ISO 19119 and OGC Service Architecture

George PERCIVALL, USA

Keywords: Geographic Information, Standards, Architecture, Services.

ABSTRACT

ISO 19119, "Geographic Information - Services," has been developed jointly with the Services Architecture SIG of the OpenGIS Consortium (OGC). As a project in ISO TC211, ISO 19119 has reached the Draft International Specification stage. OGC has adopted ISO 19119 as part of OpenGIS Abstract Specification, Topic 12 "System Architecture." The OGC Web Mapping Testbed activities provided implementation experience in the development of ISO 19119.

CONTACT

George Percivall
NASA/Global Science and Technology, Inc.
Goddard Space Flight Center, Code 710
Greenbelt, MD
USA
Tel. + 1 301 286 4073
Fax + 1 301 474 59780
E-mail: percivall@gsfc.nasa.gov
Web site: http://geo.gst.com

ISO 19119 and OGC Service Architecture

George PERCIVALL, USA

1. GEOSPATIAL SERVICES STANDARDS

The OpenGIS Consortium (OGC) and ISO TC211 have jointly developed an international standard for geospatial service architecture. ISO 19119 “Geographic Information – Services” has been adopted as part of the OGC Abstract Specification, Topic 12 “OGC Architecture.” ISO 19119 has been developed based on the implementation experience of the OGC Interoperability Program testbeds.

ISO 19119 was developed in ISO Technical Committee number 211. The scope of ISO TC211 is standardization in the field of digital geographic information (http://www.isotc211.org).

TC 211 aims to establish a structured set of standards for information concerning objects or phenomena that are directly or indirectly associated with a location relative to the Earth.

The OpenGIS Consortium is an international industry consortium of more than 220 companies, government agencies and universities participating in a consensus process to develop publicly available geoprocessing specifications (http://www.opengis.org).

A Cooperative Agreement exists between TC 211 and OGC. Both organizations are working to harmonize specification development in the areas common to the two organizations. This harmonization has been accomplished in the area of Geospatial service architecture.

The editor of ISO 19119 and the author of this paper is supported by NASA’s Geographic Interoperability Office, which is a NASA wide office for geospatial standards.

2. GOALS OF GEOGRAPHIC SERVICE ARCHITECTURE

The widespread application of computers and use of geographic information systems (GIS) have led to the increased analysis of geographic data within multiple disciplines. Based on advances in information technology, society’s reliance on such data is growing. Geographic datasets are increasingly being shared, exchanged, and used for purposes other than their producers’ intended ones. GIS, remote sensing, automated mapping and facilities management (AM/FM), traffic analysis, geopositioning systems, and other technologies for Geographic Information (GI) are entering a period of radical integration.

ISO 19119 provides a framework for developers to create software that enables users to access and process geographic data from a variety of sources across a generic computing interface within an open information technology environment.

The geographic services architecture specified in ISO 19119 has been developed to meet the following purposes:

- Provide an abstract framework to allow coordinated development of specific services,
- Enable interoperable services through interface standardization,
- Support development of a service catalogue through the definition of service metadata,
- Allow separation of data instances & service instances,
- Enable use of one provider's service on another provider's data, and
- Define an abstract framework which can be implemented in multiple ways.

ISO 19119 was developed by first considering the functionality provided by “monolithic” image processing and GIS packages. The ISO 19119 architecture provides those same functionalities and more in a distributed environment, e.g., the Internet. The concepts in the architecture were then tested and refined by implementation experience from the OGC web mapping testbeds.

3. GEOSPATIAL SERVICE ARCHITECTURE

This section describes the contents of ISO 19119 “Geographic Information – Services.” ISO 19119 is identical with the OGC Abstract Specification – Topic 12 “Services”.

3.1 Interoperability Reference Model – ISO RM-ODP

ISO 19119 is based on the Reference Model of Open Distributed Processing [ISO/IEC 10746]. Architecture is a set of components, connections and topologies defined through a series of viewpoints. Using the RM-ODP, the following viewpoints are provided in ISO 19119:

The computational viewpoint is concerned with the interaction patterns between the components (services) of the system, described through their interfaces.

The information viewpoint is concerned with the semantics of information and information processing. An information specification of an RM-ODP system is a model of theinformation that it holds and of the information processing that it carries out.

The engineering viewpoint is concerned with the design of distribution-oriented aspects, i.e., the infrastructure required to support distribution.

The technology viewpoint describes the implementation of the RM-ODP system in terms of a configuration of technology objects representing the hardware and software components of the implementation.

3.2 Computational Viewpoint: a Basis for Service Chaining

The computational viewpoint is concerned with describing the entities of a distributed system independent of implementation and semantic content.

Several terms are used extensively in ISO 19119:

- **Service**: Distinct part of the functionality that is provided by an entity through interfaces.
- **Interface**: Named set of operations that characterize the behavior of an entity.
- **Operation**: Specification of a transformation or query that an object may be called to execute. It has a name and a list of parameters.

Services are accessible through a set of interfaces that are a set of operations. The aggregation of interfaces in a service defines functionality of value to the users. Users in this context are either software agents or human users. A service provides functionality that adds value. The value is apparent to the user who invoked the service. The aggregation of operations in an interface, and the definition of interface, is for the purpose of software reusability. Interfaces are defined in order to be reusable for multiple service types. The syntax of an interface may be reused with multiple services with different semantics. Interfaces are defined through operations. An operation specifies a transformation on the state of the target object or a query that returns a value to the caller of the operation.

Key to meeting the needs of geographic uses is the combination of services to achieve results specific to a task. ISO 19119 enables this through **service chaining**. ISO 19119 enables users to combine data and services in ways that are not pre-defined by the data or service providers. This level of data/service interoperability will be achieved in stages. At first service catalogues will hold entries with tight data/service binding. Eventually the infrastructure will be available for a user to determine which data can be acted on by a loosely coupled service. This capability will be enabled by the infrastructure of the larger domain of IT.

A Service Chain is defined as a sequence of services where, for each adjacent pair of services, occurrence of the first action is necessary for the occurrence of the second action. The action of making the input of one service dependent upon another service leads to treating service chains as directed graphs, where each service is a node in the graph and references to service interactions form the edges. In some cases the directed graph structure is implicit. In other cases it is necessary to make the notion of a processing graph explicit and allow such graphs to be considered as entities in their own right. Explicit representation of a service chain allows the chain to be visually represented and passed to a chain execution service, e.g., workflow service.

There are many options for service chaining, e.g., is the chain explicit, how is the progress of the chain controlled? Different approaches reflect different priorities for different applications: user in the loop vs. user supervision. To demonstrate the breadth of the trade space defined by these variations, three architecture patterns are offered in ISO 19119:

- User defined (transparent) chaining: the Human user manages the workflow.
- Workflow-managed (translucent) chaining: in which the Human user invokes a Workflow Management service that controls the chain and the user is aware of the individual services.
- Aggregate service (opaque): in which the user invokes a service that carries out the chain, with the user having no awareness of the individual services.

In addition to the difference in visibility of the services to the user, a key distinction between these patterns is the difference in control. In transparent chaining the control is exclusively with the user. In translucent, a workflow service is present which controls the chain execution, perhaps with oversight by the human. In the aggregate pattern, the aggregate service exclusively performs the control function with no visibility by the user.

The three chaining patterns discussed could be combined in a variety of ways. Each of the lowest level services could in turn implement a chain. This is recursive composition of services supported by the opaque pattern. A service chain can become a new service. The ability to define recursive composition of services provides scalability and support for top-down progressive refinement as well as for bottom-up aggregation.

The patterns could be used to define how a library of chains is constructed. A knowledgeable user could build chains using the transparent pattern. Through iterative use of the transparent pattern a chain is constructed that produces valid results. Chains are then made available for wider use following the translucent pattern. Certain chains may become routinely used and an aggregate service is built as an interface.

An example need for a translucent or opaque chaining pattern occurs in decision support. The Decision-Maker is an individual using decision support aids to help make a decision. An example of decision support aids is a service chain. The Decision Support Aid Developer is an individual who “integrates” chains of services into decision support aids.

So far we have discussed how to construct chains and not addressed if the results of a chain are semantically valid. It is assumed that human users will determine semantic validity of the results of a service chain. Several factors to consider in the semantic evaluation of a chain result are listed in ISO 19119.

- Appropriateness of starting data: are the based datasets suited to the subsequent processing? For example, accuracy and resolution of the data, thematic values are relevant.
- Affect of services on data: how do the individual services effect the data, e.g., error sources and propagation.
- Sequence of the services: how does the order of the chain affect the results? For example, should a spatial operation, e.g., orthorectification, be performed before or after a thematic operation, e.g., resampling the attribute values?

To evaluate the fitness for use of a service in a specific context, users will review a description of the service. These service descriptions are also called **service metadata**.

Service metadata records can be managed and searched using a catalogue service as is done for dataset metadata. In order to provide a catalogue for discovering services, a schema for describing a service is needed. ISO 19119 defines a metadata model for service instances. In order to place the Service Metadata in context three types of entities need to be described:

- Service Instance: a service instance is the service itself, hosted on a specific set of hardware and accessible over a network.
- Service Metadata: a service metadata record describes a service instance including a description of the services operations and an "address" to access the specific service instance.
- Service Type: in some cases a service metadata record will describe a service instance which is of a "well known type". By well-known type, it is meant that the service conforms to a published definition of a service type, i.e., a platform-specific service specification. Some clients will be able access only services of well-known type. A user could search the service metadata catalogue to find instances of a specific well-known service type. A service registry is defined to be the service that provides details on service types.

A service instance may be tightly-coupled with a dataset instance, or it may be un-associated with specific data instances, i.e. loosely-coupled. Loosely-coupled services may have an association with data types through the service type definition. In the tightly-coupled case, the service metadata shall describe both the service and the geographic dataset, the latter being defined in accordance with ISO 19115. For the loosely-coupled case, dataset metadata need not be provided in the service metadata.

The structure of service metadata includes three major classes: a section of basic service metadata that provides a general description of the service and sections that describe the operations and data available from a particular service.

Simple service architecture is defined by a set of simplifying assumptions that are relevant to implementing a message-based architecture to support service chaining. Systems compliant to this set of guidelines shall be referred to as instances of simple service architectures.

- Message-operations. The operations consist of a request and response pairs.
- Separation of control and data. Operations of an interface separate the control of the service from the access to the data resulting from a service.
- Stateful vs. stateless service. When possible service invocations are composed of a single request-response pair with no dependence on past or future interactions.
- Known service type. All service instances are of specific service types and the client knows the type prior to runtime.
- Adequate hardware. Hardware assignment is transparent to user.

3.3 Information Viewpoint: A Basis for Semantic Interoperability

The information viewpoint is concerned with the semantics of the information processing. Each particular service will need to define its syntactical interfaces through operations and its semantics through description of the meaning of the operations and their legal sequencing. This section contains a description of a taxonomy of various services. There exist multiple possible taxonomies for services, based on various classification dimensions. The one that is used here is based on the extended OSE model.

The basis of the information viewpoint is the more general Open System Environment model [ISO/IEC TR 14252]. Six classes of information technology services are used to categorize geographic services.

- Human interaction services are services for management of user interfaces, graphics, multimedia, and for presentation of compound documents.
- Model/Information management services are services for management of the development, manipulation, and storage of metadata, conceptual schemas, and datasets.
- Workflow/Task services are services for support of specific tasks or work-related activities conducted by humans. These services support use of resources and development of products involving a sequence of activities or steps that may be conducted by different persons.
- Processing services are services that perform large-scale computations involving substantial amounts of data. Examples include services for providing the time of day, spelling checkers, and services that perform coordinate transformations (e.g., that accept a set of coordinates expressed using one reference system and converting them to a set of coordinates in a different reference system). A processing service does not include capabilities for providing persistent storage of data or transfer of data over networks.
- Communication services are services for encoding and transfer of data across communications networks.
- System management services are services for the management of system components, applications, and networks. These services also include management of user accounts and user access privileges.

The taxonomy within the Processing Services category is further detailed based on the General Feature Model as presented in ISO 19109. Processing Services modify the attributes of Features; therefore Processing Services categories are based on the property types for features given by the General Feature Model.

Table 1 – Geographic services taxonomy

- Geographic human interaction services
- Geographic model/information management services
- Geographic workflow/task management services
- Geographic processing services
  - Geographic processing services – spatial
  - Geographic processing services – thematic
  - Geographic processing services – temporal
  - Geographic processing services – metadata
- Geographic communication services
- Geographic system management services

The following provides example geographic services within the geographic services taxonomy. It is not required that a system provide any service listed in these sub-clauses. It is required that if a system provides a service named in these sub-clauses that the service shall be categorized as defined in these sub-clauses. A service catalogue compliant with ISO 19119 categorizes service metadata instances in the categories of the geographic service taxonomy.

Geographic human interaction service examples:

- Catalogue viewer. Client service that allows a user to interact with a catalogue to locate, browse, and manage metadata about geographic data or geographic services.
- Geographic viewer. Client service that allows a user to view one or more feature collections or coverages. This viewer allows a user to interact with map data.
- Geographic feature editor. Geographic viewer that allows a user to interact with feature data, e.g., displaying, querying. Supports feature annotation.

Geographic model/information management services:

- Feature access service. Service that provides a client access to and management of a feature store.
- Map access service. Service that provides a client access to a geographic graphics
- Coverage access service. Service that provides a client access to and management of a coverage store.
- Catalogue service. Service that provides discovery and management services on a store of metadata about instances.

Geographic workflow/task management services:

- Chain definition service. Service to define a chain and to enable it to be executed by the workflow enactment service.
- Workflow enactment service. Interprets a chain and controls the instantiation of services and sequencing of activities.

Geographic processing services – spatial:

- Coordinate conversion service. Service to change coordinates from one coordinate system to another coordinate system that is related to the same datum.
- Rectification service. Service for transforming an image into a perpendicular parallel projection and therefore a constant scale.
- Subsetting service. Service that extracts data from an input in a continuous spatial region either by geographic location or by grid coordinates.
- Route determination service. Service to determine the optimal path between two specified points.

Geographic processing services – thematic:

- Thematic classification service. Service to classify regions of geographic data based on thematic attributes.
- Feature generalization service. Service that generalizes feature types in a feature collection to increase the effectiveness of communication.
- Subsetting service. Service that extracts data from an input based on parameter values.

Geographic processing services – temporal:

- Temporal reference system transformation service. Service to change the values of temporal instances from one temporal reference system to another.
- Sampling service. Service that extracts data from an input using a consistent sampling scheme based on temporal position values.

Geographic processing services – metadata:

- Statistical calculation service. Service to calculate the statistics of a data set, e.g., mean, median, mode, and standard deviation.

Geographic communication services:

- Encoding service. Service that provides implementation of an encoding rule and provides an interface to encoding and decoding functionality.

Geographic system management services:

- No geographic-specific system management services have been identified.

Only a subset of available services is applicable to a specific situation, e.g., image analysis. A **service organizer folder** is an aid for users in finding services applicable to their situation. A user may construct a SOF and then make that SOF available to other users performing tasks in a similar situation. A services organizer folder (SOF) is a data structure that shall contain references to a set of services that are applicable to a given situation. The SOF need not contain service chains but may contain just individual services.

3.4 Engineering Viewpoint: A Basis for Distribution

The engineering viewpoint focuses on mechanisms for distribution of services across networks. The approach of ISO 19119 is to distribute services using a multi-tier architecture model. To support flexible deployment, IT architectures are structured as multi-tiered distributed architectures. As a reference model, a logical 4-tier architecture is presented with discussion on variations in different physical architectures. The logical architecture is the arrangement of services and associated interfaces that are present in the system (see Figure 1). The physical architecture is the arrangement of components and associated interfaces that implement the services. The components are hosted on hardware computing resources or nodes. The logical architecture can be mapped to multiple physical architectures. All tiers could be mapped into one monolithic application or could be mapped using different physical client-server architectures. As shown in Figure 1, a thick user interface client architecture will typically contain a larger part of the functionality in the user service. A thin user interface client (typically a web browser) will mostly contain user dialogue and presentation code. A Web browser client is a user interface client that interacts with a Web server, using the Internet HTTP protocol with content represented in HTML and/or XML.

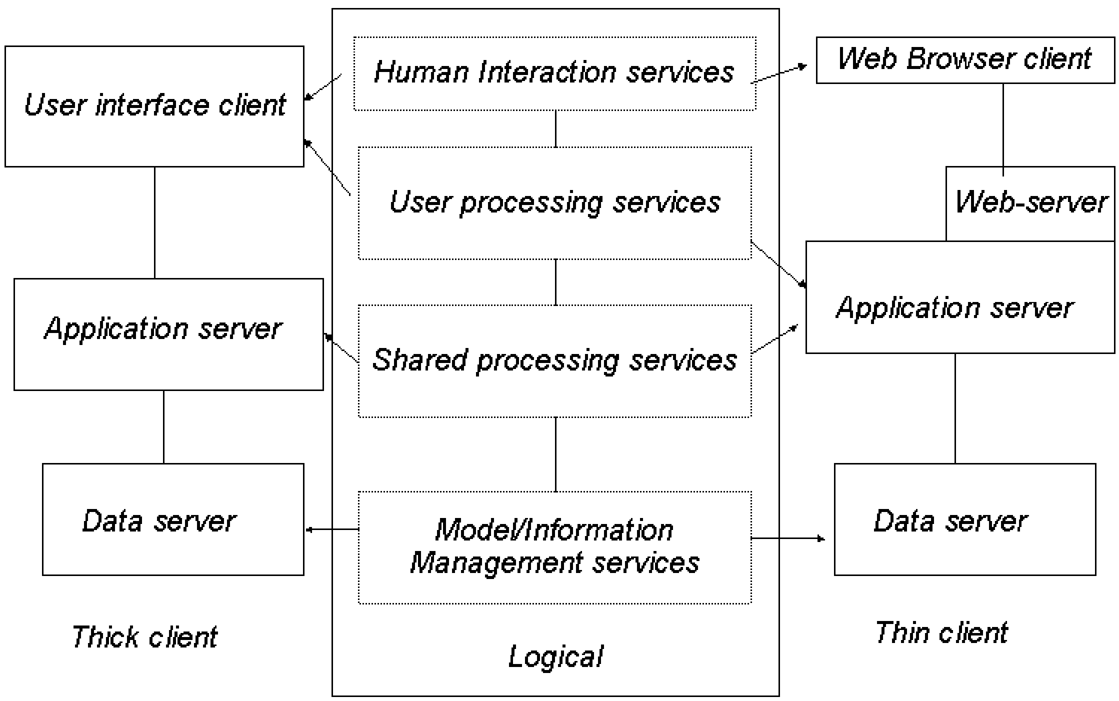

Figure 1 — Mapping logical 4 tier to thick and thin clients

3.4.1 Technology Viewpoint: a Basis for Cross Platform Interoperability

The technology viewpoint of ISO RM-ODP is concerned with the underlying infrastructure in a distributed system. To achieve interoperability in the technology viewpoint, an infrastructure that allows the components of a distributed system to interoperate is needed. This infrastructure, which may be provided by a Distributed Computing Platform (DCP), allows objects to interoperate across computer networks, hardware platforms, operating systems and programming languages. The communication between the objects in a DCP is handled by a “communication service”. The communication service allows the components in a distributed system to interoperate. In the two systems depicted in Figure 2, Communication service A allows the components of System A (Client A, Geodata Server A and GIS Service Component A) to interoperate, while Communication service B allows the components of System B to interoperate. To allow System A to interoperate with System B, the components of System A must be able to request services from the components of System B and vice versa.

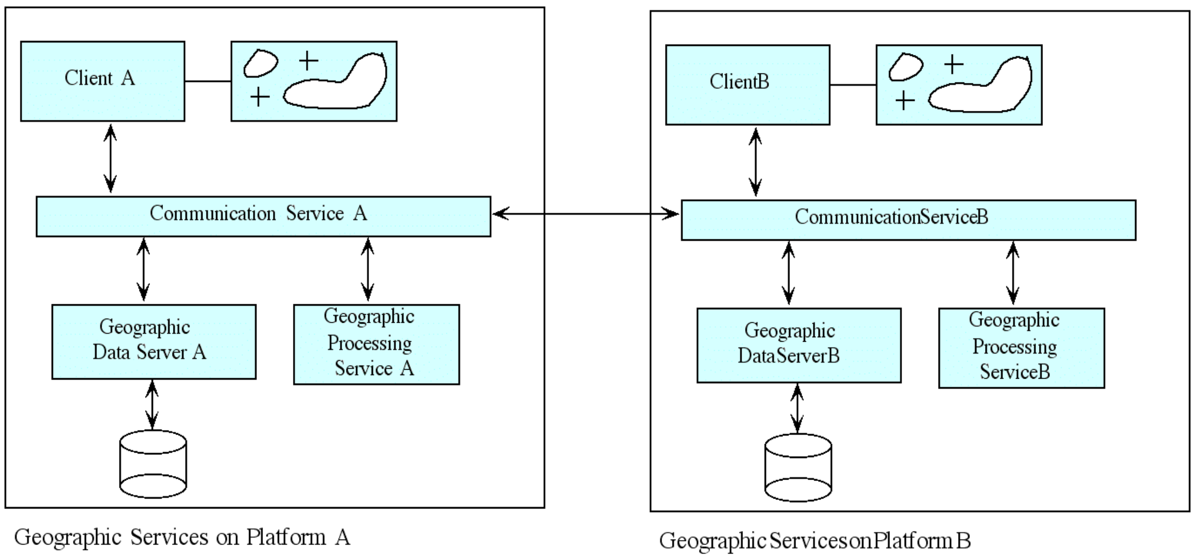

Figure 2 — Technology viewpoint model of the Interoperability Reference Model

The approach of ISO 19119 to cross platform interoperability is through a platform-neutral service specifications as the basis for multiple platform-specific service specifications. Multiple platform-specific specifications are necessary because of the variety of DCP's and the differences in the way in which they support the functional requirements. One platform-neutral service specification is needed to support interoperability of multiple platform-specific specifications. Platform-neutral models are described in UML. Platform-neutral service specifications define both static and dynamic models. Static models define objects including the attributes and operations for each object. Dynamic models capture the interaction patterns between objects and state modeling. Platform-specific models may be described in UML, together with a description of their mapping to the corresponding platform-neutral models. Development of service specifications may proceed from platform-specific to platform-neutral or from platform-neutral to platform-specific. In either case, a service specification shall not be consider complete until it has a platform-neutral model and at least one platform-specific model.

4. GEOSPATIAL SERVICE ARCHITECTURE IMPLEMENTATION

ISO 19119 has been approved by ISO TC211 as a Draft International Standard in November 2001. The concepts of ISO 19119 were based in part on the first two OGC web mapping testbeds. The third and current OGC web mapping testbed, OGC Web Services", is implementing, refining, and extending the concepts of ISO 19119.

ISO 19119 has also been used in the development of the CeoNET architecture in Canada, by CNES in France, and in the EOSDIS ClearingHOuse (ECHO) by NASA.

REFERENCES

ISO 19119, "Geographic Information - Services," ISO TC211 document number N1203, November 2001, as sent to ISO for Central Secretariat for registration as DIS, (http://www.isotc211.org).

OpenGIS Abstract Specification, Topic 12 "System Architecture", OpenGIS Consortium, version 4.2, October 2001, (http://www.opengis.org/techno/abstract.htm).

BIOGRAPHICAL NOTES

George Percivall has been a technical leader in multiple technologies and business environments. He is a recognized expert for several International Standards activities in Geographic Information Systems. He served as Chief Engineer for NASA's EOS Data and Information (EOSDIS) Core System (ECS). He was a systems engineer and manger in the GM Systems Engineering program; has performed systems engineering on the GOES/GMS geosynchronous weather satellites for Hughes Aircraft; and conducted engineering research for the US Army Corps of Engineers. He received a B.S. degree in Engineering Physics and an M.S. Degree in Electrical Engineering with Control Systems as the field of specialization. Both degrees were awarded by the University of Illinois at Urbana-Champaign. Mr. Percivall is currently the Director of GST- Geospatial Interoperability Group (http://geo.gst.com).