INTERNATIONAL STANDARD

ISO
23950

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

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.

iTeh STANDARD PREVIEW

(standards.iteh.ai)

ISO 23950:1998
https://standards.iteh.ai/catalog/standards/sist/8e606f3d-fb42-4fb2-9798-
046cc0414359/iso-23950-1998

Contents

Page

Annexes

List of Figures and Tables

iTeh STANDARD PREVIEW
(standards.iteh.ai)

ISO 23950:1998
https://standards.iteh.ai/catalog/standards/sist/8e606f3d-fb42-4fb2-9798-
046cc0414359/iso-23950-1998

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:

- n 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.

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

1. Introduction

This standard is one of a set of standards produced to facilitate the interconnection of computer systems. It is positioned with respect to other related standards by the Open Systems Interconnection (OSI) basic reference model (ISO 7498). This standard defines a protocol within the application layer of the reference model, and is concerned in particular with the search and retrieval of information in databases.

1.1 Scope and Field of Application

This standard defines the Information Retrieval Application Service (section 3) and specifies the Information Retrieval Application Protocol (section 4). The service definition describes services that support capabilities within an application; the services are in turn supported by the Z39.50 protocol. The description neither specifies nor constrains the implementation within a computer system. The protocol specification includes the definition of the protocol control information, the rules for exchanging this information, and the conformance requirements to be met by implementation of this protocol.

Intended for systems supporting information retrieval services, and for organizations such as information services, universities, libraries, and union catalogue centers, this standard addresses connection-oriented, program-to-program communication. It does not address interchange of information with terminals or via other physical media.

1.2 Version

There have been three publications of Z39.50: Z39.50-1988, Z39.50-1992, and Z39.50-1995; and there has been one publication of the Search and Retrieve Protocol, ISO 10163-1:1993. The three publications: Z39.50-1992, ISO 10163-1:1993, and Z39.50-1995 (but not Z39.50-1988) each incorporate the concept of a protocol version, and three protocol versions are defined: version 1, version 2, and version 3. ISO 10163-1:1993 is based on protocol version 1; Z39.50-1992 is based on protocol version 2; Z39.50-1995 is based on protocol version 2 as well as protocol version 3. (There is no protocol version associated with Z39.50-1988.)

This International Standard, ISO 23950, is based on version 2 and version 3. It assumes that version 1 and version 2 are identical, thus implementations that support version 2 automatically support version 1 (otherwise, version 1 is not explicitly mentioned elsewhere in this standard). Procedures within this standard that apply specifically to version 2 or version 3 are noted as such.

1.3 Referenced Standards

The following normative documents contain provisions which, through reference in this text, constitute provisions of this International Standard. For dated references, subsequent amendments to, or revisions of, any of these publications do not apply. However, parties to agreements based on this International Standard are encouraged to investigate the possibility of applying the most recent editions of the normative documents indicated below. For undated references, the latest edition of the normative document referred to applies. Members of ISO and IEC maintain registers of currently valid International Standards.

ANSI/NISO Z39.53-1994, Codes for the Representation of Languages for Information Interchange.
ANSI/NISO Z39.58-1992, Common Command Language for Online Interactive Information Retrieval.
ISO 2709:1996 — Information and documentation - Format for information interchange.
ISO 4217:1990 — Codes for the representation of currencies and funds.
ISO 7498:1984 — Information processing systems - Open systems interconnection - Basic reference model.
ISO 8649:1987 — Information processing systems - Open systems interconnection - Service definition for the association control service element.

ISO 8650:1987 — Information processing systems - Open systems interconnection - Protocol specification for the association control service element.

ISO 8777:1993 — Information and documentation - Commands for interactive text searching.

ISO 8822:1988 — Information processing systems - Open systems interconnection - Connection oriented presentation service definition.

ISO 8824:1990 — Information processing systems - Open systems interconnection - Specification of Abstract Syntax Notation One (ASN.1).

ISO 8825:1990 — Information processing systems - Open systems interconnection - Specification of basic encoding rules for Abstract Syntax Notation One (ASN.1).

ISO 10160:1997 — Information and documentation - Open systems interconnection - Interlibrary loan application service definition.

ISO 10161-1:1997 — Information and documentation - Open system interconnection - Interlibrary loan application protocol specification - Part 1: Protocol specification.

ISO 10163-1:1993 — Information and documentation - Open systems interconnection - Search and retrieve application protocol specification - Part 1: Protocol specification.

Note: Although this ISO standard supercedes ISO 10163-1, there are provisions within this standard intended to maintain compatibility with ISO 10163-1, because of existing implementations of that standard.

ISO — International register of coded character sets.

2. Definitions

For the purposes of this International Standard the following definitions apply.

A-association—See Application association.

Abstract database record—An abstract representation of the information in a database record. An abstract database record may be formed by the application of an abstract record structure (defined by a schema) to the database record. An element specification may be applied to an abstract database record forming another instance of the abstract database record.

Abstract record structure—The primary component of a database schema. An abstract record structure applied to a database record results in an abstract database record.

Abstract syntax—A description of a particular data type using an abstract syntax notation. It can be referenced by an OID (object identifier).

Abstract syntax notation—A language that allows the denotation of data types in a representation-independent manner. ASN.1 is an example.

Access point—A unique or non-unique key that can be specified either singly or in combination with other access points in a search for records. An access point may be equivalent to an element (defined by an abstract syntax), derived from a set of one or more elements, or unrelated to any element.

Access point clause—An operand of a type-1 query (informal).

Aggregate present response—Segment requests (if any) together with the Present response, for a Present operation.

APDU—See **Application Protocol Data Unit**.

Application association—A communication session between a database user and a database provider. It may consist of one or more consecutive Z-associations.

Application Protocol—The rules governing the format and exchange of information between an origin and target.

Application Protocol Control Information —Information conveyed by an application protocol data unit.

Application Protocol Data Unit—A unit of 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.

AppliedVariant—One of three usages for a variant specification. The applied variant is the variant specification that the target applies to an element included in a retrieval record. See also **variantRequest** and **supportedVariant**.

ARS—See **Abstract record structure.**

ASN.1—Abstract Syntax Notation One, as specified in ISO 8824 and ISO 8825.

Attribute—A characteristic of a search term, or one of several characteristic components which together form a characteristic of a search term.

Attribute element—An attribute represented by a pair of components: an attribute type and a value of that type.

Attribute list—A set of attribute elements and the attribute set id to which it belongs. An attribute list is combined with a search term to form an operand of a type-1 query. Usually, one of the attribute elements from the set corresponds to a normalized access point, against which the term (as qualified by the other attribute elements) is matched.

Attribute set—A set of attribute types, and for each, a list of attribute values. Each type is represented by an integer, unique within that set (as identified by its attribute set id), and each value for a given type is unique within that type.

Attribute set id—An OID that identifies an attribute set, to which an attribute element (within an attribute list) belongs.

Attribute type—A component of an attribute element. An attribute set defines one or more attribute types and assigns an integer to each type (it also defines values specific to each type). For example, bib-1 assigns the integer 1 for the attribute type "Use."

Attribute value—A component of an attribute element. An attribute set defines one or more values for each attribute type that it defines. For example, bib-1 defines the Use attribute "personal name."

Client—The application that includes the origin; the database user.

Client system—The system on which the client resides.

Composition specification—A specification that may be included in a Present request to indicate the desired composition (elements and record syntax) of the retrieval records. It includes a schema identifier, element specification, and record syntax identifier.

Conditionally confirmed service—A service that may be invoked as confirmed or non-confirmed. It is defined in terms of a request (from the origin or target) followed possibly by a response (from the peer). For example, Resource-control is a conditionally confirmed service, initiated by the target. See also **Non-confirmed service** and **Confirmed service**.

Confirmed service—A service that 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; Access-control is a confirmed service initiated by the target. See also **Non-confirmed service** and **Conditionally-confirmed service**.

Database—A collection of information units containing related information. Each unit is a database record.

Database record—A local data structure representing an information unit in a database.

Database schema—A common understanding shared by the origin and target of the information contained in the records of the database, which allows the subsequent selection of portions of that information via an element specification. A schema defines an abstract record structure, which, when applied to a database record, results in an abstract database record.

Data element—See **Element.**

Element—A unit of information defined by a schema.

ElementRequest—A request, included with an element specification, for the retrieval of a specific element. The element request may include a variantRequest, indicating the desired variant form of the element.

Element set name—An element specification in the form of a primitive name.

Element specification—An instance of an element specification format, or an element set name. An element specification transforms an abstract database record into another instance of the abstract database record (this may be a null transformation). The element specification selects elements from the abstract database record, and possibly also specifies variant forms for those elements.

Element specification format—A structure used to express an element specification.

Element specification identifier—The object identifier of an element specification format, or an element set name.

Exceptional record size—The maximum size of the record that may be included in a Present response, in the special case when a single, exceptionally large record (i.e., larger than preferred-message-size) is requested.

Facility—A logical group of Z39.50 services; in some cases, a single service. For example, the Retrieval facility consists of the Present service and the Segment service; the Search facility consists of the Search service. Alternatively, a facility might not consist of services, but instead might use services of other facilities. For example, the Explain facility does not define any services, but uses the Search and Present services.

Final fragment—A fragment that ends at the end of a record. See **Fragment.**

Fragment—A proper substring of a record. (This definition is meaningful only in the context of level-2 segmentation, described in section 3.3.; within that section, a record is considered to be a string of bytes.)
GRS—Generic Record Syntax.

Initiating request—A request that initiates an operation.

Intermediate fragment—A fragment that neither starts at the beginning nor ends at the end of a record. See **Fragment**.

IR—Information Retrieval.

Item—(1) A result set item. (2) A bibliographic item; see ISO 10160.

Maximum segment size—The largest allowable segment of an aggregate Present response (when segmentation is in effect).

Name—A linguistic construct, expressed in some language, that corresponds to an object. A name denotes (i.e., identifies) the object to which it is bound.

Non-confirmed service—A service that is defined in terms of a request from the origin or target, with no corresponding response. For example, Segment is a non-confirmed service initiated by the target. See also **Confirmed service**.

Object identifier—An unambiguous, globally-recognized, registered identifier for a data object, assigned by a registration authority.

OID—See **Object identifier**.

Operation—An initiating request and the corresponding terminating response, along with intervening related messages. For example, a Search operation always includes a Search request and Search response, and may also include access control and resource control messages. Multiple concurrent operations may occur within a Z-association.

Operation type—The name of an initiating request. For example a Search request initiates an operation whose type is "search."

Origin—The entity that initiates a Z-association and initiates operations during the Z-association.

Origin service-user—That portion of a client that makes requests upon the origin. See **Service-user**.

OSI—Open Systems Interconnection.

P-context—See **Presentation context**.

Preferred message size—The maximum size of a Search response or Present response when no segmentation is in effect. It is expressed in terms of the sum of the sizes (in bytes) of the response records, not including protocol control information.

Presentation context—The pairing of an abstract syntax with a transfer syntax, negotiated by the presentation layer, in order for that abstract syntax to be used during the application association.

Primitive—See **Service primitive**.

Primitive name—A name whose internal structure is not required to be understood or have significance to users of the name. Note: **primitive name** is not related to **primitive**.

Record syntax—An abstract syntax requested by the origin or used by the target to represent retrieval records. For a complete definition, see section 3.6.3.

Response record—A retrieval record or surrogate diagnostic record representing a database record in a Search response or (aggregate) Present response.

Result set—A local data structure used as a selection mechanism for the transfer of records, identified by a query. Its logical structure is a named, ordered list of result set items, and, possibly, unspecified information that may be used as a surrogate for the search that created the result set.

Result set item—A database name, a pointer to a record within the database, and possibly additional, unspecified information associated with the record.

Result set record—An idiomatic expression referring to the database record represented by a result set item. See **Result set**.

Retrieval record—The exportable structure defined by the application of a record syntax to an abstract database record.

RPN query—A search query represented in Reverse Polish Notation (RPN) format.

Schema—See **Database schema**.

Segment—A message that is sent (or is in preparation for transmission) by the target as part of an aggregate Present response, i.e., a Segment request or Present response.

Server—The application that includes the target; the database provider.

Server system—The system on which the server resides.

Service—(1) A Z39.50 service, as in the "search" service; (2) an extended service, as in the "persistent result set extended service"; (3) the service-provider.

Service primitive—An abstract, implementation-independent representation of an interaction between the service-user and the service-provider. The four types of service primitives are: Request, Indication, Response, and Confirmation.

Service-provider—An abstraction of the totality of those entities (the origin and target) that provide a service to peer service-users. The concept of service-provider is employed to facilitate the specification of protocol procedures. It is used only in section 4.2.2 to describe the protocol model.
Note: the service-provider is not related to the database provider or to the provider of telecommunication services.

Service-user—An origin service-user or a target service-user. That portion of a client or server that makes requests upon the origin or target respectively. The concept of service-user is employed to facilitate the specification of protocol procedures. It is used only in 4.2.2 to describe the protocol model.
Note: The service-user is not related to the database user.