Grid-enabled mediation services for geospatial information

I. Zaslavsky, C. Baru, K. Bhatia, A. Memon, P. Velikhov & V. Veytser
San Diego Supercomputer Center, University of California San Diego

ABSTRACT: The Open Grid Services Architecture (OGSA) is a convenient framework for information mediation, a popular information integration strategy that relies on middleware component called the *mediator* for rewriting user queries against heterogeneous distributed data sources. This paper explores OGSA-based information mediation for geospatial data sources and services and reports on our initial experience implementing Grid-Enabled Mediation Services (GEMS). GEMS is a collection of *grid services* being developed at the San Diego Supercomputer Center, which support source registration, dataset replication, data transfer, and query rewriting and execution. Within the services-based infrastructure, the information sources, as well as the mediator itself, represent grid services that have WSDL descriptions and follow grid service protocols for message interchange. We describe the general service architecture, as well as rationale and organization of core GEMS services supporting mediation of spatial information, and show their application in a geologic map integration scenario.

1 INTRODUCTION

Geospatial data interoperability has many facets, including: standards and specifications, infrastructure models and information integration strategies, metadata and data quality descriptions, data format and type conversion techniques, authorization, security and privacy, information assurance and business arrangements. Recent progress in all these components—in particular advances in XML-based standards for describing, serving, exchanging and rendering spatial data (OGC's WMS and WFS specifications, GML (OGC 2000, 2001, 2002), SVG (W3C 2001))—create new opportunities for seamless spatial data federation. Of the emerging infrastructure models, the *Grid Services* approach, as expressed in the *Open Grid Services Architecture* (OGSA) (Foster et al. 2001, 2002, GLOBUS 2003) and being standardized through the efforts of the Global Grid Forum (GGF 2003), is of particular interest since it provides standard mechanisms for managing security, service deployment, invocation and other interfaces, and supports complex object exchanges. The recently proposed Web Services Resource Framework (WSRF 2004) further expands the standardization effort by revising the interfaces specified in OGSI (Open Grid Services Infrastructure) while upholding the core principles of grid services architecture.

Not surprisingly, web and grid services became a component of implementation strategy within several large federally-funded projects focused on developing community *cyberinfrastructure* for a variety of scientific disciplines, e.g. the GeoSciences Network (GEON), the Biomedical Informatics Research Network (BIRN), and the Grid Physics Network (GriPhyN). Computer scientists and domain experts working together in these projects define common agreed-upon data models, dataset registration mechanisms, query templates, presentation interfaces, workflows and integrated views over distributed computational and data resources –

which together form a scalable standards-based information infrastructure that enables collaborative research and discovery.

Information mediation middleware works in conjunction with *source wrappers:* mediators are responsible for planning and orchestrating the execution of user queries across distributed resources while wrappers translate requests and responses from the common language and data model of the mediator to the language of individual sources (Wiederhold 1992). The mediation approach is becoming increasingly popular in geospatial applications (e.g., DeVogele et al. 1998, Gupta et al. 1999, Shimada and Fukui 1999, Boucelma et al. 2002) due to extreme heterogeneity (system, representational, structural, syntactic, semantic) and distributed organization of geographic data. However, scalability, security and methods for assembling mediation results into composite maps, remain serious challenges not previously addressed in the literature. This is the area where the grid services model offers a comprehensive solution, owing to systematic handling of service descriptions and interfaces for security and service lifetime management.

In this paper, we focus on services enabling geospatial mediation in grid environments. The goal is three-fold: (1) to outline grid-enabled mediation infrastructure for geospatial data, based on independent peer data nodes and supporting dataset caching/replication, and metadata propagation, (2) to describe mediator-level services that we consider necessary for geospatial data mediation, including registration, query rewriting, and spatial results assembly services, and (3) to demonstrate an application of grid-enabled mediation infrastructure to geologic map integration, within the GEON (www.geongrid.org) project. The three main sections of the paper reflect these three objectives.

2 GRID-BASED INFRASTRUCTURE FOR GEOSPATIAL INFORMATION MEDIATION

Web or Grid "Services" represent language- and system-independent re-usable functional components that are described using Web Services Description Language (WSDL – W3C 2003a) and invoked via Simple Object Access Protocol (SOAP – W3C 2003b). Grid services, in particular, emphasize security, authentication, message integrity and persistent state and lifetime management mechanisms critical for many scientific computing applications.

The SDSC GEMS system is being developed to support wide-area collaboration of scientific communities. It contains grid services for information mediation, and employs other standard services including authentication via Grid Security Infrastructure (GSI), Data Access and Integration (DAI), replication with Replica Location Service (RLS), and monitoring via NWS (Network Weather Service). All data sources and processing services in the mediation system are represented as *grid services,* i.e. they expose a standard interface to the mediator built around WSDL source descriptions, and exchange information using virtual XML documents wrapped in SOAP messages and accompanied by grid security authentication certificates (Fig. 1).

GEMS rely on Declarative Integrated Views expressed in XQuery to specify how the data sources should be combined and transformed. Compared to existing grid-based query processing services that rely on a single central server (e.g., *Grid Service Handle* server of OGSA Distributed Query Processor – DQP 2003), GEMS follow a P2P-like model of loosely-connected federation of resources that are not subject to centralized control, to ensure quality of service and scalability in terms of users, resources and supported data models. Decreasing system reliance on a single server, by making all Point-Of-Presence (POP) nodes completely independent, with propagation of registration information across heterogeneous nodes and caching/replicating datasets as needed, are important challenges that have not been sufficiently addressed, especially in the context of spatial data integration.

The general composition of the GEMS system is shown in Figure 2. Our discussion below focuses on services that are specific to mediation of geospatial information, rather than on generic (core) GEMS services related to authentication, caching, and replication.

2.1 *The bottom tier: Geospatial data nodes*

The geospatial data tier in GEMS is comprised of a set of grid data nodes that are commonly based on PostgreSQL, Oracle and DB2 with spatial options, ArcIMS and WMS servers, as well as shapefile collections and pure XML (GML) sources. In addition to datasets "hosted" by a

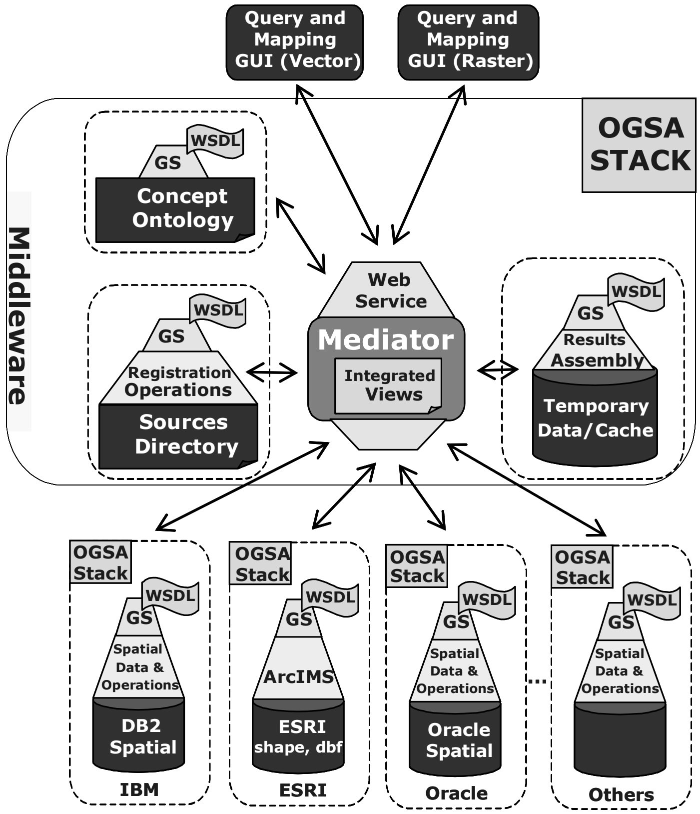

Figure 1. Grid services based mediation architecture for geospatial information.

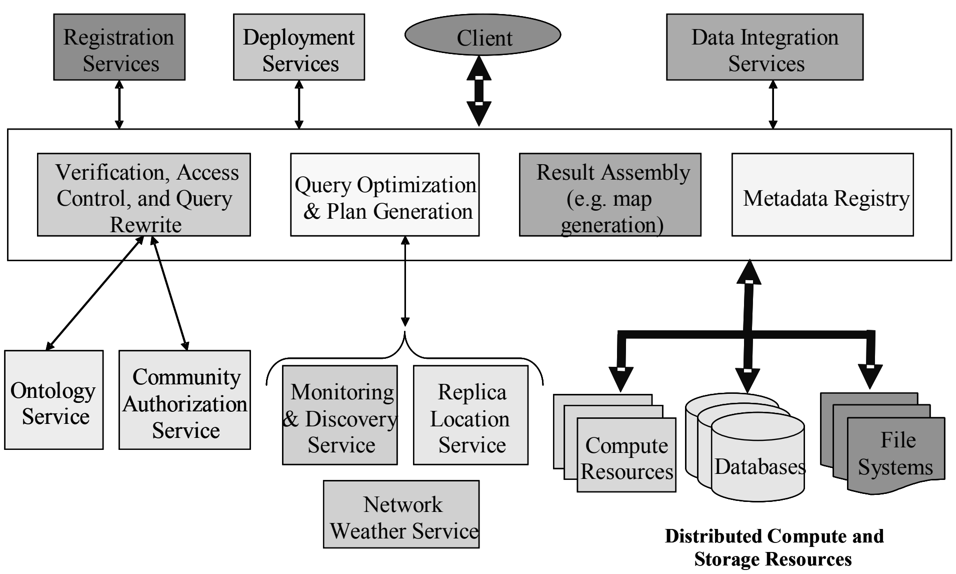

Figure 2. General composition of GEMS.

community infrastructure (i.e. managed by GEMS), mediators access various external (i.e. "nonhosted") datasets. While hosted nodes support the entire range of GEMS services (storage and replication, access control and logging, versioning, and querying), the external nodes (which, in case of GEON, include publicly accessible ArcIMS servers at *geographynetwork.com* and *gis-data.usgs.net*) only support access control and query services. All hosted services are exposed as XML (GML) sources, with an accompanying set of functions, which depend on the capabilities of the wrapper and the underlying source. For pure GML sources, no additional operations are exported, and source capabilities are limited to simple *getCoordinates* and similar requests, i.e. those retrievable by pure XQuery, while relational sources also export spatial SQL operations and stored procedures. WMS sources export the standard *GetMap*, *GetCapabilities* and *GetFeatureInfo* requests, while ArcIMS sources export a similar set of ArcXML-formatted requests (listed below in this Section). In this fashion, the task of integrating spatial information is simplified to formulating queries against a homogeneous data model with a corresponding set of operations. The task of decomposing the query into query fragments targeted to individual source capabilities is performed by the mediator. The wrappers accept SOAP messages from the mediator, convert the content of the requests into native source queries (for example, into ArcXML requests for ArcIMS servers), and add a SOAP envelope to outgoing query results.

The GEMS Data Registration Service enables users to register community data resources and processing services, external services and data sets, and *integrated views*, and make them available for discovery, access and query through a distributed metadata catalog. The key challenges in registering data include determining the appropriate content of the associated metadata, efficient indexing and replication, and propagation of metadata across the grid nodes. The availability of detailed source metadata at the mediator can improve the quality of sub-queries generated by the mediator. At the same time, the source metadata uploaded to the registry must also conform to a common source model and be fairly stable and concise. Our registry-level model of a geospatial source, for mediation purposes, generally follows the ADN (ADEPT – DLESE – NASA: ADN 2004) metadata framework, extending the required fields to support mediation and thus including:

- *Index metadata*: various metadata for indexing a data set within the grid. Data sets are typically indexed using ontologies[1] (i.e, the "semantic" dimension), which are represented using the Ontology Web Language (OWL). The registration system allows users to associate dataset schemas to an existing or user-provided ontology. Spatial datasets are also indexed by spatial (bounding rectangle) and temporal dimensions.
- *Hosted*: whether the dataset is hosted within the Grid environment, or remains external
- *Schema*: for relational data sources, includes schemas and exported function names (the schema extraction is described in (Gupta et al. 2002)). For XML (GML) sources, it includes XML schema and supported operations. For ArcIMS sources, for example, the registry keeps, beyond schema elements in ArcCatalog-generated layer metadata, records of grid services at each source that implement ArcXML's GET_SERVICE_INFO, GET_FEATURE_COUNT, GET_IMAGE, GET_FEATURES and GET_EXTRACT requests used by the map assembly services (Zaslavsky et al. 2003).
- *Access*: access mechanism for this data set, describing a local JDBC connection for databases, Web Service or OGSA service, OGSA-DAI service, etc.
- *Permissions*: a set of access control restrictions based on GEMS role-based authentication mechanism.

The Data/Service Registration process is illustrated in Figure 3 (circled numbers in the figure correspond with the registration steps outlined below). To register a data set, the owner logs into a portal, authenticates with the service (1), and enters metadata (2) providing the information described above (depending on the dataset format, some metadata can be "scraped" from the source: for example shapefile's spatial extent, schema, etc. are extracted using the free *shapelib* library (2004). Next, the GEMS registration service completes registration by:

[1] Our usage of the word "ontology" refers to a system of domain-specific concepts and their relationships, which can be represented and queried as a graph, and can be mapped to schema elements of information sources. The difference between this notion, and understanding of ontology as a philosophic category, in the GIS context, is discussed in (Smith & Mark 2001).

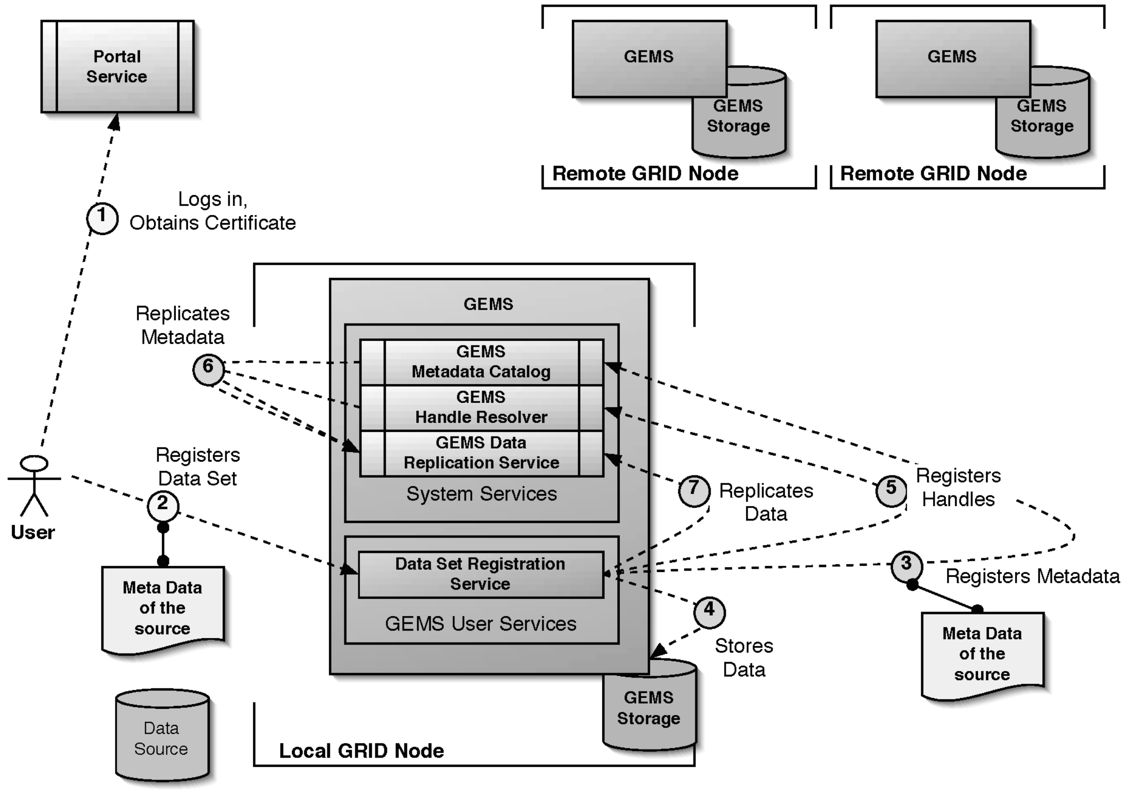

Figure 3. Data set registration in GEMS.

- placing the source metadata into a GEMS Registry (3),
- storing the data set at the local GEMS node, if the data set is hosted (4),
- replicating metadata across other Grid nodes (5) to improve availability and performance,
- archiving and replicating source data, if the source was registered as a hosted source (6) and registering references to the replicated data in the GEMS Registry Service (7).

2.2 *The middle tier: Geospatial mediation services*

As mentioned before, the mediator layer consists of a set of services for data conversion, query rewriting and execution, and results assembly, described in more detail in Section 3. As in the rest of the system, all components are loosely-coupled grid services that can be independently invoked. The middle tier also includes view definition and registration services that enable authorized users to define and publish integrated views over distributed sources. An XQuery-based view can be specified using registered source schemas, and metadata obtained through the GEMS *Discovery Service*, and published along with its XQuery view specification, description, permissions, and optional schema and ontology metadata.

2.3 *Geospatial clients*

While a detailed discussion of geospatial clients is generally beyond the scope of the paper, the architecture outline would be incomplete without a brief sketch. In a grid environment there can be a high degree of heterogeneity in mapping clients. From the mediation perspective, we are interested in client's ability to maintain state, in its rendering model (one or several server-generated images, shapefiles, SVG, etc.), and in its ability to translate user actions and state elements into mediator queries. Another aspect is that the queries issued by map clients must return meaningful maps rather than simply a query result set. A map, as a reflection of geographic "millieu" within the map spatial extent, should include layers, objects, relationships, and other pertinent components of spatial context not explicitly requested in the query (e.g. additional background layers, graticule lines and labels, place names, and highlights). While generation of a composite map is managed by GEMS map assembly service described below, a mapping client may support some limited client-side integration (e.g., displaying a stack of images produced by individual services). In an interactive Web environment, mapping is often additionally enhanced with various tooltips and information windows anchored on map elements.

We have experimented with several common clients, including a desktop GIS (ArcGIS) and three Web clients with different capabilities: a custom ArcIMS HTML viewer, an SVG viewer (a version of AxioMap, see Zaslavsky 2000), and a MapObjects-Java viewer. These clients are preconfigured to display elements of map context and support different interactivity levels. In (Zaslavsky & Memon 2004) we demonstrated how a mediator's response is translated into a presentation plan, which is then converted into map configuration files for HTML-based and SVG-based clients.

3 GEOSPATIAL DATA INTEGRATION IN GEMS

3.1 *Ontology-based rewriting*

Different geospatial sources follow different and often incompatible classification schemes and resolution standards (e.g., different land use, soil, geologic classifications, with values recorded at different hierarchical levels). Generating mappings between pairs of schemas for schema translation does not lead to a scalable solution, while associating database schema elements with concepts in a domain ontology (at the registration phase) allows querying across multiple semantically-different databases. The ontology-based rewriting services use formal ontologies materialized as OOWL files, to adjust user queries to individual source schemas and to realities of value assignments in different datasets. The include, minimally: (1) *concept expansion* service that extracts all sub-concepts of the queried term from the global ontology, and rewrites the initial query in terms of the sub-concepts; and (2) c*oncept resolution* service that extracts a set of unique values used at each source for each of the queried terms, and rewrites user query in terms of these values (at source wrappers). An example of the application of these services is shown in section 4.

3.2 *Data quality-based rewriting and evaluation service*

Spatial data are always available to certain accuracy, whether explicitly modeled in the source metadata or not. The accuracy-based rewriting service rewrites user queries against sources with known accuracy and error models, to provide *definite* and *possible* results, and evaluates the accuracy of the output map. Depending on the available accuracy specification (feature-level accuracy, or layer-level accuracy), queries can be rewritten with or without subsequent "pruning" step. In (Manpuria et al. 2003) we showed how a template query:

SELECT * from *layer1, layer2,…*
WHERE {*definitely*|*possibly*|*probably*} Aggr(spatial_condition (*layer1.geom, layer2.geom,…*))

can be rewritten if data quality information is available for input layers. The rewriting is based on a collection of error propagation templates, which provide rewriting instructions for each operation specified inside *spatial_condition.* Once a matching error template is discovered, the certainty predicate inside the WHERE clause is removed and the *spatial_condition* phrase is rewritten to reflect the semantics of the accuracy predicate. For example, in the trivial case of distance-based spatial operation and *ε-band* (Perkal 1966) certainty descriptions ($\varepsilon 1$ and $\varepsilon 2$ for the two layers respectively), a sample query:

SELECT * from *layer1, layer2* WHERE *definitely* distance(*layer1.geom*, *layer2.geom*) < D,

is rewritten as:

SELECT * from *layer1, layer2* WHERE distance(*layer1.geom*, *layer2.geom*) < [D- ε1- ε2].

3.3 *Spatial results assembly services*

Merging query results from individual sources into a composite response is an inherent mediator component. In our previous work (Baru et al. 1999, Gupta et al. 1999), the mediator received XML-formatted results from source wrappers and combined them into a single XML tree, using the instructions for stitching together result fragments specified in the initial XMAS query. This is insufficient for geospatial mediation, because:

- Query results, though returned to the mediator as virtual XML documents, typically contain or reference fragments of different types (pure XML/GML, various known text and binary vector and raster formats, and combinations of the above) with possibly different projections, schemas, referenced ontologies, spatial extents.
- The format of the output map is not completely specified by the initial query but rather determined at runtime based on the combination of client rendering capabilities, and output capabilities of sources.
- Geographic query results must be placed in spatial context not explicitly requested in the query, and the output map must comply with cartographic design principles. For this, query results are generally superimposed on a set of relevant geographic layers (perhaps retrieved from other services), and additional map requisites may be included (scale bar, north arrow, graticule grid, etc.)
- Since creating such a composite map from multiple sources is compute-intensive, the composite map should be able to support additional requests without re-querying individual data services.

To produce a composite result, map images or features retrieved from individual sources (on GET_IMAGE, GET_FEATURES, or similar requests) may be merged at the mediator or sent to the client for rendering. We have implemented a range of services that support client-level overlay or mediator-level merge of partly transparent map images from individual sources, and vector rendering of coordinate information from each source at the client or at the mediator. However, the most complete map assembly solution is accomplished by *dynamically generating* an ArcIMS image service at the mediator which can integrate both raster and vector result fragments, and generate answers to subsequent requests without re-querying the sources. The mediator-level ArcIMS image service is a typical transient grid service that is created via the grid service *Factory* interface, and supports lifetime management via *SoftStateDestruction* and *ExplicitDestruction* interfaces. The latter are invoked when additional user requests exceed the capabilities of the service and it needs to be re-initialized, or after a specified period of inactivity. The *Factory* interface creates a new grid service instance and returns a *Grid Service Handle*, which in turn can be used to retrieve the service WSDL description from the *Grid Service Reference* for subsequent querying.

A generic map assembly service has the following components (Fig. 4):

- *File Transfer Service:* This service is used to transfer selected large datasets from data source wrappers to the *staging area* at the map assembly service, using an HTTP channel, a GridFTP Web service, or any other transport service.
- *Uncompress Service:* To minimize network load, data sets are compressed at source wrappers. The *Uncompress Service* uses standard libraries (*zlib* and *Xceed*) to uncompress the data entering map assembly.
- *Image Assembly Service:* This core map assembly service combines vector and raster data fragments from individual sources, into a single ArcIMS image service, by generating a service configuration file and making the service available for querying. The newly generated service is then used to serve the resultant map image to the client (and, possibly, respond to follow-up user requests without regeneration), until it is explicitly destroyed.
- *Image Fusion Service:* This service is designed to combine raster images generated by different sources, into a composite map image (should the client be capable of displaying a single map image, eg ArcIMS HTML client)
- *Query Service:* enables querying the dynamically created map service.
- *Data Conversion Service:* Since different sources generate raster and vector data in different formats, this service is an essential part of results assembly.

– *Command Module:* represents an extensible collection of *map assembly templates* (stored as *command.xml* at the service) which bind together the processing components into a map assembly workflow.

4 GEOLOGIC MAP INTEGRATION: GRID-BASED MEDIATION IN PRACTICE

The services described above are used to resolve the following GIS request, viz. "select and map geologic formations whose geologic age is "Tertiary", in the 8-state area of Rocky Mountains, for which geologic maps are served by 9 spatial data nodes, and data sets have different database schemas and subscribe to different ontologies. The details of this system and a working demo can be accessed from GEON portal at http://www.geongrid.org.

The query is processed in the following steps:

1. GEMS *concept expansion* service parses the registered global *geologic age ontology* and returns a fragment that includes all descendents of the entered concept. For concept "Tertiary", this returned tree includes concepts "Neogene", "Pliocene", Placenzian", "Miocene", etc. (a total of 23 descendants.)
2. GEMS *concept resolution* service queries each source to return a set of actual values used to reference the initial term and its 23 descendants at that source. For term "Tertiary", these may include "Tertiary", "Quaternary/Tertiary", "Tertiary/Cretaceous", "Tertiary/Jurassic" (as in the Nevada state geologic map). This information is used to rewrite the WHERE clause of the initial query.
3. GEMS *mediator* passes the rewritten queries on to each source wrapper, and directs the output to the map assembly service. The current version of the mediator is based on X-Mediator described in (Papakonstantinou & Vassalos 2001).
4. GEMS *map assembly* service receives the initial query expression, map extent, and handles to result fragments generated by each data service (including path and data type of each fragments), and organizes then into a composite map configuration document. Then the data fragments (compressed shape files, images, or GML), are retrieved, via the *File Transfer Service*, to a local staging area and transformed as necessary, so that the map configuration can be converted into a valid ArcXML configuration file. This file is then used to create a transient grid service based on ArcIMS image service, to return the resultant map to the user.

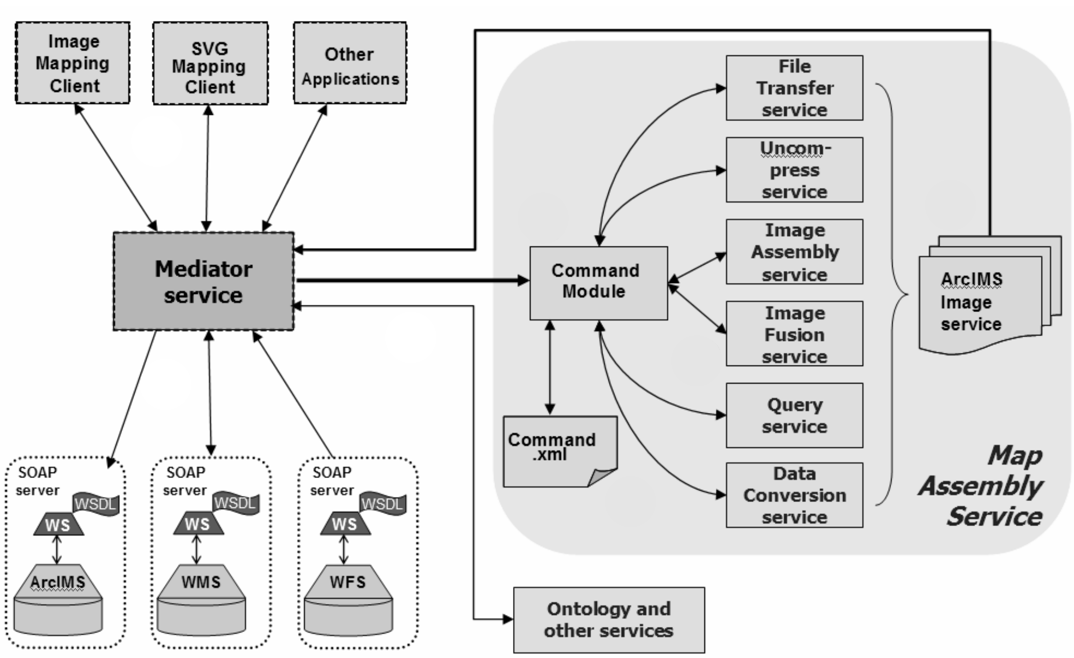

Figure 4. Internal organization of GEMS map assembly services.

5 CONCLUSION

Grid service-based approach to information integration is a promising strategy in applications that require on-demand secure query-based access to large amounts of distributed spatial data, because it provides mechanisms for dynamic resource discovery, allocation and monitoring, addresses security, authentication and authorization challenges, and is standards-conformant. This paper outlined our experience implementing the OGSA model of grid computing for registering distributed heterogeneous spatial sources on the grid and processing spatial queries against these sources, as part of GEMS (Grid-Enabled Mediation Services). While embracing the GEMS architecture, core services supporting geospatial mediation described here, implement additional functions made necessary by properties of geographic data and the need to generate composite maps as query results. Incorporating additional Grid services into the same architecture and performance tuning represent directions of our future work.

REFERENCES

ADN metadata framework, 2004. (http://www.dlese.org/Metadata/adn-item/index.htm).

Baru, C., Gupta, A., Ludäscher, B., Marciano, R., Papakonstantinou, Y., Velikhov, P. & Chu, V., 1999. XML Based Information Mediation with MIX. In *Proc. of the ACM SIGMOD 1999*, pp. 597-599.

BIRN: Biomedical Informatics Research Network, 2004 (http://www.nbirn.net)

Boucelma, O., Esid, M. & Lacroiz, Z., 2002. A WFS-based Mediation System for GIS Interoperability. In *Tenth ACM International Symposium on Advances in GIS*, pp. 23-28.

DeVogele, T., Parent, C., & Spaccapietra, S. 1998. On Spatial Database Integration. *IJGIS*, 12(4):335-352.

DQP: Distributed Query Processing on the Grid, 2003 (http://www.neresc.ac.uk/projects/OGSA-DQP/)

Foster, I., Kesselman, C. & Tuecke, S., 2001. The Anatomy of the Grid: Enabling Scalable Virtual Organizations. *International J. Supercomputer Applications* 15(3).

Foster, I., Kesselman, C., Nick, J. &Tuecke, S., 2002. The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration (http://www.globus.org/research/papers/ogsa.pdf).

GGF, 2003. Global Grid Forum (http://www.gridforum.org/)

GEON: the Geosciences Network, 2004 (http://www.geongrid.org)

GLOBUS, 2003. The GLOBUS Project (http://www.globus.org/)

GriPhyN: the Grid Physics Network, 2004 (http://www.griphyn.org)

Gupta, A., Marciano, R., Zaslavsky, I. & Baru, C., 1999. Integrating GIS and Imagery through XML-Based Information Mediation. In P. Agouris and A. Stefanidis (eds.) *Integrated Spatial Databases: Digital Images and GIS, Lecture Notes in Computer Science*, Vol. 1737, pp. 211-234.

Gupta, A., Ludäscher, B. & Martone, M., 2002. Registering Scientific Information Sources for Semantic Mediation. In *21st International Conference on Conceptual Modeling, (ER), Tampere, Finland*, pp. 182-198.

Manpuria, V., Zaslavsky, I. & Baru, C., 2003. Web services for accuracy-based spatial query rewriting in a wrapper-mediator system. In *Proc of W2GIS, Rome, December 2003* (in press).

OGC, 2000. OpenGIS Web Map Server Interfaces Implementation Specification.

OGC, 2001. Geography Markup Language (GML) 2.0.

OGC, 2002. OpenGIS Web Feature Service Implementation Specification.

Papakonstantinou, Y. & Vassalos, V., 2001. The Enosys Markets data integration platform:lessons from the trenches. In *Proceedings of the tenth international conference on Information and knowledge management*, pp. 538–540.

Perkal, J., 1966. An attempt at objective generalization. In John Nystuen (ed), *Michigan Inter-University Community of Mathematical Geographers, Discussion Paper 10*. Ann Arbor: University of Michigan.

Shapelib, 2004. C API for reading/writing ESRI Shapefiles (http://shapelib.maptools.org/)

Shimada, S. & Fukui, H., 1999. Geospatial mediator functions and container-based fast transfer interfaces in Si3CO Test-bed. *Lecture Notes in Computer Science,* Vol. 1580, pp. 265-276.

Smith, B. & Mark, D., 2001. Geographic Categories: An Ontological Investigation. *IJGIS* 15 (7): 591-612.

Wiederhold G., 1992. Mediators in the Architecture of Future Information Systems. *IEEE Computer* 25(3): 38-49.

W3C, 2001. Scalable Vector Graphics (SVG) 1.0 Specification, W3C Recommendation, 04 September 2001.

W3C, 2003a. Web Services Description Language (WSDL) Version 1.2. W3C Working Draft 24 January 2003

W3C, 2003b. Simple Object Access Protocol, W3C Proposed Recommendation, 07 May 2003.

WSRF, 2004. Web Service Resource Framework (http://www.globus.org/wsrf/)

Zaslavsky, I., 2000. A New Technology for Interactive Online Mapping with Vector Markup and XML. *Cartographic Perspectives* 37: 65-77.

Zaslavsky, I., Memon, A., Petropoulos, M. & Baru, C., 2003. Online Querying of Heterogeneous Distributed Spatial Data on a Grid. In *Proceedings of Digital Earth'2003 Conference, Brno, September 2003*.

Zaslavsky, I. & Memon, A., 2004. GEON: Assembling Maps on Demand From Heterogeneous Grid Sources. In *Proceedings of ESRI Users Conference, San Diego, CA, August 2004*.

ACKNOWLEDGMENTS

Partial support under US National Science Foundation grants #0121269 "ITR/IM: Enabling the Creation and Use of GeoGrids for Next Generation Geospatial Information", #0205049 "ITR: GEON: The Geosciences Network: A Research Project to Develop Cyberinfrastructure for the Geosciences" are gratefully acknowledged.