



Search Web Services - searchRetrieve Operation: Binding for SRU 1.2 Version 1.0

Committee Draft 01

30 June 2008

Specification URIs:

This Version:

<http://docs.oasis-open.org/search-ws/june08releases/sru-1-2-V1.0-cd-01.doc> (Authoritative)
<http://docs.oasis-open.org/search-ws/june08releases/sru-1-2-V1.0-cd-01.pdf>
<http://docs.oasis-open.org/search-ws/june08releases/sru-1-2-V1.0-cd-01.html>

Latest Version:

<http://docs.oasis-open.org/search-ws/v1.0/sru-1-2-V1.0.doc>
<http://docs.oasis-open.org/search-ws/v1.0/sru-1-2-V1.0.pdf>
<http://docs.oasis-open.org/search-ws/v1.0/sru-1-2-V1.0.html>

Technical Committee:

OASIS Search Web Services TC

Chair(s):

Ray Denenberg <rden@loc.gov>
Matthew Dovey <m.dovey@jisc.ac.uk>

Editor(s):

Ray Denenberg rden@loc.gov
Larry Dixon ldix@loc.gov
Matthew Dovey m.dovey@jisc.ac.uk
Janifer Gatenby Janifer.Gatenby@oclc.org
Ralph LeVan levan@oclc.org
Ashley Sanders a.sanders@MANCHESTER.AC.UK
Rob Sanderson azaroth@liverpool.ac.uk

Related work:

This specification is related to:

- [Search Retrieve via URL \(SRU\)](#)

Abstract:

This is a binding of the Search Web Services - searchRetrieve operation – Abstract Protocol Definition. This binding is the specification of SRU 1.2.

Status:

This document was last revised or approved by the [OASIS Search Web Services TC](#) on the above date. The level of approval is also listed above. Check the “Latest Version” or “Latest Approved Version” location noted above for possible later revisions of this document.

Technical Committee members should send comments on this specification to the Technical Committee's email list. Others should send comments to the Technical Committee by using the

“Send A Comment” button on the Technical Committee’s web page at <http://www.oasis-open.org/committees/search-ws>

For information on whether any patents have been disclosed that may be essential to implementing this specification, and any offers of patent licensing terms, please refer to the Intellectual Property Rights section of the Technical Committee web page (<http://www.oasis-open.org/committees/search-ws/ipr.php>).

The non-normative errata page for this specification is located at <http://www.oasis-open.org/committees/search-ws/>.

Notices

Copyright © OASIS® 2007. All Rights Reserved.

All capitalized terms in the following text have the meanings assigned to them in the OASIS Intellectual Property Rights Policy (the "OASIS IPR Policy"). The full Policy may be found at the OASIS website.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published, and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this section are included on all such copies and derivative works. However, this document itself may not be modified in any way, including by removing the copyright notice or references to OASIS, except as needed for the purpose of developing any document or deliverable produced by an OASIS Technical Committee (in which case the rules applicable to copyrights, as set forth in the OASIS IPR Policy, must be followed) or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by OASIS or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and OASIS DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY OWNERSHIP RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

OASIS requests that any OASIS Party or any other party that believes it has patent claims that would necessarily be infringed by implementations of this OASIS Committee Specification or OASIS Standard, to notify OASIS TC Administrator and provide an indication of its willingness to grant patent licenses to such patent claims in a manner consistent with the IPR Mode of the OASIS Technical Committee that produced this specification.

OASIS invites any party to contact the OASIS TC Administrator if it is aware of a claim of ownership of any patent claims that would necessarily be infringed by implementations of this specification by a patent holder that is not willing to provide a license to such patent claims in a manner consistent with the IPR Mode of the OASIS Technical Committee that produced this specification. OASIS may include such claims on its website, but disclaims any obligation to do so.

OASIS takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on OASIS' procedures with respect to rights in any document or deliverable produced by an OASIS Technical Committee can be found on the OASIS website. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this OASIS Committee Specification or OASIS Standard, can be obtained from the OASIS TC Administrator. OASIS makes no representation that any information or list of intellectual property rights will at any time be complete, or that any claims in such list are, in fact, Essential Claims.

The names "OASIS", [insert specific trademarked names and abbreviations here] are trademarks of OASIS, the owner and developer of this specification, and should be used only to refer to the organization and its official outputs. OASIS welcomes reference to, and implementation and use of, specifications, while reserving the right to enforce its marks against misleading uses. Please see <http://www.oasis-open.org/who/trademark.php> for above guidance.

Table of Contents

1	Introduction.....	6
1.1	Terminology	6
1.2	Normative References	6
2	Overview and Model.....	7
2.1	Data model.....	7
2.2	Processing Model	7
2.3	Result Set Model	7
2.4	Diagnostic Model	8
3	Request Parameters.....	9
3.1	Request Parameters for this Binding.....	9
3.2	Relationship of Actual Parameters to Abstract Parameters	9
3.2.1	Abstract Request Parameters	9
3.2.2	Abstract/Actual Request Parameters	10
3.2.3	Abstract/Excluded Request Parameters	10
3.2.4	New Request Parameters	10
4	Response Elements	12
4.1.1	Actual Response Elements for this Binding	12
4.2	Relationship of Actual Elements to Abstract Elements	12
4.2.1	Abstract Response Elements.....	12
4.2.2	Abstract/Actual Response Elements.....	13
4.2.3	Abstract/Excluded Response Elements	13
4.2.4	New Response Elements	13
5	Parameter and Element Descriptions.....	15
5.1	maximumRecords.....	15
5.2	recordPacking.....	15
5.3	recordSchema	15
5.4	resultSetTTL and resultSetIdleTime	15
5.5	Stylesheet	16
5.6	records.....	16
5.7	diagnostics.....	16
5.8	extraRequestData, extraResponseData, and extraRecordData	16
5.9	echoedSearchRetrieveRequest.....	16
6	Response Schema for SRU 1.2	18
6.1	Structure of the <Record> Element.....	18
	Example.....	19
7	Diagnostics	20
7.1	Diagnostic List	20
7.2	Diagnostic Schema.....	20
7.3	Diagnostic Examples	20
7.3.1	Non-Surrogate Example.....	20
7.3.2	Surrogate Example.....	20
8	Extensions	22
A.	Acknowledgements	23

B. SRU/CQL 2.0 Features 24

 B.1 Background 24

 B.2 Proposed Changes for SRU 2.0 and CQL 2.0 24

 B.2.1 Allow Non-XML Record Representations 24

 B.2.2 Proximity as Boolean Modifier 24

 B.2.3 Proximity as Relation 25

 B.2.4 Parenthesized Terms 25

 B.2.5 Faceted Searching 25

 B.2.6 Result Set Size 25

 B.2.7 Multiple Query Types 25

 B.2.8 Eliminate the Version and Operation Parameters 26

1 Introduction

2 This is a binding of the OASIS SWS (Search Web Services) searchRetrieve operation – ABSTRACT
3 PROTOCOL DEFINITION.

4 **This binding is the specification of SRU 1.2.**

5 1.1 Terminology

6 The key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD
7 NOT”, “RECOMMENDED”, “NOT RECOMMENDED”, “MAY”, and “OPTIONAL” in this document are to be
8 interpreted as described in [RFC2119]. When these words are not capitalized in this document, they are
9 meant in their natural language sense.

10 1.2 Normative References

11 [RFC2119] S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*,
12 <http://www.ietf.org/rfc/rfc2119.txt>, IETF RFC 2119, March 1997.

2 Overview and Model

2.1 Data model

The data model in the Abstract Protocol Model says that a “*datastore* is a collection of units of data. Such a unit is referred to as an *abstract item...*”.

In this binding:

- A datastore is referred to as a database.
- An abstract item is referred to as an *abstract record*, or *record*.

The Abstract Protocol Model further notes that “Associated with a datastore are one or more formats that the server may apply to an abstract item, resulting in an exportable structure referred to as a *response Item*. . Such a format is referred to as a response item type or *item type...*”

In this Binding:

- A response item is referred to as a *response record*, or *record*.
Note that “abstract record” and “response record” are referred to as “record” when the meaning is clear from the context.
- An item type is referred to as a *record schema*.

A record schema is an XML schema. All records within a response, as well as the response itself, are transferred in XML. Data is not assumed to be stored in XML; records that are not natively XML must be first transformed into XML before being transferred.

2.2 Processing Model

A client sends a searchRetrieve request to a server. The client formulates the request according to a binding that describes a particular request construction. (An example of such a binding is: *Search Web Services (SWS) Auxiliary Binding for HTTP GET* which describes the construction of an http URL to encode parameter values of the form ‘*key=value*’.)

The request includes [request parameters](#) including a query to be matched against the database at the server. The server processes the query, creating a result set of records that match the query.

The request also indicates the desired number of records to be included in the response and includes the identifier of a record schema for transfer of the records in the response. The entire response (including all of the response records) is packaged in the SRU 1.2 [response schema](#).

The response includes records from the result set, diagnostic information, and a result set identifier that the client may use in a subsequent request to retrieve additional records.

2.3 Result Set Model

The result set model for SRU 1.2 is as described in the Abstract Protocol Definition, with the following additional feature:

When a result set record becomes unavailable, if a client then requests that record, the server is expected to supply a surrogate diagnostic in place of the record. For example suppose a result set of three records is created, subsequently the record at position 2 is deleted, and the client then requests records 1 through 3. The server should supply, in order: record 1, a surrogate diagnostic for record 2, record 3.

52 2.4 Diagnostic Model

53 A server supplies diagnostics in the response as appropriate. A diagnostics is *fatal* or *non-fatal*. A fatal
54 diagnostic is generated when the execution of the request cannot proceed and no entries are available.
55 For example, if the client supplied an invalid query there might be nothing that the server can do. A non-
56 fatal diagnostic is one where processing may be affected but the server can continue. For example if a
57 particular record is not available in the requested schema but others are, the server may return the ones
58 that are available rather than failing the entire request.

59 Non-fatal diagnostics are further divided into two categories: 'surrogate' and 'non-surrogate'. Surrogate
60 diagnostics take the place of a record (as described in the [Result Set Model](#)). Non-surrogate, non-fatal
61 diagnostics are diagnostics saying that while some or all the entries are available, something may have
62 gone wrong; for example the requested sorting algorithm might not be available. Or, it may be just a
63 warning.

64 To summarize: A surrogate diagnostic replaces a record; a non-surrogate diagnostic refers to the
65 response at large and is supplied in addition and external to the records. A non-surrogate diagnostic may
66 be fatal or non-fatal. So three combinations are possible:

- 67 • surrogate, non-fatal diagnostic
 - 68 • non-surrogate, non-fatal diagnostic
 - 69 • non-surrogate, fatal diagnostic
- 70 ("Fatal, surrogate" is not a valid combination.)

71 3 Request Parameters

72 3.1 Request Parameters for this Binding

73 The following table describes actual request parameters defined in this binding. It combines those that
74 are included from the Abstract Protocol Definition (table 3) and new request parameters (table 5).

75
76
77
78

79 . Table 1: Summary of Request Parameters

Actual Parameter Name	Occurence	Restrictions/Values/Description
operation	mandatory	The string: 'searchRetrieve'.
version	mandatory	The string '1.2'.
query	mandatory	A query expressed in CQL.
startRecord	optional	Positive integer. Default if omitted is 1.
maximumRecords	optional	See maximumRecords .
recordPacking	optional	'string' or 'xml'. Default is 'xml'. See recordPacking .
recordSchema	optional	See recordSchema .
resultSetTTL	optional	See resultSetTTL and resultSetIdleTime .
stylesheet	optional	See stylesheet .
extraRequestData	optional	See extraRequestData , extraResponseData , and extraRecordData

80 3.2 Relationship of Actual Parameters to Abstract Parameters

81 3.2.1 Abstract Request Parameters

82 The following table, copied from the Abstract Protocol Definition, summarizes the abstract parameters of
83 the SWS searchResponse request. Each parameter occurs either in table 3 or table 4.

84 Table 2: Abstract Request Parameters

Abstract Parameter Name	Description
responseType	e.g. 'text/html', 'application/atom+xml', 'application/x-sru'
query	The search query of the request.
startPosition	The position within the result set of the first record to be returned.

maximumItems	The number of items requested to be returned.
group	The number of the result group requested to be returned.
responseItemtype	e.g. 'string', 'jpeg', 'dc', 'iso2709'. From list provided by server.
sortOrder	The requested order of the result set.

85 3.2.2 Abstract/Actual Request Parameters

86 The following table lists those abstract request parameters that are realized as actual parameters in this
87 binding. See table 1 for their descriptions

88

89

90 *Table 3: Abstract/Actual Request Parameters*

Abstract Parameter Name	Actual Parameter Name
query	query
startPosition	startRecord
maximumItems	maximumRecords
responseItemtype	recordSchema

91 3.2.3 Abstract/Excluded Request Parameters

92 The following table summarizes the abstract request parameters that are excluded from this binding, and
93 why they are excluded.

94 *Table 4: Abstract/Excluded Request Parameters*

Abstract Parameter Name	Why Excluded
responseType	The response schema for SRU 1.2 is fixed.
group	request by group is not a feature of SRU 1.2.
sortOrder	The sort order is specified within the query.

95 3.2.4 New Request Parameters

96 The following table lists actual request parameters defined in this binding that are not defined as abstract
97 parameters. See table 1 for their descriptions.

98

99 *Table 5: New Request Parameters*

Actual Parameter Name
operation
version
recordPacking

resultSetTTL
stylesheet
extraRequestData

100

4 Response Elements

101

4.1.1 Actual Response Elements for this Binding

102

The following table describes the top-level XML elements in the response. These correspond to the actual response elements (from tables 8 and 9).

103

104 *Table 6: Summary of Actual Response Elements*

Actual Element Name	Type	Occurrence	Restrictions/Values/ Description
<version>	<i>xsd:string</i>	Mandatory, non-repeatable	The string '1.2'.
<numberOfRecords>	<i>xsd:integer</i>	Mandatory, non-repeatable	If the query fails this MUST be 0.
<resultSetId>	<i>xsd:string</i>	Optional, non-repeatable	
<resultSetIdleTime>	<i>xsd:integer</i>	Optional, non-repeatable	see resultSetTTL and resultSetIdleTime .
<record>	<record>	Optional, repeatable	See records
<nextRecordPosition>	<i>xsd:integer</i>	Optional, non-repeatable	If there are no remaining records, this field MUST be omitted.
<diagnostics>	<diagnostics> (non-surrogate)	Optional, non-repeatable	see diagnostics
<extraResponseData>	<i>structured</i>	Optional, non-repeatable	See extraRequestData , extraResponseData , and extraRecordData
<echoedSearch RetrieveRequest>	<i>structured</i>	Optional, non-repeatable	see echoedSearch RetrieveRequest

105

4.2 Relationship of Actual Elements to Abstract Elements

106

4.2.1 Abstract Response Elements

107

The following table, copied from the Abstract Protocol Definition, summarizes the abstract elements of the SWS searchResponse response.

108

109 *Table 7: Abstract Response Elements*

Abstract Element Name	Description/Reference
numberOfItems	The number of items matched by the query.
numberOfGroups	The number of result groups in the result set.

resultSetId	The identifier for a result set created through the execution of the query.
item	An individual response item (one of possibly many).
nextPosition	The next position within the result set following the final returned item.
nextGroup	The next result group following the group being returned.
diagnostics	Error message and/or diagnostics.
echoedRequest	The server may echo the request back to the client.

110 4.2.2 Abstract/Actual Response Elements

111 The following table lists those abstract response elements that are realized as actual elements in this
112 binding. See table 6 for their descriptions.

113 *Table 8: Abstract/Actual Response Elements*

Abstract Element Name	Actual Element Name
numberOfItems	numberOfRecords
resultSetId	resultSetId
item`	record
nextPosition	nextRecordPosition
diagnostics	diagnostics
echoedSearchRetrieveRequest	echoedSearchRetrieveRequest

114 4.2.3 Abstract/Excluded Response Elements

115 The following table summarizes the abstract response elements that are excluded from this binding, and
116 why they are excluded.

117 *Table 4: Abstract/Excluded Response Elements*

Abstract Element Name	Why Excluded
numberOfGroups	request by group is not a feature of SRU 1.2.
nextGroup	request by group is not a feature of SRU 1.2.

118 4.2.4 New Response Elements

119 The following table summarizes actual response elements defined in this binding that are not defined as
120 abstract elements. See table 9 for their descriptions.

121 *Table 9: New Response Elements*

Actual Element Name
version

resultSetIdleTime
extraResponseData

122 5 Parameter and Element Descriptions

123 5.1 maximumRecords

124 The request parameter maximumRecords is a non-negative integer. It may be omitted; if so the server
125 determines the default value. The server may return less than this number of records, for example if there
126 are fewer matching records than requested, but MUST NOT return more.

127 5.2 recordPacking

128 In order that records which are not well formed do not break the entire message, it is possible to request
129 that they be transferred as a single string with the <, > and & characters escaped to their entity forms.
130 Moreover some toolkits may not be able to distinguish record XML from the XML that forms the response.
131 However, some clients may prefer that the records be transferred as XML in order to manipulate them
132 directly with a stylesheet that renders the records and potentially also the user interface.

133 This distinction is made via the recordPacking parameter in the request. If the value of the parameter is
134 'string', then the server should escape the record before transferring it. If the value is 'xml', then it should
135 embed the XML directly into the response. Either way, the data is transferred within the 'recordData' field.
136 If the server cannot comply with this packing request, then it MUST return a diagnostic.

137 5.3 recordSchema

138 The requestParameter recordSchema is the XML schema of the records to be supplied in the response.
139 The value of the parameter is the short name that the server assigns to the identifier for the schema, as
140 listed in the server's explain file. The default value if not supplied is determined by the server.

141 For example, for the [MODS Schema Version 3.3](http://www.loc.gov/standards/sru/resources/schemas.html) the identifier is info:srw/schema/1/mods-v3.3, as
142 shown in the table at <http://www.loc.gov/standards/sru/resources/schemas.html> and the short name
143 might (but need not) be 'mods'. (Note: schema identifiers are not restricted to those in this table.)

144 The server MUST supply records in the requested schema only. If the schema is unknown or a record
145 cannot be rendered in that schema, then the server MUST return a diagnostic:

- 146 • If the schema is unknown, the server should supply a non-surrogate (fatal) diagnostic:
147 *info:srw/diagnostic/1/66*: "Unknown schema for retrieval".
- 148 • If an individual record cannot be rendered in the requested schema, the server should supply a
149 surrogate (non-fatal) diagnostic in place of the record: *info:srw/diagnostic/1/67*: "Record not
150 available in this schema".

151 5.4 resultSetTTL and resultSetIdleTime

152 The request parameter resultSetTTL is the time (in seconds) that the client requests that the result set
153 created should be maintained. The server MAY choose not to fulfill this request, and may respond with a
154 different value, via the response element resultSetIdleTime.

155 resultSetIdleTime is a good-faith estimate by the server of the idle time, in seconds. That is, the server
156 projects (but does not guarantee) that the result set will remain available and unchanged (both in content
157 and order) until there is a period of inactivity exceeding this idle time. The idle time must be a positive
158 integer, and should not be so small that a client cannot realistically reference the result set again. If the
159 server does not intend that the result set be referenced, it should omit the result set identifier in the
160 response.

161 The response element resultSetIdleTime may be less-than, equal-to, or greater-than the request
162 parameter resultSetTTL, and may be supplied or omitted regardless of whether resultSetTTL is supplied
163 or omitted. Thus the two (from a protocol point of view) are independent.

164 5.5 Stylesheet

165 The request parameter 'stylesheet' is a URL for a stylesheet. The client requests that the server simply
166 return this URL in the response, in the href attribute of the xml-stylesheet processing instruction before
167 the response xml. (It is likely that the type will be XSL, but not necessarily so.) If the server cannot fulfill
168 this request it must supply a [non-surrogate diagnostic](#) .

169 The purpose is to allow a thin client to turn the response XML into a natively renderable format, often
170 HTML or XHTML. This allows a web browser, or other application capable of rendering stylesheets, to act
171 as a dedicated client without requiring any further application logic.

172

173 **Example**

```
174 http://z3950.loc.gov:7090/voyager?version=1.2&operation=searchRetrieve  
175 &stylesheet=/master.xsl&query=dinosaur
```

176 This requests the server to include the following as beginning of the response:

```
177 <?xml version="1.0"?>  
178 <?xml-stylesheet type="text/xsl" href="/master.xsl"?>  
179 <sru:searchRetrieveResponse ...
```

180 5.6 records

181 The response element 'records' is a sequence of <record> elements. See [Structure of the <Record>](#)
182 [Element](#). Each contains either a record, or a surrogate diagnostic explaining why that record could not be
183 transferred. All records are transferred in XML. Records may be expressed as a single string, or as
184 embedded XML. If a record is transferred as embedded XML, it must be well formed and should be
185 validatable against the record schema.

186 5.7 diagnostics

187 The response element 'diagnostics' includes one or more non-surrogate diagnostics.

188 Note: See [Diagnostic Model](#). Non-surrogate diagnostics are distinguished from surrogate diagnostics.
189 The latter occur in the 'records' element of the response (they take the place of the record for which they
190 are a surrogate). Non-surrogate diagnostics, both fatal and non-fatal, occur in the 'diagnostics' element.

191 5.8 extraRequestData, extraResponseData, and extraRecordData

192 The request parameter 'extraRequestData'; response element <extraResponseData>, shown in
193 [Response Schema for SRU 1.2](#); and the <extraRecordData> element listed in [Structure of the Record](#)
194 [Element](#), are fields in which additional information may be provided. The mechanism is described in
195 [Extensions](#).

196 5.9 echoedSearchRetrieveRequest

197 The response element <echoedSearchRetrieveRequest> is as shown in the example below. Note the
198 two sub-elements <xQuery> and <baseUrl>.

199 <xQuery> represents an XCQL [\[reference\]](#) rendering of the query.

200 Note: This has two benefits.

- 201 • The client can use XSLT or other XML manipulation to modify the query without having a
202 CQL query parser.
- 203 • The server can return extra information specific to the clauses within the query.

204 <baseURL> allows the client to reconstruct queries by simple concatenation, or retrieve the explain
205 document to fetch additional information such as the title and description to include in the results
206 presented to the user.

207 Echoed Request Example

```
208 <echoedSearchRetrieveRequest>  
209   <version>1.2</version>  
210   <query>dc.title = dinosaur</query>  
211   <recordSchema>mods</recordSchema>  
212   <xQuery>  
213     <searchClause xmlns="http://www.loc.gov/zing/cql/xcql/">  
214       <index>dc.title</index>  
215       <relation>  
216         <value>=</value>  
217       </relation>  
218       <term>dinosaur</term>  
219     </searchClause>  
220   </xQuery>  
221   <baseUrl>http://z3950.loc.gov:7090/voyager</baseUrl>  
222 </echoedSearchRetrieveRequest>  
223  
224  
225
```

226

6 Response Schema for SRU 1.2

227

An example of a searchRetrieve response:

228

229

230

231

232

233

234

235

236

237

238

239

240

241

242

243

244

245

246

247

248

249

250

251

252

253

254

255

256

257

258

259

260

261

262

263

264

```

</searchRetrieveResponse>
  <numberOfRecords>10</numberOfRecords>
  <resultSetId>resultA</resultSetId>
  <resultSetIdleTime>180</resultSetIdleTime>
  <records>
    <record>
      record 1 ... See Example Record
    </record>
    <record>
      record 2 ....
    </record>
  </records>
  <nextRecordPosition>3</nextRecordPosition>
  <echoedSearchRetrieveRequest>
    ... see Echoed Request Example
  </echoedSearchRetrieveRequest>
  <diagnostics>
    <diagnostic>
      first non-surrogate diagnostic ... see Non Surrogate Diagnostic Example
    </diagnostic>
    <diagnostic>
      second non-surrogate diagnostic
    </diagnostic>
  </diagnostics>
  <extraResponseData>
    see Extension Example
  </extraResponseData>
</searchRetrieveResponse>

```

265

6.1 Structure of the <Record> Element

266

Each <record> element is structured into the elements shown in the following table.

267

268

269

270

Table 3: Structure of the <Record> Element

Element	Type	Occurrence	Description
<recordSchema>	<i>xsd:string</i>	mandatory	The URI identifier of the XML schema in which the record is encoded. Although the request may use the server's assigned short name, the response must always be the full URI.

<recordPacking>	<i>xsd:string</i>	mandatory	'string' or 'xml'.
<recordData>	<stringOrXmlFragment>	mandatory	The actual record.
<recordIdentifier>	<i>xsd:string</i>	optional	An identifier for the record by which it can unambiguously be retrieved in a subsequent operation. For example via the 'rec.identifier' index in CQL.
<recordPosition>	<i>xsd:positiveInteger</i>	optional	The position of the record within the result set.
<extraRecordData>	,<xmlFragment>	optional	Any additional information to be transferred with the record.

271 **Example**

272 An example record, in the simple Dublin Core schema, packed as XML:

```

273
274   <record>
275     <recordSchema>info:srw/schema/1/dc-v1.1</recordSchema>
276     <recordPacking>xml</recordPacking>
277     <recordData>
278       <srw_dc:dc xmlns:srw_dc="info:srw/schema/1/dc-v1.1">
279         <dc:title>This is a Sample Record</dc:title>
280       </srw_dc:dc>
281     </recordData>
282     <recordPosition>1</recordPosition>
283     <extraRecordData>
284       <rel:score xmlns:rel="info:srw/extensions/2/rel-1.0"> 0.965
285     </rel:score>
286     </extraRecordData>
287   </record>

```

288

289 7 Diagnostics

290 7.1 Diagnostic List

291 Diagnostics for use with SRU 1.2 are listed at [http://www.loc.gov/standards/sru/resources/diagnostics-](http://www.loc.gov/standards/sru/resources/diagnostics-list.html)
292 [list.html](http://www.loc.gov/standards/sru/resources/diagnostics-list.html). This diagnostic list has the namespace: info:srw/diagnostic/1. For example, the URI
293 info:srw/diagnostic/1/10 identifies the diagnostic “Query syntax error”.

294 Diagnostics used in SRU 1.2 need not be limited to this list, nor need this list be used exclusively for
295 SRU 1.2.

296 7.2 Diagnostic Schema

297 The diagnostic schema for SRU 1.2 has three elements, 'uri', 'details' and 'message'.

298 The required 'uri' field is a URI, identifying the particular diagnostic. The 'details' part contains information
299 specific to the diagnostic. The 'message' field contains a human readable message to be displayed. Only
300 the uri field is required, the other two are optional.

301 The identifier for the diagnostic schema is: info:srw/schema/1/diagnostics-v1.1

302 *Table 3: Elements of the Diagnostic Schema*

Element	Type	Occurrence	Description
<uri>	xsd:anyURI	Mandatory	The diagnostic's identifying URI.
<details>	xsd:string	Optional	Any supplementary information available, often in a format specified by the diagnostic
<message>	xsd:string	Optional	A human readable message to display to the end user. The language and style of this message is determined by the server, and clients should not rely on this text being appropriate for all situations.

303 7.3 Diagnostic Examples

304 7.3.1 Non-Surrogate Example

```
305 Non-surrogate, fatal diagnostic:  
306 <diagnostics>  
307   <diagnostic xmlns="http://www.loc.gov/zing/srw/diagnostic/">  
308     <uri>info:srw/diagnostic/1/38</uri>  
309     <details>10</details>  
310     <message>Too many boolean operators, the maximum is 10.  
311       Please try a less complex query.</message>  
312   </diagnostic>  
313 </diagnostics>
```

314 7.3.2 Surrogate Example

```
315 Surrogate, non-fatal diagnostic:  
316 <records>  
317   <record>  
318     <recordSchema> info:srw/schema/1/diagnostics-v1.1</recordSchema>  
319     <recordData>
```

```
320         <diagnostic xmlns="http://www.loc.gov/zing/srw/diagnostic/">
321             <uri>info:srw/diagnostic/1/65</uri>
322             <message>Record deleted by another user.</message>
323         </diagnostic>
324     </recordData>
325 </record> ...
326 </records>
```

327

8 Extensions

328 The name for an extension must begin with 'x-' (lower case x followed by hyphen. SRU 1.2 and future
329 versions will never include an official extension with a name beginning with 'x-', so this will never clash
330 with a mainstream extension name.) It is recommended that the extension name be 'x-' followed by an
331 identifier for the namespace for the extension, again followed by a hyphen, followed by the name of the
332 element within the namespace.

333

example

334

```
http://z3950.loc.gov:7090/voyager?...&x-info4-onSearchFail=scan
```

335 Note that this convention does not guarantee uniqueness since the extension name will not include a full
336 URI. The extension owner should try to make the name as unique as possible. If the namespace is
337 identified by an ['info:srw' URI](#), then the recommended convention is to name the extension "x-info/NNN-
338 XXX" where NNN is the 'info:srw' authority string, and XXX is the name of the extension. Extension
339 names MUST never be assigned with this form except by the proper authority for the given 'info'
340 namespace.

341

Response

342 The response may include `extraResponseData` with any well-formed XML, and hence servers can
343 include namespaced XML fragments within it in order to convey information back to the client. The
344 extension MUST supply a namespace and the element names with which to do this, if feedback to the
345 client is necessary.

346

example:

347

```
<sru:extraResponseData>  
  <auth:token xmlns:auth="info:srw/extension/2/auth-1.0">  
    277c6d19-3e5d-4f2d-9659-86a77fb2b7c8  
  </auth:token>  
</sru:extraResponseData>
```

348

349

350

351

352 **Semantics:** If the server does not understand a piece of information in an extension parameter, it may
353 silently ignore it. This is unlike many other request parameters, where if the server does not implement
354 that particular feature it MUST respond with a diagnostic. If the particular request requires some
355 confirmation that it has been carried out rather than ignored, then the extension designer should include a
356 field in the response. The semantics of parameters in the request may not be modified by extensions,
357 however the semantics of parts of the response may be modified by extensions. The response semantics
358 may be changed in this way only if the client specifically requests the change. Clients should in any case
359 be prepared to receive the regular semantics, as servers are at liberty to ignore extensions.

360 `ExtraResponseData` may be sent that is not directly associated with the request. For example it may
361 contain cost information regarding the query or information on the server or database supplying the
362 results. This data must, however, have been requested. As the request may be echoed, the server must
363 be able to transform the elements into their XML form. If it encounters an unrecognized element, the
364 server may either make its best guess as to how to transform the element, or simply not return it at all. It
365 should not, however, add an undefined namespace to the element as this would invalidate the response.
366 If the content of the element is an XML structure, then the extension designer should also specify how to
367 encode this structure in a URL. This may simply be to escape all of the special characters, but the
368 designer could also create a string encoding form with rules as to how to generate the XML in much the
369 same fashion as the relationship between CQL and XCQL.

370

371 **Acknowledgements**

372 The following individuals have participated in the creation of this specification and are gratefully
373 acknowledged:

374 **Participants:**

375 Kerry Blinco, Australian Department of Education, Employment and Workplace Relations
376 Ray Denenberg, Library of Congress
377 Larry Dixon, Library of Congress
378 Matthew Dovey, JISC
379 Janifer Gatenby, OCLC/PICS
380 Ralph LeVan, OCLC
381 Ashley Sanders, University of Manchester
382 Rob Sanderson, University of Liverpool

383 A. SRU/CQL 2.0 Features

384 Non-normative Annex

385 A.1 Background

386 This specification of the SRU 1.2 Binding is a straightforward rendering of and is equivalent to the existing
387 SRU 1.2 specification. It was developed as a proof of concept for the Abstract Protocol Definition (as is
388 the Open Search binding). The CQL Committee Draft is similarly a straightforward rendering of the CQL
389 1.2 specification.

390 Now the OASIS SWS Technical Committee turns its attention to substantive changes to SRU and CQL,
391 one of the reasons for the creation of the committee. This annex describes some of the changes that
392 have been contemplated. The list that follows is neither a complete list, nor have any of the proposals in
393 this list been approved at any level. The purpose of this list of proposals is to generate discussion on
394 these and other proposals.

395 A.2 Proposed Changes for SRU 2.0 and CQL 2.0

396 A.2.1 Allow Non-XML Record Representations

397 Although SRU does not require the database to store records as XML, it does require it to transfer the
398 data as XML. Many formats do not map easily into XML, for example multimedia, images and even
399 complex text formats. This proposal is to allow non-xml serialized data in the response, signaled by
400 additional values for the recordPacking parameter. Two possible values are:

- 401 • **By Value: recordPacking="base64"**
402 In order to include arbitrary data into an XML document, it can be base-64 encoded. To return
403 base-64 encoded data, the recordPacking value should be "base64", and the recordData element
404 should include the base-64 encoded data.
- 405 • **By Reference: recordPacking="uri"**
406 Instead of including the data directly into the response, the server could simply provide a link to it.
407 In this case, the recordPacking value should be "uri", and the recordData element should include
408 only the URI of the record.

409 A.2.2 Proximity as Boolean Modifier

410 The proposal is to deprecate the PROX BOOLEAN operator and instead have a BOOLEAN prox
411 modifier, with the same modifiers: 'unit', 'distance' and 'ordered'. The rationale for doing this is twofold.

- 412 • Proximity is not very frequently implemented other than as the special case of adjacency (or
413 phrase searching) so it does not warrant a keyword slot when it could instead be a modifier.
- 414 • More importantly, it is not currently possible to express negative proximity – the case when two
415 terms are not proximal.

416 Proximity would then look like:

- 417 • fish and/prox/distance=2/unit=word/ordered chips
418 (fish followed by chips 2 words away)
- 419 • fish not/prox/distance=2/unit=word/ordered chips
420 (fish not followed by chips 2 words away).

421 **A.2.3 Proximity as Relation**

422 Alternatively proximity can be represented as a searchClause: add a new relation to CQL called 'window'
423 allowing full proximity, rather than just adjacency as is currently possible. CQL.'window' would take the full
424 range of proximity modifiers: 'distance', 'unit' and 'ordered'. For example:

- 425 • `dc.title window/distance=2/unit=word/ordered "fish chips"`
426 (fish followed by chips 2 words away)

427 This proposal also allows for more than 2 terms to be part of a proximity query (however only as relates to
428 a single index):

- 429 • `dc.title window/distance<5/unit=word "fish chips salt"`
430 (fish, chips and salt all within a span of 5 words)
- 431 • `dc.title window/distance>5/unit=word "fish chips salt"`
432 (fish, chips and salt all with at least 5 words between any combination)

433 **A.2.4 Parenthesized Terms**

434 The above window relation does not allow choice. Window is the equivalent of "all" in that regard. For
435 example:

- 436 • `dc.title window/distance<5/unit=word ((fish and chips) and (salt or
437 vinegar))`
438 (fish and chips and one of salt or vinegar, in a 5 word window)

439 This is the only way this query can be expressed. And there are other applications for this style of query.

440 **A.2.5 Faceted Searching**

441 Faceted searching is commonly supported by search engines; for example, one might wish to do a
442 search on a database for books about a particular topic and then see how many records there are in
443 different time periods.

444 The proposed change is to add a new, optional response element for faceted results. In addition there
445 would be a request parameter to request faceted results.

446 **A.2.6 Result Set Size**

447 The proposal is to add a parameter allowing the client to indicate how much effort the server should take
448 to determine or estimate the number of records in the result set. Similarly, the response might include a
449 parameter indicating how accurate is the result-set-size reported.

450 The server may be able to determine the exact number of records, or provide a realistic estimate, but it
451 may be an expensive process. The server might prefer not go through that process unless the client
452 requests that it do so. Or the client might want to explicitly request that the server go through or not go
453 through that process.

454 For example, the client might want the first 10 records regardless of how many there are. In that case if
455 the server goes through the process of determining how many records there are, it may go through an
456 expensive process for nothing.

457
458 There is also the (special) case where the server cannot determine or estimate the number of records in
459 the result set. In that case it might be useful to have a special value or some way to indicate this
460 condition.

461 **A.2.7 Multiple Query Types**

462 The proposal is to support multiple query types. CQL would be one query type but there could be other
463 query types as well, for example, Parameterized Query and XQuery.

464 There are a number of possible ways to integrate support for multiple queries into the protocol.

- 465
- 466
- 467
- **Specify the query type via the parameter name.** For example "cql=" would be a cql query, "pq=...." would be a Parameterized Query, etc. The list of possible queries could be a controlled list in the standard, or the query types could be registered.
 - **A 'queryType' parameter.** The actual query would be contained within a fixed parameter, for example "query". E.g. "queryType=cql&query=.....".
 - **Via the Explain record.** Let Explain specify the name of the parameter used for each query type. So one server could say that that a CQL query is specified with the "cql" parameter and another server could say that a CQL query is specified with the "query" parameter.
- 468
- 469
- 470
- 471
- 472

473 **A.2.8 Eliminate the Version and Operation Parameters**

474 The version parameter in SRU is based on the assumption that the same base URL is to be used for
475 multiple versions of the protocol, and SRU prescribes a mechanism to allow different versions to attempt
476 to interoperate. Instead, different base URLs could be used for different versions, base URLs and
477 corresponding versions exposed via explain, and then a client would know the version of a server, so
478 there would never be a mismatch.

479 Similarly, the operation parameter could be eliminated on the same basis.

480