddressService, the AddressDataAccessService and the VirtualAddressDataService described in Chapter 4, could be deployed in a cloud, probably with some minor modifications.

3.2.5 Note on 'grid', 'Grid' and 'GRID'

In the literature different spellings are observed: 'grid' or 'data grid' (Chervenak *et al.* 2000, Zaslavsky *et al.* 2004, Grimshaw and Natrajan 2005); 'Grid' or 'Data Grid' (Foster *et al.* 2001, Foster 2002, Chervenak *et al.* 2005, Baker *et al.* 2005, OGF 2007c; Venugopal *et al.* 2006); and even 'GRID' (Rajabifard 2005). In this dissertation the term is written with an uppercase, as in 'Grid', when writing about the concept of an infrastructure for resource sharing, while using the lower case 'grid' and 'data grid' for localized or specific Grid implementations. This approach can be observed in the literature but all authors do not consequently apply it.

While the term 'Grid service' has been deprecated in the OGSA Glossary of Terms (OGF 2007c), in this dissertation the term is used to denote a web service in a grid environment.

3.3 The Grid architecture

In this section two aspects of the Grid architecture that are important for Compartimos are presented. *The Anatomy of the Grid* (Foster *et al.* 2001) and *The Physiology of the Grid* (Foster *et al.* 2002), which evolved into the Open Grid Services Architecture (OGSA) from the Open Grid Forum (2006), describe the high level architecture of the Grid. The two aspects of the architecture that are discussed here are the Grid components organized as a layered architecture and the Grid as a service-oriented architecture (SOA). The layered architecture is of interest because in Figure 36 of Chapter 4 each Compartimos component is assigned to one of these layers, thus showing the level of abstraction at which each Compartimos component operates. The SOA aspect of a Grid is of interest since Web service implementation specifications for spatial data discovery and access exist, and these are described and related to Compartimos in the section on Technology choices in Chapter 5 .

3.3.1 The Grid as a layered architecture

The Grid can be described in terms of a number of layers, each at a different level of abstraction, ranging from the fabric layer (the actual hardware) at the lowest level to the application layer (where applications operate in a virtual organization environment) at the highest level. Each layer provides services to the layer above it, and makes use of services that are provided by the layer below it. Each layer also provides a virtualization of the resources on the lower level, e.g. the differences between hard disks from different vendors are accommodated by the operating systems in the Grid fabric layer, and on the application layer a storage resource or computing resource is requested, regardless of all the intricate details of the actual device, the discovery mechanisms to locate it and the communication protocols to use it.

In Compartimos there is also abstraction and virtualization of the distributed address data sources, as well as services and protocols that determine behavior and coordination on the collective layer, i.e. for the collection of address datasets. The layers in Figure 9 below provide a reference for this allocation of Compartimos components to individual grid layers, which will be discussed in Chapter 4 . For the purposes of describing the layered architecture of the Grid, the layers presented by Venugopal *et al.* (2006) and Foster *et al.* (2001) are merged into four main layers of the Grid architecture, illustrated in Figure 9 below, and described as follows.

Application layer. On this layer domain-specific applications operate within a virtual organization (VO) environment to achieve the VO’s collaborative goal, such as climate modeling, hybrid earthquake engineering experiments and physics data analysis; or mapping and geocoding in the case of the work in this dissertation.

Grid resources and services layer. Consists of protocols for the secure negotiation, initiation, monitoring, control, accounting and payment of individual resources (the *Resource* layer according to Foster *et al.* 2001) as well as collections of resources (the *Collective* layer according to Foster *et al.* 2001). This layer provides the required level of abstraction of individual resources so that the Grid can provide a wide variety of behaviors for collections of resources that are based on these resource abstractions. Examples are resource discovery, resource brokering, job scheduling, replication and replica management, resource monitoring, and accounting services.

Connectivity and communication layer. Consists of the communication and authentication protocols that are required to interact with the resources on the Fabric layer. These protocols are mainly drawn from the TCP/IP protocol stack but are augmented with protocols incorporating specific Grid requirements such as the GridFTP protocol, which provides efficient transfer of large data files in a grid. Existing authentication protocols such as public key infrastructure (PKI) in the form of X.509-certificates are integrated and extended on this layer.

Fabric layer. Consists of the physical resources to which the Grid coordinates shared access. These are computational resources (clusters, supercomputers), storage resources (RAID disks, tape archives), data resources (databases, files) and instruments (sensors, telescope, accelerator). Each of these resources is controlled by software such as the operating system, file system and/or database management system.

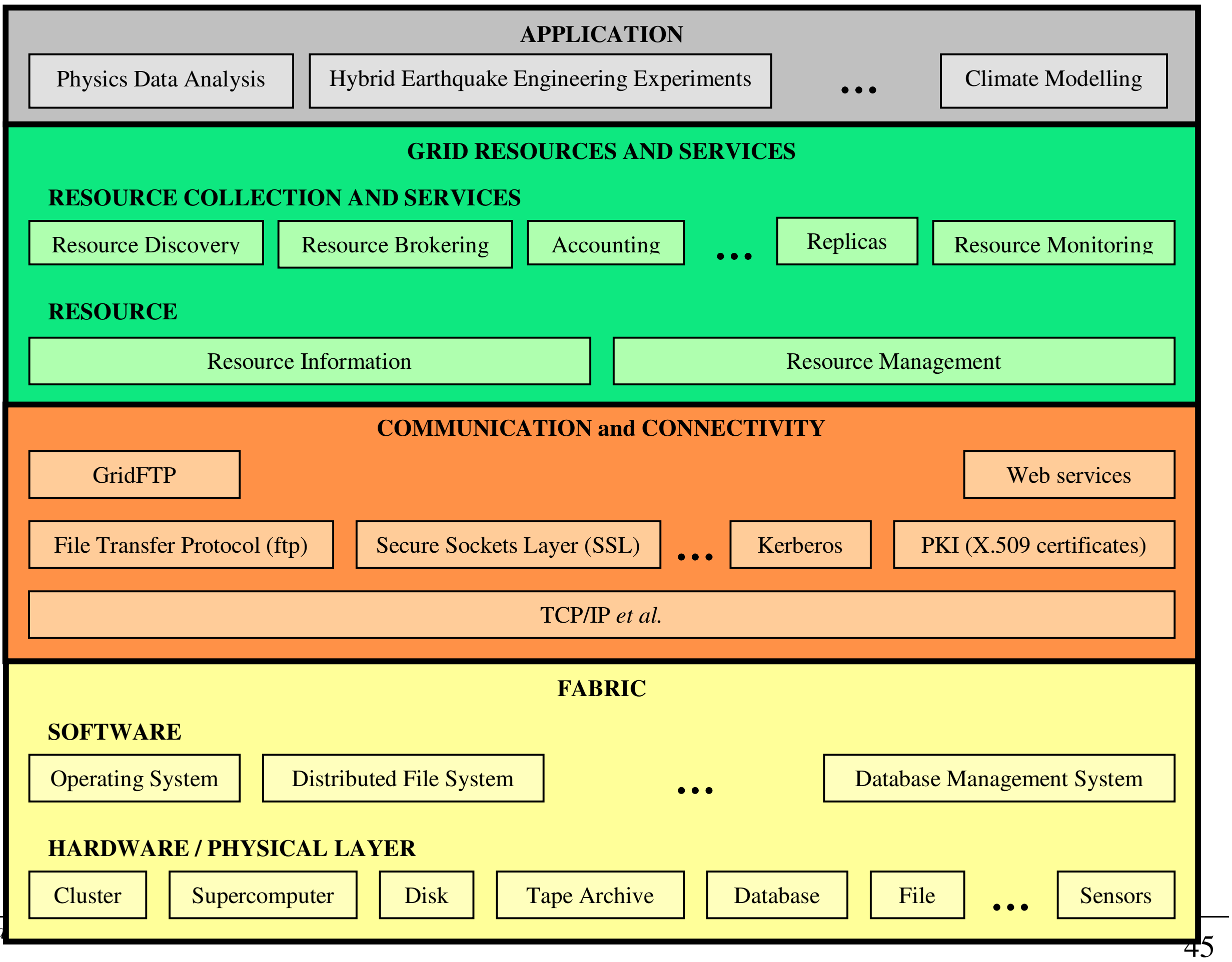

Figure 9. The four main layers of the Grid architecture

3.3.2 The Grid as a service-oriented architecture

A service-oriented architecture (SOA) refers to the specific style of building a reliable distributed system that delivers functionality as services, with the additional emphasis on loose coupling between interacting services (OGF 2007c). An SOA is typically implemented by a set of Web services that provide the capabilities and behaviors of the system. OGSA is a service-oriented architecture in which the core capabilities and behaviors are described as a set of services, the Grid services. These Grid services are loosely coupled peers that, either singly or as a part of an interacting group of services, realize the capabilities and behaviors of OGSA (OGF 2006).

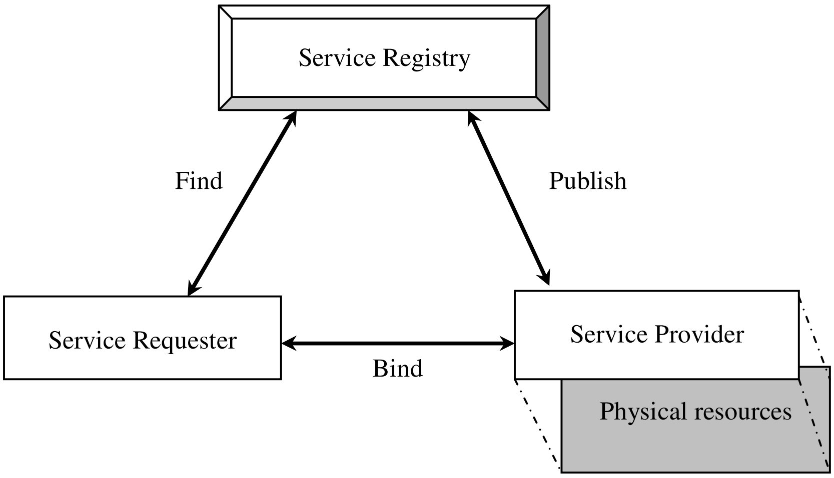

Figure 10. Service-oriented architecture

A service-oriented architecture is further based on the concept that a service provider publishes its services at a service registry. A service requester finds details about a specific service at the registry, and then proceeds to bind to the service at the service provider and starts interacting with the service at the service provider. This concept is illustrated in Figure 10. OGSA is a specific profile of the core Web Service standards that includes the Web Services Description Language (WSDL), Simple Object Access Protocol (SOAP) and Universal Description Discovery and Integration (UDDI) that have been specified by the W3C.

Web Services (WS) standards do not support and were not designed to meet all Grid requirements, but the Grid community is actively involved in the development and evolution of WS standards to revise, modify and extend existing standards and to introduce new standards where applicable. For example, Web services are stateless, but OGSA requires stateful Web services. The Web Services Resource Framework (WSRF) and WS-Management families of standards, published by OASIS, were prompted by the Grid community (Baker *et al.* 2005). While statefulness is a fundamental requirement for the Grid, other application domains, such as general-purpose Web servers, benefit from these standards as well.

Compartimos follows a service-oriented approach, i.e. there is a registry of discoverable AddressDataAccessServices, published on the address data grid by address data providers, which the VirtualAddressDataService discovers and coordinates in order to achieve the virtual address dataset. These services are presented and discussed in Chapter 4 . The SOA aspect of a Grid is further of interest since OGC Web service implementation specifications for spatial data discovery and access exist, and these are described and related to Compartimos in the section on Technology choices in Chapter 5 .

3.4 Data grids

3.4.1 What is a data grid?

A *data grid* is a special kind of Grid in which data resources are shared and coordinated. In OGSA (OGF 2007c) a *data resource* is defined as an entity (and its associated framework) that provides a data access mechanism or can act as a source or sink of data. These data resources are typically heterogeneous in terms of their syntax and semantics. Examples of data resources are flat files, tables in a relational database, sensors or data streams. Referring back to Foster's (2002) definition of a grid, in a data grid

1. the individual *data resources* that are shared on the grid live in different control domains and consist of flat files, tables in relational databases, or other sources of data;
2. these constituent *data resources* are coordinated to deliver a *data service* that is significantly greater than the sum of its parts and,
3. all of this is achieved through standard, open, general-purpose *data protocols and data interfaces.*

Or, adapting the definition of a grid in the OGSA Glossary of Terms (OGF 2007c):

A ***system***
that is concerned with the
integration, virtualization, and management of data services and data resources
in a ***distributed, heterogeneous environment***
that ***supports virtual organizations*** (collections of users and data resources)
across traditional *administrative and organizational domains (real organizations).*

Data grids are used for the sharing and integration of distributed data that are managed and administered independently, also referred to as *data federation*. Data grids are also applied in areas of science, technology and commerce where there is a need for efficient access to, and the movement and management of, large quantities of data in a distributed environment, also known as *data-intensive environments*.

Compartimos is a reference model for an address data grid in which distributed heterogeneous sources of address data are managed and administered independently of each other. The data grid provides coordinated access to these distributed heterogeneous sources of address data and is therefore an example of a data grid that is used for *data federation*.

3.4.2 OGSA-Data Access and Integration (OGSA-DAI)

OGSA-Data Access and Integration (OGSA-DAI) is a service-oriented architecture for database access over the Grid that allows for the integration of heterogeneous databases into an OGSA-type grid, and this technology is therefore relevant to Compartimos where coordinated access to distributed heterogeneous sources of address data is provided. OGSA-DAI originated from a need to include in Grid implementations data that is stored in a DBMS. Current DBMSs already provide a large range of functionality to securely store, query and maintain large volumes of data, but none of them have been OGSA Grid-enabled. Rather than to build an OGSA Grid-enabled DBMS from scratch, the principle behind OGSA-DAI is to provide the necessary middleware that will ‘OGSA Grid-enable’ existing DBMSs. As such OGSA-DAI has to reconcile DBMS implementation differences (IBM DB2, Oracle, MS SQLServer, etc.) and accommodate the variety of database paradigms (relational, object, XML, etc.). Additionally, OGSA-DAI provides distributed query functionality (OGSA-DQP) that allows a user to send a single data request to an OGSA Grid-enabled data grid and the OGSA-DQP then takes care of coordinating the request among the different data resources (Antonioletti *et al.* 2005).

Figure 11 shows the layered architecture of OGSA-DAI. The OGSA-DAI Basic Services are implemented as OGSA Grid Data Services and access the underlying databases using drivers and also provide for data formatting, data delivery and request handling. A Grid application can either use the OGSA Grid Data Services (OGSA-GDS) directly or use the OGSA-DQP to coordinate access to the multiple databases. OGSA-DAI implements its services as OGSA Grid Services, and is therefore also a service-oriented architecture.

OGSA-DAI has been used in a number of data grid implementations, including the eDiaMoND project described in section 3.5. In some ways OGSA-DAI works like Open Database Connectivity (ODBC) and Java Database Connectivity (JDBC): ODBC and JDBC provide a standard programming interface to any DBMS, and as long as a driver exists for a specific DBMS, that DBMS can be accessed through the standard ODBC/JDBC interface. OGSA-DAI provides the standard interface to the various databases (in various DBMSs and paradigms) in a data grid. In Compartimos the *AddressDataAccessService* performs the role of such a 'driver'.

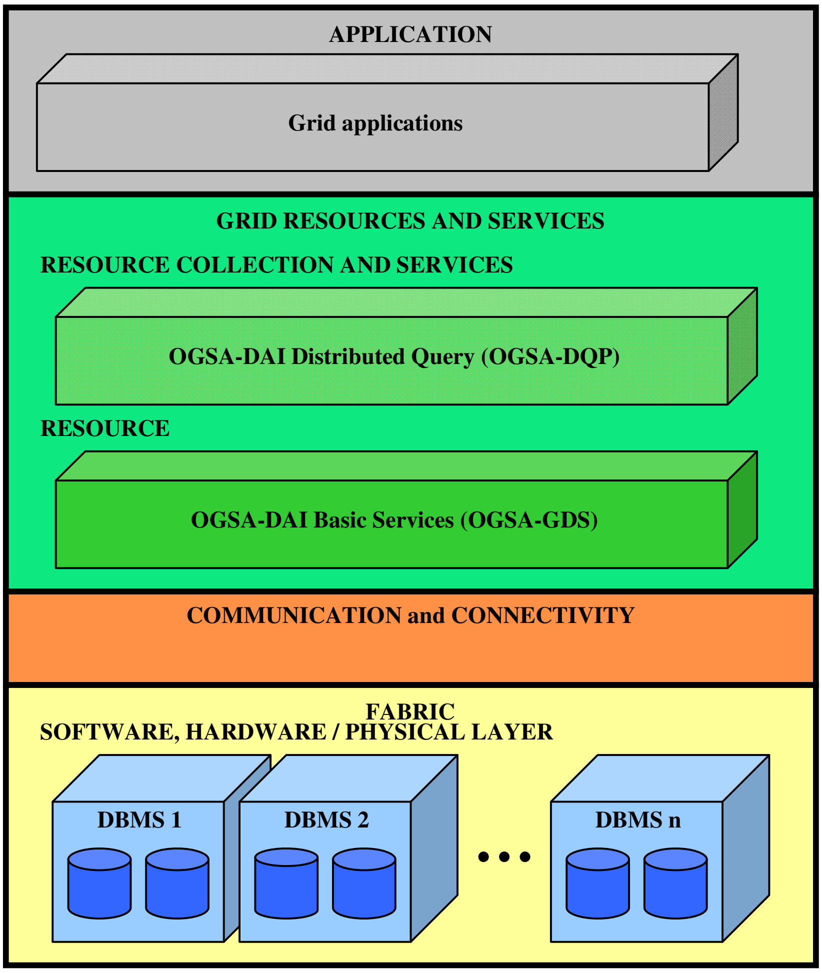

Figure 11. The OGSA-DAI layered architecture, adapted from Antonioletti (2005) to show the four main Grid layers of Figure 9

In an SDI environment address data is stored in data resources that include relational databases (such as Oracle and Microsoft SQL Server), spatial databases (such as Oracle Spatial and ArcSDE), as well as proprietary geographic files (such as ESRI .SHP and MapInfo .TAB files). In order to integrate the data from the multiple heterogeneous sources of address data described in Chapter 2, both the syntactic (Oracle Spatial, ESRI SHP, etc.), as well as the semantic (address data model) differences have to be accommodated.

The OGC has published a specification for a Web Feature Service (WFS) (OGC 2005) that provides a platform-independent data access interface to features (representations of real world objects) in a spatial dataset. A draft ISO standard, ISO 19142 (draft), *Geographic information – Web Feature Service*, is also available. One potential approach to syntactic interoperability in an address data grid is to make use of a single OGC WFS that is able to interpret different types of address data resources (Oracle Spatial, ArcSDE, etc.), but returns address data in a common format such as the standard open Geography Markup Language (GML), an XML grammar for the modeling, transfer, and storage of geographic information (ISO 19136:2007). An alternative approach is to make use of many OGC WFSs, one per type of spatial data resource, each translating between the proprietary format and a common format such as GML. The same goes for semantic interoperability: either a single service is able to interpret all kinds of address data, meaning that 'understanding' an additional address dataset requires an updated implementation of the service; or there is one service for each address data model, meaning that an additional address dataset requires the registration of an additional service. The latter approach scales better for many different formats and/or models and is therefore followed in the Compartimos reference model. Figure 12 below illustrates the two different approaches.

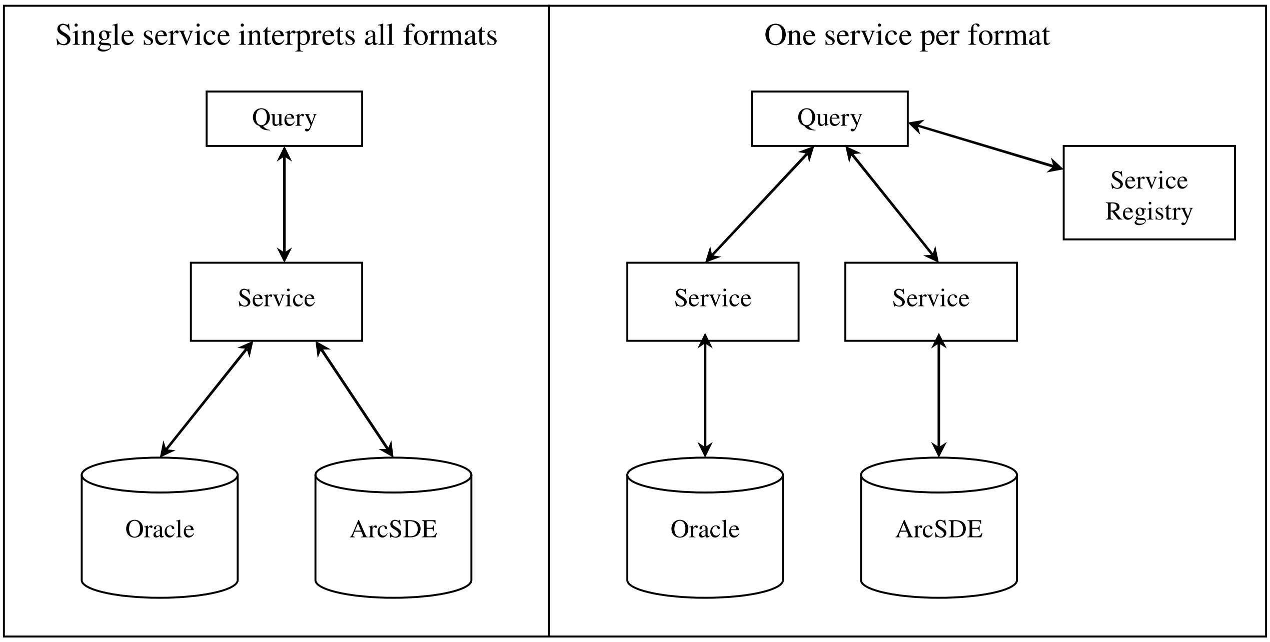

Figure 12. Two approaches to resolving syntactic heterogeneity

3.5 Examples of data grid implementations

In this section the following existing data grid implementations are described:

- the Laser Interferometer Gravitational Wave Observatory (LIGO);
- the Earth System Grid (ESG);
- the e-DiaMoND project; and
- the Geoscience network (GEON).

Table 4. Existing data grid implementations (author's summary)

	LIGO	Earth System Grid (ESG)	e-DiaMoND	GEON
Application domain	Physics and astronomy	Climate modeling	Breast cancer treatment	Earth sciences
Region	United States	United States	United Kingdom	Northern America
Number of data sites	Two	Around 10 centers and laboratories	Scalable to 90+ Breast Care Units (BCUs) in the UK	3 GEON data nodes, 15 GEON points of presence
Total data volume	One terabyte per day, ca. 365 terabytes per year	250 terabytes until 2006, ca. 70 terabytes per year	Estimated 480 terabytes per year, when fully operational	Each data node can store 4 terabytes of data
Metadata	Descriptive metadata about the data in the files (in a relational database)	Climate model metadata (in a relational database)	Patient data and metadata on image files (in a relational database)	Metadata about data made available by providers (in a relational database)
Format of data resources	Files with data from the LIGO detector	Files containing climate research data	Image files with mammography	Relational data, ESRI .SHP files, LiDAR
Size of individual data item	1-100 megabytes per file	Unknown	Estimated 75 megabytes per image file	Varies considerably, depending on what a user uploads
Number of data items	More than 40 million files	Millions of files	1000 cases	Around 4,500 (searching the portal on March 2008)
Interaction	Portal www.ligo.org	Portal www.earthsystemgrid.org	Service registry and Web services	Portal www.geongrid.org, as well as Web service registry
Software	Globus Toolkit, Lightweight Data Replicator (LDR)	Globus Toolkit, OPeNDAP-G (Grid-enabled Open-source Project for a Network Data Access Protocol)	Globus Toolkit, OGSA-DAI, IBM: DB2, Content Manager, Visual Age C++, WebSphere Application Server, Apache TomCat	Globus Toolkit, OGSA-DAI, Storage Resource Broker (SRB), ROCKS, GridSphere, PostgresSQL, IBM DB2, MySQL, ArcIMS, GRASS, ArcSDE

These implementations were selected specifically to illustrate variety in data resource types, data volumes and client interaction in order to provide a broad comparative base for Compartimos. A summary overview (prepared by the author) of some of the characteristics of these implementations is supplied in Table 4. In section 3.6 these existing data grid implementations are related to the work on Compartimos in this dissertation in order to position this work in the bigger context of data grids. In Chapter 4 the summary is repeated, this time including the requirements for an address data grid in an SDI, so as to distinguish the Compartimos requirements from other data grid implementations.

3.5.1 Laser Interferometer Gravitational Wave Observatory (LIGO)

LIGO is a facility dedicated to the detection of cosmic gravitational waves and the harnessing of these waves for scientific research. Gravitational waves are ripples in the fabric of space and time produced by violent events in the distant universe, for example by the collision of two black holes or by the cores of supernova explosions. These ripples in the space-time fabric travel to Earth, bringing with them information about their violent origins and about the nature of gravity. As such LIGO is a scientific tool to assist research in both physics and astronomy (http://www.ligo.caltech.edu/, Chervenak *et al.* 2005).

The LIGO facility includes three interferometers at two sites that generate approximately one terabyte of data every day. Refer to Figure 13 for photos of the facilities. The data generated must be scientifically analyzed for it to be of any value. The analysis is computationally intensive (some classes of astrophysical searches can require hundreds of Teraflops), and the data volume itself is huge (1 TB per day).


LIGO Hanford Observatory
Hanford, Washington, United States

Livingstone Observatory
Livingstone, Louisiana, United States

Figure 13. LIGO installations (http://www.ligo.caltech.edu/)

Nine sites within the LIGO collaboration (each operated independently) currently provide computing facilities based on commodity cluster computing. The scientists who have the expertise to perform this analysis are spread across 41 institutions on several continents, and this community is growing all the time. The key challenge for LIGO is to get the data from the LIGO detectors to the sites where analysis happens and to make those sites accessible to the participating scientists.

The data management challenge faced by LIGO is therefore to replicate approximately one TB/day of data to multiple sites securely, efficiently, robustly, and automatically; to keep track of where replicas have been made for each piece of the data; and to use the data in a multitude of independent analysis runs. The nine sites each use mass storage systems, but different systems are used at different sites. Scientists and analysts need a coherent mechanism to learn which data items are currently available, where they are, and how to access them.

The work on the Lightweight Data Replicator (LDR) was developed for the LIGO system and the knowledge gained from its development served as input to the Data Replication Service (DRS) of the Globus Toolkit, which provides a pull-based replication capability similar to that provided in the LIGO LDR system. (www.globus.org).

3.5.2 Earth System Grid (ESG)

ESG is a data grid that connects important U.S. repositories of climate model data that are geographically distributed at a number of national laboratories and research centers. In 2003 alone, climate change research sponsored by the US Department of Energy (DOE) has produced at least 72 terabytes of scientific data that is stored across several of the DOE sites. The goal of the ESG is to improve the daily management and tracking of this data by facilitating data publishing by climate modelers, and by facilitating access to the data by the worldwide climate community.

The main entry point to the ESG is through a Web portal on which users search or browse ESG catalogues to locate desired datasets, with the option of browsing both metadata about experiments as well as individual files and their contents. Users select datasets or individual files to be retrieved or request an “aggregation” – a specific set of variables subject to a spatiotemporal constraint. Selected data can be downloaded to the user’s system for analysis.

The ESG portal provides a central location for enforcing authentication, authorization and accounting (AAA) services, and brokers the formulation and submission of user data requests (transfer, download, sub-setting) among the distributed data nodes. The portal also provides the interface through which authorized data providers publish datasets into the system. Data are stored either on disk farms for faster high-performance online access to frequently requested datasets, or on deep archives for less frequently accessed datasets (Foster *et al.* 2006). Refer to Figure 14 for the ESG topology.

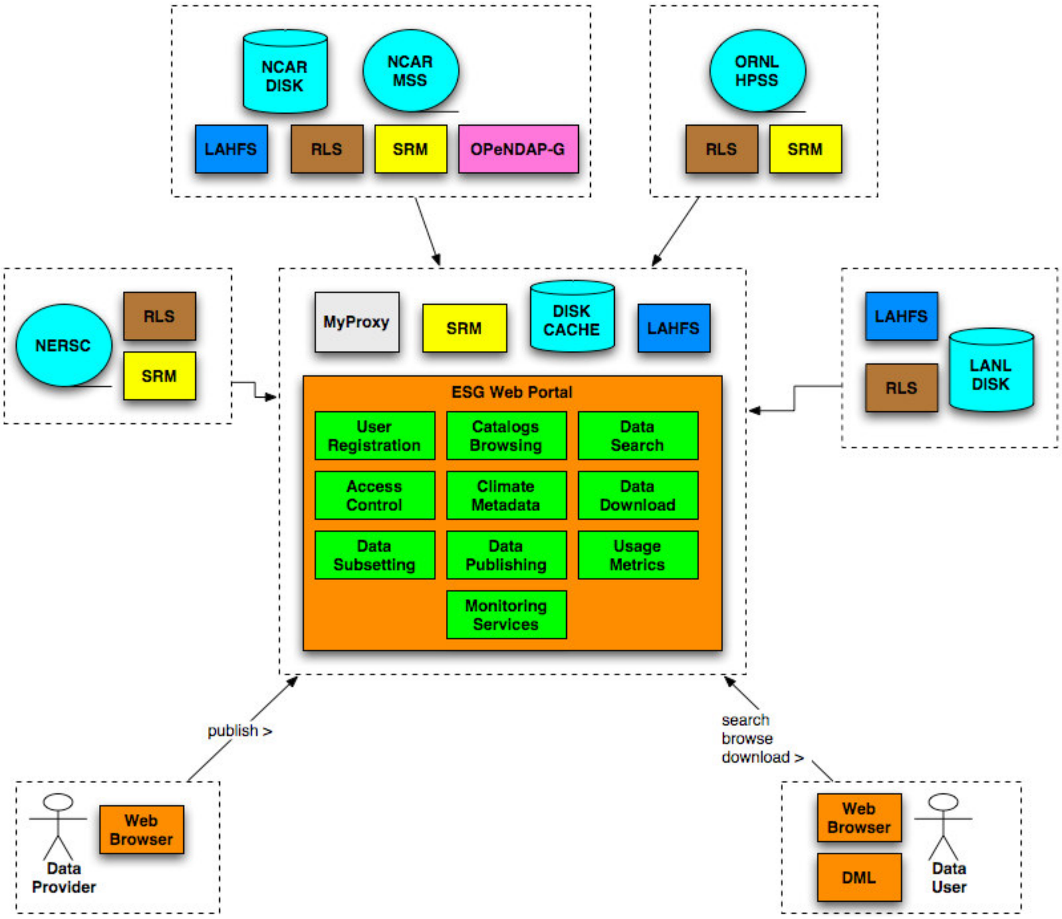

Figure 14. The ESG topology (Foster *et al.* 2006)

3.5.3 e-DiaMoND

e-DiaMoND is a British research project that developed a prototype Grid infrastructure for mammogram databases that are stored at geographically distributed independently managed breast care units (BCUs). The goal with the infrastructure was to enable the sharing of mammography archives, the easy, dynamic and real-time movement of information with high levels of security to facilitate the sharing of workloads, and collaboration of radiologists without being in the same physical location. Figure 15 shows examples of mammography.

The core e-DiaMoND system consists of middleware and a virtual image store comprising physical databases, each owned and administered by a different organization (the BCUs). The e-DiaMoND grid is formed by participating BCUs coming together as a virtual organization to unite their individual databases as a single logical resource, the virtual image store. Clients interact with the e-DiaMoND registry to discover services, and then interact with the grid through these services, which can be divided into data services—that allow each hospital to see all of the data owned by the participating nodes—and compute services—that can perform potentially complex and long-running calculations on the image data.

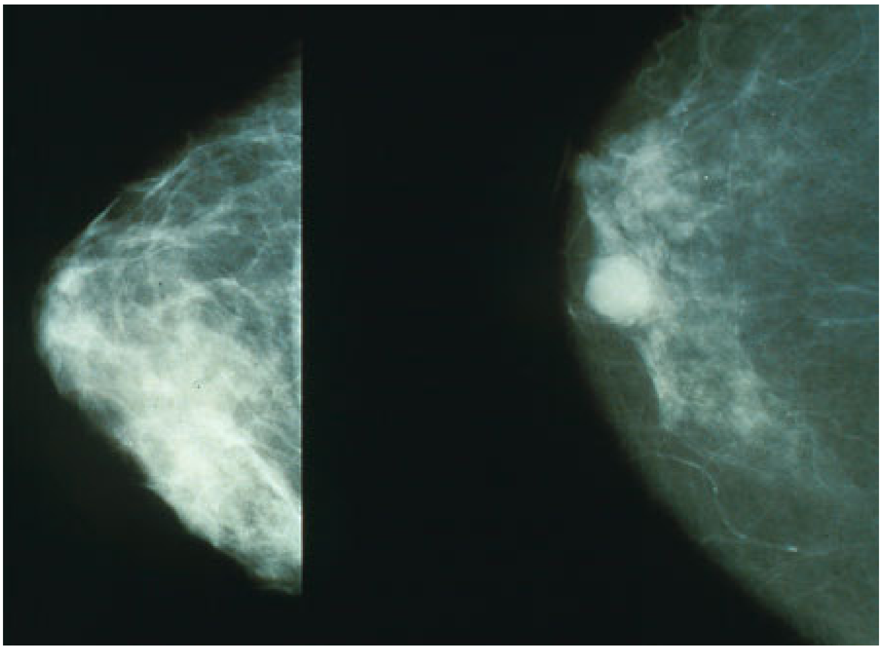

Figure 15. Example mammography image: normal (left) versus cancerous (right)

3.5.4 GEON

The Geoscience network (GEON) project aims to develop cyber-infrastructure in support of integrated research to gain a more quantitative understanding of the 4-D evolution of the North American lithosphere (the crust and uppermost mantle of the Earth). Scientists in this area are providers (computing, storage, etc.) as well as consumers (e.g. for analyses) of resources and are themselves distributed at a number of organizations. The themes that provided the initial guidelines for realizing this cyber-infrastructure illustrate the need to accommodate a very wide range of data variety (www.geongrid.org):

(1) gravity modeling of 3-D geological features such as plutons, using semantic integration of (igneous) rock and gravity databases, and other geological and geophysical data;

(2) study of active tectonics via integration of LiDAR data sets, data on distribution of faults and earthquakes, and geodynamics models; and

(3) study of lithospheric structure and properties across diverse tectonic environments via the integration of geophysical, petrologic, geochronologic, and structural data and models.

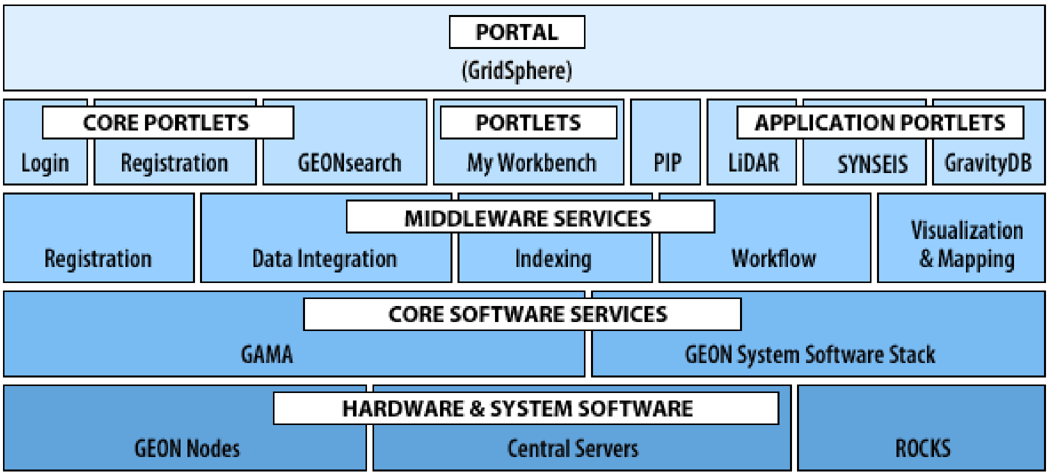

Figure 16. The various components of the GEON system

The GEON architecture consists of a single GEON portal and a number of distributed GEON sites and GEON service providers. The portal provides a single point of entry into the distributed GEON environment. Refer to Figure 16. GEON is based on a service-oriented architecture (SOA) in which remote sites and service providers communicate via Web services. All servers, at the central site as well as remote sites, run the same software stack, which is configured and deployed using ROCKS, a server clustering toolkit.

A site deploys a Point of Presence (PoP) node to host data, tools, and other resources contributed into the GEON network. Alternatively data, tools and other resources are deployed ("hosted") on the central site and hosted from the central location. Resources (data and tools) may also be contributed into the GEON network by providing Web service interfaces to the same. In this case, there is no requirement to install a PoP node at the remote site. The provider only needs to ensure that their Web services are able to interoperate with the GEON SOA environment.

3.6 Related work

3.6.1 Compartimos in relation to the examples of data grid implementations

In this section the work on Compartimos is compared to the examples of data grid implementations that were described in the previous section. Compartimos itself is presented in detail in Chapter 4 and this section serves to relate work on Compartimos to other existing data grid implementations. This related work falls mainly into the Computer Science discipline and only some of it overlaps with GISc. Other related work in the GISc discipline specifically was discussed in Chapter 2.

Compartimos, a reference model for an address data grid in an SDI, falls into the *application domain* of Geographic Information Science (GISc) but it is not in the scientific or research domain; rather it attempts to solve a real-world problem. ESG and GEON focus on data relating to the physical aspects of the Earth, while addresses are assigned by people according to certain rules and conventions, often in relation to (invisible) administrative boundaries and man-made features. Address data is influenced by the authorities that define these boundaries, as well as construct and maintain address-related infrastructure such as streets, houses and buildings.

The *regions* listed for the data grid implementations in Table 4 in the previous section span a single country or continent, but could be expanded globally should the need arise. The need for address data is firstly on a national, and later on an international scale. Consider the European situation where international seamlessly integrated address data for the whole of Europe is required in terms of the INSPIRE Directive. For the Compartimos implementation sample data of different address types (street addresses, site addresses, etc.) from multiple local sources in South Africa was used and the resulting consolidated dataset spans a single country.

Many of the early Grid implementations, including the data grid implementations described in the previous section, were built with the *purpose* to fulfill the demand for high performance computing and data-intensive applications in scientific research. This is a natural development since experts initiating the Grid computing concept were centered at universities and research centers and were probably in a better position to understand the scientific requirements. However, the initial concept of the Grid has expanded to the wider area of resource sharing among members of virtual organizations and is not restricted to scientific applications anymore. The OGF has passed some of the components that were identified for standardization in OGSA onto the W3C and OASIS (an industry consortium), from which one can conclude that Grid computing is progressing on its way to reach a wider audience. Refer also to current hype about cloud computing, discussed earlier in section 3.2.3. The main *purpose* of Compartimos is not high performance computing, but data integration and federation, although high performance computing could be a side benefit from the grid.

The *number of data sites* in the data grid implementations that are described in the previous section are limited, ranging from two to ninety. Address data typically resides at the local authorities in a country, the number of which in a country differs, but is very often more than a hundred. South Africa, for example, has around 260; in the UK there are 376 in England and Wales, and 32 in Scotland (Coetzee *et al.* 2008b); Australia has 750 (Jacoby 2002); and Denmark had 271 that were recently rationalized to 98 (Lind 2004). In this regard Compartimos is more like a sensor grid that has hundreds or thousands of sources of data, the sensors. However, in a sensor grid new data is continuously generated, while in Compartimos each local authority shares a comparatively stable database of addresses.

The *total data volumes* of the LIGO, ESG and eDiaMoND data grids suggest that they are data intensive, whereas the volumes of address data are usually smaller: even the conservative estimate of assuming that there is a spatial address for every person on earth, each taking up 5-10K bytes (extremely conservative estimate: a typical address record will take up even less space) of storage space results in (only) a total of 32TB to 66TB of data. While this data might change over time, there is not a continuous stream of new address data. Thus comparatively, address data volumes are small.

The address data grid stores *metadata* about data providers and the content of their data. This is stored in a relational database and is available through the *CatalogueService*. Metadata for each individual address record is stored at the data provider. In this regard Compartimos is similar to the GEON grid but differs from the LIGO and eDiaMoND data grids. In Compartimos the design of the data model for the metadata is based on the ISO 19115:2003, *Geographic information – Metadata* standard.

Although the data in the ESG is linked to locations on the Earth, it is limited to climate research data, which is often better represented by raster-type data. The spatial simplicity of an address feature – very often represented as a point feature – on the other hand is best represented by vector data. The *format of data resources* in Compartimos compares well to the GEON grid where data is stored in proprietary GIS files, such as ESRI .SHP files, of which the format is similar to the type of address data that one finds at a local authority. Both ESG and GEON store and provide geographic data, but in relation to Compartimos, it is a different kind of geographic data. ESG and GEON focus on data relating to the physical aspects of the Earth, while address data in Compartimos is assigned by people according to certain rules and conventions, often in relation to administrative boundaries and man-made features. Thus address data is influenced by the authorities that define these boundaries and construct and maintain address-related infrastructure such as streets, houses and buildings.

The size of an address record is in the range of 5-10K, which is a lot smaller than the *size of individual data items* in the data grids described in the previous section: 75MB and 100MB data files in the LIGO and eDiaMoND grids respectively.

Interaction with Compartimos is through a portal, similar to the LIGO, ESG and GEON grids, but Grid services can also be registered in a service registry from where third party organizations can discover them in order to use them, similar to the GEON and eDiaMoND data grid. Thus Compartimos has a service-oriented architecture.

3.6.2 Other related work

In this section data grid related work that is similar to Compartimos is described. The work confirms that the research described in this dissertation is relevant and novel.

Zaslavsky *et al.* (2004) describe *Smart Atlas*, a GIS-based atlas environment enabling users to discover, access, visualize and query heterogeneous brain images and image markup. These brain images are organized into atlases of spatial data with 2-D and 3-D visualization techniques. The spatial data sources for Smart Atlas are web-enabled and include ArcIMS feature and image services and distributed Grid sources. The Smart Atlas client can be invoked from the Biomedical Informatics Research Network (BIRN) portal at www.nbirn.net. Even though Smart Atlas does not represent geographic data as such, it is relevant to Compartimos from the point of view that it includes ArcIMS data sources, which are a common format for storing address data.

The *multi-node and multi-data source grid GIS (MMG^{2}IS)* is a Grid system that was developed by integrating Grid technology with networked GIS systems (Wang *et al.* 2004). Grid technology that was used includes the Globus Toolkit and Grid development tools such as the Java Commodity Grid Kit and the IBM Grid application framework for Java. MMG^{2}IS employs virtualization through a system architecture that is based on a virtual machine: the different underlying machines are virtualized into a single virtual machine representing a single Grid-enabled operating system. Compartimos makes use of the concept of virtualization of address data sources, albeit at different levels and for different purposes. The MMG^{2}IS is further relevant because it integrates spatial data. It was, however, developed for a data-intensive environment whereas the purpose of Compartimos is data integration and federation.

Zhao *et al.* (2004) present a geospatial registry approach in which the OGC WRS (Web Registry Service), a de facto standard that supports the publishing of, and run-time access to, geospatial resources, as a wrapper, is used to extend the capabilities of the conventional Grid Metadata Catalogue Service (MCS) to the processing of geospatial queries against multiple heterogeneous spatial data sources and services. The implementation of this approach was used in the NASA Grid Data Service environment. This report is relevant to the Compartimos catalogue service.

The Distributed Earth Observation System Information Service (DEOSIS) (Aloisio *et al.* 2005b) was developed by the University of Lecce and aims at managing and accessing earth observation and geospatial heterogeneous data sources in a Grid environment. A Grid-based architecture provides a secure, scalable and pervasive environment for earth observation and geospatial data management among several virtual organizations. In this grid the metadata model is based on ISO 19115:2003, *Geographic information – Metadata,* which is interesting because the Compartimos catalogue data model is also based on ISO 19115.

The geographic data for which the DEOSIS data grid was developed, is mostly Earth observation data (similar to the ESG and GEON grids), which differs from address data: Earth observation data is often gathered with an 'eye-in-the-sky', be it aerial photography or satellite imagery. Addresses on the other hand have a close link to land administration and legal processes, and are therefore usually assigned by people and maintained by the relevant authorities.

Xue *et al.* (2008) write about the remote sensing information grid node (RSIN), which is a tool for dealing with climate change and quantitative environmental monitoring. In their case data-intensive retrieval of remote sensing information is required. They have developed a failure management strategy using a throughput-estimated model for data nodes. The failure management strategy is of general interest but this report is of specific interest to the work in this dissertation because spatial data is used in a grid but the application is for data intensive operations, and not for data federation as in the case of Compartimos.

Hua *et al.* (2005) present a design and small-scale implementation of a spatial data grid for a few hundred MBs as a proof of concept that grid computing technology can be applied to share distributed spatial data. They propose their spatial data grid as the GIS infrastructure for sharing spatial data comprehensively and thereby eliminating information islands. In this regard, their research supports the proposed data grid approach to sharing address data in an SDI, discussed in Chapter 6. In conclusion, Hua *et al.* recommend that more research should be done on the combination of grid computing with GIS. Compartimos is an example of such multi-disciplinary research that aims to apply grid computing to the sharing of distributed spatial data.

Ghimire *et al.* (2005) propose the integration of geographic information Web services with mobile agents and the grid in order to deliver large size spatial datasets over limited bandwidths. They recommend that the GI community should reconsider its adherence to the HTTP POST and GET communication patterns and move towards a SOAP based communication, which is a necessary and urgent step that would pave the road for a number of important performance issues, and from the point of view of Compartimos it would also enable seamless integration into the OGSA architecture. Aydin et al. (2008) report on an OGC WFS implementation that supports the three mandatory operations of the WFS implementation specification through a WSDL interface. Work on this recommendation is also now in progress at OGC, as reported in the *Summary of the OGC Web Services, Phase 5 (OWS-5) Interoperability Testbed* (OGC 2008) that SOAP and WSDL interfaces have been developed for four foundation OGC interface standards, WMS, WFS-T, WCS-T, and WPS, allowing these services to be integrated into industry standard service chaining tools.

In 2007, the OGC released a Request For Quotation and Call for Participation (RFQ/CFP) to solicit proposals for the *Canadian Geospatial Data Infrastructure (CGDI) Interoperability Pilot* project (http://www.opengeospatial.org/standards/requests/38). One of the principles of the CGDI is the widespread dissemination of data, which is at the same time managed at or near its source. Data users require authoritative geospatial information, accessible directly from as close as possible to its source, in order to make timely and effective decisions. Evolving the Canadian GeoBase portal to operate in a more distributed fashion and making maintenance transactions more efficient, will help to meet those user requirements.

The CGDI project focuses on three vector-based data themes: geographic name, national road network and administrative boundaries. The functional scope of this project includes investigations in the following areas: access by users to closest-to-source data, transactional updates exchanged between data suppliers and GeoBase, and the use of distributed services architecture to support end-user online applications. The project includes participants from provincial and federal agencies and from the private sector. Agencies participate at different levels, some work with private sector partners while others will use their existing infrastructures to provide access to manage and disseminate data. While this project, strictly speaking, does not involve Grid computing, according to the RFQ/CFP it does aim to take Web services one step further by introducing transactions into the Web Feature Server. Like Compartimos, this project works with vector data in an SDI environment. A live demonstration of the project, available online at http://www.ogcnetwork.net/cgdi, showed the benefits of access to a CGDI distributed network of federal, provincial and territorial data servers (14 in total) and highlighted:

1. access to place names, roads and municipal boundary data from a distributed network of federal, provincial & territorial servers (14 in total);
2. direct updates of data in provincial servers; and
3. a demonstration of an emergency response scenario

In 2007 the OGF and OGC signed a memorandum of understanding (MoU) to collaborate (GridToday 2007). The OGC Web Processing Service (WPS) was chosen as a starting point with the following initial goals:

- Integrate the WPS with a range of “back-end” processing environments to enable large scale processing as an application driver for both grid and data interoperability issues.
- Integration of the WPS with workflow management tools.
- Integration of OGC catalogues and data repositories with grid data movement tools such as GridFTP.

The goal is to enhance operational hurricane forecasting, location-based services and anything to do with data on a map, which naturally includes address data. The two organizations hope that from the mutual understanding of technical requirements and approaches, other opportunities and capabilities will emerge. The announcement of the MoU was published in November 2007 and thus the collaboration is still in its very initial stages. However, the collaboration proves that on an international level, the geospatial community is increasingly interested in utilizing grid technology as solution to its problems, while the grid community has found another community that can benefit from its technology. The results from the work in this dissertation will provide valuable input towards the mutual understanding of requirements and approaches in the Grid and geospatial communities, which the OGF/OGC collaboration seeks to build.

Rajabifard *et al.* (2005) acknowledge that new developments in database management software, including Grid computing technologies will change the way in which data is stored and maintained. However information about actual application of grid technology in SDIs is limited. One such project that incorporates Grid computing technology into an SDI is the recently launched *Geodateninfrastruktur-Grid* (GDI-Grid) project (http://www.d-grid.de/index.php?id=398&L=1), which is part of D-Grid, a long-term German strategic initiative in Grid computing. *Geodateninfrastruktur* is the German word for spatial data infrastructure and the project aims to find solutions for the efficient integration and processing of geographic data within geographic information systems (GIS) and spatial data infrastructures. As such it is the aim of GDI-Grid to develop standardized SDI services that can be used by a wide range of users and/or applications and thereby opening up access to existing SDI resources. The project was launched in July 2007 and to date there no publications have been listed on the project's website. From an OGF-22 workshop report (OGF 2008b) it is evident that members of the GDI-Grid are involved in the OGC-OGF Collaboration. GDI-Grid operates in an SDI environment and from that point of view is similar to Compartimos, but from the information available on the project, the GDI-Grid does not seem to include address data specifically.

52°North Initiative for Geospatial Open Source Software GmbH (http://52north.org/) is an international research and development company with the mission to promote the conception, development and application of free open source geo-software for research, education, training and practical use. One of their focus areas is geo-processing, which allows for the standardized and web-based processing of geographic data. They have developed a number of web services based on the OGC's Web Processing Service (WPS) interface standard. First reports on the grid-enablement of WPS, for example, Baranski (2008), have been published on their website. *52°North* are also involved in the OGC-OGF Collaboration according to the OGF-22 workshop report (OGF 2008b).

From the work that is discussed in this section, it is clear that the research community, as well as industry, recognizes the importance of grid computing for SDIs and geospatial data in general. The 52°North Initiative and the GDI-Grid project, as well as the MoU between the OGF and the OGC, are proof of this recognition, but the projects are still in their initial stages and results are not yet available. While ESG and GEON are examples of existing spatial data grid implementations, their purpose is slightly different to that of a data grid in an SDI environment that incorporates multiple local authorities.

In summary, the related work that was discussed here confirms that the approach in this dissertation of the data grid as enabler for SDI data sharing is innovative and new, and it proves that the work is extremely relevant at this point in time in both the Computer Science and the GISc disciplines. The work is also unique because Compartimos is designed for address data.