First edition
1998-07-15

Information and documentation — Information retrieval (Z39.50) — Application service definition and protocol specification

Information et documentation — Recherche d'information (Z39.50) — Définition du service de l'application et spécification du protocole

Reference number
ISO 23950:1998(E)

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

Draft International Standards adopted by the technical committees are circulated to the member bodies for voting. Publication as an International Standard requires approval by at least 75 % of the member bodies casting a vote.

International Standard ISO 23950 was prepared by ANSI/NISO (as ANSI/NISO Z39.50-1995) and was adopted, under a special “fast-track procedure”, by Technical Committee ISO/TC 46, *Information and documentation*, in parallel with its approval by the ISO member bodies.

Annexes 1 to 12 form an integral part of this International Standard. Annexes 13 to 16 are for information only.

International Organization for Standardization
Case postale 56 • CH-1211 Genève 20 • Switzerland
Internet iso@iso.ch

Printed in Switzerland

Contents

Page

Annexes

List of Figures and Tables

Foreword

(Informative)

ISO 23950 is identical in text to ANSI/NISO Z39.50-1995 (except for certain style discrepancies between ISO and ANSI standards, for example, ISO standards have "Annexes" where ANSI standards have "Appendices"), with identical name: *Information Retrieval (Z39.50) Application Service Definition and Protocol Specification*. Note that "Z39.50" is explicitly incorporated into the name of both standards, in order to avoid any possible confusion that these are different standards, and because "Z39.50" is commonly used to refer to the service and protocol defined by this standard. Note that related standards ISO 10162 and ISO 10163 are superseded by the approval of this standard, ISO 23950. Throughout the remainder of this foreword, references to "Z39.50-1995" pertain to ANSI/NISO Z39.50-1995, which is identical to ISO 23950. References to Z39.50-1988, Z39.50-1992, and Z39.50-1994, refer to earlier versions, not identical to ISO 23950.

ANSI/NISO Z39.50-1995, Information Retrieval (Z39.50) Application Service Definition and Protocol Specification, was a revision of ANSI/NISO Z39.50-1992. Because draft versions of Z39.50-1995 were referred to as Z39.50-1994, implementors should take note that any draft referred to as Z39.50-1994 is not the latest version of this standard. Z39.50-1995 is the final, approved version of the standard which was preceded by various drafts referred to as Z39.50-1994.

This protocol was originally proposed in 1984 for use with bibliographic information. Interest in Z39.50 broadened and in 1990 the Z39.50 Implementors Group (ZIG) was established. Members of the ZIG include manufacturers, vendors, consultants, information providers, and universities who wish to access or provide access to various types of information, including bibliographic, text, image, financial, public utility, chemical, and news. ZIG membership is open to all interested parties.

In 1989 the Z39.50 Maintenance Agency was formed and administered at the Library of Congress. It was assigned to revise Z39.50-1988 to achieve bit-compatibility with the International Standard, ISO 10162/10163, Search and Retrieve, SR. At that time, various enhancements to support a wide range of information retrieval activities had been proposed for the 1992 version. However, many of these features were not fully developed, and their incorporation into the 1992 standard would have caused significant delay. Consequently the maintenance agency deferred the proposed new features with a commitment to implementors that development of the required features would proceed, and that the next version would be a compatible superset of the 1992 standard. Z39.50-1992 replaced and superseded Z39.50-1988, and became a compatible superset of SR.

In 1992 the maintenance agency conducted a formal survey of Z39.50 implementors to determine the relative importance of proposed new features. The survey's purposes were (a) to begin to narrow the list to a manageable set, (b) to determine whether the proposed features were adequately specified and understood, and (c) to gauge their perceived cost and complexity. The survey results revealed certain features to be indispensable, and that certain others features could be eliminated from further consideration. For a third set of features, the survey was inconclusive and the disposition of those features eventually was determined by consensus.

Development of Z39.50-1995 began in late 1991. For each meeting of the ZIG, from December 1991 through April 1994, the maintenance agency developed a revised draft. Implementors carefully scrutinized each draft and discussed them at length both over the ZIG Internet mail list, and at the ZIG meeting. Comments and discussion for each draft and the agreements reached at each ZIG meeting were incorporated into a subsequent draft. In April 1994, the ZIG recommended that the draft be put into final form.

The 1992 version came to be known as "version 2," and the 1995 version, "version 3." However, although these version designations do have specific *protocol* significance, they do not refer to versions of the *standard*. Z39.50-1992 specifies protocol version 2; Z39.50-1995 specifies protocol versions 2 and 3.

Although Z39.50-1992 replaced and superseded Z39.50-1988 (which is obsolete) the relationship between Z39.50-1992 and Z39.50-1995 is quite different: Z39.50-1995 is a compatible superset of the 1992 version. An implementor may obtain complete details of version 2 from the Z39.50-1995 document, and build an implementation compatible with Z39.50-1992.

Basics 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 database and identify records that meet specified criteria, and (b) retrieve some or all of the identified records.

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

Z39.50-1992 provides the following basic capabilities, all of which are supported in this standard as well. The client may send a search, indicating one or more databases, and including a query as well as 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 may then retrieve selected records. The client assumes that records selected by the search form a "result set" (an ordered set, order determined by the 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. For example, the client might specify "If 5 or fewer records are identified, transmit 'full' records; if more than 5 records are found, transmit 'brief' records."
- The client may indicate a *preferred syntax* for response records, for example, USMARC.
- 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.

Query Formulation

This standard fully specifies and mandates support of the *type-1* query, expressed by individual search terms, each with a set of attributes, specifying, for example, type of term (subject, name, etc.), whether it is truncated, and its structure. The server is responsible for mapping attributes to the logical design of the database. Terms may be combined in a type-1 query, linked by Boolean operators. Terms and operators are expressed in Reverse Polish Notation.

Attribute Sets

The attributes associated with a search term belong to a particular attribute set, whose definition is *registered*, that is, assigned a unique and globally recognized *attribute-set-id*, an *Object Identifier*, which is included within the query.

Annex 3, ATR, defines and registers the attribute-set *bib-1*, which specifies various attributes useful for bibliographic queries. The bib-1 attribute set was developed by the bibliographic community; it is intended that attribute sets will be developed and registered as needed by other communities. Additional attribute sets may be registered outside of the standard.

Response Records

The protocol distinguishes two types of records that may occur in response messages from the server: database and diagnostic records.

Annex 5, REC, registers object identifiers for various MARC formats, including USMARC, UKMARC, Norway MARC and CANMARC; these object identifiers accompany database records returned by the server. The appendix defines several other types of record formats, and provides for registration of additional record formats.

Diagnostic records are similarly accompanied by an object identifier which identifies their format. Annex 4, ERR, defines and registers two diagnostic record formats (one of which was defined in Z39.50-1992) that include various diagnostic codes useful for bibliographic applications. Additional diagnostic record formats may be registered.

New Features

Provided below is a summary of the enhancements in Z39.50-1995 (versus the 1992 version). The designations "version 2" and "version 3" refer to protocol version; "Z39.50-1992" and "Z39.50-1995" refer to the respective standards. Thus where a particular feature is described as "new in Z39.50-1995," that generally means it applies in either protocol version. An example is Scan: an implementor may add the Scan service to an existing implementation of Z39.50-1992 without incorporating any other new features.

The enhancements described below fall into four categories: search, retrieval, new services and facilities, and miscellaneous enhancements.

Search

Attributes. A number of enhancements pertain to attributes and attribute sets. In version 3, attributes may be combined from different attribute sets within a single query (even for a single search term). This presents two advantages: First, it is useful when searching multiple databases (although version 2 supports multiple-database searches, all attributes within a query must belong to a single attribute set, which inhibits the ability to search multiple databases, unless those databases are similar). Second, new attribute sets may now be defined with less replication.

Version 3 provides two further enhancements allowing flexibility in the definition of attribute sets. First, new data types for attribute values are defined (in version 2 only numeric values are allowed). Second, an attribute set definition may now list alternative sets of evaluation rules (for example, whether the server is allowed to substitute an attribute that it thinks is more appropriate), and the query may select one of the alternatives. The enhanced bib-1 attribute set definition exploits this new feature.

The bib-1 definition in Z39.50-1995 also included many new attributes (as well as all of the attributes in Z39.50-1992).

Extended Result Set Model. The basic model of a result set is developed in Z39.50-1992; the 1995 version describes an "extended result set model," which supports extended proximity searching.

The extended model also supports a new version 3 search function, *restriction*, which is (in effect) an operation on a result set. It permits selection of records from a result set, based on specified attributes.

Search Term. The search term for a query may take on a variety of data types in version 3. (In version 2 a search term is binary and thus essentially has no data type, so the type is often described by a structure attribute.) This enhancement will simplify queries (as well as attribute set definitions) by reducing the need for structure attributes.

Intermediate Results. In Z39.50-1995 the server may provide information per query *component* (i.e., per sub-query, per database), as part of the Search response (version 3 only), or as part of resource-control when the server reports on the progress of the search. The server may also create and provide access to a result set for individual query components.

Retrieval

Segmentation. In version 2, a retrieval response is limited to a single message; the server attempts to fit the requested records into the message, and if it cannot, it simply fits as many as it can. The client might want to retrieve, for example, ten thousand records, knowing it cannot retrieve them in a single message. Typically the client will request all ten thousand records, wait for the response, determine how many records are retrieved, and then send another request for the remaining records. This works well in many environments but is unacceptably slow for high-speed networks. The server must await a request before sending each set of records, which introduces a delay; the delay may be negligible for conventional networks, but is intolerable for high-speed networks. In version 3 a server may respond to a retrieval request with multiple consecutive response messages without intervening requests.

A more serious segmentation problem occurs when a *single* record is too large to fit in a single message. Version 3 thus introduces a second level of segmentation: an individual record may span response messages. A client or server may choose to support either level of segmentation, or no segmentation (in which case version 2 rules apply).

Retrieval Tools. The ZIG worked intensively over two years to develop an extensive model and suite of tools for a wide range of retrieval functions to support various retrieval applications, in particular, document retrieval. The model is detailed in Annex 14, RET. Several new object classes were designated in Z39.50-1995 (schemas, tagSets, variants) and specific objects from these and other classes are defined. Annex RET provides detailed semantics for these objects and describes how they are used together to provide a variety of document retrieval capabilities. Following are a few examples:

- A single database record might include a number of documents. The client may discover and retrieve a specific document, rather than the entire database record.
- The client may retrieve a specific portion of a document, logical or physical, for example, specific pages, a specific chapter, a specific caption, all captions, or all images. The client might retrieve just headings, for example, all chapter or section headings.
- A document might be available in a wide variety of formats (e.g., PostScript, SGML), languages, presentation parameter (e.g., line length, lines per page, columns), and other variants. The client may discover what variants are supported for a document, as well as information associated with a particular variant form: for example the cost to retrieve the document according to a specific variant, or its size. Finally, the client may then retrieve the document (or specific portion) according to the desired variant.
- Associated with a document, for a given search, may be *hits*: pointers to terms (within the document) relevant to the search. The client might retrieve hits along with a document to quickly locate the satisfying portions. Or the client might retrieve only the hits (ranked in order of importance), and subsequently retrieve only the indicated satisfying portions.

New Services and Facilities

Scan and Sort. Scan and Sort were new services in Z39.50-1995. These are used respectively to scan terms in a list or index, and to sort a result set.

Scan is currently the only service in the Z39.50 Browse facility, but it is intended that various other browse capabilities will be added in future versions.

Extended Services. Extended Services was a new facility in Z39.50-1995. It includes a new Z39.50 service, the *Extended Services service*, used to initiate a specific extended service task, which is executed outside of the Z39.50 session and whose progress may be monitored using Z39.50 services. Specific extended services include: save a result set, set a periodic query schedule, export a document, order a document, and update a database.

Explain. The new Explain facility allows a client to retrieve details of the server implementation: general features (description, contact information, hours of operation, restrictions, usage cost, etc.) databases available for searching, indexes, attribute sets, attribute details, schemas, record syntaxes, sort capabilities, and extended services. The server maintains Explain information in a special database that may be accessed by the client using the Z39.50 search and retrieval facilities. The format of the Explain information is detailed in the standard.

Some Explain information is transparent to the client, intended for direct display to the client-user, and is so designated (e.g., “general features”). Some Explain information is intended to be shared by client and user. For example, the client may retrieve a list of searchable databases; for each database in the list the client might display an *informal* name, an icon, and a brief description. Meanwhile the client would retain the *actual* database name to be used in a protocol message, which probably would not be displayed. Some Explain information may be completely transparent to the user. For example, the client may retrieve information about attributes supported for a database and use that information when formulating a query (when converting a user-supplied query to a Z39.50 type-1 query).

Miscellaneous Enhancements

Termination and Re-initialization. Version 3 includes a more flexible approach to termination of a Z39.50 session, to allow, in effect, re-initialization without taking down the network connection.

Concurrent Operations. Multiple concurrent operations are allowed in version 3. In version 2, operations are strictly serial.

Diagnostics. Most Z39.50 services include diagnostic capability. In version 2 a diagnostic must conform to a specific format defined within the standard. In version 3, diagnostic formats may be externally defined and registered. One such (new) format is defined, along with a comprehensive set of diagnostics.

Access Control Formats. Z39.50-1992 provides access control, but does not define any access control formats. Z39.50-1995 defined formats for encryption and authentication, and a format allowing the server to prompt the client for arbitrary information.

Character Set Support. A new data type, “International String,” has been introduced for character strings. Its definition allows greater flexibility for a client and server to agree to the use of a particular language and one or more character sets during a session.

Units. New data types are introduced for support of units. These definitions allow standard representations to be used to represent unit type and unit. For example, unit type might be “mass,” and unit, “kilogram.”

Extensibility and Negotiation. Version 3 provides a powerful extensibility feature. Each protocol message includes a field designated for information whose format is to be defined externally. These externally defined formats will be registered and maintained by the Z39.50 Maintenance Agency as provisional extensions to the standard and for experimental use and possible consolidation into a subsequent version.

In Z39.50-1995 the concept of a “negotiation record” was introduced. The client may include a negotiation record within the initialization message to propose that some condition be in effect for the session (for example, the use of a particular language and one or more character sets). The server may respond, indicating whether the proposal is accepted, or indicate a counter-proposal.

The negotiation record is an application of the new extensibility feature. Negotiation records will be defined externally and maintained by the Z39.50 Maintenance Agency.