

General description of the CGDI Geodata Discovery Service - Z39.50

1. Identification

1.1 Service Name:

Geodata Discovery Service - Z39.50

1.2 Purpose:

This service provides a mechanism to:

- Search catalogues¹ that describe geodata resources¹;
- retrieve detailed information (metadata) that describes the geodata resource, in order to evaluate its suitability to meet a need;

If an organization has a collection¹ of geodata, implementing this service allows the collection to be searched by consumers within the CGDI.

To facilitate geodata discovery, the CGDI will contain one or more geodata registries. A registry is a catalogue that holds descriptions of other geodata collections within the CGDI (i.e. a catalogue of catalogues). Geodata registries provide the ability to search through multiple geodata catalogues simultaneously.

This service has the same purpose as the CGDI Geodata Search Service - HTTP service. Implementing either of these services may fulfill the purpose, or both may be implemented for increased interoperability.

1.3 Scope:

This service is applicable to **any** catalogue of geodata resources. A catalogue is any server that provides this interface. A catalogue contains metadata that describes a collection of related geodata resources. Geodata resources are any geospatial data entities. See Figure 1 for a graphic representation of these relationships. Example geodata resources are:

- a satellite image;
- a collection of related geographic features, such as the road network of British Columbia;
- a collection of geodata products, such as a collection of Landsat-5 images over Canada.

To support the definition of the catalogue itself through "simple" metadata (i.e., the definition of the catalogue in CGDI geodata registries), all geodata resources described within a single catalogue should be related in some way, with all resources being at a similar aggregation level. A single catalogue should **not**, for example, contain one entry that defines a geodata collection and another that defines a geodata "product".

¹ See definition in Section 1.3.

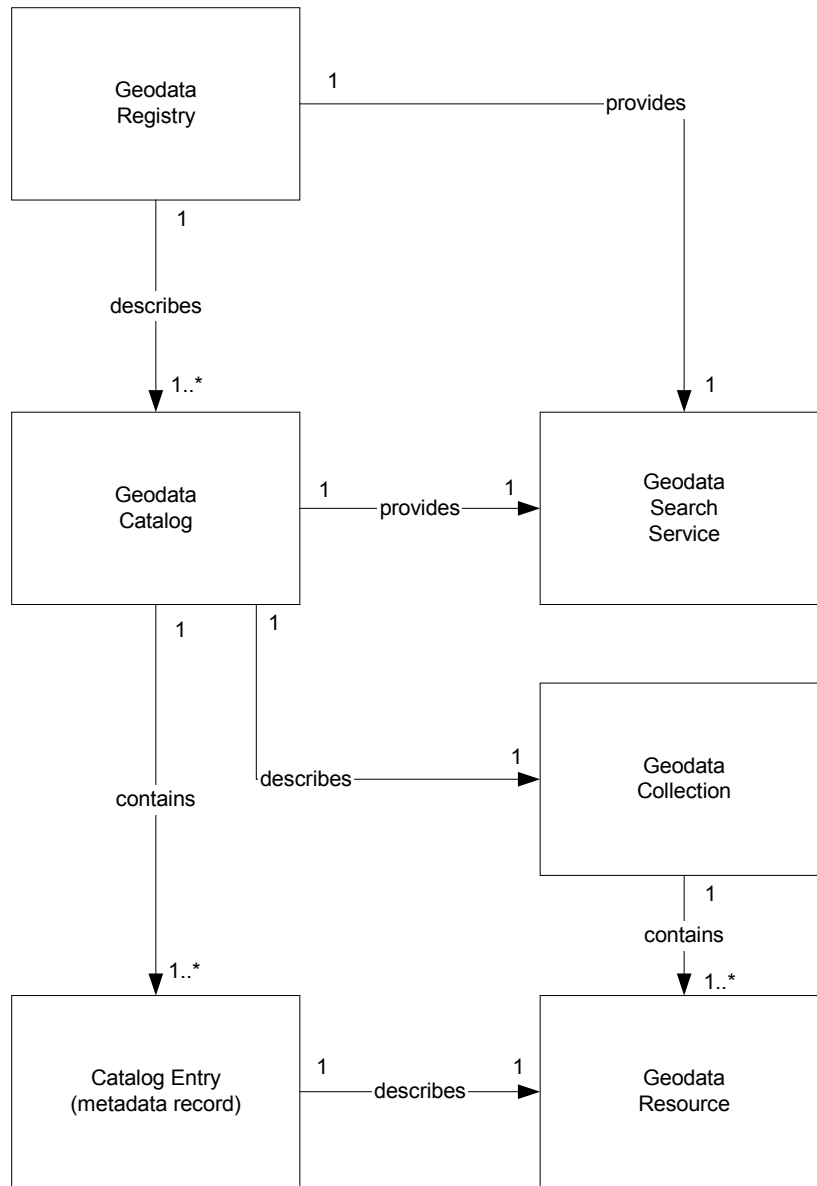


Figure 1 Entity Relationship Diagram

This service provides an interface to search for and retrieve metadata records. It does **not** provide an interface to populate the catalogue. Though this service may provide information on how to access geodata resources (within the metadata), it does **not** provide interfaces to perform the access (i.e. no interface is provided to, for example, download an image listed in the catalogue).

1.4 Service Version:

This specification defines version 1.0 of the CGDI Geodata Discovery Service - Z39.50.

2. Description

2.1 Functional Description:

This service is based on the Z39.50 protocol. It defines an interface whereby a catalogue can be searched, and selected metadata records returned. The following messages are to be supported by a CGDI geodata search server (note that these are a subset of the full set of Z39.50 messages):

- `initRequest` - initiates a connection between the client and server.
- `initResponse` - server's response to an `initRequest`, which completes the connection between the client and server;
- `searchRequest` - transmits a query from the client to the server. This message also specifies the conditions under which catalogue entries are included in the `searchResponse`;
- `searchResponse` - server's response to an `searchRequest`. The set of catalogue entries selected by the `searchRequest` is called the "result set". The `searchResponse` returns the number of entries in the result set, and if the result set is small, it also returns the catalogue entries. The server maintains the result set for subsequent `presentRequests` (see below). The entries in the result set are ordered and are referenced by the position (beginning with one) within the set;
- `presentRequest` - a request from the client to retrieve some, or all, of the results from a previously executed search. The request also indicates the required format of the results, both the level of detail of the results and the record format;
- `presentResponse` - server's response to a `presentRequest` which contains the requested results;
- `releaseRequest` - initiate termination of the connection;
- `releaseResponse` - completes connection termination;
- `abort` - aborts the connection.

The Z39.50 protocol is stateful. The client issues an `initRequest`, followed by a `searchRequest` (which selects a result set), and then one or more `presentRequests`. The client may send a `releaseRequest` message at any time to terminate the connection. A simplified interaction diagram is provided in Figure 2, and a simplified server state transition diagram is provided in Figure 3.

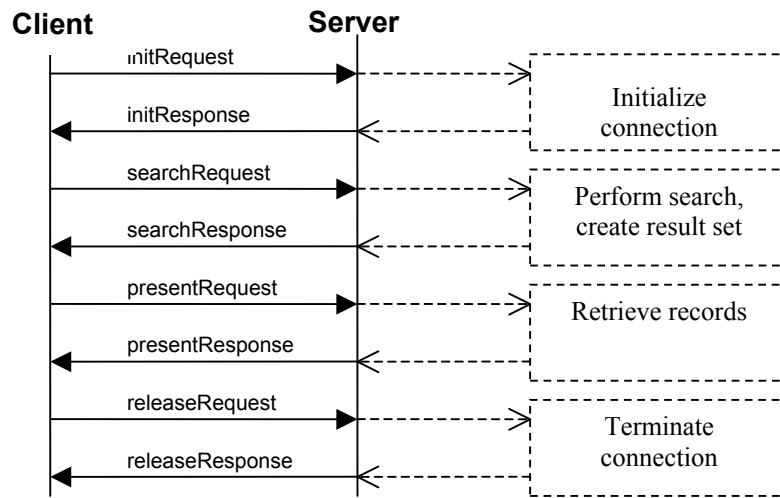


Figure 2 Example Geodata Discovery Service Sequence Diagram

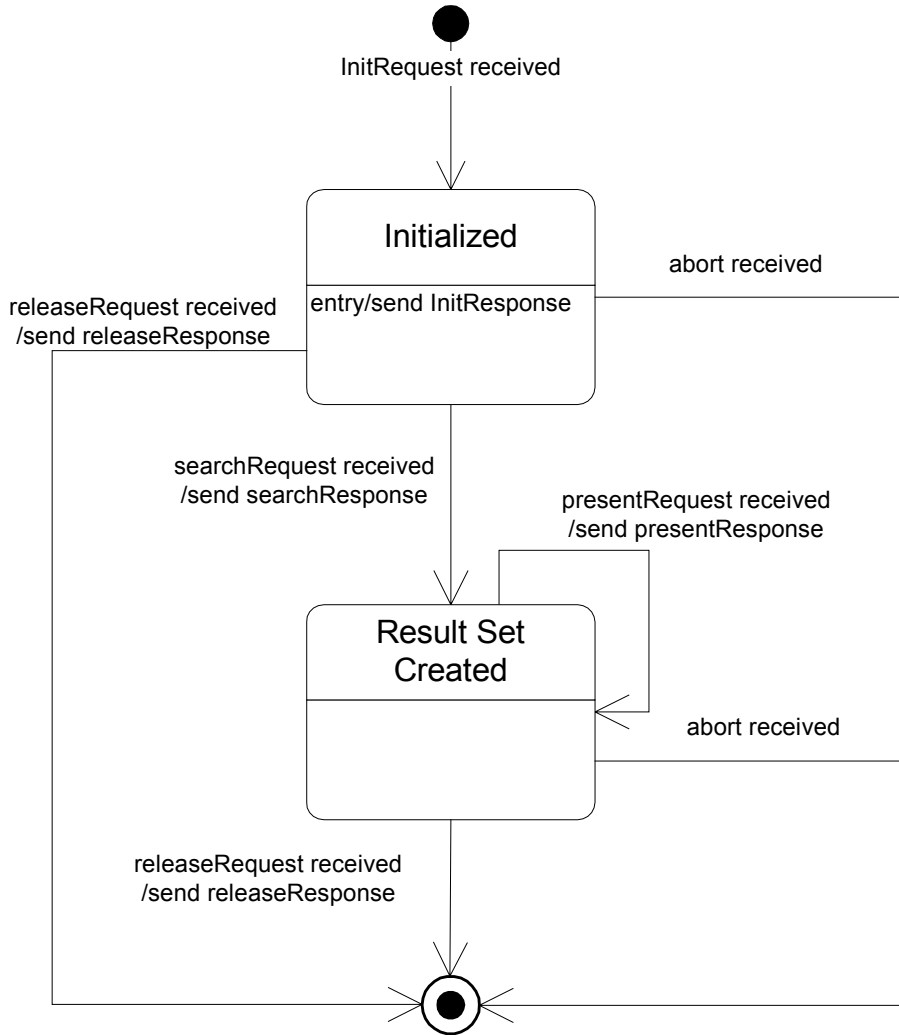


Figure 3 Simplified Geodata Discovery Server State Transition Diagram

2.2 Reference Documents

[1] Application Service Definition and Protocol Specification for Open Systems Interconnection, ANSI/NISO Z39.50-1992
 (<http://www.cni.org/pub/NISO/docs/Z39.50-1992/>)

[2] Z39.50 Application Profile for Geospatial Metadata or "GEO", Version 2.2, 27 May 2000 (<http://www.blueangeltech.com/standards/GeoProfile/geo22.htm>)

[3] ISO/DIS 19115. Geographic Information - Metadata

3. Data Requirements

3.1 Input data

The critical input data for this service is the query, which is part of the searchRequest message. The query used is the Z-39.50 type-1 query, which uses Reverse Polish Notation (RPN) formatting.

The query consists of operands, and operators (AND, OR or AND-NOT) that operate on the operands. A conceptual example query would be:

(operand1 AND operand2) OR operand3.

formatted in RPN, this query would be:

operand1 operand2 AND operand3 OR

Each operand represents a component of the query that selects those catalogue entries that have a particular value (or set of values) within a particular metadata "field". An operand conceptually represents a statement such as:

Title field contains the word Landsat

In Z39.50, the operand consists of an attribute list and a term. The term is the "value" to be matched ("Landsat" in the above example). The attribute list specifies how the match is made. The attribute list contains the following attributes:

- A Use Attribute - identifies a "field" in the metadata to be matched against ("Title" in the above example);
- A Relationship Attribute - specifies a matching relationship such as "contains", "equals" or "less than";
- A Structure Attribute - specifies, for example, if the match is at the word or phrase level.
- An optional Truncation Attribute - specifies whether one or more characters may be omitted in matching the search term.

3.2 Reference data

To implement this service, the provider must have a digital database of metadata (catalogue) that describes geodata resources. This service operates on the digital catalogue.

4. User Interface Requirements

None.

5. Organizational Requirements/Recommendations

None.

6. Access Requirements and Constraints

6.1 Access Requirements:

The server providing this interface requires a TCP/IP network connection to all the clients to be served. If the server lies behind a firewall, the firewall must be configured to allow traffic from clients through to the server port.

6.2 Access Constraints:

The Z39.50 protocol does provide an authentication mechanism in the `initRequest` message. This mechanism requires that the client and server agree, prior to initiation of the connection (i.e. outside the scope of this service), on the access information that will be provided. It is recommended that CGDI servers providing this interface do not enforce user authentication.

7. Standardization

7.1 Reference Standards

This section identifies applicable standards, profiles and functional standards that constrain the design and implementation of the service in terms of its functionality and external interfaces.

This service is based on the Z39.50 protocol. Z39.50 is a generic protocol that enables a client to request that a server search a database to identify records that meet a specified criteria, and to retrieve some or all of the identified records. To use Z39.50, one needs to define a set of attributes that can be searched, and the format of the results. The attribute sets and result formats are particular to each application domain, and are defined in an Application Profile. For the CGDI the currently mandated Application Profiles are one of:

- The *Z39.50 Application Profile for Geospatial Metadata or "GEO", Version 2.2*. GEO is based on the Content Standard for Digital Geospatial Metadata (CSDGM) that was developed by the Federal Geographic Data Committee. The CSDGM standard defines the data schema of a geodata catalogue entry. The GEO profile defines attribute numbers that map one-to-one to CSDGM field. In addition it defines a number of attribute numbers that match against multiple fields (for example there is an attribute number to match against all fields).
- In the future, a profile based on the ISO geographic metadata standard [3].

7.2 Reference Implementation

This section identifies an existing implementation of the service that can be used as a testing environment, or examined to assist in the development of another implementation.

The reference implementation, within the CGDI, for this service is provided by CompuSult's MetaManager product. This product is available at <http://www.fgdctoolkit.com/html/download.html>.

The CGDI Discovery portal (formally CEONet) provides the reference client for this service. Through the portal data search interface, any geodata search server that has been previously registered can be searched. A Web API is provided, through which any instance of a Z39.50 geodata catalogue service (including those not registered) can be invoked (follow the Web API links from the CGDI Discovery home page, which is accessed through <http://ceonet.gc.ca>). A "wizard" facilitates the construction of a URL that executes a search on the search server and returns the results to the web browser.

7.3 Other Related Implementations

Other implementations of Z39.50 server software are listed at <http://www.niso.org/z3950.html#software>. The more popular products for the geospatial domain are:

- Isite (<http://clearinghouse2.fgdc.gov/ftp/>)
- YAZ toolkit (<http://www.indexdata.dk/yaz/>)
- Blue Angel Metastar (<http://www.blueangeltech.com/products.htm>)
- SMMS GeoConnect™ Geodata Management Server (<http://www.rtseusa.com/enter/geoconnect/>).

8. Implementation Specification

8.1 Statefull

Yes

8.2 Methods:

8.2.1 Method 1:

8.2.1.1 Method Name

initRequest / initResponse

8.2.1.2 Method Purpose

Initialize the connection between client and server and negotiate the protocol version, and message size to be used for future messages.

8.2.1.3 Logical Input

The *initRequest* message contains:

- authentication information (optional);
- the preferred and maximum message size;
- the list of functions (search, present, etc.) that are to be used in the session.

8.2.1.4 Logical Output

The *initResponse* message contains:

- the preferred and maximum message size;
- the list of functions (search, present, etc.) that are supported;
- a result indicating acceptance or rejection of the connection.

8.2.2 Method 2:

8.2.2.1 Method Name

searchRequest / searchResponse

8.2.2.2 Method Purpose

To select a result set that matches a client supplied query, and optionally to return the result records.

8.2.2.3 Logical Input

The *searchRequest* message contains:

- a query;
- a list of databases to be searched. Traditionally, within the CGDI, only one database is specified.

- A name to be given to the result set and a hint on how to handle result set name clashes.
- The boundaries that divide the possible result set sizes into three ranges: "small", "medium" and "large". When the result set size fall in to the "small" or "medium" range, the searchResponse will actually include results, otherwise the searchResponse includes only the number of results.
- The set of elements (metadata fields) to be included in the search result. Two element sets are specified, one to be used when the result set size is "small", and the other to use if the result set size is "medium". GEO defines 3 element sets: brief, summary and full.
- The preferred record syntax for the results. GEO requires support of 4 record syntaxes: GRS-1, HTML, SUTRS and XML. The CGDI does not require support of GRS-1 or SUTRS (check this??).

8.2.2.4 Logical Output

The searchResponse message contains:

- the status of the search.
- the size of the result set;
- the number of result records included within the searchResponse message. This will be zero if the result set size falls in the "large" range, in which case the results must be retrieved through a presentRequest.
- optionally, if the result set size falls into the "small" or "medium" range, the result records that match the query. Each record contains the metadata fields defined by the element set specified in the searchRequest (note the fields may be different depending on whether the result set size is "small" or "medium" - see the searchRequest specification above). Note that if any of the result records cannot be retrieved, it is replaced with a diagnostic record. Also note that fewer records may be returned if returning all records would exceed the negotiated maximum message size.
- the position, within the result set, of the next record following those returned (or zero if the last record in the result set is being returned).

8.2.3 Method 3:

8.2.3.1 Method Name

presentRequest / presentResponse

8.2.3.2 Method Purpose

To retrieve a number of records from a previously created result set.

8.2.3.3 Logical Input

The presentRequest message contains:

- the result set from which to retrieve records;
- the number of records to retrieve, and the position within the result set from which to retrieve;
- the set of elements (metadata fields) to be included in the search results. GEO defines 3 element sets: brief, summary and full.
- The preferred record syntax for the results. GEO requires support of 4 record syntaxes: GRS-1, HTML, SUTRS and XML. The CGDI does not require support of GRS-1 or SUTRS (check this??).

8.2.3.4 Logical Output

The presentResponse message contains:

- the status if the present operation.
- the number of results included within the presentResponse message.
- the result records that were requested. Each record contains the metadata fields defined by the element set specified in the presentRequest. Note that if any of the result records cannot be retrieved, it is replaced with a diagnostic record. Also note that fewer records may be returned, if returning all records would exceed the negotiated maximum message size.
- the position, within the result set, of the next record following those returned (or zero if the last record in the result set is being returned).

8.2.4 Method 4:

8.2.4.1 Method Name

releaseRequest / releaseResponse

8.2.4.2 Method Purpose

To terminate the connection between client and server.

8.2.4.3 Logical Input

The releaseRequest message has no content .

8.2.4.4 Logical Output

The releaseResponse message has no content.

8.2.5 Method 5:

8.2.5.1 Method Name

abort

8.2.5.2 Method Purpose

To abort the connection between client and server.

8.2.5.3 Logical Input

The abort message has no content.

8.2.5.4 Logical Output

There is no response to an abort message.