

Information Retrieval Z39.50 for Protocol Specification

Engr. Faizullah Mahar
Department of Electrical Engineering and Computer Science,
Balochistan University of Engineering and Technology, Khuzdar, Pakistan

Abstract: The information retrieval service describes an activity between two applications: an initiating application, the client and a responding application, the server. The server is associated with one or more databases. Communication between client and server is carried out by the Z39.50 protocol. The protocol specifies format and procedures governing the exchange of messages between client and server. Thus enabling a client to request that the server search a data base and identify records that meets specified criteria and retrieve some or all of the identified records. This paper describes the protocol procedures, protocol model, rules of extensibility, general and specific conformance requirements and detail requirements for the protocol addresses communication between the client and server.

Key words: Protocol, retrieval, state table, conformance, syntax

Basic of the protocol

The protocol specifies formats and procedures governing the exchange of messages between a client and server, thus enabling the client to (a) request that the server search a data base and identify records that meet specified criteria and (b) retrieve some or all of the identified records (Faizullah, 1996).

The client may initiate requests on behalf of a user; the protocol addresses communication between the client server (which may reside on different computers); it does not address interaction between the client and user.

The client may send a search, indicating one or more databases and including a query as well as the parameters which determine whether records identified by the search should be returned as part of the response. The server responds with a count of records identified and possibly some or all of the records. The client assumes that records selected by the search form a 0result set0 (an ordered set, order determined by server) and records may be referenced by position within the set. Optional capabilities include:

- ! The client may specify an element set indicating data elements to retrieve in cases where the client does not wish to receive complete database records.
- ! The client may indicate a preferred syntax for response records.
- ! The client may name a result set for subsequent reference.
- ! The client may delete a named result set.

- ! The server may impose access control restrictions on the client by demanding authentication before processing a request.
- ! The server may provide resource control by sending an unsolicited or solicited status report; the server may suspend processing and allow the client to indicate whether to continue.

Information retrieval service

The Information Retrieval service describes an activity between two applications: an initiating application, the client and responding application, the server. The server is associated with one or more databases.

Communication between client and server is carried out by the Z39.50 protocol, the information retrieval application protocol the formats and procedures governing the transfer of information between a Z39.50 origin /target pair the specification is logically divided into procedures pertaining to client and procedures pertaining to server. The portions of the client and server that carry out the Z39.50 protocol procedures are referred to respectively as the Z39.50 origin and the Z39.50 target.

Model and characteristics of the information retrieval service

Communication between origin and target is via a Z39.50-Association (Z-association) within an application association. A Z-association is explicitly established by the origin and may be explicitly terminated by either origin or target, or implicitly terminated by termination of the A-association.

There may be multiple consecutive Z-associations within an A-association. There may be multiple consecutive as well as concurrent operations within a Z-association.

The role of origin and target may not be reversed within a Z-association. A Z-association can not be restarted, thus once a Z-association is terminated no status information is retained except information that is explicitly saved.

Z39.50 Services

Z39.50 services are carried out by the exchange of messages between the origin and target. A message is a request or a response. Services are defined to be confirmed, non-confirmed, or conditionally confirmed.

A confirmed service is defined in terms of a request (from the origin or target) followed by a response (from the peer). For example, search is a confirmed service initiated by the origin, the Search service is defined in terms of a search request from the origin followed by a search response from the target. Access-control is an example of a confirmed service initiated by target.

Anon Confirmed Service is defined in terms of a request from the origin or target, with no corresponding response. For example, Trigger Resource Control is non-confirmed service initiated by origin; Segment is a non-confirmed service initiated by the target.

A Conditionally Confirmed service is a service that may be invoked as either a confirmed or non-confirmed service. It is defined in terms of a request (from the origin or target) followed possibly by a response (from the peer). For example, Resources-control is a conditionally confirmed service initiated by the target (Thoms, 1991)

Z39.50 Operations

This standard describes eight operation types: Init, Search, Present, Delete, Scan, Sort, Resources-report and Extended services.

A request from the origin of a particular operation type initiates an operation of that type (for example a search request initiates a search operation) which is termed by the respective response from the target. Only the origin may initiate an operation and not all origin requests do so.

A request that initiates an operation is called an initiating request and response that ends an operation is called a terminating response. From the origin perspective, an operation begins when it issues the initiating request and ends when it receives the terminating response. From the target perspective, the operation begins when it receives the initiating request and ends when it sends the terminating response. An operation consists of the initiating request and the terminating response, along with any intervening related messages (Thoms, 1991).

Protocol procedures

Presentation and association control services

The information retrieval protocol may be used in conjunction with the presentation layer and the association control service element (ACSE)

Service provided by the presentation layer

Z39.50 may use the presentation services to provide a presentation connection for communication between a Z39.50 origin/target pair. The communication service that supports in communication oriented service.

A Z39.50 origin establishes application-associations as necessary with the target. The Z39.50 application service element (ASE) may then used to provide a connection-oriented interaction between Z39.50 systems.

Association control services

The complete application service may include ACSE and one or more specific application service.

ACSE, is used to establish an A-association and provides association management. The life of an A-association has three distinct phases: establishment, information transfer and termination. ACSE provides services for the establishment and termination phases, including the selection of an application context, specifying information including the set of service elements that are valid during the information transfer phase. Prior to the exchange of Z39.50 APDUs, the information retrieval service user invokes the association control services required to establish an association

with an application context encompassing the information retrieval service (Halsall, 1985).

Protocol model

To specify protocol procedure the abstract implementation-independent concepts of service -user, service provider and service primitive are used.

A service provider provides a communication path between two service users. In this model, the service provider is analogous to the application layer composed at the Z39.50 origin/target pair. The client is modeled as a service-user together with a target. The two service users are referred to as the origin service-user and target service-user.

A service primitive is an element of interaction between a service-user and the service provider. There are four types of service primitives: Request, Indication, Response and Communication.

For a confirmed service initiate by the origin (i.e., for Z39.50: Init, Search, Present, Delete, Resource-report, Sort, Scan, Extended-services) they are used as follows:

- ! Request- A primitive issued by the origin to the service-provider in order to invoke some procedure
- ! Indication-A primitive issued by the service provider to the target service-user to indicate that a procedure has been invoked by its peer
- ! Response -A primitive issued by service-user to the service provider at the completion of the procedure previously invoked by an indication
- ! Confirmation- A primitive issued by the service-provider to the origin service-user to complete the procedure previously invoked by a request

Notes

- ! For confirmed service initiated by the target (i.e., for Z39.50: Access-control and Resources-control) the role of origin and target are reversed.
- ! For non-confirmed service (i.e., for Z39.50: Segment, Trigger-resource-control, Close) only the Request and Indicative primitives are used.

Primitives are conceptual and their use neither specifies nor precludes any specific implementation of a service. Only primitives that correspond to some element of the service involving the exchange of information between system are defined.

From the perspective of the service-user, the service provider is system-independent. For the exchange of the protocol however, a distinction is drawn between the portion of the service-provider residing on the client and the portion of the service-provider residing on server (respectively, the origin and the target). The sequence of interactions for a confirmed service initiated by the origin is:

- ! Request Primitive from origin service-user to service-provider.
- ! Protocol message from origin to target.

- ! Indication Primitive from service-provider to target service-user.
- ! Response primitive from target service-user to service provider.
- ! Protocol Message from target to origin.
- ! Confirmation Primitive from service-provider to origin service user.

Notes

- ! For a confirmed service initiated by the target, the role of origin and target are reversed.
- ! For a non-confirmed service, only step-1 through 3 applies.

The following illustrates the sequence of interactions that occur for search operation:

- ! Search request from origin service-user to service-provider.
- ! Search APDU from origin to target.
- ! Search indication from service-provider to target service-user
- ! Search response from target service-user to service-provider.
- ! Search-response APDU from target to origin.
- ! Search confirms from service-provider to origin service-user.

The interactions between service-user and service-provider, as represented by step 1 and 6 for the client and by step 3 and 4 for the server, are described solely to facilitate the specification of protocols. These steps do not represent inter system communication and the means by which they are implemented are not constrained by this affection. For example, in an actual implementation the target service-user and service provider must be combined in a single program and step 3 and 4 might not have any real physical manifestation

State table

For both origin and target, there three protocol machines defined, one for the Z-Association (called the "Z. -machine") and two for 39.50 operations called "o-machines"). One O-machine is for a present operation and one O-machine is for any other type operation, excluding Init, which is included in the Z-machine.

There is one instance of the Z-machine (within a given application association) each for the origin and target; there may be multiple concurrent instance of the O-machines.

Each state table shows the interrelationship between the state of an operation or Z-Association, the incoming events that occur in the protocol, the action taken and, finally, the resulting state. The state tables do not constitute a formal definition of the IRMP. They are included to provide a more precise specification of the protocol procedures. The following conventions are used in the state tables (Satton and Mc Gill, 1991).

State table cells

The intersection of an incoming event (row) and a state (column) forms a cell. A blank cell represents the combination of an incoming event and a state that is not defined for the IRMP.

A non-blank cell represents an incoming event and state that is defined for the IRPM. Such a cell contains one or more actions, separated by semicolon (;). The last such action specified is always a transition to the resulting state, in parenthesis.

Invalid intersections

Blank cells indicate an invalid intersection of an incoming event and state. The state table defines correct operation only. They do not specify actions to be taken in response to incorrect operation (for example, erroneous protocol control information, incorrect protocol actions, etc.). Such actions are not within the scope of the specification, although implementations must consider them.

Predicates

Some actions are predicated on a certain condition, or "predicate. The notation for these actions taken one of the following forms:

:`[predicate]` actions: or :`[predicate]` actions else actions:

where "action" is either single action or multiple actions separated by semicolons. The following predicates are defined:

predicate	Meaning
resp	"Response required" on Resource control PDU.
noResp	"No response required" on a resource control PDU
conc	Concurrent operations in effect
noOps	No active operations.

Variables

The following variables are used:

Variable	Meaning
<op>	An operation type (other than Init): search, present, delete, scan, sort, resources-report, extended services.
opCnt	Number of active operations.
ResSt	Return state. An Integer; the action "(resSt)" means " go to the state whose value is reSt."

Rules of extensibility

All syntactical errors in received APDUs are considered to be protocol errors except for the following case: Unknown data elements and unknown options within the Options data element, will be ignored are received Init APDUs

Conformance

A system claiming to implement the procedures in this standard shall comply with the conformance requirements

The system shall:

- (a) Act in the role of origin or target.
- (b) Support the Init, Search and present services.

A system must support the Init, Search, and present services.

This means that an origin must be capable of sending Init, Search, and present requests and receiving the respective responses. A target must respond properly to Init, Search and Present requests with respective responses.

An origin may indicate (via option bits) during initialization that it does not intend to utilize the present service during the Z-association; this does not constitute non-conformance. If, however, an origin indicates that it does not intend to utilize the Present service and the target refuse, this does constitute non-conformance on the part of the target.

- (b) Support the syntax
- (c) Support the syntax
- (d) Support the type-1 Query

An origin must be capable of formulating a type-1 query within a search request and a target should expect to receive a type-1 query.

An origin or target may support other query types. If the origin fails to send a type-1 query during a Z-Association, this does not constitute non-conformance on the part of the origin. If, however, the origin does not send a type-1 query and the targeted responds with a diagnostic indicating " query type not supported" this does constitute non-conformance on the part of the target.

- (e) Support(at minimum) version2 of the protocol
- (f) Follow the procedures specified
- (g) Assign values of APDU data elements according the procedures

This standard is one of a set of standards to facilitate the interconnection of computer Systems. This standard defines protocol within the application layer of the reference model and is concerned in particular with the search and retrieval of information in databases (Ahuja, 1990).

This standard specifies a client/ server based protocol for information retrieval. It specifies procedures and structures for a client to search a database provided by server, retrieve data base records identified by a search, scan a term list and sort a result set. It also specifies access control, resource control, extended services and a "help" facility.

The protocol address communication between the client and server (which may reside on different computers), it does not address interaction between the client and the end user.

Definitions and nomenclature

APDU	Application Protocol Data Unit A unit information, transferred between origin and target, whose format is specified by the Z39.50 protocol, consisting of application-protocol-information and possibly application-user data.
ASC	Application Service Element
ACSE	Association Control Service Element
Client	The application that includes the origin; the data base user.
Data base	A collection of information units containing related information. Each unit is data base record
Initiating request	A request that initiates an operation.
Server	The application that includes the target, the data base provider.
Starting fragment	A fragment that starts at the beginning of the record.
Final fragment	A fragment that ends at the end of a record.

References

Ahuja, 1990. Design and Analysis of Computer Communication Networks, McGraw-Hill Computer Science Series, 1990.

Faizullah, M., 1996. Design aspects of cost effective local and long distance computer communication system, paper presented in 36th convention of IEP held in May, 1996.

Halsall, F., 1985. Introduction to data communication and computer net works, Addison-Wesley, Wokingham, England, 1985.

Salton and McGill, 1991. Introduction to modern Information Retrieval, McGraw-Hill Computer Science Series, 1991.

Thomas, K.J., 1991. Computer Applications for Engineers, John Wiley and Sons, 1991.