The Design and Implementation of a Grid-enabled Catalogue Service

Yaxing Wei[1, 2], Liping Di[2], Baohua Zhao[1], Guangxuan Liao[3], Aijun Chen[2], Yuqi Bai[2], Yang Liu[2]
[1] Department of Computer Science, University of Science and Technology of China (USTC), Hefei, Anhui, P.R. China 230027
[2] Laboratory for Advanced Information Technology and Standards (LAITS), George Mason University (GMU), 9801 Greenbelt Road, Suite 316-317, Lanham, MD 20706, USA
[3] State Key Laboratory of Fire Science (SKLFS), University of Science and Technology of China (USTC), Hefei, Anhui, P.R. China 230027

Abstract – **Catalogue service is crucial to the sharing and utilizing of geographic information. This paper provides a new architecture and its implementation of a geographic information catalogue service that can effectively manage distributed geographic data and services. The catalogue service is based on the Open Geospatial Consortium (OGC) Catalogue Service for Web Specification and the ebXML Registry Information Model (ebRIM). The ebRIM is extended with ISO 19115 (including draft part 2) and ISO 19119 international standards for better description of geographic information. The adoption of the ebRIM and international standards will greatly advance the interoperability of the catalogue service and the sharing of the abundant geographic information. The catalogue service is Grid-enabled to facilitate Grid's abilities of on-demand, ubiquitous access to distributed computing, data, and services resources. The Grid-enabled catalogue service is GSI-supported which enables it to interoperate with other Grid Services in a secure way. By being integrated with RLS, MDS and other Grid-enabled OGC Services, this Grid-enabled catalogue service prototype provides secure discovery, retrieval, and management of geographic data and services.**

I. INTRODUCTION

Open Geospatial Consortium (OGC) is an international organization promoting the interoperability and sharing of geospatial resources and services in the distributed environment through the development of consensus-based implementation specifications. OGC specifications are widely used by geospatial communities for sharing data and resources and some of the specifications are becoming ISO standards. OGC web services specifications allow seamless access to geospatial data in a distributed environment, regardless the format, projection, resolution, and the archive location. The fundamental ones include Web Coverage Services (WCS), Web Feature Services (WFS), Web Map Services (WMS) and Catalogue Service for Web (CSW). WCS, WFS and WMS define the interfaces for access to diverse geospatial datasets. CSW specifies the interfaces, HTTP protocol bindings, and a framework for defining application profiles required to publish and access digital catalogues of metadata for geographic data, services, and related resource information [1].

The Grid technology is a rapid developing technology, which enables large-scale sharing of resources within formal or informal consortia of individuals and/or institutions, sometimes referred to as virtual organizations [2]. The Grid technology can greatly promote the sharing, management, discovery, and utilization of huge amount of distributed geographic data and services. We make the OGC Catalogue Service Grid-enabled by introducing Grid Services into the catalogue service and enabling it interoperable with other useful Grid Services to facilitate the sharing of geographic information.

By utilizing Grid technology and integrating OGC CSW specification with ebRIM and ISO 19115/19119 standards, we present a new architecture and its implementation of a Grid-enabled geospatial catalogue service that can effectively manage distributed geographic data and services.

In this paper we first introduce the information model of this Grid-enabled catalogue service, followed by a brief description of its interfaces and operations. The technologies for the backend repository and the Grid-enabled geospatial catalogue service are discussed in a greater detail. The implementation of this Grid-enabled catalogue service and its collaboration with other Grid Services in a realistic geospatial Grid testbed are presented. The future research directions are also discussed in the concluding section.

II. DESIGN OF THE GRID-ENABLED GEOSPATIAL CATALOGUE SERVICE

A. EbRIM-Based Information Model

The catalogue information model is a conceptual model that specifies how geographic information metadata are organized within the catalogue. A well-designed information model can greatly improve the content utilization and interoperability of the catalogue. The ebXML Registry Information Model (ebRIM) version 2.5 is chosen to be the base of the catalogue information model and is extended with ISO 19115 and ISO 19119 standards to facilitate the description of geographic information metadata.

The ebRIM is a widely adapted information model that defines what types of objects are stored in a registry as well as the relationships among these objects classes [3]. The classes defined in the ebRIM are organized as a tree structure with its root being the RegistryObject class, which is the ancestor of most other classes.

The ebRIM is a general information model and is not sufficient to describe the geospatial metadata. We extended it with two international geographic standards: ISO 19115 (including draft part 2) and ISO 19119. ISO 19115 defines the abstract models required for describing geographic information and provides information about the identification, the extent, the quality, the spatial and temporal schema, spatial reference, and distribution of digital geographic data [4]. ISO 19115 draft part 2 is an extension of ISO 19115 to provide additional metadata needed for describing imagery and gridded data [5]. ISO 19119 identifies and defines the architecture patterns for service interfaces used for geographic information and the relationships to the Open Systems Environment model [6].

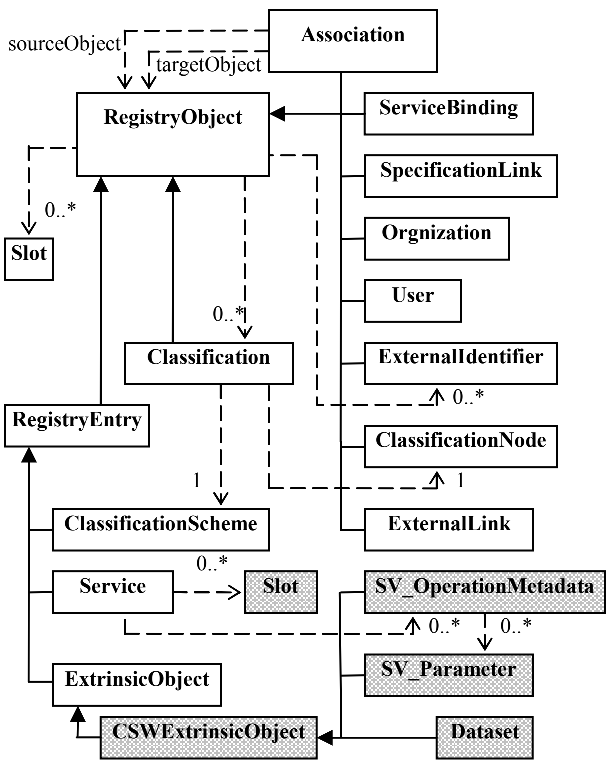

Figure 1. High-level view of catalogue information model

The ebRIM is extended with ISO 19115 and ISO 19119 in two ways. The first is to derive new metadata classes from existing ebRIM classes by importing new classes into the ebRIM class tree structure. For example, class CSWExtrinsicObject is derived from an existing class in the ebRIM – ExtrinsicObject – to represent all the metadata objects describing objects that may not be intrinsic to the catalogue. The new attribute in class CSWExtrinsicObject, repositoryItem, is used to specify the location of the object described by a CSWExtrinsicObject object. Class Dataset is derived from CSWExtrinsicObject in order to describe geographic datasets. Many new attributes are added to the Dataset class based on ISO 19115 and its draft part 2. The inherited repositoryItem attribute may be used to specify the location of the described geographic dataset. Two other classes, SV_OperationMetadata and SV_Parameter, are derived from class CSWExtrinsicObject. These two new classes both come from ISO 19119 standard to provide more detailed description of geographic service interfaces.

The second way to extend ebRIM is to use Slots to extend an existing class. Every class extended from class RegistryObject has the capability to add Slots into itself. The Service class included in the ebRIM can be used to describe geographic service but the available attributes in the class Service are not sufficient. Thus, new attributes derived from ISO 19119 are added to the Service class through Slots.

Figure 1 shows a high-level view of the catalogue information model. It also provides a simple description of the classes (including ebRIM classes and extended classes) and their relationships.

B. External Interfaces of the Catalogue Service

The external interfaces specify how clients interact with the catalogue service. They also define how the contents in the catalogue are managed and utilized.

The interfaces design of this catalogue service is based on CSW version 2.0, which defines the formal HTTP protocol binding interfaces and operations a web-based catalogue service should comply. It also specifies the request, response and exception messages for each operation [1]. Three interfaces, OGC_Service, Discovery, and Manager, are chosen from the CSW specification. Interface OGC_Service has only one operation, getCapabilities, which provides clients an OGC compatible XML document describing the capabilities and related information of the catalogue service. Interface Discovery provides three operations that allow clients to discover and obtain the resources registered in the catalogue. Interface Manager allows a client to update catalogue content using either a 'push' (through operation transaction) or a 'pull' (through operation havestRecords) style of publication.

Table 1 lists the interfaces of the catalogue service and the operations in each interface.

Interface	Operation
OGC_Service	getCapabilities()
Discovery	getRecords()
	describeRecords()
	getRegistryObjectById()
Manager	transaction()
	harvestRecords()

Table 1. Catalogue Service interfaces and operations

C. Grid-enabling of the Catalogue Service

Grid-enabling of a catalogue service is to transform the service from a Web Service to a Grid Service while preserving the interfaces, request and response messages defined in the CSW specification. A Grid Service is a Web Service that conforms to a set of conventions (interfaces and behaviors) defined in Open Grid Services Infrastructure (OGSI). The OGSI is built on top of Web Services standards and is the foundation for Grid services. It defines mechanisms for creating, managing, and exchanging information among Grid services [7].

According to the OGSI specification, each Grid service instance is associated with some transient state and has properties that control its lifetime. Additionally, the specification defines the operations that must be supported and the way the state and lifetime-related information can be accessed. Similar to Web services, a GWSDL document - with few non-compliant extensions to the WSDL standard - is used to describe the operations, messages, and portTypes (interfaces) of a Grid Service instance [8].

Security is a very important issue in a collaborative Grid computing environment. The Grid-enabled catalogue service works in a secure Grid computing environment through the supporting of Grid Security Infrastructure (GSI). The GSI uses public key cryptography, specifically public/private keys and X.509 certificates, as the basis for creating secure Grids.

III. IMPLEMENTATION OF THE GRID-ENABLED GEOSPATIAL CATALGOUE SERVICCE

A. The Backend Database

We use an open source relational database – MySQL 4.1.7 – as the backend repository to store the metadata of geographic data and services. The Grid-enabled catalogue service is XML-based, that is, all the classes in the catalogue service information model appear in the service requests and responses in XML form. Therefore, the mapping between XML objects and relational database tables has to be done. Each type of complex XML element (an element with attributes or sub elements) or simple element with multiple occurrences is mapped to a table. Every simple element with single occurrence or attribute is mapped to a field of a table which is mapped from that simple element's or attribute's parent element. The relationship between an element and its sub-elements is represented by the relationship between the tables mapped from them.

B. Implementation of the Grid-enabled Catalogue Service

The Grid-enabled catalogue service is implemented on Globus Toolkit 3.2 (GT3.2), an open source implementation of OGSI version 1.0 developed by the Globus project. GT3.2 provides libraries and tools to implement a Grid Service in Java programming language. It also provides the implementation of the building blocks defined in Open Grid Services Architecture (OGSA) as Grid Services to enable the resource monitoring (MDS), discovery (RLS), and management (GRAM), plus security (GSI and CAS) and file management (RFT, GridFTP) in a Grid environment. GT3.2 offers authentication, authorization, and secure communications through its implementation of GSI.

Implementation of the Grid-enabled catalogue service is similar to that of a pure web-based catalogue service. The major differences include the followings:

1. The abstract GWSDL document that describes the interfaces and messages of the Grid-enabled catalogue service. The types and messages definitions are the same as their definitions in the CSW specification [1]. But all the three portTypes (interfaces) chosen from CSW should be extended from ogsi:GridService portType to inherit the mandatory operations a Grid Service should contain. The binding and service definitions are generated from this abstract GWSDL document by GT3.2.

2. New parameters in the Web Service Deployment Descriptor (WSDD) file. These new parameters enable the Grid-enabled catalogue service to support GSI. For example, the securityConfig and authorization parameters are added to the WSDD file to enable communication encryption and server side authorization based on a gridmap file, as the followings:

<parameter name="securityConfig" value="org/globus/ogsa/impl/security/descriptor/gsi-security-config.xml"/>

<parameter name="authorization" value="gridmap"/>

C. A GSI-based Geospatial Data Grid

We set up a GSI-based geospatial data Grid by combining the Grid-enabled catalogue service (CSW) and Grid-enabled Web Coverage Services (WCS) [9], with iGSM and Grid Services (RLS, MDS, RFT) of Globus Toolkit 3.2. The Grid-enabled catalogue service plays an important role in such a geospatial data Grid. The architecture of the Grid is shown in figure 2.

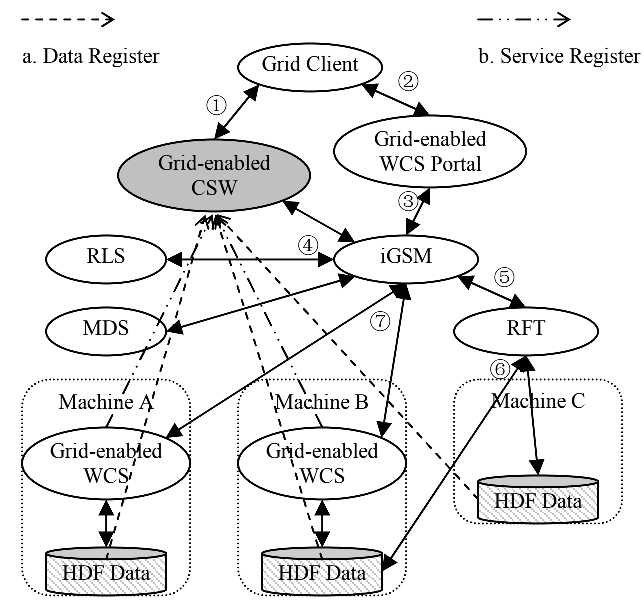

Figure 2. Architecture of the GSI-based Geospatial Data Grid

In Figure 2, the Grid-enabled WCS and the intelligent Grid Service Mediator (iGSM) are Grid Services implemented by this project. iGSM is a Grid Service that connects all the components in this geospatial Grid. Reliable File Transfer (RFT), Replication Location Service (RLS), Monitoring and Discovering Service (MDS) are all GT3.2 components. RFT handles data transferring between two machines in a reliable way. RLS manages the mappings between logical file names and physical file names. MDS can monitor the status (CPU status, Network bandwidth, etc.) of any machine in a Grid system. This Grid system is GSI-based and all the components communicate with each other in a secure way.

This geospatial data Grid currently manages 10TB of geospatial data, primarily NASA's EOS data products. Each granule of data logically contains one or multiple geospatial coverages. All these coverages are registered as Dataset objects into the Grid-enabled CSW in logical level (represented by the dashed line a), that is, all the replicas of a logical coverage are registered only once. All the mappings between a logical coverage and its replicas are registered into the RLS. The datasets are accessed through Grid-enabled WCS. All Grid-enabled WCSs are registered in the Grid-enabled CSW as Service objects (represented by the dashed line b) and available to all members in the Grid. A particular machine in the Grid may or may not have a Grid-enabled WCS installed.

If an authorized Grid client needs a coverage, it will first issue a Grid request containing the description of the coverages it needs to the Grid-enabled CSW. The Grid-enabled CSW returns to the client a list of logical coverages that meet the client-specified requirements. The client selects one dataset from the list and issues a Grid WCS request to the Grid-enabled WCS portal. The portal does nothing but delivers the request to iGSM. iGSM will query RLS and MDS to choose an optimal replica of the logical coverage dataset based on various criteria such as computing capability and network bandwidth. iGSM then check whether a Grid-enabled WCS is available on the machine where the coverage replica resides. If the Grid-enabled WCS is available, iGSM will generate a Grid WCS request to generate a concrete coverage dataset and send it to the requesting client. Otherwise, several processes will be automatically activated. iGSM will first query the Grid-enabled CSW and MDS to select the most available Grid-enabled WCS. It then requests RFT to transfer the coverage replica from its original machine to the machine on which the selected Grid-enabled WCS resides. With data replica transferred, iGSM will issue a Grid WCS request to the selected Grid-enabled WCS to get the concrete coverage. Finally the generated coverage will be returned to the Grid client through iGSM and the Grid-enabled WCS portal.

The above-described Grid system connects geographically distributed resources (data, services, CPUs, storages, etc.) into an integrated virtual organization within which seamless management and sharing of the heterogeneous geographic information is feasible and possible. Any authorized Grid client can discover and access the geographic datasets managed by the virtual organization without having to know their physical locations, and most likely, many other properties such as data format and map projection, due to the abundance of services usually available in such an organization.

IV. CONCLUSION AND FUTURE WORK

The ebRIM information model extended with ISO 19115 (including draft part 2) and ISO 19119 is suitable to describe the metadata of geospatial data and services. A Grid-enabled and GSI-supported CSW based on such an information model can act as the central registry of distributed geographic information in a large-scale secure geospatial Grid. Integration with other Grid services and useful Globus components will greatly promote the management and sharing of distributed geographic information.

We plan to serve virtual geographic products based on the geographic ontology's and workflow technologies. The Grid-enabled CSW is also a key component for production of workflow-based virtual geographic datasets. The information model of the Grid-enabled CSW should be enriched and more Grid components should be integrated into the system to achieve that goal.

V. ACKNOWLEDGEMENT

This project was supported by grants from the NASA Earth Science Data and Information System Project (ESDISP) (NCC5-645, PI: Dr. Liping Di) and NASA Earth Science Technology Office (ESTO) (NAG-13409, PI: Dr. Liping Di). Additional funding was provided by the Open Geospatial Consortium (OGC) for the development of Web Coverage Server, Web Map Server and Catalogue Service for Web as a part of OGC WMT II and OWS-II.

VI. REFERENCE

[1] Open Geospatial Consortium. OpenGIS® Catalogue Services Specification (04-021r2), version 2.0, http://portal. opengeospatial.org/files/files/?artifact_id=5929, May 11, 2004.

[2] K. Czajkowski, S. Fitzgerald, I. Foster, C. Kesselman, Grid Information Services for Distributed Resource Sharing. *Proceedings of the Tenth IEEE International Symposium on High-Performance Distributed Computing (HPDC-10), IEEE Press*, August 01, 2001.

[3] OASIS/ebXML Registry Information Model v2.5 – Committee Approved Specification, http://www.oasis-open.org/committees/regrep/documents /2.5/specs/ebrim-2.5.pdf, June 2003.

[4] ISO. Geographic Information – Metadata, ISO 19115:2003, http://www.isotc211.org/protdoc/FDIS/ISO_FDIS_19115_(E).pdf, Dec.01, 2002.

[5] ISO draft. Geographic Information – Metadata – Part 2: Extensions for imagery and gridded data, ISO/WD 19115-2.2, http://www.isotc211.org/protdoc/211n1337/211n1337.pdf, Oct. 13, 2003.

[6] ISO draft. Geographic Information – Services, ISO/DIS 19119.

[7] GGF. Open Grid Services Infrastructure (OGSI) Version 1.0, http://www.ggf.org/documents/GWD-R/GFD-R.015.pdf, June 27, 2003.

[8] S. Parastatidis, Introducing Open Grid Services, http://www.sys-con.com/webservices/articleprint.cfm?id=619.

[9] Open Geospatial Consortium. Web Coverage Service (WCS) version 1.0.0, https://portal.opengeospatial.org/ files/?artifact _ id=3837, Aug. 27, 2003.