

Open Geospatial Consortium Inc.

Date: 2007-02-23

Reference number of this document: OGC 07-006r1

Version 2.0.2, Corrigendum 2 Release

Category: OpenGIS® Implementation Specification

Editors: Douglas Nebert, Arliss Whiteside, Panagiotis (Peter) Vretanos

OpenGIS® Catalogue Services Specification

Copyright © 2007 Open Geospatial Consortium, Inc. All Rights Reserved.
To obtain additional rights of use, visit <http://www.opengeospatial.org/legal/>.

Recipients of this document are invited to submit, with their comments, notification of any relevant patent rights of which they are aware and to provide supporting documentation.

Document type:	OpenGIS® Publicly Available Standard
Document subtype:	Implementation Specification
Document stage:	Proposed Revision
Document language:	English

License Agreement

Permission is hereby granted by the Open Geospatial Consortium, ("Licensor"), free of charge and subject to the terms set forth below, to any person obtaining a copy of this Intellectual Property and any associated documentation, to deal in the Intellectual Property without restriction (except as set forth below), including without limitation the rights to implement, use, copy, modify, merge, publish, distribute, and/or sublicense copies of the Intellectual Property, and to permit persons to whom the Intellectual Property is furnished to do so, provided that all copyright notices on the intellectual property are retained intact and that each person to whom the Intellectual Property is furnished agrees to the terms of this Agreement.

If you modify the Intellectual Property, all copies of the modified Intellectual Property must include, in addition to the above copyright notice, a notice that the Intellectual Property includes modifications that have not been approved or adopted by LICENSOR.

THIS LICENSE IS A COPYRIGHT LICENSE ONLY, AND DOES NOT CONVEY ANY RIGHTS UNDER ANY PATENTS THAT MAY BE IN FORCE ANYWHERE IN THE WORLD.

THE INTELLECTUAL PROPERTY IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT OF THIRD PARTY RIGHTS. THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE DO NOT WARRANT THAT THE FUNCTIONS CONTAINED IN THE INTELLECTUAL PROPERTY WILL MEET YOUR REQUIREMENTS OR THAT THE OPERATION OF THE INTELLECTUAL PROPERTY WILL BE UNINTERRUPTED OR ERROR FREE. ANY USE OF THE INTELLECTUAL PROPERTY SHALL BE MADE ENTIRELY AT THE USER'S OWN RISK. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR ANY CONTRIBUTOR OF INTELLECTUAL PROPERTY RIGHTS TO THE INTELLECTUAL PROPERTY BE LIABLE FOR ANY CLAIM, OR ANY DIRECT, SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM ANY ALLEGED INFRINGEMENT OR ANY LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR UNDER ANY OTHER LEGAL THEORY, ARISING OUT OF OR IN CONNECTION WITH THE IMPLEMENTATION, USE, COMMERCIALIZATION OR PERFORMANCE OF THIS INTELLECTUAL PROPERTY.

This license is effective until terminated. You may terminate it at any time by destroying the Intellectual Property together with all copies in any form. The license will also terminate if you fail to comply with any term or condition of this Agreement. Except as provided in the following sentence, no such termination of this license shall require the termination of any third party end-user sublicense to the Intellectual Property which is in force as of the date of notice of such termination. In addition, should the Intellectual Property, or the operation of the Intellectual Property, infringe, or in LICENSOR's sole opinion be likely to infringe, any patent, copyright, trademark or other right of a third party, you agree that LICENSOR, in its sole discretion, may terminate this license without any compensation or liability to you, your licensees or any other party. You agree upon termination of any kind to destroy or cause to be destroyed the Intellectual Property together with all copies in any form, whether held by you or by any third party.

Except as contained in this notice, the name of LICENSOR or of any other holder of a copyright in all or part of the Intellectual Property shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Intellectual Property without prior written authorization of LICENSOR or such copyright holder. LICENSOR is and shall at all times be the sole entity that may authorize you or any third party to use certification marks, trademarks or other special designations to indicate compliance with any LICENSOR standards or specifications.

This Agreement is governed by the laws of the Commonwealth of Massachusetts. The application to this Agreement of the United Nations Convention on Contracts for the International Sale of Goods is hereby expressly excluded. In the event any provision of this Agreement shall be deemed unenforceable, void or invalid, such provision shall be modified so as to make it valid and enforceable, and as so modified the entire Agreement shall remain in full force and effect. No decision, action or inaction by LICENSOR shall be construed to be a waiver of any rights or remedies available to it.

None of the Intellectual Property or underlying information or technology may be downloaded or otherwise exported or reexported in violation of U.S. export laws and regulations. In addition, you are responsible for complying with any local laws in your jurisdiction which may impact your right to import, export or use the Intellectual Property, and you represent that you have complied with any regulations or registration procedures required by applicable law to make this license enforceable

Contents	Page
i. Preface.....	x
ii. Submitting organizations	x
iii. Document contributor contact points.....	xi
iv. Revision history	xi
v. Changes to the OGC Abstract Specification.....	xii
vi. Future work.....	xii
Foreword.....	xiii
Introduction.....	xiv
1 Scope.....	1
2 Conformance.....	1
3 Normative references.....	1
4 Terms and definitions	3
5 Conventions	7
5.1 Symbols (and abbreviated terms).....	7
5.2 UML notation.....	7
5.3 Document terms and definitions.....	9
6 Catalogue abstract information model.....	9
6.1 Introduction.....	9
6.2 Query language support	10
6.2.1 Introduction.....	10
6.2.2 OGC_Common catalogue query language	10
6.2.3 Extending the Common Catalogue Query Language	16
6.2.4 Query language realization	18
6.3 Core catalogue schema	18
6.3.1 Introduction.....	18
6.3.2 Core queryable properties.....	19
6.3.3 Core returnable properties.....	21
6.3.4 Information structure and semantics.....	24
7 General catalogue interface model.....	25
7.1 Introduction.....	25
7.2 Interface definitions	26
7.2.1 Overview.....	26
7.2.2 Catalogue Service class.....	29
7.2.3 OGC_Service class	29
7.2.4 Discovery class	31
7.2.5 Session class.....	44
7.2.6 Manager class.....	51

7.2.7	Brokered Access class.....	56
7.3	Protocol, interface and operation specializations.....	58
7.4	Dynamic model.....	58
7.4.1	Introduction.....	58
7.4.2	UML state diagram notation.....	59
7.4.3	Catalogue server state machine.....	59
7.4.4	Discovery state.....	60
7.4.5	Access state diagram.....	62
7.4.6	Management state.....	65
7.4.7	Explain state diagram.....	65
8	Z39.50 protocol binding.....	66
8.1	Architecture.....	66
8.1.1	Introduction.....	66
8.1.2	Supported services.....	67
8.1.3	Core queryable elements.....	67
8.2	General model to Z39.50 protocol binding operations mapping.....	67
8.3	Z39.50 BER implementation notes.....	69
8.3.1	Message encoding.....	70
8.3.2	Additional search info.....	70
8.3.3	Order extended service.....	71
8.4	Search/Retrieve Web Service (SRW/SRU) implementation notes.....	83
9	CORBA/IIOP protocol binding.....	85
9.1	Architecture.....	85
9.1.1	Supported services.....	85
9.1.2	Core queryable elements.....	85
9.2	Content types (Catalogue entry types).....	85
9.3	Supported query languages.....	85
9.4	Result set encodings.....	86
9.4.1	XML.....	86
9.4.2	Name-Value pairs.....	86
9.5	General model to CORBA protocol binding operations mapping.....	87
9.6	Interface definition - IDL.....	88
9.6.1	Introduction.....	88
9.6.2	Enumerations.....	89
9.6.3	Structures and unions.....	89
9.6.4	Definitions for brokered access.....	92
9.6.5	Capabilities.....	93
9.6.6	General messages.....	94
9.6.7	Discovery messages.....	96
9.6.8	Management messages.....	98
9.6.9	Access messages.....	99
9.6.10	Exceptions.....	100
9.6.11	Catalogue Service interfaces.....	101
9.6.12	Basic interfaces.....	102
9.6.13	Complete IDL.....	103
10	HTTP protocol binding (Catalogue Services for the Web, CSW).....	115
10.1	Architectural principles.....	115

10.2	The HTTP protocol	115
10.2.1	Overview	115
10.2.2	Message headers	116
10.2.3	Content encoding	117
10.2.4	Request methods	117
10.2.5	Message payload	118
10.3	Operation request and response encoding	123
10.3.1	Introduction	123
10.3.2	SOAP	123
10.3.3	Namespaces	125
10.3.4	Predicate languages	126
10.3.5	General model message mapping	127
10.3.6	Common request parameters	128
10.3.7	Exception reporting	129
10.4	Operations overview	129
10.5	GetCapabilities operation	131
10.5.1	Introduction	131
10.5.2	Operation request	131
10.5.3	Operation response	132
10.5.4	OperationsMetadata section standard contents	132
10.5.5	Examples	134
10.6	DescribeRecord operation	134
10.6.1	Introduction	134
10.6.2	KVP encoding	135
10.6.3	XML encoding	135
10.6.4	Parameter descriptions	136
10.6.5	Response	137
10.6.6	Examples	138
10.7	GetDomain operation	138
10.7.1	Introduction	138
10.7.2	KVP encoding	138
10.7.3	XML encoding	139
10.7.4	Parameter descriptions	140
10.7.5	Response	140
10.7.6	Examples	141
10.8	GetRecords operation	142
10.8.1	Introduction	142
10.8.2	KVP encoding	142
10.8.3	XML encoding	145
10.8.4	Parameter descriptions	147
10.8.5	Response	157
10.8.6	Examples	159
10.9	GetRecordById operation	160
10.9.1	Introduction	160
10.9.2	KVP encoding	160
10.9.3	XML encoding	161
10.9.4	Parameter descriptions	162
10.9.5	Response	163

10.9.6	Examples.....	164
10.10	Record locking.....	165
10.11	Transaction operation.....	165
10.11.1	Introduction.....	165
10.11.2	KVP encoding.....	165
10.11.3	XML encoding.....	166
10.11.4	Response.....	168
10.12	Harvest operation.....	169
10.12.1	Introduction.....	169
10.12.2	KVP encoding.....	170
10.12.3	XML encoding.....	172
10.12.4	Parameter descriptions.....	172
10.12.5	Response.....	174
10.12.6	Examples.....	175
10.13	XML Schemas.....	175
11	Specializing general model through protocol bindings and profiles.....	176
11.1	Introduction.....	176
11.2	Interface definitions.....	177
11.3	Query model components.....	178
11.3.1	Query language/model.....	178
11.3.2	Common search and retrieval elements.....	178
11.4	Catalogue Application Profiles.....	179
11.5	Structure and format.....	181
11.5.1	Introduction.....	181
11.5.2	System context.....	182
11.5.3	Information models.....	182
11.5.4	External interfaces.....	183
11.6	Compliance.....	184
Annex A (normative)	Abstract conformance test suite.....	185
Annex B (informative)	Description of Distributed Search.....	186
Annex C (informative)	Details of order operation.....	189
C.1	Introduction.....	189
C.2	BrokeredAccessRequest.....	189
C.3	BrokeredAccessResponse.....	190
C.4	Parameter Type Definitions.....	192
Annex D (informative)	Sample CSW capabilities document.....	198
Annex E (Normative)	Minimum XPath subset.....	203
Bibliography	204

Figures	Page
Figure 1 — UML notations.....	8
Figure 2 — Information discovery continuum	9
Figure 3 — Reference model architecture	26
Figure 4 — General OGC catalogue UML static model	27
Figure 5 — getCapabilities operation UML static model.....	30
Figure 6 — “query” operation UML static model	33
Figure 7 — “present” operation UML static model.....	38
Figure 8 — describeRecordType operation UML static model.....	41
Figure 9 — getDomain operation UML static model	43
Figure 10 — “initialize” operation UML static model	45
Figure 11 — “close” operation UML static model.....	46
Figure 12 — “status” operation UML static model	48
Figure 13 — “cancel” operation UML static model.....	50
Figure 14 — “transaction” operation UML static model.....	52
Figure 15 — harvestResource operation UML static model	54
Figure 16 — “order” operation UML static model.....	57
Figure 17 — UML state diagram notation.....	59
Figure 18 — Catalogue session state diagram	60
Figure 19 — Discovery state diagram (without Status and Cancel).....	61
Figure 20 — Discovery state diagram (complete).....	62
Figure 21 — Access state diagram	63
Figure 22 — Order estimation state diagram.....	64
Figure 23 — Order submit state diagram	64
Figure 24 — Management state diagram.....	65
Figure 25 — Explain state diagram	66
Figure 26 — Z39.50 Protocol binding Sequence Diagram.....	69
Figure 27 — Name-value pair results	87
Figure 28 — Catalogue service web.....	115
Figure 29 — Protocol sequence diagram.....	130
Figure 30 — Conceptual architecture	131

Figure 31 — Relationship of general model, protocol binding, and application profile	177
Figure 32 — Application profiles specify concrete catalogue services	179
Figure 33 — Relationships between base specs, profiles, and implementations	180
Figure B.1 — Query network topology resulting in duplicates	187
Figure B.2 — Query network topology resulting in a loop	187

Tables	Page
Table 1 — Common queryable elements	20
Table 2 — Composition of compound element “BoundingBox”	20
Table 3 — Composition of compound element “Association”	21
Table 4 — List of common returnable properties	22
Table 5 — Attribute of Catalogue Service class	29
Table 6 — Definition of getCapabilities operation	29
Table 7 — UML attribute in getCapabilities operation request	30
Table 8 — UML attributes in getCapabilities operation normal response	31
Table 9 — Definition of “query” operation	32
Table 10 — UML attributes and roles in “query” operation request	34
Table 11 — UML attributes in SortSpec data type	36
Table 12 — UML attributes in SessionInfo data type	36
Table 13 — UML attributes in QueryExpression data type	36
Table 14 — UML attributes and role “query” operation normal response	37
Table 15 — Definition of “present” operation	38
Table 16 — UML attributes and role in “present” operation request	39
Table 17 — UML attributes and role in “present” operation normal response	40
Table 18 — Definition of describeRecordType operation	40
Table 19 — UML attributes and role in describeRecordType operation request	41
Table 20 — UML attributes and role in describeRecordType operation normal response	42
Table 21 — Definition of getValueDomain operation	42
Table 22 — UML attribute in getDomain operation request	43
Table 23 — UML attributes and role in getValueDomain operation normal response	43
Table 24 — Definition of “initiaize” operation	44
Table 25 — UML role name in “initialize” operation request	45

Table 26 — UML role name in “initiaize” operation normal response.....	45
Table 27 — Definition of “close” operation.....	46
Table 28 — UML role name in “close” operation request.....	46
Table 29 — UML role name in “close” operation normal response	47
Table 30 — Definition of “status” operation.....	47
Table 31 — UML attribute and role in “status” operation request.....	48
Table 32 — UML attribute and role in “status” operation normal response	49
Table 33 — Definition of “cancel” operation.....	49
Table 34 — UML attributes and role in “cancel” operation request	50
Table 35 — UML attributes and role in “cancel” operation normal response	51
Table 36 — Definition of “transaction” operation	52
Table 37 — UML attributes in “transaction” operation request.....	53
Table 38 — UML attributes in “transaction” operation normal response	53
Table 39 — harvestResource operation.....	54
Table 40 — UML attributes in harvestResource operation request	55
Table 41 — UML attributes in harvestResource operation normal response.....	55
Table 42 — Definition of “order” operation.....	56
Table 43 — UML attributes in “order” operation request.....	57
Table 44 — UML attributes in “order” operation normal response	58
Table 45 — Correspondence of Z39.50 Attributes to general model equivalents.....	67
Table 46 — General Model to Z39.50 protocol binding operations mapping and obligation.....	68
Table 47 — Order extended service	72
Table 48 — SearchRetrieve request parameters.....	84
Table 49 — SearchRetrieve response parameters	84
Table 50 — General Model to CORBA protocol binding operations mapping and obligation.....	88
Table 51 — HTTP message elements.....	116
Table 52 — Selected HTTP Request Methods.....	117
Table 53 — Mapping of Dublin Core names to XML element names.....	119
Table 54 — HTTP method bindings.....	123
Table 55 — KVP encoding for constraints.....	127
Table 56 — General model to CSW mapping.....	128
Table 57 — KVP encoding of common operation request parameters	128

Table 58 — Additional section name value and meaning	132
Table 59 — Section names and contents	132
Table 60 — Required values of the OperationsMetadata section attributes.....	133
Table 61 — Optional values of the OperationsMetadata section attributes	133
Table 62 — KVP encoding for DescribeRecord operation request.....	135
Table 63 — KVP encoding for GetDomain operation request.....	139
Table 64 — Interface parameters that may be interrogated using GetDomain operation	140
Table 65 — KVP encoding for GetRecords operation request.....	143
Table 66 — <SearchResults> Parameters	159
Table 67 — KVP encoding for GetRecordById operation request	161
Table 68 — KVP encoding for Harvest operation request.....	171
Table 69 — URIs for well known metadata standards	173
Table 70 — Structure of an application profile	181
Table 71 — System context: required subclauses	182
Table 72 — Information models: required subclauses	183
Table 73 — Public interfaces: required subclauses	183
Table C.1 — Brokered Access Request Types.....	193
Table C.2 — Order Status Codes.....	194

i. Preface

This document explains how Catalogue Services version 2.0 are organised and implemented for the discovery and retrieval of spatial data and services metadata. The prior public version of this specification was 1.1.1. Catalogue Services version 2.0 supersedes and deprecates version 1.1.1.

This revision of this document has been significantly improved, largely based on the comments on document 04-021 received by the Revision Working Group. The changes made in this second corrigendum are summarized in the Corrigendum description document [OGC 07-010].

ii. Submitting organizations

The following organizations submitted the original document or its revisions to the Open GIS Consortium, Inc. in response to the OGC Request 6, Core Task Force, Catalogue Working Group, A Request for Proposals: OpenGIS[®] Catalogue Interface (OpenGIS[®] Project Document Number 98-001r2):

BAE SYSTEMS Mission Solutions (formerly Marconi Integrated Systems, Inc.)

Blue Angel Technologies, Inc.

Environmental Systems Research Institute (ESRI)

Geomatics Canada (Canada Centre for Remote Sensing (CCRS))

Intergraph Corporation

MITRE

Oracle Corporation

U.S. Federal Geographic Data Committee (FGDC)

U.S. National Aeronautics and Space Administration (NASA)

U.S. National Imagery and Mapping Agency (NIMA)

Contributing Entities

The submitting entities were grateful for the contributions from the following companies in the development and revision of this Interface Specification:

Compusult, Limited

Con terra GmbH

Cubewerx Inc.

Galdos Systems, Inc

GEODAN IT bv
 Hammon, Jensen, Wallen & Associates, Inc (HJW)
 Ionic Software, sa
 JRC (Joint Research Centre), European Commission
 SICAD GEOMATICS
 Traverse Technologies

iii. Document contributor contact points

All questions regarding this document should be directed to the editor or the contributors:

Name	Organization
Doug Nebert	U.S. Federal Geographic Data Committee
Arliss Whiteside	BAE SYSTEMS Mission Solutions
Panagiotis (Peter) Vretanos	CubeWerx, Inc.
Louis Reich	NASA (Computer Sciences Corp)
Richard Martell	Galdos Systems Inc.
Uwe Voges	con terra GmbH

iv. Revision history

Date	Release	Editor	Primary clauses modified	Description
12Aug1999	1.0	Nebert	N/A	Original Specification entitled "Catalogue Interface Implementation Specification" OGC Document 00-034
28Mar2001	1.1	Nebert	Made fine-grain CORBA and OLE/COM Annexes to Informative, added abstract conformance test suite, fixed coarse-grain CORBA IDL	Document only made available to OGC membership pending passage of Version 2.0. (OGC Document 01-040)
11Nov2002	1.1.1	Nebert, Katz,	State diagram changes, renamed specification and changed WWW Profile to Z39.50 Profile, added introductory words as required for new format	Document primarily reflects conversion to newer OGC/ISO document format
22Dec2003	2.0.0	Nebert	All	Reorganised and largely rewrote document.
6Mar2004	2.0.0	Nebert	Clauses 6,7, 9,10,11	Edited CORBA, Z39.50, and HTTP to reflect changes in General Model, other minor revisions to document
29Mar2004	2.0.0	Whiteside	All	

14Apr2004	2.0.0	Whiteside	All	
11May2004	2.0.0	Nebert	Merge updates on Clauses 1-5, 6-8, 9, 10, and 11 from separate editors	Responded to all RWG review comments.

v. **Changes to the OGC Abstract Specification**

The OpenGIS[®] Abstract Specification requires minor changes to accommodate the technical contents of this document to reflect the use of catalogue services to search for and retrieve any type of information object (data, service instance, service type, schema, style description, etc.) based on its properties described in “metadata.”

vi. **Future work**

Future work may include expansion of definitions of ordering capabilities, standing orders or queries, and transactional interfaces. Individual Application Profiles (see Clause 11) will be submitted as separate but dependent specifications.

Foreword

This version deprecates and replaces Catalogue Service Specification 1.1.1 (OGC Document 02-087r3).

This document, through its implementation profiles, references several external standards and specifications as dependencies:

- a) Common Object Request Broker Architecture (CORBA/IIOP), Version 2.X, The Object Management Group (OMG): <http://www.omg.org>
- b) Information and documentation -- Information retrieval (Z39.50) -- Application service definition and protocol specification:
<http://www.iso.ch/iso/en/CatalogDetailPage.CatalogDetail?CSNUMBER=27446&ICS1=35&ICS2=240&ICS3=30>
- c) ISO/IEC TR 10000-1:1998. *Information Technology – Framework and taxonomy of International Standardised Profiles – Part 1: General principles and documentation framework*. Technical Report, JTC 1. Fourth edition. Available [online]:
<[http://standards.iso.org/ittf/PubliclyAvailableStandards/c030726_ISO_IEC_TR_10000-1_1998\(E\).zip](http://standards.iso.org/ittf/PubliclyAvailableStandards/c030726_ISO_IEC_TR_10000-1_1998(E).zip)>.
- d) ISO/IEC 10746-2:1996. *Information Technology – Open Distributed Processing – Reference Model: Foundations*. Common text with ITU-T Recommendation X.902. Available [online]:
<[http://standards.iso.org/ittf/PubliclyAvailableStandards/s018836_ISO_IEC_10746-2_1996\(E\).zip](http://standards.iso.org/ittf/PubliclyAvailableStandards/s018836_ISO_IEC_10746-2_1996(E).zip)>.
- e) Unified Modeling Language (UML) Version 1.3, The Object Management Group (OMG): <http://www.omg.org/cgi-bin/doc?formal/00-03-01>
- f) The Extensible Markup Language (XML), World Wide Web Consortium, <http://www.w3.org/TR/1998/REC-xml-19980210>

Annex A, the Abstract Conformance Test Suite, is normative to this specification and shall be implemented when a computing environment requires catalogue services. All other annexes are informative and provide background information, such as terminology and alternative implementation approaches.

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. The Open GIS Consortium, Inc. shall not be held responsible for identifying any or all such patent rights.

Introduction

This document specifies the interfaces between clients and catalogue services, through the presentation of abstract and implementation-specific models. Catalogue services support the ability to publish and search collections of descriptive information (metadata) for data, services, and related information objects. Metadata in catalogues represent resource characteristics that can be queried and presented for evaluation and further processing by both humans and software. Catalogue services are required to support the discovery and binding to registered information resources within an information community.

OpenGIS® Catalogue Services Specification

1 Scope

This OGC™ document specifies the interfaces, bindings, and a framework for defining application profiles required to publish and access digital catalogues of metadata for geospatial data, services, and related resource information. Metadata act as generalised properties that can be queried and returned through catalogue services for resource evaluation and, in many cases, invocation or retrieval of the referenced resource. Catalogue services support the use of one of several identified query languages to find and return results using well-known content models (metadata schemas) and encodings. This OpenGIS® document is applicable to the implementation of interfaces on catalogues of a variety of information resources.

The target audience for this specification is the community of software developers who are:

- a) Implementers of OGC compliant Catalogue servers
- b) Implementers of OGC compliant Catalogue clients

2 Conformance

Abstract conformance to the mandatory catalogue service interfaces is described in Annex A. In a given community, a test suite should include test metadata records with a variety of element values and a series of queries that would return correct and properly formatted results. Test data and queries may be included in associated Application Profiles.

3 Normative references

The following normative documents contain provisions that, through reference in this text, constitute provisions of this part of OGC 03-108. For dated references, subsequent amendments to, or revisions of, any of these publications do not apply. For undated references, the latest edition of the normative document referred to applies.

ANSI/NISO Z39.50-2003, *Information Retrieval (Z39.50): Application Service Definition and Protocol Specification* (revision of Z39.50-1995)
<http://www.iso.ch/iso/en/CatalogDetailPage.CatalogDetail?CSNUMBER=27446&ICS1=35&ICS2=240&ICS3=30>

IETF RFC 2045 (November 1996), *Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies*, Freed, N. and Borenstein N., eds., <http://www.ietf.org/rfc/rfc2045.txt>

IETF RFC 2141 (May 1997), *URN Syntax*, R. Moats, <http://www.ietf.org/rfc/rfc2141.txt>

IETF RFC 2396 (August 1998), *Uniform Resource Identifiers (URI): Generic Syntax*, Berners-Lee, T., Fielding, N., and Masinter, L., eds., <http://www.ietf.org/rfc/rfc2396.txt>

IETF RFC 2616 (June 1999), *Hypertext Transfer Protocol – HTTP/1.1*, Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and Berners-Lee, T., eds., <http://www.ietf.org/rfc/rfc2616.txt>

IANA, Internet Assigned Numbers Authority, *MIME Media Types*, available at <http://www.iana.org/assignments/media-types/>

CORBA/IIOP, *Common Object Request Broker Architecture*, Version 2.X, The Object Management Group (OMG): <http://www.omg.org>

ISO/IEC TR 10000-1:1998. *Information Technology – Framework and taxonomy of International Standardised Profiles – Part 1: General principles and documentation framework*. Technical Report, JTC 1. Fourth edition, Available [online]: [http://standards.iso.org/ittf/PubliclyAvailableStandards/c030726_ISO_IEC_TR_10000-1_1998\(E\).zip](http://standards.iso.org/ittf/PubliclyAvailableStandards/c030726_ISO_IEC_TR_10000-1_1998(E).zip).

ISO/IEC 10746-2:1996. *Information Technology – Open Distributed Processing – Reference Model: Foundations*. Common text with ITU-T Recommendation X.902, Available [online]: [http://standards.iso.org/ittf/PubliclyAvailableStandards/s018836_ISO_IEC_10746-2_1996\(E\).zip](http://standards.iso.org/ittf/PubliclyAvailableStandards/s018836_ISO_IEC_10746-2_1996(E).zip).

ISO 4217:2001, *Codes for the representation of currencies and funds*

ISO 8601:2000(E), *Data elements and interchange formats - Information interchange - Representation of dates and times*

ISO 11180, *Postal addressing*

ISO/IEC 14977:1996, *Information technology – Syntactic metalanguage – BNF*

ISO 19106:2003, *Geographic Information – Profiles*

ISO 19115:2003, *Geographic Information – Metadata*

ISO/DIS 19119, *Geographic Information – Services*

ISO/TS 19139, *Geographic Information – Metadata -Implementation Specification*

OASIS/ebXML Registry Services Specification v3.0

OGC 99-054, *OpenGIS® Simple Features Specification for CORBA*

OGC 99-113, *OGC Abstract Specification Topic 13: Catalogue Services*

OGC 02-112, *OGC Abstract Specification Topic 12: OpenGIS Service Architecture*

OGC 04-095, *Filter Encoding Implementation Specification*

OGC 05-008c1, *OWS Common Implementation Specification*

NOTE This OWS Common Implementation Specification contains a list of normative references that are also applicable to this Implementation Specification.

OMG UML, *Unified Modeling Language, Version 1.3*, The Object Management Group (OMG): <http://www.omg.org/cgi-bin/doc?formal/00-03-01>

W3C Recommendation January 1999, *Namespaces In XML*,
<http://www.w3.org/TR/2000/REC-xml-names>

W3C Recommendation 6 October 2000, *Extensible Markup Language (XML) 1.0* (Second Edition), <http://www.w3.org/TR/REC-xml>

W3C Recommendation 2 May 2001: *XML Schema Part 0: Primer*,
<http://www.w3.org/TR/2001/REC-xmlschema-0-20010502/>

W3C Recommendation 2 May 2001: *XML Schema Part 1: Structures*,
<http://www.w3.org/TR/2001/REC-xmlschema-1-20010502/>

W3C Recommendation 2 May 2001: *XML Schema Part 2: Datatypes*,
<http://www.w3.org/TR/2001/REC-xmlschema-2-20010502/>

W3C Recommendation (24 June 2003): *SOAP Version 1.2 Part 1: Messaging Framework*, <http://www.w3.org/TR/SOAP/>

W3C Recommendation (16 November 1999): *XML Path Language (XPath) Version 1.0*,
<http://www.w3.org/TR/xpath.html>

In addition to this document, this specification includes several normative XML Schema files. These are posted online at the URL <http://schemas.opengis.net/>. These XML Schema files are also bundled with this document. In the event of a discrepancy between the bundled and online versions of the XML Schema files, the online files shall be considered authoritative.

4 Terms and definitions

For the purposes of this document, the following terms and definitions apply:

4.1

client

software component that can invoke an **operation** from a **server**

4.2

data clearinghouse

collection of institutions providing digital data, which can be searched through a single interface using a common metadata standard [ISO 19115]

4.3

data level

stratum within a set of layered levels in which data is recorded that conforms to definitions of types found at the application model level [ISO 19101]

4.4

dataset series

collection of datasets sharing the same product specification [ISO 19113, ISO 19114, ISO 19115]

4.5

feature catalogue

catalogue containing definitions and descriptions of the feature types, feature attributes, and feature relationships occurring in one or more sets of geographic data, together with any feature operations that may be applied [ISO 19101, ISO 19110]

4.6

geographic dataset

dataset with a spatial aspect [ISO 19115]

4.7

geographic information

information concerning phenomena implicitly or explicitly associated with a location relative to the Earth [ISO 19128 draft]

4.8

identifier

a character string that may be composed of numbers and characters that is exchanged between the client and the server with respect to a specific identity of a resource

4.9

interface

named set of operations that characterize the behaviour of an entity [ISO 19119]

4.10

metadata dataset

metadata describing a specific dataset [ISO 19101]

4.11**metadata entity**

group of metadata elements and other metadata entities describing the same aspect of data

NOTE 1 A metadata entity may contain one or more metadata entities.

NOTE 2 A metadata entity is equivalent to a class in UML terminology [ISO 19115].

4.12**metadata schema**

conceptual schema describing metadata

NOTE ISO 19115 describes a standard for a metadata schema. [ISO 19101]

4.13**metadata section**

subset of metadata that defines a collection of related metadata entities and elements [ISO 19115]

4.14**operation**

specification of a transformation or query that an object may be called to execute [ISO 19119]

4.15**parameter**

variable whose name and value are included in an operation **request** or **response**

4.16**profile**

set of one or more base standards and - where applicable - the identification of chosen clauses, classes, subsets, options and parameters of those base standards that are necessary for accomplishing a particular function [ISO 19101, ISO 19106]

4.17**qualified name**

name that is prefixed with its naming context

EXAMPLE The qualified name for the road no attribute in class Road defined in the Roadmap schema is RoadMap.Road.road_no. [ISO 19118]

4.18**reporting group**

data with common characteristics forming a subset of a dataset

NOTE 1 Common characteristics can include belonging to an identified feature type, feature attribute or feature relationship; sharing data collection criteria; sharing original source; or being within a specified geographic or temporal extent.

NOTE 2 A reporting group can be as small as a feature instance, an attribute value, or a single feature relationship. [ISO 19109, ISO 19113].

4.19

request

invocation of an **operation** by a **client**

4.20

response

result of an **operation**, returned from a **server** to a **client**

4.21

resource

an object or artefact that is described by a record in the information model of a catalogue

4.22

schema

formal description of a model [ISO 19101, ISO 19103, ISO 19109, ISO 19118]

4.23

server

service instance

a particular instance of a **service** [ISO 19119 edited]

4.24

service

distinct part of the functionality that is provided by an entity through interfaces [ISO 19119]

capability which a service provider entity makes available to a service user entity at the interface between those entities [ISO 19104 terms repository]

4.25

service interface

shared boundary between an automated system or human being and another automated system or human being [ISO 19101]

4.26

service metadata

metadata describing the **operations** and **geographic information** available at a **server** [ISO 19128 draft]

4.27

state

condition that persists for a period

NOTE The value of a particular feature attribute describes a condition of the feature [ISO 19108].

4.28

transfer protocol

common set of rules for defining interactions between distributed systems [ISO 19118]

4.29 version

version of an Implementation Specification (document) and XML Schemas to which the requested operation conforms

NOTE An OWS Implementation Specification version may specify XML Schemas against which an XML encoded operation request or response shall conform and should be validated.

5 Conventions

5.1 Symbols (and abbreviated terms)

Some frequently used abbreviated terms:

CORBA	Common Object Request Broker Architecture
DCP	Distributed Computing Platform
HTTP	Hypertext Transfer Protocol
IDL	Interface Definition Language
ISO	International Organization for Standardization
KVP	Keyword Value Pair
MIME	Multipurpose Internet Mail Extensions
OGC	Open GIS Consortium, also referred to as OpenGIS [®]
TBD	To Be Determined
TBR	To Be Reviewed
UML	Unified Modeling Language
XML	Extensible Markup Language
Z39.50	Service definition for information search and retrieval, also known as ISO 23950

5.2 UML notation

Some of the diagrams in this document are presented using the Unified Modeling Language (UML) static structure diagram. The UML notations used in this document are described in Figure 1, below.

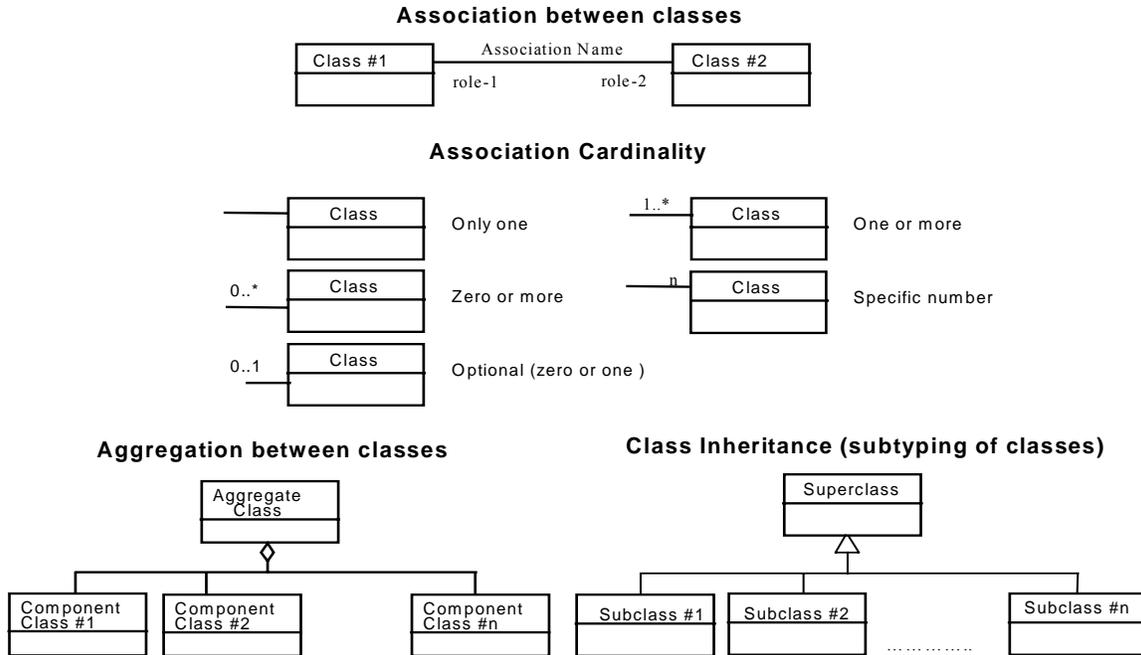


Figure 1 — UML notations

In these UML class diagrams, the class boxes with a light background are the primary classes being shown in this diagram, often the classes from one UML package. The class boxes with a gray background are other classes used by these primary classes, usually classes from other packages.

In this diagram, the following stereotypes of UML classes are used:

- a) <<Interface>> A definition of a set of operations that is supported by objects having this interface. An Interface class cannot contain any attributes.
- b) <<Type>> A stereotyped class used for specification of a domain of instances (objects), together with the operations applicable to the objects. A Type class may have attributes and associations.
- c) <<DataType>> A descriptor of a set of values that lack identity (independent existence and the possibility of side effects). A DataType is a class with no operations whose primary purpose is to hold the information.
- d) <<CodeList>> A flexible enumeration that uses string values for expressing a list of potential values. If the list alternatives are completely known, an enumeration shall be used; if the only likely alternatives are known, a code list shall be used.
- e) <<Enumeration>> A data type whose instances form a list of alternative literal values. Enumeration means a short list of well-understood potential values within a class.

In this document, the following standard data types are used:

- a) `CharacterString` – A sequence of characters
- b) `Boolean` – A value specifying TRUE or FALSE
- c) `Integer` – An integer number
- d) `Identifier` – Unique identifier of an object
- e) `URI` – An identifier of a resource that provides more information
- f) `URL` – An identifier of an on-line resource that can be electronically accessed

5.3 Document terms and definitions

This document uses the specification terms defined in Subclause 5.3 of [OGC 05-008c1].

6 Catalogue abstract information model

6.1 Introduction

The abstract information model specifies a BNF grammar for a minimal query language, a set of core queryable attributes (names, definitions, conceptual datatypes), and a common record format that defines the minimal set of elements that should be returned in the brief and summary element sets.

The geospatial community is a very broad-based community that works in many different operational environments, as shown in the information discovery continuum in Figure 2. On one extreme there are tightly coupled systems dedicated to well defined functions in a tightly controlled environment. At the other extreme are Web based services that know nothing about the client. This document provides a specification that is applicable to the full range of catalogue operating environments.

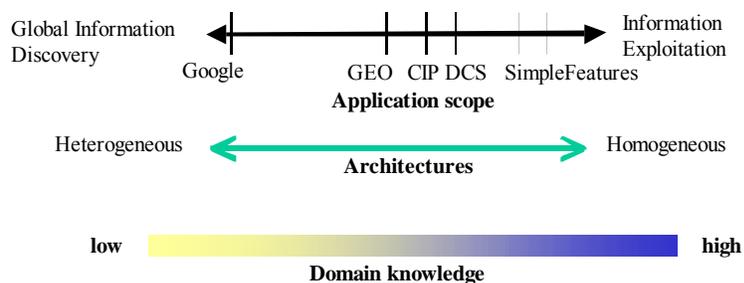


Figure 2 — Information discovery continuum

6.2 Query language support

6.2.1 Introduction

The query capabilities of the OGC General Catalogue Model provide a minimum set of data types and query operations that can be assumed of OGC Compliant Catalogue implementations. In addition, these Query Capabilities provide a high degree of flexibility enabling alternate styles of query, result presentation, and the potential support of any geo-enabled query language. This flexibility is provided by the query operation that contains the parameters needed to select the query result presentation style and to provide a query expression which includes the actual query with an identification of the query language used. The query operation, query expression, and other related operations are further discussed in Clause 7.2.4.

The interoperability goal is supported by the specification of a minimal abstract query (predicate) language, which shall be supported by all compliant OpenGIS Catalogue Services. This query language supports nested Boolean queries, text matching operations, temporal data types, and geospatial operators. The minimal query language syntax is based on the SQL WHERE clause in the SQL SELECT statement. Implementations of query languages that are transformable to the OGC_Common Catalogue Query Language are the OGC Filter Specification and the CIP and GEO profiles of Z39.50 Type-1 queries.

The minimal query language assists the consumer in the discovery of datasets of interest at all sites supporting the OpenGIS Catalogue Services. The ability to specify alternative query languages allows for evolution and higher levels of interoperability among more tightly coupled communities of Catalogue Service Providers and Consumers.

6.2.2 OGC_Common catalogue query language

This subclause defines the BNF for the OGC_Common Catalogue Query Language. OGC_Common is the query language to be supported by all OGC Catalogue Interfaces in order to support search interoperability.

Assumptions made during the development of OGC_Common Query Language:

- a) The query will have a syntax similar to the SQL “Where Clause.”
- b) The expressiveness of the query will not require extensions to various current query systems used in geospatial catalogue queries other than the implementation of some geo operators.
- c) The query language is extensible.
- d) OGC_Common supports both tight and loose queries. A tight query is defined where if a catalogue doesn’t support an attribute/column specified in the query, no entity/row can match the query and the null set is returned. In a loose query, if an attribute is undefined, it is assumed to match.

BNF definition of OGC_Common Query Language:

```

<SQL terminal character> ::= <SQL language character>
<SQL language character> ::= <simple Latin letter>
                               | <digit>
                               | <SQL special character>
<simple Latin letter> ::= <simple Latin upper case letter>
                          | <simple Latin lower case letter>
<simple Latin upper case letter> ::=
    A | B | C | D | E | F | G | H | I | J | K | L | M | N | O
    | P | Q | R | S | T | U | V | W | X | Y | Z
<simple Latin lower case letter> ::=
    a | b | c | d | e | f | g | h | i | j | k | l | m | n | o
    | p | q | r | s | t | u | v | w | x | y | z
<digit> ::=
    0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9
<SQL special character> ::= <space>
                             | <double quote>
                             | <percent>
                             | <ampersand>
                             | <quote>
                             | <left paren>
                             | <right paren>
                             | <asterisk>
                             | <plus sign>
                             | <comma>
                             | <minus sign>
                             | <period>
                             | <solidus>
                             | <colon>
                             | <semicolon>
                             | <less than operator>
                             | <equals operator>
                             | <greater than operator>
                             | <question mark>
                             | <left bracket>
                             | <right bracket>
                             | <circumflex>
                             | <underscore>
                             | <vertical bar>
                             | <left brace>
                             | <right brace>
<space> ::= /*space character in character set in use
             In ASCII it would be 40*/
<double quote> ::= "
<percent> ::= %
<ampersand> ::= &
<quote> ::= '
<left paren> ::= (
<right paren> ::= )
<asterisk> ::= *
<plus sign> ::= +
<comma> ::= ,
<minus sign> ::= -
<period> ::= .
<solidus> ::= /

```

```

<colon> ::= :
<semicolon> ::= ;
<less than operator> ::= <
<equals operator> ::= =
<greater than operator> ::= >
<question mark> ::= ?
<left bracket> ::= [
<right bracket> ::= ]
<circumflex> ::= ^
<underscore> ::= _
<vertical bar> ::= |
<left brace> ::= {
<right brace> ::= }
<separator> ::= { <comment> | <space> | <newline> }
/* The next section of the BNF defines the tokens available to the
   language. I have deleted the concepts of bit string, hex string and
   national character string literal. Keywords have been added to support
   the geo literals. */
<token> ::= <nondelimiter token>
           | <delimiter token>
<nondelimiter token> ::= <regular identifier>
                       | <key word>
                       | <unsigned numeric literal>
<regular identifier> ::= <identifier body>
<identifier body> ::=
<identifier start> [ { <underscore> | <identifier part> } ]
<identifier start> ::= <simple latin letter>
<identifier part> ::= <identifier start>
                    | <digit>
<key word> ::= <reserved word>
<reserved word> ::= AND | POINT | LINESTRING | POLYGON | MULTIPOINT
                  | MULTILINESTRING | MULTIPOLYGON | EMPTY | DATE
                  | TIME | TIMESTAMP | FALSE | TRUE | UNKNOWN | LIKE
                  | MINUTE | MONTH | NOT | NULL
<unsigned numeric literal> ::= <exact numeric literal>
                              | <approximate numeric literal>
<exact numeric literal> ::= <unsigned integer>
                           [ <period>[<unsigned integer>]]
                           | <period> <unsigned integer>
<unsigned integer> ::= {<digit>}
<approximate numeric literal> ::= <mantissa> E <exponent>
<mantissa> ::= <exact numeric literal>
<exponent> ::= <signed integer>
<signed integer> ::= [ <sign> ] <unsigned integer>
<sign> ::= <plus sign> | <minus sign>
<character string literal> ::=
    <quote> [ {<character representation>} ] <quote>
<character representation> ::= <nonquote character> | <quote symbol>
<quote symbol> ::= <quote><quote>
/*End of non delimiter tokens*/
/* I have limited the delimiter tokens by eliminating, interval strings
and delimited identifiers BNF and simplifying the legal character set
to the characters to a single set so no identification of character set
would be needed decision. */
<delimiter token> ::= <character string literal>
                    | <SQL special character>

```

```

        | <not equals operator>
        | <greater than or equals operator>
        | <less than or equals operator>
        | <concatenation operator>
        | <double greater than operator>
        | <right arrow>
        | <left bracket>
        | <right bracket>
<character string literal> ::=
    <quote> [ {<character representation>} ] <quote>
<character representation> ::= <nonquote character> | <quote symbol>
<quote symbol> ::= <quote><quote>
<not equals operator> ::= <>
<greater than or equals operator> ::= >=
<less than or equals operator> ::= <=
/*The following section is intended to give context for identifier and
namespaces. It assumes that the default namespace is specified in the
query request and does not allow any overrides of the namespace */
<identifier> ::=
    <identifier start> [ { <colon> | <identifier part> } ]
<identifier start> ::= <simple Latin letter>
<identifier part> ::= <simple Latin letter> | <digit>
<attribute name> ::= <simple attribute name> | <compound attribute
name>
<simple attribute name> ::= <identifier>
<compound attribute name> ::= <identifier><period>
                               [ {<identifier><period>}... ]
                               <simple attribute name>

/*The rest of the BNF is the real BNF for the query capabilities.*/
<search condition> ::= <boolean value expression>
<boolean value expression> ::= <boolean term>
                               | <boolean value expression> OR <boolean term>
<boolean term> ::= <boolean factor>
                  | <boolean term> AND <boolean factor>
<boolean factor> ::= [ NOT ] <boolean primary>
<boolean primary> ::= <predicate> |
                    <routine invocation> |
                    <left paren> <search condition> <right paren>
<predicate> ::= <comparison predicate>
               | <text predicate>
               | <>null predicate>
               | <temporal predicate>
               | <classification predicate>
               | <existence predicate>

/* This set of productions define temporal predicates */
<temporal predicate> ::= <attribute_name> BEFORE <date-time expression>
                       | <attribute_name> BEFORE OR DURING <period>
                       | <attribute_name> DURING <period>
                       | <attribute_name> DURING OR AFTER <period>
                       | <attribute_name> AFTER <date-time expression>
<date-time expression> ::= <date-time> | <period>

```

/* This set of productions enables loose or tight queries. For example the predicate "cloudcover EXISTS" evaluates as true for all record instances where the attribute cloudcover is a member of the record schema. Similarly, the predicate "cloudcover DOESNOTEXIST" evaluates as true for all record instances where the attribute cloudcover is not a member of the record schema.*/

```

<existence_predicate> ::= <attribute_name> EXISTS
                        | <attribute_name> DOES-NOT-EXIST

<comparison predicate> ::= <attribute name> <comp op> <literal>
<text predicate> ::= <attribute name> [ NOT ] LIKE <character pattern>
<null predicate> ::= <attribute name> IS [ NOT ] NULL
<character pattern> ::= <character string literal>
    /* In a character pattern the character percent is used as a
    wildcard to represent an arbitrary string. This allows LIKE to
    implement the effect of many characters matching operations, such as:
    contains, begins with, ends with, not contains, not begins with, not
    ends with, and so forth. For example:
        attribute like '%contains_this%'
        attribute like '%begins_with_this%'
        attribute like '%ends_with_this%'
        attribute like '%d_ve' will match 'dave' or 'dove'
        attribute not like '%will_not_contain_this%'
        attribute not like '%will_not_begin_with_this%'
        attribute not like '%will_not_end_with_this'    */
<comp op> ::= <equals operator>
              | <not equals operator>
              | <less than operator>
              | <greater than operator>
              | <less than or equals operator>
              | <greater than or equals operator>
<literal> ::= <signed numeric literal>
              | <general literal>
<signed numeric literal> ::= [ <sign> ] <unsigned numeric literal>
<general literal> ::= <character string literal>
                    | <datetime literal>
                    | <boolean literal>
                    | <geography literal>
<boolean literal> ::= TRUE
                    | FALSE
                    | UNKNOWN
<routine invocation> ::= | <geoop name><georoutine argument list>
                        | <relgeoop name><relgeoop argument list>
                        | <routine name><argument list>
<routine name> ::= < attribute name>
<geoop name> ::= EQUALS | DISJOINT | INTERSECTS | TOUCHES | CROSSES
                | WITHIN | CONTAINS | OVERLAPS | RELATE
<relgeoop name> ::= DWITHIN | BEYOND
<argument list> ::=
    <left paren> [ <positional arguments> ] <right paren>
<positional arguments> ::=
    <argument> [ { <comma> <argument> } ]
<argument> ::= <literal> | <attribute name>
<georoutine argument list> ::=
<left paren><attribute name><comma><geometry literal><right paren>

```

```

<relgeoop argument list> ::= <left paren>
    <attribute name><comma><geometry literal><comma><tolerance>
    <right paren>
<tolerance> ::= <unsigned numeric literal><comma><distance units>
<distance units> ::= = "feet" | "meters" | "statute miles" |
    "nautical miles" | "kilometers"
/*this set of units is just an example. The real list of distance unit
must be developed*/
<geometry literal> := <Point Tagged Text>
    | <LineString Tagged Text>
    | <Polygon Tagged Text>
    | <MultiPoint Tagged Text>
    | <MultiLineString Tagged Text>
    | <MultiPolygon Tagged Text>
    | <GeometryCollection Tagged Text>
    | <Envelope Tagged Text>
<Point Tagged Text> := POINT <Point Text>
<LineString Tagged Text> := LINESTRING <LineString Text>
<Polygon Tagged Text> := POLYGON <Polygon Text>
<MultiPoint Tagged Text> := MULTIPOINT <Multipoint Text>
<MultiLineString Tagged Text> := MULTILINESTRING <MultiLineString Text>
<MultiPolygon Tagged Text> := MULTIPOLYGON <MultiPolygon Text>
<GeometryCollection Tagged Text> :=
    GEOMETRYCOLLECTION <GeometryCollection Text>
<Point Text> := EMPTY | <left paren> <Point> <right paren>
<Point> := <x><space><y>
<x> := numeric literal
<y> := numeric literal
<LineString Text> := EMPTY
    | <left paren>
    | <Point>
    | {<comma><Point >}
    | <right paren>
<Polygon Text> := EMPTY
    | <left paren>
    | <LineString Text>
    | {<comma><LineString Text>}
    | <right paren>
<Multipoint Text> := EMPTY
    | <left paren>
    | <Point Text>
    | {<comma><Point Text >}
    | <right paren>
<MultiLineString Text> := EMPTY
    | <left paren>
    | <LineString Text>
    | {<comma><LineString Text>}
    | <right paren>
<MultiPolygon Text> := EMPTY
    | <left paren>
    | <Polygon Text>
    | {<comma><Polygon Text>}
    | <right paren>
<GeometryCollection Text> := EMPTY
    | <left paren>
    | <geometry literal>

```

```

                                {<comma><geometry literal>}
                                <right paren>
<Envelope Tagged Text> ::= ENVELOPE <Envelope Text>
<Envelope Text> ::= EMPTY
                    | <left paren>
                      <WestBoundLongitude><comma>
                      <EastBoundLongitude><comma>
                      <NorthBoundLatitude><comma>
                      <SouthBoundLatitude>
                      <right paren>
<WestBoundLongitude> ::= <signed numeric literal>
<EastBoundLongitude> ::= <signed numeric literal>
<NorthBoundLatitude> ::= <signed numeric literal>
<SouthBoundLatitude> ::= <signed numeric literal>

<date-time> ::= <full-date> "T" <UTC-time>

<full_date> ::= <date-year> "-" <date-month> "-" <date-day>
<date-year> ::= <digit><digit><digit><digit>
<date-month> ::= <digit><digit>
<date-day> ::= <digit><digit>

<UTC-time> ::= <time-hour> ":" <time-minute> ":" <time-second> "Z"
<time-hour> ::= <digit><digit>
<time-minute> ::= <digit><digit>
<time-second> ::= <digit><digit>[. {<digit>}]

<duration> ::= "P" (<dur-date> | <dur-time>)
<dur-date> ::= <dur-day> | <dur-month> | <dur-year> [<dur-time>]
<dur-day> ::= {<digit>} "D"
<dur-month> ::= {<digit>} "M" [<dur-day>]
<dur-year> ::= {<digit>} "Y" [<dur-month>]

<dur-time> ::= "T" (<dur-hour> | <dur-minute> | <dur-second>)
<dur-hour> ::= {<digit>} "H" [<dur-minute>]
<dur-minute> ::= {<digit>} "M" [<dur-second>]
<dur-second> ::= {<digit>} "S"

<period> ::= <date-time> "/" <date-time>
           | <date-time> "/" <duration>
           | <duration> "/" <date-time>

```

6.2.3 Extending the Common Catalogue Query Language

The Common Catalogue Query Language BNF can be extended by adding new predicates, operations, and datatypes. The following discussion is an example of extending the BNF to include a CLASSIFIED-AS operator using the patterns identified in OASIS/ebXML Registry Services Specification v2.5. This extension could appear in a protocol binding or an Application Profile.

This specification makes no assumptions about how taxonomies are maintained in a catalogue, or how records are classified according to those taxonomies. Instead, this specification defines a routine, CLASSIFIED-AS, in order to support classification queries based on taxonomies.

The CLASSIFIED-AS routine takes three arguments. The first argument is the abstract entry point whose classification is being checked. The second argument is the key name string that represents a path expression in the taxonomy. The last argument is the key value string that represents the corresponding path expression containing key values that are the targets of the query. In both cases, the first element of the path expression for the key name argument and key value arguments shall be the name of the taxonomy being used. The normal wildcard matching characters, ‘_’ for a single character and ‘%’ for zero or more characters, may be used in the key value expression which is the last argument of the CLASSIFIED_AS routine.

The following set of productions defines the CLASSIFIED-AS routine.

```
/* The following example: */
/* */
/* RECORD CLASSIFIED AS CLASSIFICATIONSCHEME='GeoClass' */
/*      ='/GeoClass/North America/%/Ontario' */
/* */
/* Will find all records in all the Ontario's in North America. */
```

The following are the required BNF specializations:

```
<classop argument list> ::= <left paren> <entry_point> <comma>
    <Classification Scheme> <comma><Classification Node> <right
paren>

<entry_point> ::= <identifier>
<Classification Scheme> ::= <identifier>
<classop name> ::= CLASSIFIED_AS

<Classification Node> ::= <identifier> | <solidus><path
element>[<solidus><path element>]...
<path element> ::= <character pattern>

<routine invocation> ::= | <geoop name><georoutine argument list>
    | <relgeoop name><relgeoop argument list>
    | <routine name><argument list>
    | <classop><classop argument list>
```

Consider the following example:

```
CLASSIFIED_AS('RECORD', 'GeoClass', 'GeoClass/NorthAmerica/%/Ontario')
```

In this example, we are searching records classified according to the **GeoClass** taxonomy. Specifically, we are looking for all catalogue records classified as *Continent=NorthAmerica*, *Country=**any country*** and *State=Ontario*. Notice how the wildcard character ‘%’ is used to search for any Country node.

Here is the same example encoded using XML:

```
<ogc:Filter xmlns:ogc="http://http://www.opengis.net/ogc">
  <ogc:ClassifiedAs>
    <ogc:TypeName>csw:Record</ogc:TypeName>
    <ogc:ClassificationScheme>GetClass</ogc:ClassificationScheme>
```

```
<ogc:ClassificationNode>/GeoClass/NorthAmerica/%/Ontario</ogc:ClassificationNode>
  </ogc:ClassifiedAs>
</ogc:Filter>
```

In order for catalogue clients to be able to determine which taxonomies are available, a catalogue implementation should advertise the list of available taxonomies in its capabilities document. If a query is executed against a non-existent taxonomy, then an exception should be raised.

6.2.4 Query language realization

Many OGC service operations have the requirement to pass and process a query as a structure to perform a request. There are several query languages and messaging mechanisms identified within OGC specifications. Application Profiles should be explicit about the selected query languages and any features peculiar to a scope of application. The following items should be addressed in the preparation of an Application Profile with respect to query language support:

- a) Support for “abstract” query against well-known queryable entry points (e.g. OGC Core). Some specifications promote or require the exposure of well-known field-like objects as common search targets (queryables), allowing interrogation of a service without prior negotiation on information content. The mandatory queryable attributes which shall be recognised by all OGC Catalogue Services are discussed in Subclause 6.3.2.
- b) Selection of a query language. Some specifications describe one or more query languages that can be supported. Identify the name and version of required query language(s) anticipated by this Application Profile for use.
- c) Supported data types (e.g. character, integer, coordinate, date, polygon) and operator types (e.g. inequality, proximity, partial string, spatial, temporal). Query languages may be restricted in their implementation or extended with functions not described in the base specification. This narrative should provide lists or reference documents with the enumerated data types and operator types required by this Application Profile. In addition, any description of special techniques (e.g. supporting joins or associations) that are expected by an Application Profile should be described.

6.3 Core catalogue schema

6.3.1 Introduction

Metadata structures, relationships, and definitions -- known as conceptual schemas -- exist for multiple information communities. For the purposes of interchange of information within an information community, a metadata schema may be defined that provides a common vocabulary which supports search, retrieval, display, and association between the description and the object being described. Although this specification does not require the use of a specific schema, the adoption of a given schema within an

information-sharing community ensures the ability to communicate and discover information.

The geomatics standardization activity under ISO Technical Committee 211 includes a formal schema for geospatial metadata that is intended to apply to all types of information. This metadata standard, ISO 19115:2003 includes a proposal for core metadata elements in common use. ISO Draft Technical Specification 19139 defines a formal encoding and structure of ISO metadata for exchange. Where a catalogue service advertises such application schemas, catalogues that handle geographic dataset descriptions should conform to published metadata standards and encodings, e.g. ISO 19115:2003, and support XML encoding per ISO 19139 or profiles thereof. Service metadata elements should be consistent with ISO 19119.

6.3.2 Core queryable properties

The goal of defining core queryable properties is query interoperability among catalogues that implement the same protocol binding and query compatibility among catalogues that implement different protocol bindings perhaps through the use of “bridges” or protocol adapters. Defining a set of core queryable properties also enables simple cross-profile discovery, where the same queries can be executed against any catalogue service without modification and without detailed knowledge of the catalogue's information model. This requires a set of general metadata properties that can be used to characterize any resource.

Tables 1, 2 and 3 define a set of abstract queryables that binding protocols shall realize in their core queryable schemas. Binding protocols shall further specify a record identifier (ID) based on the native platform ID types. Binding protocols shall also specify how the values of core queryable properties shall be encoded in service requests. Binding protocols may choose to use a single comma-separated list for a compound datatypes or may label each sub-element for clarity and order flexibility. Application profiles may further modify or redefine the realization of the core queryables and how their values are encoded.

Table 1 — Common queryable elements

Name	Definition	Data type
Subject ^a	The topic of the content of the resource ^b	CharacterString
Title ^a	A name given to the resource	CharacterString
Abstract ^a	A summary of the content of the resource	CharacterString
AnyText	A target for full-text search of character data types in a catalogue	CharacterString
Format ^a	The physical or digital manifestation of the resource	CharacterString
Identifier ^a	An unique reference to the record within the catalogue	Identifier
Modified ^c	Date on which the record was created or updated within the catalogue	Date-8601
Type ^a	The nature or genre of the content of the resource. Type can include general categories, genres or aggregation levels of content.	CodeList ^f
BoundingBox ^d	A bounding box for identifying a geographic area of interest	BoundingBox, See Table 2
CRS ^e	Geographic Coordinate Reference System (Authority and ID) for the BoundingBox	Identifier
Association	Complete statement of a one-to-one relationship	Association, See Table 3

a Names, but not necessarily the identical definition, are derived from the Dublin Core Metadata Element Set, version 1.1:ISO Standard 15836-2003 (February 2003)

b Typically, a Subject will be expressed as keywords, key phrases or classification codes that describe a topic of the resource. Recommended best practice is to select a value from a controlled vocabulary or formal classification scheme.

c DCMI metadata term <<http://dublincore.org/documents/dcmi-terms/>>.

d Same semantics as EX_GeographicBoundingBoxclass in ISO 19115.

e If not supplied, the BoundingBox CRS is a Geographic CRS with the Greenwich prime meridian.

f A “CodeList” is a CharacterString taken from an authoritative list of CharacterStrings or Identifiers. The authority may optionally be identified in the value.

Table 2 — Composition of compound element “BoundingBox”

Name	Definition	Data type
WestBoundLongitude	Western-most coordinate of the limit of the resource's extent, expressed in longitude in decimal degrees (positive east)	numeric
SouthBoundLatitude	Southern-most coordinate of the limit of the resource's extent, expressed in latitude in decimal degrees (positive north)	numeric
EastBoundLongitude	Eastern-most coordinate of the limit of the resource's extent, expressed in longitude in decimal degrees (positive east)	numeric
NorthBoundLatitude	Northern-most, coordinate of the limit of the resource's extent, expressed in latitude in decimal degrees (positive north)	numeric

Table 3 — Composition of compound element “Association”

Name	Definition	Data type
Target	Referenced resource	Identifier
Source	Referencing resource	Identifier
Relation	The name of the description of the relationship	CodeList or Identifier

All realizations of the core queryable properties in a binding protocol shall include all the properties listed in tables 1,2 or 3 even if the underlying information model does not include information that can be mapped into all properties. Core properties that cannot have a value assigned to them because the information is not available in the information model of the catalogue shall be considered as having a value of NULL.

The properties “Title”, “Identifier” and the pseudo-property “AnyText” shall be supported as mandatory queryables in all implementations. Protocol bindings shall describe mechanisms to identify and elaborate on the queryables and operations supported by a given catalogue service.

6.3.3 Core returnable properties

A set of core properties returned from a metadata search is encouraged to permit the minimal implementation of a catalogue service independent of a companion application profile, and to permit the use of metadata returned from different systems and protocol bindings. The core metadata is returned as a request for the Common Element Set. The Common Element Set is a new group of public metadata elements, expressed using the nomenclature and syntax of Dublin Core Metadata, ISO 15836. Table 4 provides some interpretation of Dublin Core elements in the context of metadata for geospatial data and services.

Table 4 — List of common returnable properties

Dublin Core element name	Term used in OGC queryables	Definition	Data type
title	Title	A name given to the resource. Also known as "Name".	CharacterString
creator		An entity primarily responsible for making the content of the resource.	CharacterString
subject	Subject	A topic of the content of the resource. This is a place where a Topic Category or other taxonomy could be applied.	CharacterString
description	Abstract	An account of the content of the resource. This is also known as the "Abstract" in other aspects of OGC, FGDC, and ISO metadata.	CharacterString
publisher		An entity responsible for making the resource available. This would equate to the Distributor in ISO and FGDC metadata.	CharacterString
contributor		An entity responsible for making contributions to the content of the resource.	CharacterString
date	Modified	The date of a creation or update event of the catalogue record.	ISO-8601 date
type	Type	The nature or genre of the content of the resource.	CodeList
format	Format	The physical or digital manifestation of the resource.	CharacterString
identifier	Identifier	A unique reference to the record within the catalogue.	Identifier
source	Source	A reference to a resource from which the present resource is derived.	CharacterString
language		A language of the intellectual content of the catalogue record.	CharacterString
relation	Association	The name of the relationship that exists between the resource described by this record and a related resource referenced using the <i>Source</i> or <i>dc:source</i> property.	CodeList
coverage	BoundingBox	The spatial extent or scope of the content of the resource.	
rights		Information about rights held in and over the resource.	CharacterString

The core elements are recommended for a response but do not need to be populated. The support for a common syntax for the returnable properties as a "common" Summary Element Set is defined in the protocol binding clauses.

EXAMPLE The following shows an example query, using the GetRecords operation described in Clause 10. The operation references the concrete realization of the core queryables, using Dublin Core Metadata Elements, also described in Clause 10.

```
<?xml version="1.0" encoding="UTF-8"?>
```

```

<GetRecords
  service="CSW"
  version="2.0.2"
  maxRecords="5"
  startPosition="1"
  resultType="results"
  outputFormat="application/xml"
  outputSchema="http://www.opengis.net/cat/csw/2.0.2"
  xmlns="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:csw="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:ogc="http://www.opengis.net/ogc"
  xmlns:ows="http://www.opengis.net/ows"
  xmlns:dc="http://purl.org/dc/elements/1.1/"
  xmlns:dct="http://purl.org/dc/terms/"
  xmlns:gml="http://www.opengis.net/gml"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.opengis.net/cat/csw/2.0.2
    ../../csw/2.0.2/CSW-discovery.xsd"
  <Query typeNames="csw:Record">
    <ElementSetName typeNames="csw:Record">full</ElementSetName>
    <Constraint version="1.1.0">
      <ogc:Filter>
        <ogc:And>
          <ogc:PropertyIsLike escapeChar="\ " singleChar="?"
wildCard="*">
            <ogc:PropertyName>dc:title</ogc:PropertyName>
            <ogc:Literal>*Elevation*</ogc:Literal>
          </ogc:PropertyIsLike>
          <ogc:PropertyIsEqualTo>
            <ogc:PropertyName>dc:type</ogc:PropertyName>
            <ogc:Literal>Service</ogc:Literal>
          </ogc:PropertyIsEqualTo>
          <ogc:PropertyIsGreaterThanOrEqualTo>
            <ogc:PropertyName>dct:modified</ogc:PropertyName>
            <ogc:Literal>2004-03-01</ogc:Literal>
          </ogc:PropertyIsGreaterThanOrEqualTo>
          <ogc:Intersects>
            <ogc:PropertyName>ows:BoundingBox</ogc:PropertyName>
            <gml:Envelope>
              <gml:lowerCorner>14.05 46.46</gml:lowerCorner>
              <gml:upperCorner>17.24 48.42</gml:upperCorner>
            </gml:Envelope>
          </ogc:Intersects>
        </ogc:And>
      </ogc:Filter>
    </Constraint>
  </Query>
</GetRecords>

```

The response to such a query, might be:

```

<?xml version="1.0" encoding="UTF-8"?>
<csw:Record
  xmlns:csw="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:dc="http://purl.org/dc/elements/1.1/"
  xmlns:dct="http://purl.org/dc/terms/"

```

```

xmlns:ows="http://www.opengis.net/ows"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://www.opengis.net/cat/csw/2.0.2
    ../../../../csw/2.0.2/record.xsd">
  <dc:creator>U.S. Geological Survey</dc:creator>
  <dc:contributor>State of Texas</dc:contributor>
  <dc:publisher>U.S. Geological Survey</dc:publisher>
  <dc:subject>Elevation, Hypsography, and Contours</dc:subject>
  <dc:subject>elevation</dc:subject>
  <dct:abstract>Elevation data collected for the National Elevation
Dataset (NED) based on 30m horizontal and 15m vertical
accuracy.</dct:abstract>
  <dc:identifier>ac522ef2-89a6-11db-91b1-7eea55d89593</dc:identifier>
  <dc:relation>OfferedBy</dc:relation>
  <dc:source>ddl1b2ce7-0722-4642-8cd4-6f885f132777</dc:source>
  <dc:rights>Copyright © 2004, State of Texas</dc:rights>
  <dc:type>Service</dc:type>
  <dc:title>National Elevation Mapping Service for Texas</dc:title>
  <dct:modified>2004-03-01</dct:modified>
    <dc:language>en</dc:language>
  <ows:BoundingBox>
    <ows:LowerCorner>-108.44 28.229</ows:LowerCorner>
    <ows:UpperCorner>-96.223 34.353</ows:UpperCorner>
  </ows:BoundingBox>
</csw:Record>

```

6.3.4 Information structure and semantics

Some services that implement OGC Implementation Specifications expect a rigid syntax for the information resources to be returned, whereas others do not. This subclause allows an Application Profile to be specific about what information content, syntax, and semantics are to be communicated over the service. The following items should be addressed in an Application Profile.

- a) Identify information resource types that can be requested. In the case of a catalogue service, the information resources being described by the metadata may include geographic data, imagery, services, controlled vocabularies, or schemas among a wide variety of possible types. This subclause allows the community to specify or generalise the resource types being described in metadata for their scope of application.
- b) Identify a public reference for the information being returned by the service (e.g. ISO 19115:2003 “Geographic Information – Metadata “). Include any semantic resources including data content model, dictionary, feature type catalog, code lists, authorities, taxonomies, etc.
- c) Identify named groups of properties (element sets) that may be requested of the service (e.g. “brief,” “summary,” or “full”) and the valid format (syntax) for each element set. Identify valid schema(s) with respect to a given format to assist in the validation of response messages.

- d) Specialise the core queryable properties list by making some optional queryable attributes mandatory, deleting other optional attributes and adding queryable attributes that should be standard across all profile users
- e) Optional mapping of queryable and retrievable properties against other public metadata models or tags.
- f) Expected response/results syntax and content Message syntax and schemas (e.g. brief/full, individual elements).

7 General catalogue interface model

7.1 Introduction

The General Catalogue Interface Model provides a set of abstract service interfaces that support the discovery, access, maintenance and organization of catalogues of geospatial information and related resources. The interfaces specified are intended to allow users or application software to find information that exists in multiple distributed computing environments, including the World Wide Web (WWW) environment. All behaviour requiring sessions is expressed by a dynamic model of conversation state and state transitions. The model expresses the states and messages that trigger the changes in state.

Implementation design guidance is included in the protocol binding clauses of this specification. Each protocol binding includes a mapping from the general interfaces, operations, and parameters specified in this clause to the constructs available in a chosen protocol. Application profiles are intended to further document implementation choices.

An Application Profile is predicated on the existence of one protocol binding in the base specification. In the case of the Catalogue Services Specification, a profile could reference CORBA, Z39.50, or HTTP protocol bindings. In most, but not all, protocol bindings, there may be restrictions or refinements on implementation of the General Model agreed within an implementation community. This subclause provides an overview of the portions of the General Catalogue Model that are realised by Protocol Bindings and Application Profiles.

Figure 3 shows the Reference Architecture assumed for development of the OGC Catalogue Interface. The architecture is a multi-tier arrangement of clients and servers. To provide a context, the architecture shows more than just catalogue interfaces. The bold lines illustrate the scope of OGC Catalogue and Features interfaces.

The Application Client shown in Figure 3 interfaces with the Catalogue Service using the OGC Catalogue Interface. The Catalogue Service may draw on one of three sources to respond to the Catalogue Service request: a Metadata Repository local to the Catalogue Service, a Resource service, or another Catalogue Service,. The interface to the local Metadata Repository is internal to the Catalogue Service. The interface to the Resource service can be a private or OGC Interface. The interface between Catalogue Services is the OGC Catalogue Interface. In this case, a Catalogue Service is acting as both a client and server. Data returned from an OGC Catalogue Service query is processed by the

requesting Catalogue Service to return the data appropriate to the original Catalogue request. See Annex B for more about Distributed Searching.

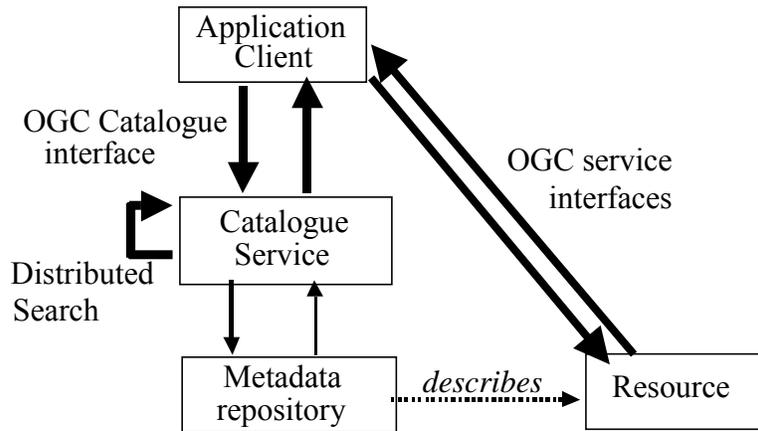


Figure 3 — Reference model architecture

7.2 Interface definitions

7.2.1 Overview

Figure 4 is a general UML model of OGC catalogue service interfaces, in the form of a class diagram. Operation signatures have been suppressed in this figure for simplicity but are described in detail below. This model shows the Catalogue Service class plus five other classes with which that class are associated. A Catalogue Service is a realization of an OGC Service. Each instance of the Catalogue Service class is associated with two or more of these other classes, depending on the abilities included in that service instance. Each of these other classes defines one or several related operations that can be included in a Catalogue Service class instance. The Catalogue Service class directly includes only the serviceTypeID attribute, with a fixed value for the service type.

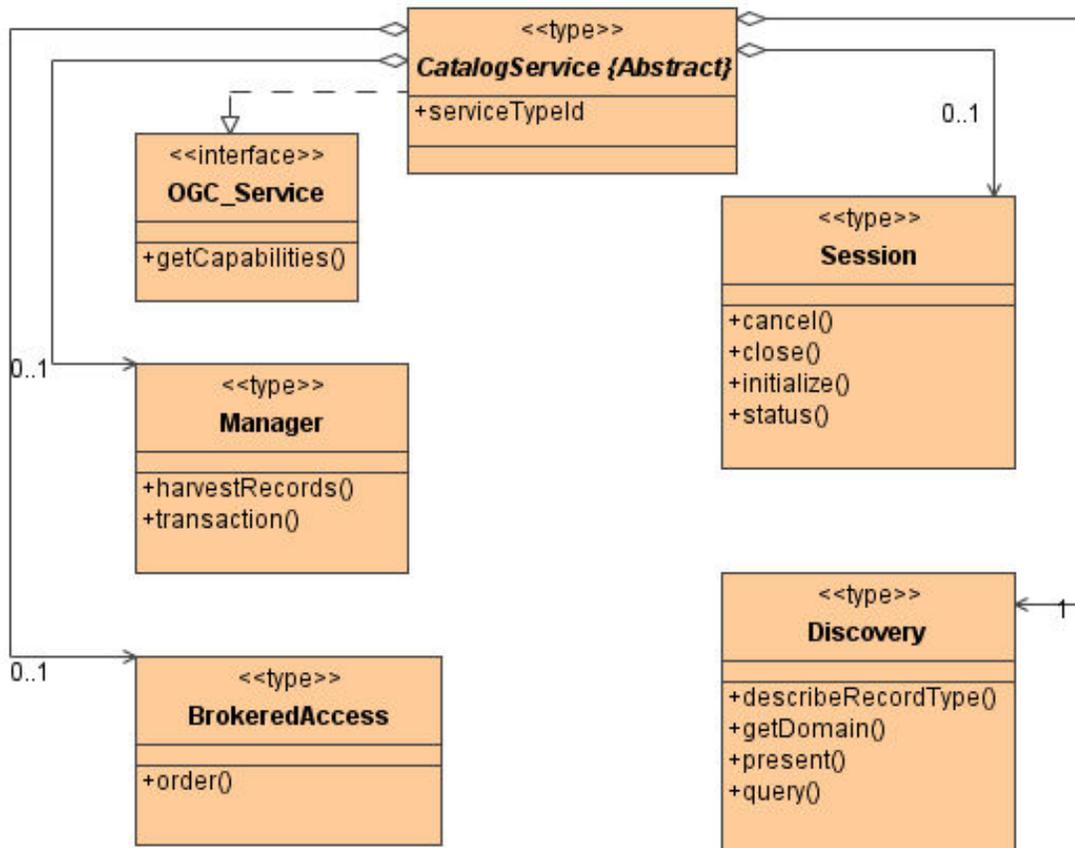


Figure 4 — General OGC catalogue UML static model

In Figure 4 an instance of the `CatalogService` type is a composite object that is a high-level characterization of a catalogue service. Its constituent objects are themselves components that provide functional behaviours to address particular areas of concern. A protocol binding may realise specific configurations of these components to serve different purposes (e.g. a read-only catalogue for discovery, a transactional catalogue for discovery and publication, or a ‘stateful’ catalogue that supports session management).

The associated classes shown in this figure are mandatory or optional for implementation as indicated by the association multiplicity in the UML diagram. Therefore, a compliant catalogue service shall implement the `OGC_Service`, `CatalogService`, and `Discovery` classes. An application profile or protocol binding can implement additional classes associated with the Catalogue Service class. A catalogue implementation shall recognise all operations defined within each included class, and shall generate a message indicating when a particular operation is not implemented.

The protocol binding clauses of this specification provide more detail on the implementation of these conceptual interfaces. For example, the names of the classes and operations in this general UML model are changed in some of the protocol bindings. The names of some operation parameters are also changed in some protocol bindings.

Application Profiles may further specialise the implementation of these interfaces and their operations, including adding classes. In general, however, the interfaces and operations described here shall have the same semantics and granularity of interaction regardless of the protocol binding used.

The Catalogue Service class can be associated with the:

- a) OGC_Service class, which provides the getCapabilities operation that retrieves catalogue service metadata. This class is always realised by the Catalogue Service class, and is thus always implemented by a Catalogue Service implementation.
- b) Discovery class, which provides four operations for client discovery of resources registered in a catalogue. This class has a required association from the Catalogue Service class, and is thus always implemented by a Catalogue Service implementation. The “query” operation searches the catalogued metadata and produces a result set containing references to all the resources that satisfy the query. This operation optionally returns metadata for some or all of the found result set. The “present” operation returns selected metadata for some or all of the resources referenced in a specific previous result set. The describeRecordType operation retrieves the type definition used by metadata of one or more registered resource types. The getDomain operation retrieves information about the valid values of one or more named metadata properties.
- c) Session class, which provides four operations for interactive sessions between a server and a client. This class has an optional association from the Catalogue Service class; this interface is implemented by the Catalogue Service implementation. The “initialise” operation initiates an interactive session, and the “close” operation terminates a session. The “status” operation retrieves the current status of a previously initiated operation, and the “cancel” operation terminates a previously initiated operation.
- d) Manager class, which provides two operations for inserting, updating, and deleting the metadata by which resources are registered in a catalogue. This class has an optional association from the Catalogue Service class; this interface is implemented by the Catalogue Service implementation. The transaction operation performs a specified set of “insert”, “update”, and “delete” actions on metadata items stored by a Catalogue Service implementation—this enables a “push” style of publication. The harvestRecords operation requests the Catalogue Service to retrieve resource metadata from a specified location, often on a regular basis—this behaviour reflects a ‘pull’ style of publication.
- e) Brokered Access class, which provides the “order” operation for ordering an identified resource that is registered in a catalogue but is not directly accessible to the client. This class has an optional association from the Catalogue Service class; this interface is implemented by the Catalogue Service implementation.

The five classes that can be associated with the Catalogue Service class allow different OGC catalogue services to provide significantly different abilities. A particular protocol

binding is used by each Application Profile and a particular set of these catalogue service classes is specified by each Application Profile.

Each of the catalogue classes is described further in the following subclauses. These subclauses discuss the operations and parameters of each operation in this general model. Specific protocol bindings or application profiles can define additional parameters. For example, the HTTP Protocol Binding adds the Service, Request, and Version parameters to all operation requests to be consistent with other OGC Web Services.

7.2.2 Catalogue Service class

The Catalogue Service class provides the foundation for an OGC catalogue service. The Catalogue Service class directly includes only the serviceTypeID attribute, as specified in Table 5. In most cases, this attribute will not be directly visible to catalogue clients.

Table 5 — Attribute of Catalogue Service class

Name	Definition	Data type	Multiplicity
serviceTypeID	Identification of catalogue service type	URI, as specified in IETF RFC 2396	One (Mandatory)

7.2.3 OGC_Service class

7.2.3.1 Introduction

The OGC_Service class allows clients to retrieve service metadata by providing the getCapabilities operation. This class is always realised by the Catalogue Service class, and is thus always implemented by a Catalogue Service instance.

NOTE This getCapabilities operation corresponds to CatalogueService.explainServer operation in OGC Catalogue version 1.1.1.

7.2.3.2 getCapabilities operation

The getCapabilities operation is more completely specified in Table 6.

Table 6 — Definition of getCapabilities operation

Definition	Allows clients to retrieve service metadata describing Catalogue Service instance
Receives	Optional identifier(s) of requested parts of the complete service metadata document
Returns	Service metadata document for Catalogue Service instance. Some document contents depend on the set of classes that are associated with the Catalogue Service class, as defined by the specific protocol binding, and on other details of that protocol binding. Other document contents depend on the types of data defined by the specific application profile, and on other details of that profile.
Exceptions	Invalid Parameter Value, Missing Parameter Value
Pre-conditions	None
Post-conditions	Service metadata document returned to requesting client, either complete or including selected parts

Figure 5 provides a UML model of the OGC_Service class that shows the complete signature of the getCapabilities operation, plus classes for the getCapabilities operation request and the ServiceMetadata operation response. The abstract GetCapabilities and Service Metadata classes are specialised by each service that uses the OGC_Service class. The detailed contents of both the CatalogGetCapabilities and CatalogServiceMetadata classes depend on the protocol binding, and perhaps also on the Application Profile, and are thus not detailed here.

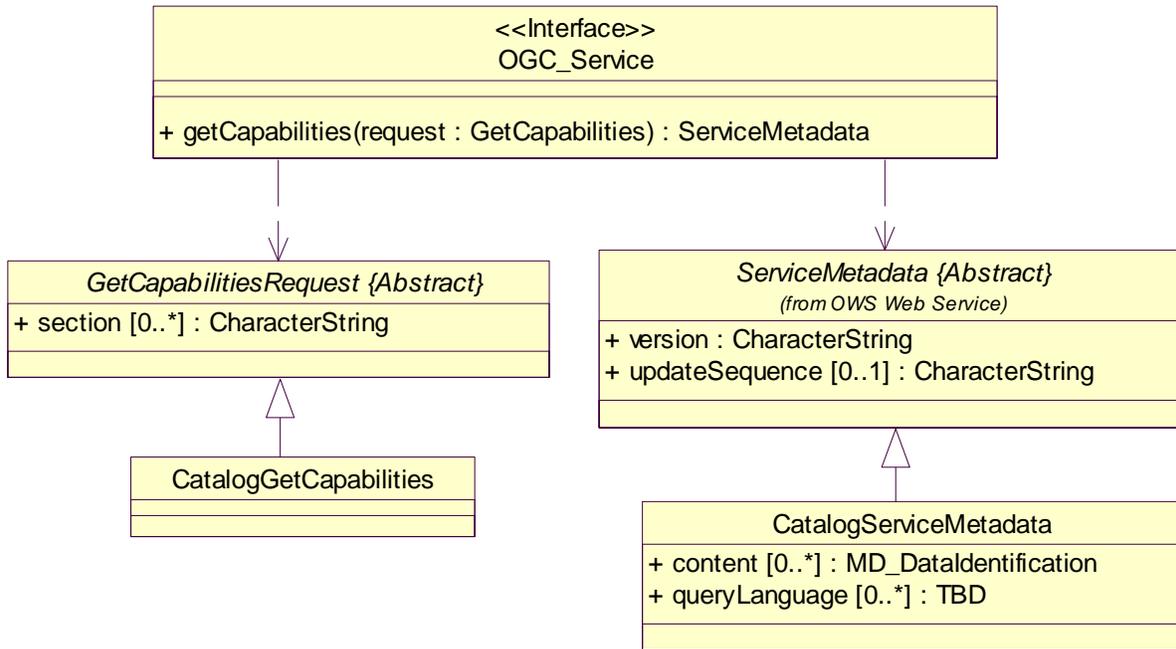


Figure 5 — getCapabilities operation UML static model

The GetCapabilities operation request includes one “section” attribute listed and defined in Table 6.

Table 7 — UML attribute in getCapabilities operation request

Name	Definition	Data type and value	Optionality and use
section	Name of requested section in complete service metadata document	Character String type, not empty Allowed values specified by each Application Profile	Zero or more (Optional) Return complete service metadata document when omitted

The normal GetCapabilities operation response is a service metadata document that includes the “section” attributes listed and defined in Table 8, as selected by the “section” attribute in the operation request.

Table 8 — UML attributes in getCapabilities operation normal response

Name	Definition	Data type	Optionality and use
ServiceIdentification	Metadata about this specific server	SV_ServiceIdentification in ISO 19119	Zero or one (Optional) Include when requested
ServiceProvider	Metadata about the organization operating this server	SV_ServiceProvider in ISO 19119	Zero or one (Optional) Include when requested
OperationMetadata	Metadata about an operations specified by this service, including the URL(s) for operation requests	SV_OperationMetadata in ISO 19119	Zero or more (Optional) Include when requested Repeated for each operation implemented by this server
Content	Metadata about a collection or type of resource catalogued by this server	MD_DataIdentification in ISO 19115 (adapted)	Zero or more (Optional) Include when requested Repeated for each collection and type of resources catalogued
QueryLanguage	Metadata about a query language supported by this server, specifying the query abilities implemented	Character string	Zero or more (Optional) Include when requested Repeated for each query language implemented by this server

NOTE 1 The term “Capabilities XML” document was previously used for what is here called “service metadata” document. The term “service metadata” is now used because it is more descriptive and is compliant with OGC Abstract Specification Topic 12 (ISO 19119).

NOTE 2 This general model assumes that operation failure will be signalled to the client in a manner specified by each protocol binding.

7.2.4 Discovery class

7.2.4.1 Introduction

The Discovery class allows clients to discover resources registered in a catalogue, by providing four operations named “query”, “present”, describeRecordType, and getDomain. This class has a required association from the Catalogue Service class, and is thus always implemented by all Catalogue Service implementations. The Session class can be included with the Discovery class, in associations with the Catalogue Service class.

The “query” and “present” operations may be executed in a session or stateful context. If a session context exists, the dynamic model in Subclause 7.4.3 defines the states of the session and the allowed transitions between states. When the “query” and “present” state does not include a session between a server and a client, any “memory or shared information” between the client and the server may be based on private understandings or features available in the protocol binding. The describeRecordType and getDomain

operations do not require a session context. If a session context exists, the simple request-response dynamic model shown in Subclause 7.4.4 would apply.

7.2.4.2 “query” operation

The “query” operation is more completely specified in Table 9. Figure 6 provides a UML model of the “query” operation that shows the complete Discovery class with the QueryRequest and QueryResponse classes and the classes they use. This UML diagram does not detail the <<CodeList>> stereotyped classes named ResourceType, QueryScope, ResultType, ElementSet, ResponseSchema, ReturnFormat, SortOrder, and QueryLanguage. The operation request includes the attributes and association role names listed and defined in Table 10 through Table 13. The normal operation response includes the attributes and association role names listed and defined in Table 14.

NOTE The query operation corresponds to Discovery.query in OGC Catalogue version 1.1.1.

Table 9 — Definition of “query” operation

Definition	Allows clients to ask a catalogue to execute a query that searches the catalogued metadata and produces a result set containing (zero or more) references to all the registered resources that satisfy the query. The server may maintain the result set for subsequent retrieval requests.
Receives	Specifications of query constraints and of metadata to be returned
Returns	Number of items in result set, and/or selected metadata for some or all of the result set. The client can specify the maximum number of records for which metadata is returned. When metadata return is requested, the service implementation shall first sort the result set as specified by the client. Most of the metadata returned depends on the metadata requested and on the types of data defined by the specific Application Profile.
Exceptions	Missing Parameter Value, Invalid Parameter Value, Nonexistent collection or type
Pre-conditions	The client knows the schema of the information model that the catalogue supports and can thus form valid query expressions.
Post-conditions	Response returned to requesting client, containing number of items in result set and/or selected metadata for some or all of result set

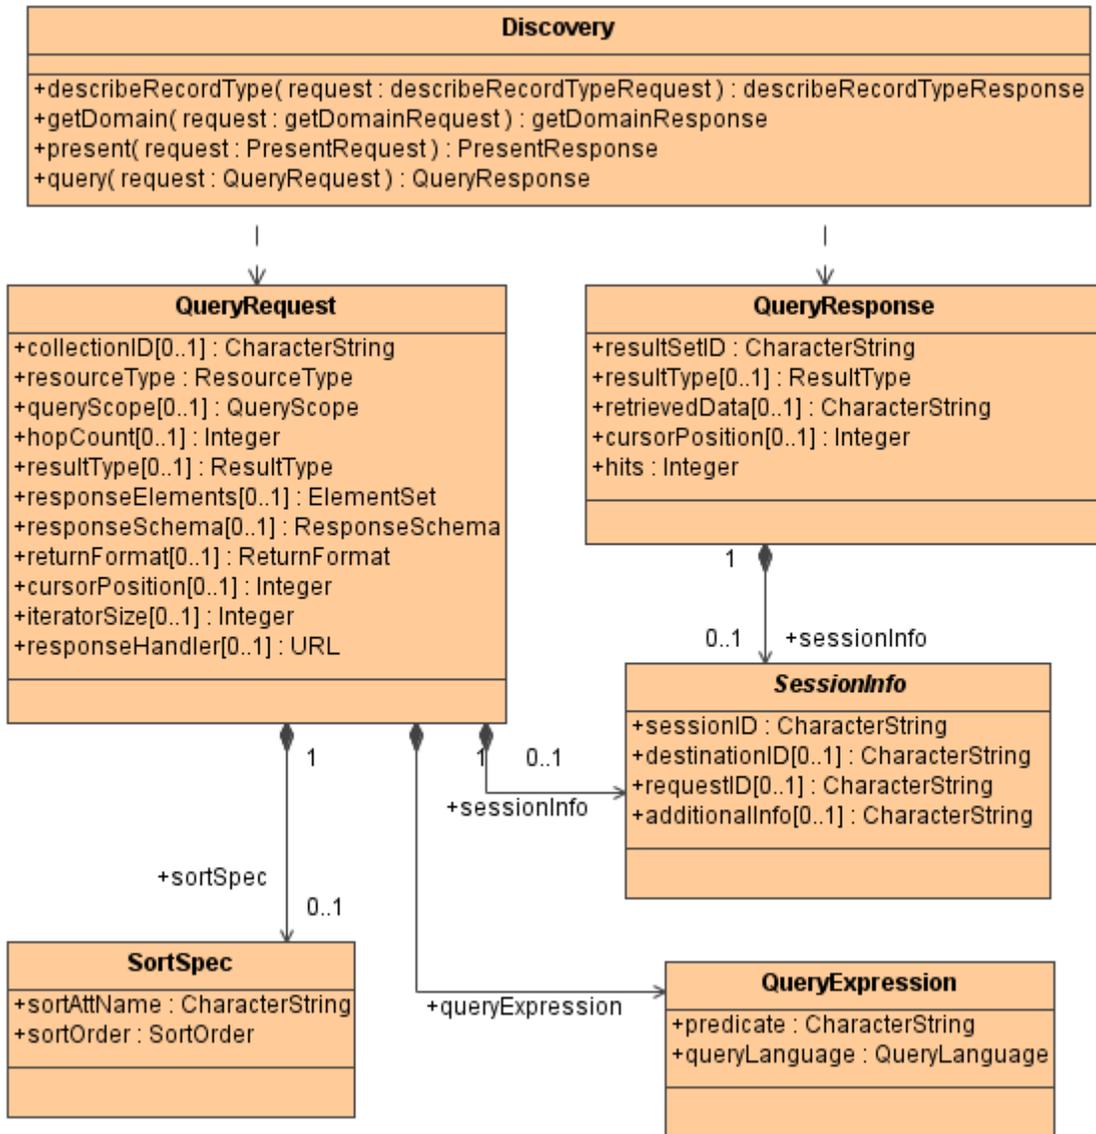


Figure 6 — “query” operation UML static model

Table 10 — UML attributes and roles in “query” operation request

Name	Definition	Data type and value	Optionality and use
sessionInfo	The core set of parameters required in each message exchanged between a client and server operating in a session context, where these parameters support message routing and session management	SessionInfo, See Table 12	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent
queryExpression	The query language and predicate expressing query constraints	QueryExpression, See Table 13	One (Mandatory)
collectionID	Specifies the search space for this query. Search space can be all catalogue holdings, a previously named result set, or a named subspace of the catalogue holdings	Character String type, not empty Specific values that may be referenced are application profile or protocol binding dependent	Zero or one (Conditional) Include when required by protocol binding, otherwise optional
resourceType	A catalogue may contain references to several different resource types. This parameter provides for the selection of one of those types for processing	CodeList type ^a	One (Mandatory)
queryScope	Scope of this query.	CodeList type with allowed values of “local” and “distributed”	Zero or one (Optional) Default value is “local”
hopCount	Maximum number of message hops before distributed search is terminated. Each catalogue decrements value by one when request is received, and does not forward request if hopCount=0.	Non-negative integer	Zero or one (optional) Default value is “2” Included only when queryScope has value “distributed”
resultType	Specifies how client wants result set presented and the behaviour of the catalogue as to when a response is sent	CodeList Type ^b	Zero or one (Conditional) Default values specified by protocol binding or application profile
ResponseElements ^c	Specifies set name or list of metadata elements to be returned in the context of a specific metadata structure	Either a list of elements as name/type pairs OR CodeList type named ElementSet with allowed values of “brief,” “summary” “full” and “browse”	Zero or one (Optional) Default value is “brief”

Name	Definition	Data type and value	Optionality and use
ResponseSchema ^C	The name of the “well-known” or advertised (in the capabilities) schema of the response	Code List type with one mandatory value of “OGCCORE” that represents the core catalogue schema. Other values may be defined by application profiles. Examples of such values might be: “FGDC”, “ISO-19119”, ISO-19139”, ANZLIC	Zero or one (Optional). If the parameter is not specified then the default value is “OGCCORE”.
sortSpec	Sorting information to the server for formatting data returned to the client	SortSpec, See Table 11	Zero or one (Optional) Default is sorted on ID in descending order
returnFormat	Specifies format (MIME or Internet media type) for returning result set metadata	CodeList type XML HTML TXT	Zero or one (Optional) Default is “XML” Include when results to be returned
cursorPosition	First result set resource to be returned for this operation request	Positive integer	Zero or one (Optional) Default is “1” Include when results to be returned
iteratorSize	Specifies maximum number of result set resources to be returned	Non-negative integer	Zero or one (Optional) Default is “10” Include when results to be returned
responseHandler	Network location to which the response will be forwarded when operation has been completed, for asynchronous requests	URL	Zero or one (Optional) If not included, process request synchronously
<p>a Values and definitions of resourceType codes:</p> <p>Data set – the lowest level packaging of Features that have been catalogued</p> <p>Data set collection – a grouping of data sets that have commonality (ISO 19115: data set series)</p> <p>Service – a set of interfaces that provide access to or operations on data (e.g. catalogue service)</p> <p>b Values and definition of resultType codes and behaviours in session based environments:</p> <p>validate - the QueryResponse is returned as soon as QueryRequest has been determined to be valid. Query processing continues after the QueryResponse is returned.. Reasons for failure are provided in the diagnostic of QueryResponse.</p> <p>resultSetID - the QueryResponse is returned as soon as the resultSetID is available and the query has completed processing.</p> <p>hits- the QueryResponse is returned as soon as the query has completed processing and the number of hits has been determined. Metadata records are not returned in the QueryResponse</p> <p>results - the QueryResponse is returned as soon as the query has completed processing and the results have been formatted for return. Metadata records are returned in the QueryResponse</p> <p>c The information model of this specification is the core catalogue schema defined in Subclause 6.3. It represents the common part of the information model which all application profiles shall support. This specification only supports 'OGCCORE' as the value of the 'responseSchema' parameter and a value of “brief”, “summary” or “full” for the value of the 'responseElements' parameter. Additional values for the responseSchema and responseElements parameters may be defined by application profiles.</p>			

Table 11 — UML attributes in SortSpec data type

Name	Definition	Data type and value	Optionality and use
sortAttName	Identifies the result set attribute type to be sorted on	Character String	Zero or one (Optional) Default is ID as defined by server
sortOrder	How the attributes are to be ordered by the sort	Code List type with allowed values of “ascending” and “descending”	Zero or one (Optional) Default is descending

Table 12 — UML attributes in SessionInfo data type

Name	Definition	Data type and value	Optionality and use
sessionID	Unique identifier for this client/server session. The value is assigned in response to an “initiate” request. All further messages within that session will contain that identifier in the sessionID parameter.	CharacterString	One (Mandatory)
destinationID	Identifies the target for this message. It can identify a server, service, or a process within a service, or a list of services to which messages may be sent in a distributed environment.	CharacterString	Zero or one (Conditional) As required by protocol binding or distributed search environment
requestID	Uniquely identifies this message. In the case of a request message, this identifier can be used to monitor and control the processing resulting from the request message.	CharacterString	Zero or one (Optional)
additionalInfo	This parameter provides a means of passing additional data that may only be relevant within the context of a specific message exchange.	CharacterString	Zero or one (Optional)

Table 13 — UML attributes in QueryExpression data type

Name	Definition	Data type and value	Optionality
queryLanguage	Specifies the predicate language and version used in a query expression	Code List, known values of “OGC_Common”, “Filter_Type-1”	One (Mandatory)
predicate	The constraint expression for selecting entries from a catalogue	CharacterString	One (Mandatory)

Table 14 — UML attributes and role “query” operation normal response

Name	Definition	Data type and value	Optionality and use
sessionInfo	The core set of parameters required in each message exchanged between a client and server operating in a session context, where these parameters support message routing and session management	SessionInfo, See Table 12	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent
resultSetID	Identifier of the Result Set generated for the query. Further query, present and cancel requests for this Result Set will supply this value through the collectionID parameter	Character String as defined by protocol binding and may be further defined by application profile	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent
resultType	How the server responded to the query request.	CodeList type with allowed values of “dataset”, “datasetcollection” and “service”	Zero or one (Optional)
retrievedData	A subset of the results of this query request, organised and formatted as specified in the presentation, messageFormat, and sortField parameters.	ReturnData Set of resource descriptions/records	Zero or one (Conditional) Include when resultType = Results
cursorPosition	Last result set resource returned for this operation request.	Positive integer	Zero or one (Conditional) Include when results returned
hits	Number of entries in the result set.	Non-negative integer	One (Mandatory)

NOTE This general model assumes that operation failure will be signalled to the client in a manner specified by each protocol binding.

7.2.4.3 “present” operation

The “present” operation is more completely specified in Table 15. provides a UML model of the “present” operation that shows the complete Discovery class with the PresentRequest and PresentResponse classes and the class they use. This UML diagram does not detail the <<CodeList>> stereotyped classes named PresentationSetType, ResponseSchema, ReturnFormat, and SortOrder. The operation request includes the attributes and association role name listed and defined in Table 16. The normal operation response includes the attributes and association role name listed and defined in Table 17.

Table 15 — Definition of “present” operation

Definition	Allows clients to retrieve selected metadata for some or all of the resources referenced in a specific previous result set or a list of resource identifiers. This operation can be used repetitively to retrieve more of the result set, each time retrieving metadata for a maximum number of the resources listed, starting at a specified position.
Receives	Specifications of sorting and of metadata to be returned, optionally including maximum number of records for which metadata is to be returned
Returns	Metadata document containing selected metadata for some or all of the specific result set, after it is sorted as specified by the client. Most of the metadata returned depends on the metadata requested and on the types of data defined by the specific Application Profile
Exceptions	Missing Parameter Value, Invalid Parameter Value , Unrecognized collection identifier.
Pre-conditions	Client has previously performed “search” operation, and the server has provided a result set identifier that the client can use to perform the present operation.
Post-conditions	Metadata document returned to requesting client, containing selected metadata for some or all of sorted result set

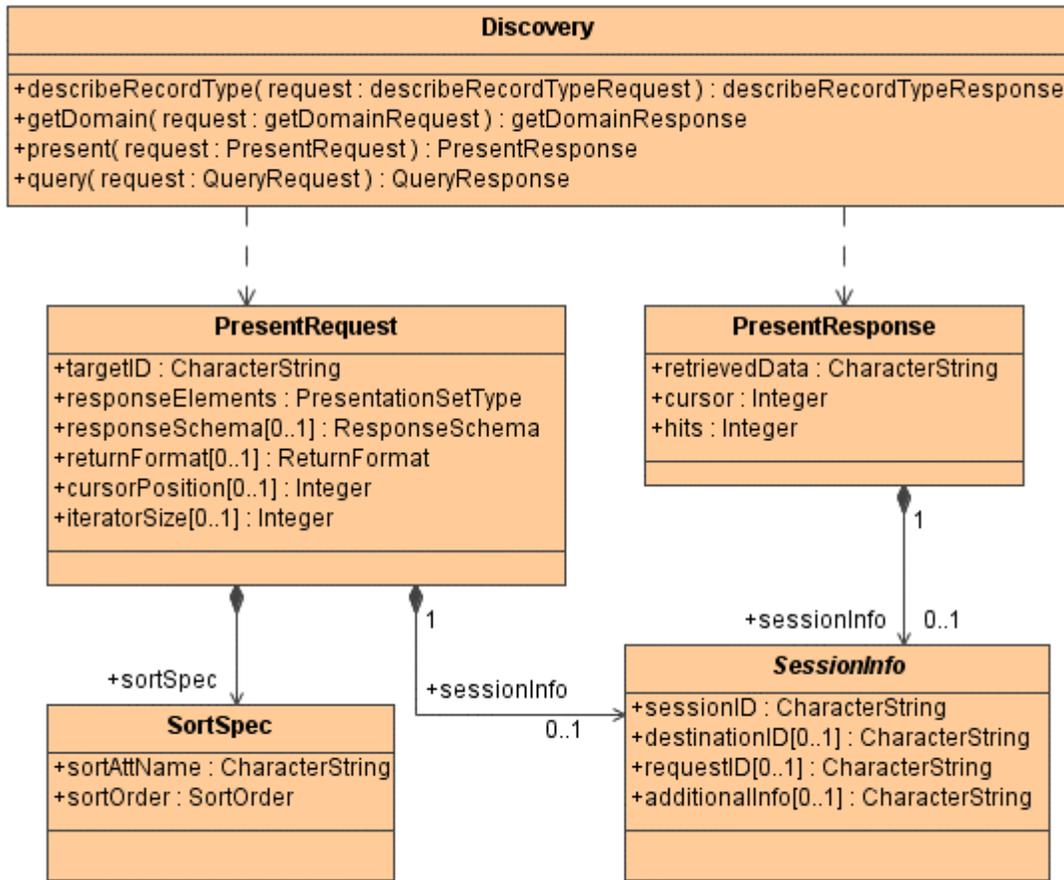


Figure 7 — “present” operation UML static model

Table 16 — UML attributes and role in “present” operation request

Name	Definition	Data type and value	Optionality and use
sessionInfo	The core set of parameters required in each message exchanged between a client and server operating in a session context, where these parameters support message routing and session management	SessionInfo, See Table 12	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent
targetID	Unique identifier of result set from which resources to be returned or of resources to be returned	Either resultSetID OR list of resource IDs	One (Mandatory)
responseElements	Specifies set name or list of metadata elements to be returned in the context of a specific metadata structure.	Either a list of elements as name/type pairs OR Codelist type named PresentationSetType with allowed values of “brief,” “summary” “full” and “browse”	Zero or one (Optional) Default is “Brief” Include when results returned
responseSchema	The name of the “well-known” or advertised (in the capabilities) schema of the response	Code List type with allowed values of “FGDC”, “ISO-19115”, ISO19139”, ANZLIC	Zero or one (Optional)
sortSpec	SortField provides sorting information to the server for formatting data returned to the client	SortSpec , See Table 11	Zero or one (Optional) Default is sorted on ID in descending order
returnFormat	Specifies format (MIME or Internet media type) for returning result set metadata.	CodeList type with allowed values of XML HTML TXT	Zero or one (Optional) Default is “XML”
cursorPosition	First result set resource to be returned for this operation request.	Positive integer	Zero or one (Optional) Default is “1”
iteratorSize	Specifies maximum number of result set resources to be returned.	Non-negative integer	Zero or one (Optional) Default is “10”

Table 17 — UML attributes and role in “present” operation normal response

Name	Definition	Data type and value	Optionality and use
sessionInfo	The core set of parameters required in each message exchanged between a client and server operating in a session context, where these parameters support message routing and session management	SessionInfo, See Table 12	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent
retrievedData	A subset of the results of this query request, organised and formatted as specified in the presentation, messageFormat, and sortField parameters.	Record type for the catalogue or collection.	One (Mandatory)
cursor	Last result set resource returned for this operation request.	Non-negative integer	One (Mandatory)
hits	Number of entries in the result set.	Non-negative integer	One (Mandatory)?

NOTE This general model assumes that operation failure will be signalled to the client in a manner specified by each protocol binding

7.2.4.4 describeRecordType operation

The describeRecordType operation is more completely specified in Table 18. provides a UML model of the describeRecordType operation that shows the complete Discovery class with the DescribeRecordTypeRequest and DescribeRecordTypeResponse classes and the class they use. The operation request includes the attributes and association role name listed and defined in Table 19. The normal operation response includes the attributes and association role name listed and defined in Table 20.

NOTE The describeRecordType operation corresponds to CG_Discovery.explainCollection operation in OGC Catalogue version 1.1.1.

Table 18 — Definition of describeRecordType operation

Definition	Allows clients to retrieve type definition(s) used by metadata of one or more registered resource types
Receives	Optional identifications of requested record type(s) and of desired format
Returns	Type definition document containing definition(s) of type(s) used by the metadata of one or more registered resource types. This type definition shall include the structure (schema), queryables, element sets, and formats of the metadata used for one or more registered resource types. The contents of the result of this operation depend on the types of metadata that can currently be used by registered resources.
Exceptions	Missing Parameter Value, Invalid Parameter Value, Nonexistent type
Pre-conditions	None
Post-conditions	Type definition document returned to requesting client, containing definition(s) of type(s) used by the metadata of one or more registered resource types

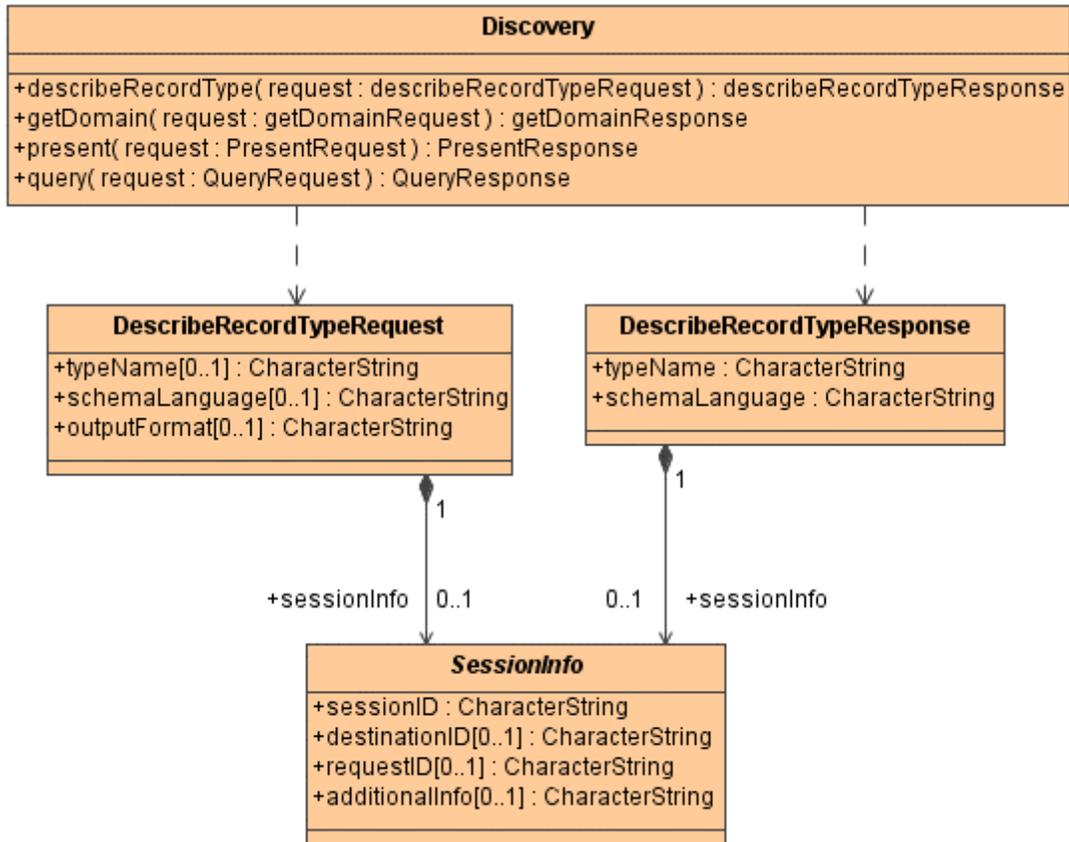


Figure 8 — describeRecordType operation UML static model

Table 19 — UML attributes and role in describeRecordType operation request

Name	Definition	Data type and value	Optionality and use
sessionInfo	The core set of parameters required in each message exchanged between a client and server operating in a session context, where these parameters support message routing and session management	SessionInfo, See Table 12	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent
typeName	Name of metadata record type(s) for which type information is to be returned	Character String type Values specified by protocol binding	Zero or more (Optional) Return all types when omitted
schemaLanguage	The schema language of the response message	Character String type Values specified by protocol binding	Zero or one (Optional) Use XML Schema when omitted
outputFormat	Document format for output	Character String type Value is MIME type	Zero or one (Optional) Use application/xml when omitted.

Table 20 — UML attributes and role in describeRecordType operation normal response

Name	Definition	Data type and value	Optionality and use
sessionInfo	The core set of parameters required in each message exchanged between a client and server operating in a session context, where these parameters support message routing and session management	SessionInfo, See Table 12	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent
typeName	Name of metadata record type for which type information is returned	Character String Values are names of metadata record types	One or more (Mandatory) Include one for each record type to be returned
schemaLanguage	The schema language used to describe the type	Character String. Non-empty Values specified by protocol binding	One (Mandatory).

7.2.4.5 getDomain operation

The getDomain operation is more completely specified in Table 21. provides a UML model of the getDomain operation that shows the complete Discovery class with the GetDomainRequest and GetDomainResponse classes and the class they use. The operation request includes the attributes listed and defined in Table 22. The normal operation response includes the attributes and association role name listed and defined in Table 23.

Table 21 — Definition of getValueDomain operation

Definition	Allows clients to retrieve the domain (allowed values) of a metadata property or request parameter at the time the request is invoked. The returned information may be static domain information, but may also be dynamic in that the allowed values are determined at runtime. The operation does a <i>best attempt</i> at returning information about a metadata property or request parameter.
Receives	Names of one or more requested metadata properties or request parameters.
Returns	Descriptions of domains of one or more requested metadata properties or request parameters
Exceptions	Missing Parameter Value, Invalid Parameter Name
Pre-conditions	None
Post-conditions	Descriptions of domains returned to requesting client, containing the domain descriptions for all the identified metadata properties or request parameters.

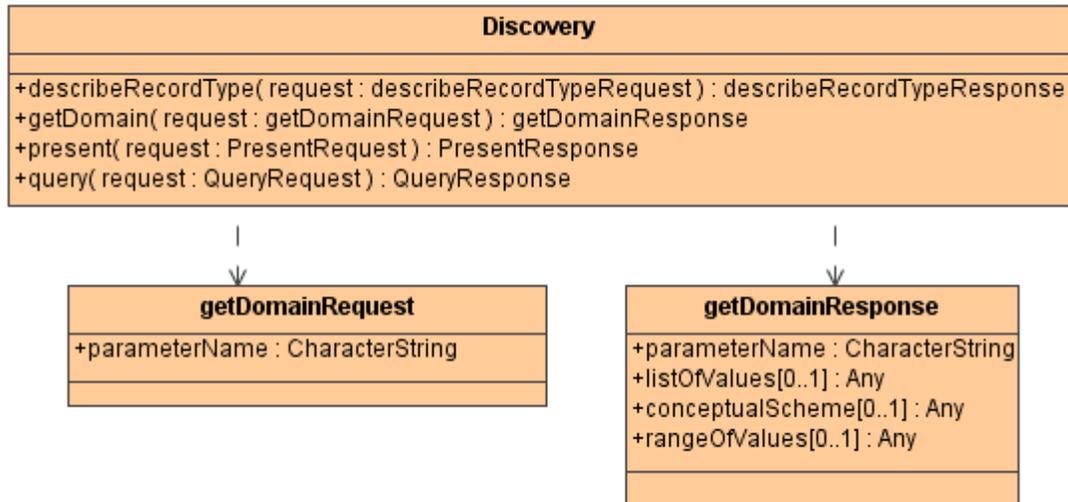


Figure 9 — getDomain operation UML static model

Table 22 — UML attribute in getDomain operation request

Name	Definition	Data type and value	Optionality
parameterName	The name of a metadata property or request parameter	Character string. Non-empty Allowed values specified by protocol binding	One (Mandatory)

Table 23 — UML attributes and role in getValueDomain operation normal response

Name	Definition	Data type and value	Optionality and use
parameterName	Name or identifier of metadata property or request parameter	Character String type, not empty	One (Mandatory)
listOfValues	Unordered list of domain values	Data type of list elements depends on the data type of the parameter whose domain is being described	Zero or one (Optional) ^a
conceptualScheme	Reference to an authoritative list of domain values for the specified parameter	Data type of list of values in the authoritative list depends on the data type of the parameter whose domain is being described	Zero or one (Optional) ^a
rangeOfValues	Range of domain values expressed by specifying a minimum and maximum value	Data type of the minimum and maximum values depends on the data type of the parameter whose domain is being described	Zero or one (Optional) ^a

^a For any single parameter, only one of listOfValues, conceptualScheme or rangeOfValues should be used to describe the value domain.

7.2.5 Session class

7.2.5.1 Introduction

The Session class allows use of interactive sessions between a client and a server, by providing four stateful operations named “initiate”, “close”, “status”, and “cancel”. This class encapsulates the operations pertaining to session management. This class has an optional association from the Catalogue Service class, in which case this class is implemented by the Catalogue Service implementation.

NOTE The four Session operations are patterned after similar Z39.50 services.

7.2.5.2 “initialize” operation

The “initialize” operation is used to establish a session context with a Catalogue Service and is more completely specified in Table 24. provides a UML model of the “initialize” operation that shows the complete Session class with the InitiateRequest and InitiateResponse classes and the class they use. The operation request includes the association role name listed and defined in Table 25. The normal operation response includes the association role name listed and defined in Table 26.

Table 24 — Definition of “initiaize” operation

Definition	Allows clients to initiate an interactive session with a server, and generates a unique identifier used to track the session
Receives	Identifier of operation request
Returns	Data describing success or failure of this operation, plus session identifier when successful
Exceptions	Missing Parameter Value, Invalid Parameter Value
Pre-conditions	Existence of a Catalogue Service instance that supports sessions
Post-conditions	Result document returned to requesting client

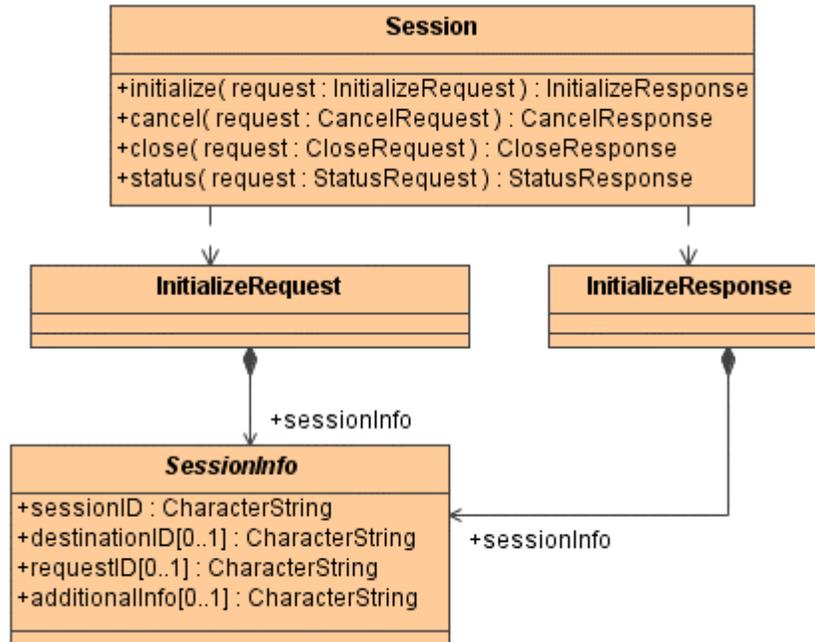


Figure 10 — “initialize” operation UML static model

Table 25 — UML role name in “initialize” operation request

Name	Definition	Data type	Optionality and use
sessionInfo	The core set of parameters required in each message exchanged between a client and server operating in a session context, where these parameters support message routing and session management	SessionInfo, See Table 12	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent

Table 26 — UML role name in “initiaize” operation normal response

Name	Definition	Data type	Optionality and use
sessionInfo	The core set of parameters required in each message exchanged between a client and server operating in a session context, where these parameters support message routing and session management	SessionInfo, See Table 12	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent

NOTE This general model assumes that operation failure will be signalled to the client in a manner specified by each protocol binding.

7.2.5.3 “close” operation

The “close” operation is more completely specified in Table 27. provides a UML model of the “close” operation that shows the complete Session class with the CloseRequest and CloseResponse classes and the class they use. The operation request includes the

association role name listed and defined in Table 28. The normal operation response includes the association role name listed and defined in Table 29.

Table 27 — Definition of “close” operation

Definition	Allows clients to terminate an interactive session with a server
Receives	Identifiers of session and of operation request
Returns	Optional identities of session or request in acknowledgement.
Exceptions	Missing Parameter Value, Invalid Parameter Value
Pre-conditions	Client has previously initiated identified session
Post-conditions	Optional result document returned to requesting client , and results sets created during the session are deleted and other resources are released.

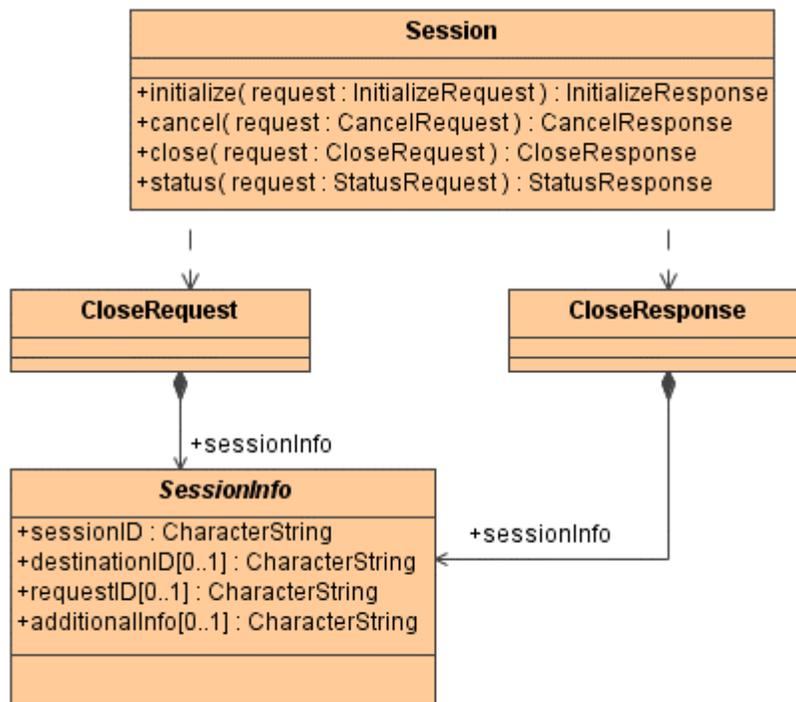


Figure 11 — “close” operation UML static model

Table 28 — UML role name in “close” operation request

Name	Definition	Data type	Optionality and use
sessionInfo	The core set of parameters required in each message exchanged between a client and server operating in a session context, where these parameters support message routing and session management	SessionInfo, See Table 12	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent

Table 29 — UML role name in “close” operation normal response

Name	Definition	Data type	Optionality and use
sessionInfo	The core set of parameters required in each message exchanged between a client and server operating in a session context, where these parameters support message routing and session management	SessionInfo, See Table 12	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent

NOTE This general model assumes that operation failure will be signalled to the client in a manner specified by each protocol binding.

7.2.5.4 “status” operation

The “status” operation is more completely specified in Table 30. provides a UML model of the “status” operation that shows the complete Session class with the StatusRequest and StatusResponse classes and the class they use. The operation request includes the attribute and association role name listed and defined in Table 31. The normal operation response includes the attribute and association role name listed and defined in Table 32.

Table 30 — Definition of “status” operation

Definition	Allows clients to retrieve current status of specified previously initiated operation in a session, either currently running or completed
Receives	Identifiers of session, previous operation request, and this operation request
Returns	Session or request IDs
Exceptions	Missing Parameter Value, Invalid Parameter Value
Pre-conditions	Client has previously initiated identified session and operation
Post-conditions	Result document returned to requesting client

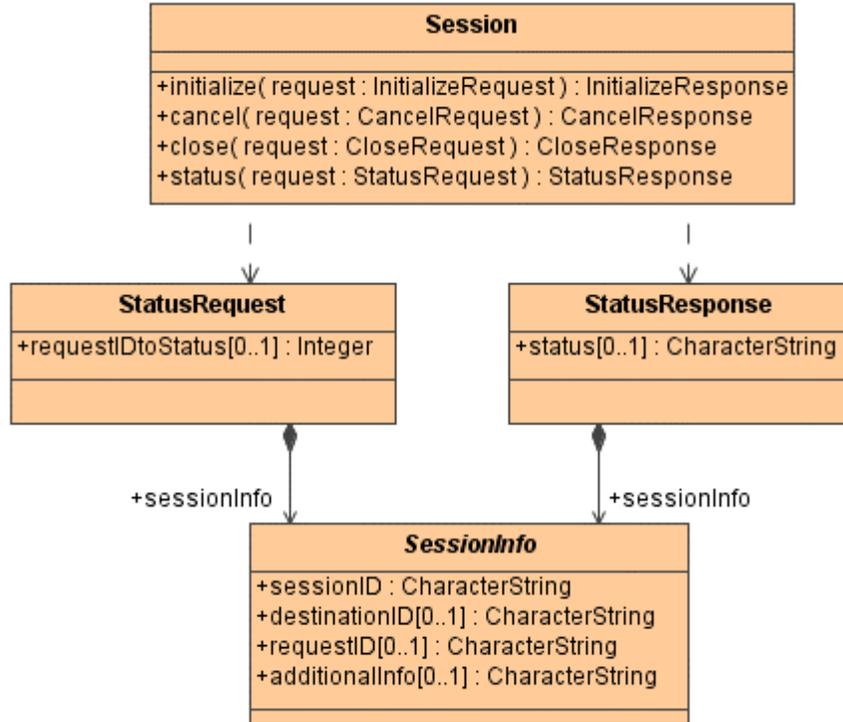


Figure 12 — “status” operation UML static model

Table 31 — UML attribute and role in “status” operation request

Name	Definition	Data type and value	Optionality and use
sessionInfo	The core set of parameters required in each message exchanged between a client and server operating in a session context, where these parameters support message routing and session management	SessionInfo, See Table 12	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent
requestIDtoStatus	Unique requestID for operation execution to be obtain status of	Integer Value of request to status	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent

Table 32 — UML attribute and role in “status” operation normal response

Name	Definition	Data type and value	Optionality and use
sessionInfo	The core set of parameters required in each message exchanged between a client and server operating in a session context, where these parameters support message routing and session management	SessionInfo, See Table 12	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent
status	Message containing the status disposition	CharacterString Values TBD	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent

NOTE This general model assumes that operation failure will be signalled to the client in a manner specified by each protocol binding.

7.2.5.5 “cancel” operation

The “cancel” operation is more completely specified in Table 33. provides a UML model of the “cancel” operation that shows the complete Session class with the CancelRequest and CancelResponse classes and the class they use. The operation request includes the attributes and association role name listed and defined in Table 34. The normal operation response includes the attributes and association role name listed and defined in Table 35.

Table 33 — Definition of “cancel” operation

Definition	Allows clients to cancel a previously initiated operation in a session, either currently running or completed. Any partial or completed result set from that operation is discarded.
Receives	Identifiers of session, previous operation request, and this operation request
Returns	Data describing success or failure of this operation
Exceptions	Missing Parameter Value, Invalid Parameter Value
Pre-conditions	Client has previously initiated identified session and operation
Post-conditions	Search request is cancelled, current result set erased (resources freed) unless otherwise specified

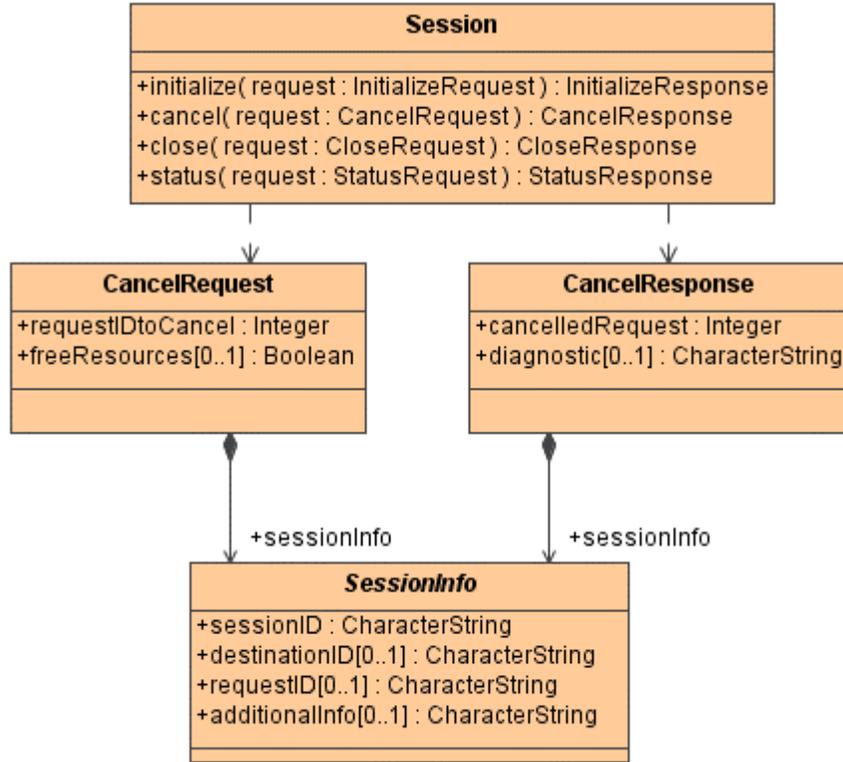


Figure 13 — “cancel” operation UML static model

Table 34 — UML attributes and role in “cancel” operation request

Name	Definition	Data type and value	Optionality and use
sessionInfo	The core set of parameters required in each message exchanged between a client and server operating in a session context, where these parameters support message routing and session management	SessionInfo, See Table 12	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent
requestIDtoCancel	Unique requestID for operation execution to be cancelled	Integer ID of request to cancel	One (Mandatory)
freeResources	If set to FALSE, any partial result set is not deleted until the client terminates the session.	Boolean	Zero or one (Optional) Free resources if omitted

Table 35 — UML attributes and role in “cancel” operation normal response

Name	Definition	Data type and value	Optionality and use
sessionInfo	The core set of parameters required in each message exchanged between a client and server operating in a session context, where these parameters support message routing and session management	SessionInfo, See Table 12	Zero or one (Conditional) Include when sessions are supported/required by protocol binding and/or result sets are persistent
cancelledRequest	Unique requestID for operation that was the target of the cancel request	Integer ID of request cancelled	One (Mandatory)
diagnostic	Text message describing the result of the cancel request	CharacterString Values TBD	Zero or one (Optional)

NOTE This general model assumes that operation failure will be signalled to the client in a manner specified by each protocol binding.

7.2.6 Manager class

7.2.6.1 Introduction

The Manager class allows a client to insert, update and/or delete catalogue content. This class has an optional association from the CatalogueService class; it is not required that a catalogue service implement publishing functionality. Two operations are provided: “transaction” and harvestResource.

The “transaction” operation allows a client to formulate a transaction, and send it to the catalogue to be processed. The transaction may contain metadata records and elements of the information model that the catalogue understands. To use the transaction operation, the client must know something about the information model that the catalogue implements.

The “harvestResource” operation, on the other hand, directs the catalogue to retrieve an accessible metadata record and processes it for inclusion in the catalogue, perhaps periodically re-fetching the metadata records to refresh the information in the catalogue. The client does not need to be aware of the information model of the catalogue when using the “harvestResource” operation, since the catalogue itself is doing the work required to process the information. The client is simply pointing to where the metadata resource to be harvested is.

7.2.6.2 “transaction” operation

The “transaction” operation is more completely specified in Table 36. Figure 14 provides a UML model of the “transaction” operation that shows the complete Manager class with the TransactionRequest and TransactionResponse classes and the classes they use. The operation request includes the attributes listed and defined in Table 37. The normal operation response includes the attributes listed and defined in Table 38.

Table 36 — Definition of “transaction” operation

Definition	Allows clients to request a specified set of “insert”, “update”, and “delete” actions on the content managed by a Catalogue Service instance.
Receives	Specification of set of “insert”, “update”, and “delete” actions, plus an optional identifier. At least one action shall be included.
Returns	A summary of the transaction results that identifies newly created entries when applicable. Most contents of the result depend on the types of data defined by the specific protocol binding and Application Profile.
Exceptions	Missing Parameter Value, Invalid Parameter Value, Transaction Failed
Pre-conditions	User is authorized to modify catalogue contents
Post-conditions	Catalogue entries are inserted, updated, and/or deleted as requested, and the integrity and consistency of catalogue contents are preserved..

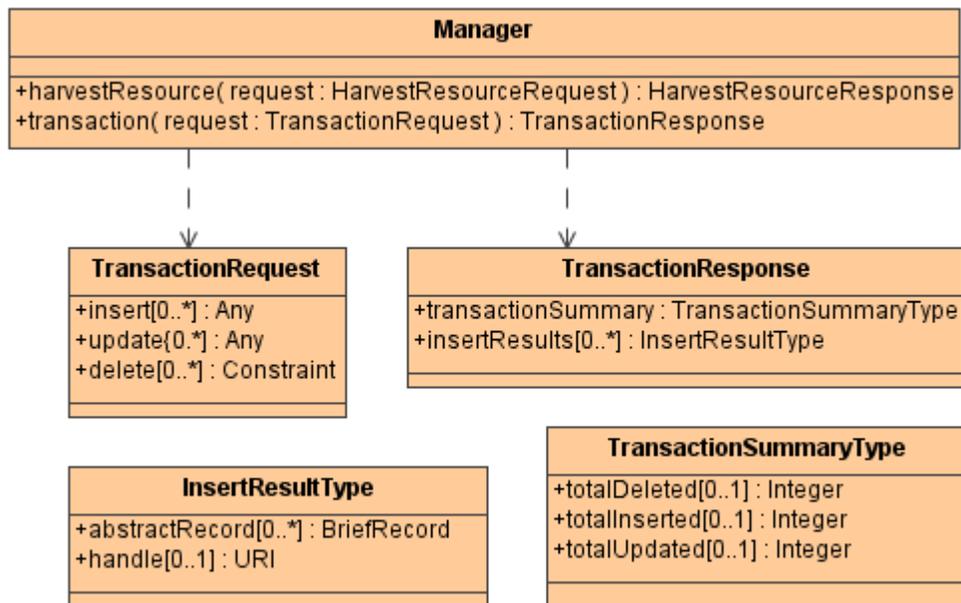


Figure 14 — “transaction” operation UML static model

Table 37 — UML attributes in “transaction” operation request

Name	Definition	Data type and value	Optionality and use
insert	The insert action is used to create new metadata records in a catalogue. Each insert action may contain one or more new metadata record instances that are to be inserted into the catalogue.	Any, a container for one or more metadata record instances The schema for metadata records is defined in the protocol binding and may be extended or redefined in an Application Profile	Zero or more (Optional) Include when client wishes to insert one or more new catalogue records
update	The update action is used to modify existing records in the catalogue. The update action contains a single new metadata record instance and a predicate that defines the set of catalogue records that will be modified. The predicate may identify zero or more records that are to be modified by the update action. The encoding of the predicate is specified in the protocol binding and may be further qualified or extended in an Application Profile.	Any, contains one instance of a metadata record that will be used to update existing records in catalog The schema of the record is defined in the protocol binding and may be extended or redefined in an Application Profile	Zero or more (Optional) Include when client wishes to modify one or more existing catalogue records
delete	The delete action is used to remove one or more records from a catalogue. The records to be removed are identified by specifying a predicate with the operation. The predicate may identify zero or more records that are to be removed from the catalogue by the delete action. The encoding of the predicate is specified in the protocol binding and may be further qualified or extended in an Application Profile.	The delete action requires a constraint predicate that identifies the records in the catalogue to be removed	Zero or one (Optional) Include when client wishes to delete one or more existing records from a catalogue

Table 38 — UML attributes in “transaction” operation normal response

Name	Definition	Data type and value	Optionality
transactionSummary	Summary of transaction results that includes the numbers of records inserted, updated, and deleted by the actions specified in the transaction	TransactionSummaryType Total number of records inserted, updated, and deleted (Integer)	One (Mandatory)
insertResults	Brief representation of a record created by the transaction, which shall include the record identifier May contain a handle that relates newly created record with the insert action that created it	InsertResultType Structure composed of brief record type (application profile or protocol binding dependent) and an optional handle	Zero or more (Optional) Include one for each record created

7.2.6.3 harvestResource operation

The harvestResource operation facilitates the retrieval of remote resources from a designated location and provides for optional transactions on the local catalogue. The

harvestResource operation is described in Table 39. Figure 15 provides a UML model of the “harvestResource” operation that shows the complete Manager class with the HarvestResourceRequest and HarvestResourceResponse classes. The operation request includes the attributes listed and defined in Table 40. The normal operation response includes the attributes listed and defined in Table 41.

Table 39 — harvestResource operation

Definition	Allows a user to request that a catalogue service attempt to retrieve a resource from a specified location, and to optionally create one or more entries for that resource. A harvest attempt may occur periodically if an interval is specified.
Receives	A request message containing the source of the resource to be harvested
Returns	An acknowledgement that a harvestRequest has been received and validated (if a responseHandler is specified) or a summary of the harvest results that identifies newly harvested records (if a responseHandler is not specified). Most contents of the result depend on the types of data defined by the specific protocol binding and Application Profile.
Exceptions	InvalidRequest, ResourceNotFound
Pre-conditions	The user is permitted to modify catalogue contents, unless the scope of the harvest does not include an insert or update transaction
Post-conditions	One or more records are harvested from a remote system and optionally new catalogue entries are created or existing entries are updated, and the integrity and consistency of the catalogue contents are preserved

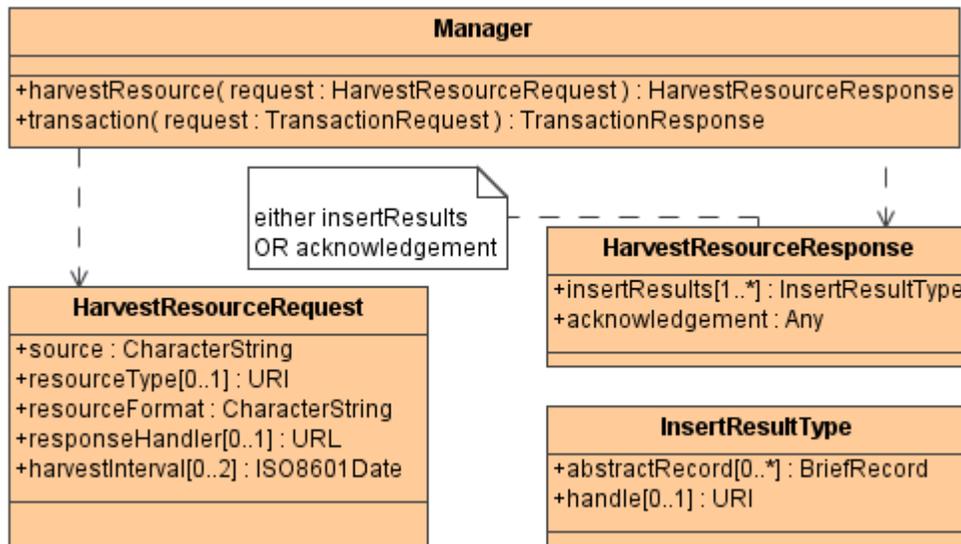


Figure 15 — harvestResource operation UML static model

Table 40 — UML attributes in harvestResource operation request

Name	Definition	Data type and value	Optionality and use
source	Location from which resource to be retrieved	URL	One (Mandatory)
resourceType	Identifier of type of resource to be harvested, if known	URI	Zero or one (Optional) If the parameter is not specified then the catalogue should determine the resourceType from the content of the message
resourceFormat	Identifier of media type indicating the format of resource to be harvested	CharacterString Value shall be a media type supported by catalogue	One (Mandatory)
responseHandler	Network location to which the response will be forwarded when operation has been completed, for asynchronous requests	URL	Zero or one (Optional) If not included, process request synchronously
harvestInterval	Time interval between harvest attempts	Period Using ISO 8601 Period syntax (e.g., P6M indicates an interval of six months)	Zero or one (Optional) If the parameter is not specified then the catalogue should harvest the resource once in response to the request.

Table 41 — UML attributes in harvestResource operation normal response

Name	Definition	Data type and value	Optionality and use
acknowledgement	Summary of transaction results, with contents depending on the protocol binding and Application Profile (e.g. total records affected by each action)	Any	One (Mandatory)
insertResults	Brief representation of a record created by the transaction, which shall include the record identifier May contain a handle that relates newly created record with the insert statement that created it	InsertResultType A structure composed of the brief record type (application profile or protocol binding dependant) and an optional handle	One or more (Mandatory) Include one for each new record created in catalogue

NOTE This general model assumes that operation failure will be signalled to the client in a manner specified by each protocol binding.

7.2.7 Brokered Access class

7.2.7.1 Introduction

The Brokered Access class allows clients to place an order for an identified registered resource, for use when that resource is a data product that is not directly accessible to clients. This class has an optional association from the Catalogue Service class, in which case this interface is implemented by the Catalogue Service implementation.

Not all resources can be accessed directly. Brokered access provides for accessing resources that are controlled. Controlled resources might include those for which one or more of the following applies:

- a) A fee is charged
- b) Have security limitations
- c) Require additional processing
- d) Are not available electronically

NOTE This class is included partially for backwards compatibility. This class may be deprecated in the future to instead use a general framework for ordering more than catalogued data sets.

7.2.7.2 “order” operation

The single “order” operation is more completely specified in Table 42. Figure 16 provides a UML model of the “order” operation that shows the complete BrokeredAccess class with the OrderRequest and OrderResponse classes. The operation request includes the attributes listed and defined in Table 43. The normal operation response includes the attributes listed and defined in Table 44.

Table 42 — Definition of “order” operation

Definition	Allows clients to order a specified product or resource, and negotiate order price and other factors with ordering service
Receives	Identifiers of desired product or resource of this order , user billing information and type of order
Returns	Order modifications, order estimates, order status
Exceptions	Missing Parameter Value, Invalid Parameter Value, ResourceNotFound, InvalidUserID
Pre-conditions	User registered with order service, resource identifiers known
Post-conditions	Response returned to requesting client

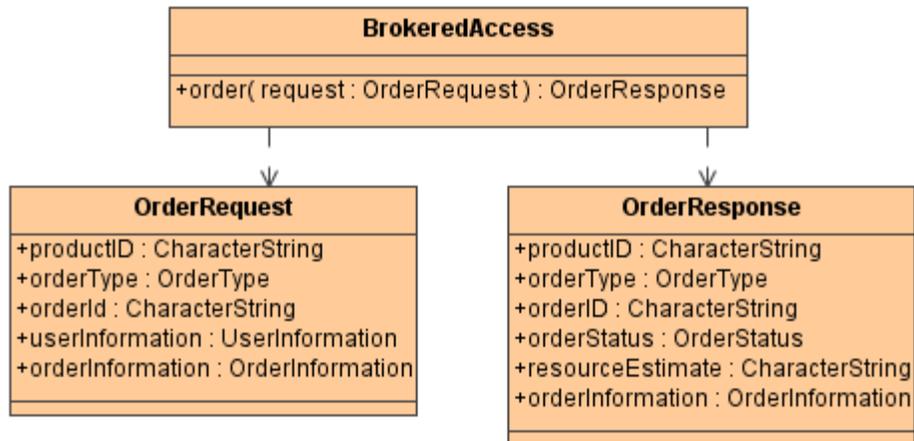


Figure 16 — “order” operation UML static model

Table 43 — UML attributes in “order” operation request

Name	Definition	Data type and value	Optionality
productID	Unique identifier for specific resource being ordered, taken from catalogue metadata	Character String type, not empty	One (Mandatory)
orderType	Type of order request	OrderType (CodeList), with allowed values of: orderEstimate, orderQuoteAndSubmit, orderMonitor, and orderCancel	One (Mandatory)
orderId	Unique identifier for an order	Character String type, not empty	One (Mandatory)
orderInformation	Specification of current order as provided by the client	OrderSpecification data type, see Annex C	One (Mandatory)
userInformation	Needed requester identification, for billing and delivery	UserInformation data type, see Annex C	One (Mandatory)

Table 44 — UML attributes in “order” operation normal response

Name	Definition	Data type and value	Optionality
productID	Unique identifier for specific resource being ordered	Character String type, not empty Values from catalogue metadata	One (Mandatory)
orderType	Type of order request	OrderType (CodeList), with allowed values of: orderEstimate, orderQuoteAndSubmit, orderMonitor, and orderCancel	One (Mandatory)
orderID	Unique identifier for this order	Character String type, not empty Value assigned by client	One (Mandatory)
orderStatus	Current status of the order	CodeList data type ^a	One (Mandatory)
resourceEstimate	Estimate of the resources needed to process and/or deliver the requested resource. Examples of these resources are time until delivery and cost.	Character String type, not empty Values TBD	One (Mandatory)
orderInformation	Specification of current order, as provided by client or modified by server during resource estimation	OrderSpecification data type, see Annex C	One (Mandatory)
<small>a Possible orderStatus values are orderBeingEstimated, orderEstimated, orderBeingQuoted, orderBeingProcessed, orderCompleted, orderNotValid, and orderCancelled.</small>			

7.3 Protocol, interface and operation specializations

The catalogue service specification includes recognized Protocol Bindings, formerly known as Implementation Profiles. Clause 11 defines the rules by which an Application Profile as a dependent specification can be written. Protocol Bindings shall implement the features present in the General Model, described in Clause 7 following the optionality expressed there. Protocol Bindings interpret the general model in the referenced implementation environment. These artifacts are discussed in detail in Subclause 7.2 and the Protocol Binding clauses of this document.

7.4 Dynamic model

7.4.1 Introduction

The Catalogue Interface defines a stateful session (a stateless interface will be added in future versions of this Implementation Specification). This subclause defines the states of the session and the allowed transitions between the states. All other state transitions are disallowed and are considered errors if exhibited by a server.

A physical server may support more than one session. Each of the sessions is independent when viewed from the interface defined by this specification.

In the state models below, a transition is typically triggered by a request. Following the messaging model introduced earlier, a Request is paired with a Response. Generally, a

transaction in this model is bounded by a request-response pair. Note that a transaction can be stasured or cancelled while it is active, i.e., before a response is issued. Once the server has sent a Response, the server treats the receipt of a StatusRequest (or CancelRequest) as an error, to which it responds gracefully. Gracefully means that the server should respond with a StatusResponse (or a CancelResponse) with a diagnostic indicating that the RequestIDtoStatus (or the RequestIDtoCancel) is not recognised. The server shall not change state in response to a StatusRequest (or CancelRequest) when the transaction is complete, i.e., a Response has been sent.

7.4.2 UML state diagram notation

The state diagrams in the following subclauses use the UML notation. Figure 17 provides a summary of the UML notation used in the following subclauses. Transitions are the paths between states. A transition will occur if the event occurs and the guard condition is true. If a transition occurs, the Action is completed prior to entering the next state.

Composite states contain multiple sub-states. Both the Sequential Composite State and the Concurrent Composite State types are used in model for the Catalogue Interface. In a Sequential Composite State only one sub-state is active at any given time. UML defines that when a transition enters Concurrent Composite State all of the sub-states are active, although some of the sub-states may remain in the Initial State. When exiting a composite state, all sub-states are exited as well.

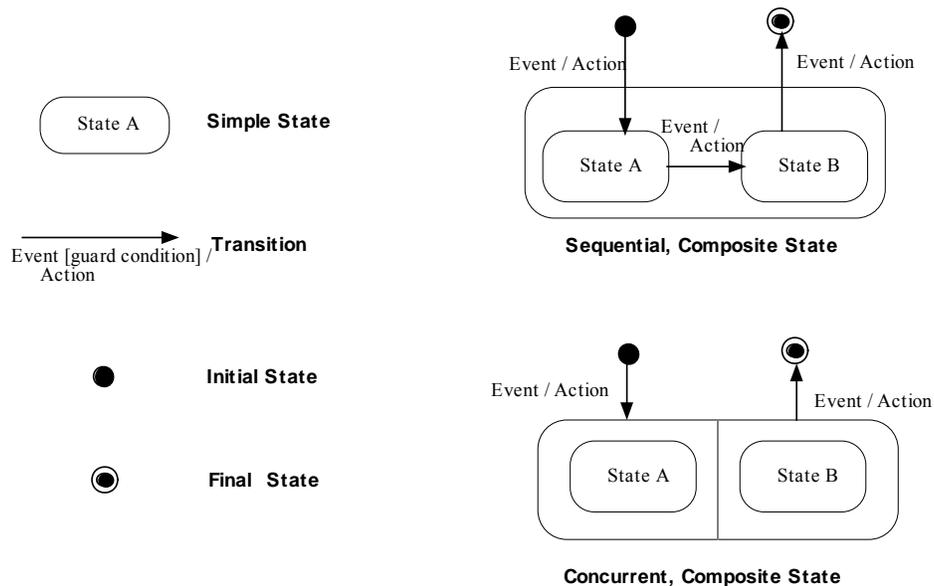


Figure 17 — UML state diagram notation

7.4.3 Catalogue server state machine

The top-level state diagram for the Catalogue Interface is shown in Figure 18. After a successful initialization, the session will be in the Main state. The Main state is a concurrent, composite state, consisting of four substates: Discovery, Access, Management, and Explain. While in the Main state, Requests (other than

TerminateRequest) may cause transitions internal to the substates. To determine what transition occurs for the various Requests, the internals of the substates must be examined. (If a server does not support interfaces associated with a substate, the substate is not present for sessions with that server. For example, if the server does not support access, then the Access Substate is not present.)

When a TerminateRequest is received, the session will transition from the Main state to the end state, ending all processing associated with the substates of Main. The Catalogue Session state diagram allows the server to end a session after a designated, configurable duration, i.e., timeout. When a session times-out, the server closes the session without notification to the client. The server shall be prepared to respond to client requests for a session that has timed out by returning the paired response containing a diagnostic indicating that the session does not exist.

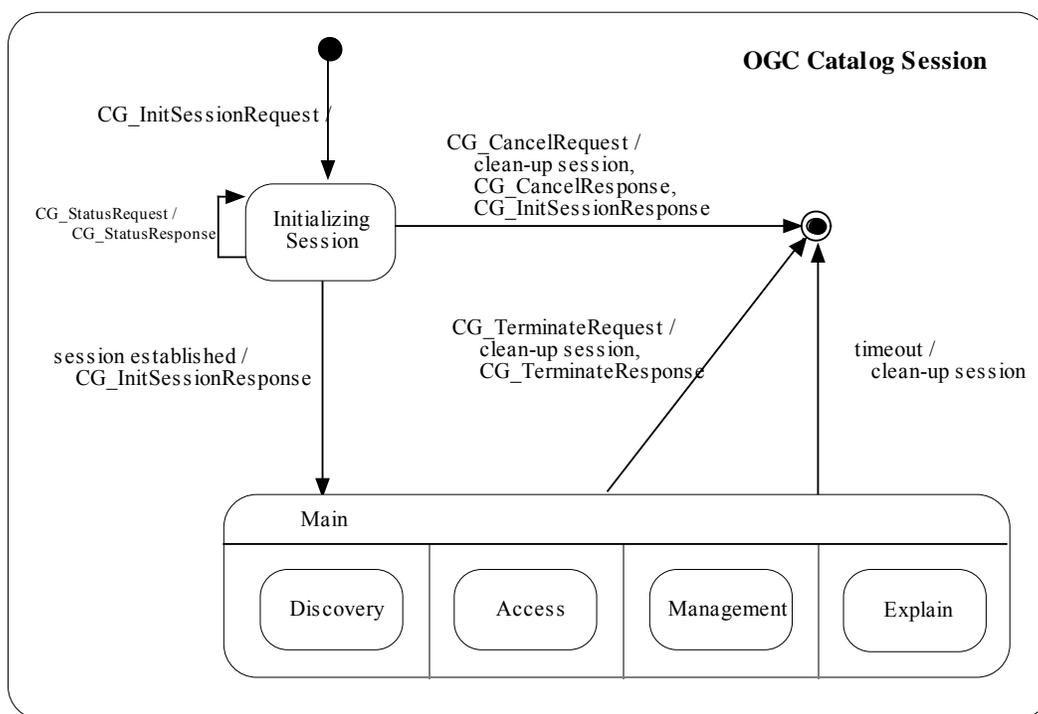


Figure 18 — Catalogue session state diagram

7.4.4 Discovery state

Two views of the Discovery State diagram are provided: Figure 19 shows an abbreviated state diagram, Figure 20 shows the complete Discovery state diagram. The abbreviated version is only provided to assist the reader in understanding the complete diagram.

A session can be in the Discovery substate, once a successful initialization has occurred at which time the Discovery substate will be in the initial state. Upon receiving any QueryRequest, the Discovery state will transition to the Processing Query state. Transitions leaving the Processing Query state are dependent upon the resultType that was requested in the QueryRequest that caused entry into the Processing Query state. The four potential values for resultType are Validate, Result Set ID, Hits, Results. If a

PresentRequest is sent by the client prior to the query completing, the session will transition to the Processing Query and Formatting Results state. The formatting of records and a PresentResponse shall occur causing a transition to the Processing Query state, prior to completing the query and sending a QueryResponse, if necessary.

When the query completes and the resultType was not Results, the state will transition to the Idle state, sending a QueryResponse unless the resultType was Validate, in which case a response has already been sent. When the resultType was Results, the state will pass to the Formatting Records for Query state, until the results are ready and a QueryResponse containing the records can be sent. While in the Idle state, a PresentRequest may be sent by the client, in which case, if a result set is present, the state will transition to the Formatting Records state, until the results are ready and a QueryResponse containing the records can be sent. As will be seen in the next diagram, there need not be a result set when the Discovery substate is Idle. If no result set is present while in the Idle state and a PresentRequest is received, the state will not transition and a PresentResponse will be returned with a diagnostic.

If a QueryRequest is received while in the Idle state, the result set for the session, if present, will be reset, and the state will transition to the Processing Query state, creating a new result set. A catalogue session can only have a single result set. (Future enhancements of the Catalogue Interface may allow multiple result sets to exist in a session.) The result set is also deleted when a TerminateRequest is received and the Catalogue Interface state, which includes the Discovery substate, transitions from Main to the end state.

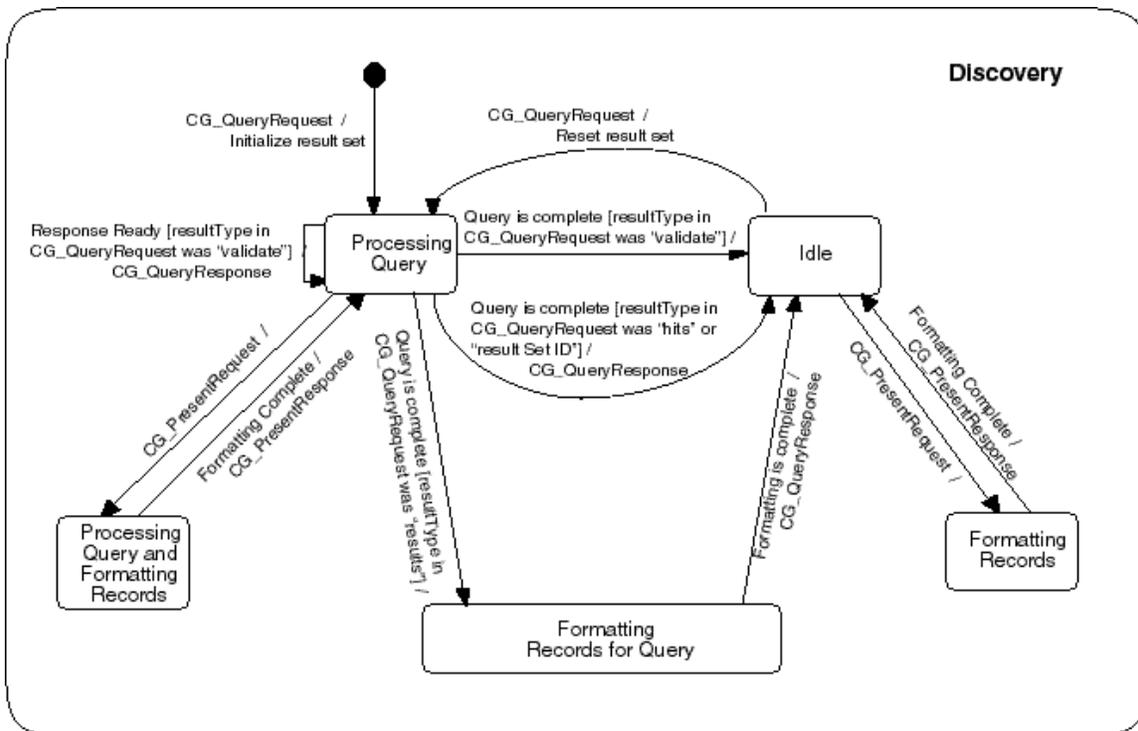


Figure 19 — Discovery state diagram (without Status and Cancel)

The complete Discovery state diagram adds StatusRequest and CancelRequest. The substates of Discovery remain the same, but additional transitions are present. If a CancelRequest is received while in the Processing Query state, the session will transition to the Idle state. Depending upon the value of the freeResources parameter in the CancelRequest, a result set may or may not exist once in the Idle State. Note that because the client sets the request ID in a request, the client knows the ID that is used in a status or cancel request.

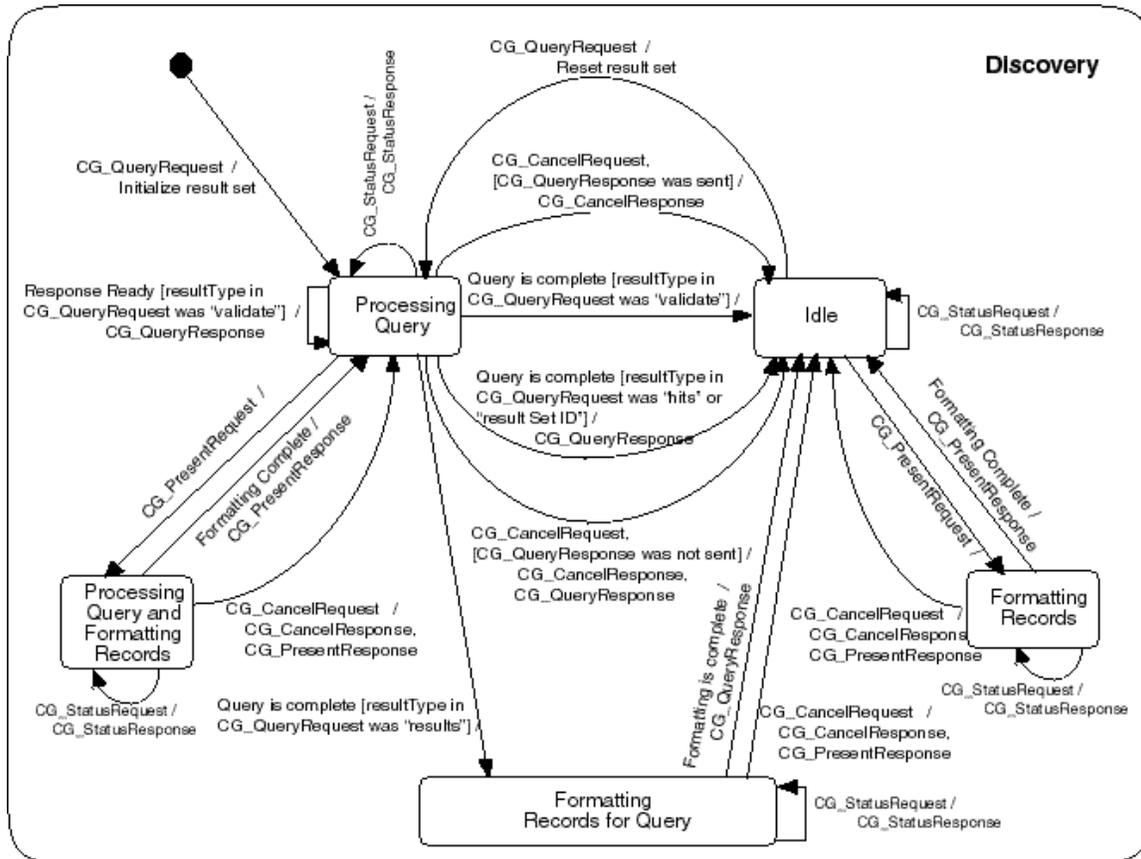


Figure 20 — Discovery state diagram (complete)

7.4.5 Access state diagram

The Access State Diagram is shown in Figure 21. A session can be in the Access substate, once a successful initialization has occurred at which time the Access substate will be in the initial state. Upon receiving a BrokeredAccessRequest, the Access state will transition to the Processing Request State. During the Processing Request State, the state of an Order may be modified based on the contents of the BrokeredAccessRequest. The state of the Order is a separate state machine; see Figure 22 and Figure 23 where order estimation and order submission are diagrammed. Transitions in the Order state may occur independent of OGC Catalogue Interface requests, e.g., order fulfilled is a transition that occurs without a BrokeredAccessRequest. The server may delete orders. The server shall be prepared to respond to client request for an order that has been deleted

by returning the paired response containing a diagnostic indicating that the order does not exist.

Once the processing of a BrokeredAccessRequest has completed a response is sent and the state transitions to Idle. Transition out of the Idle state occurs upon the client sending a BrokeredAccessRequest in which case the state transitions to Processing Request. When a TerminateRequest is received, the Catalogue Interface state, which includes the Access substate, transitions from Main to the end state also closing the Access state.

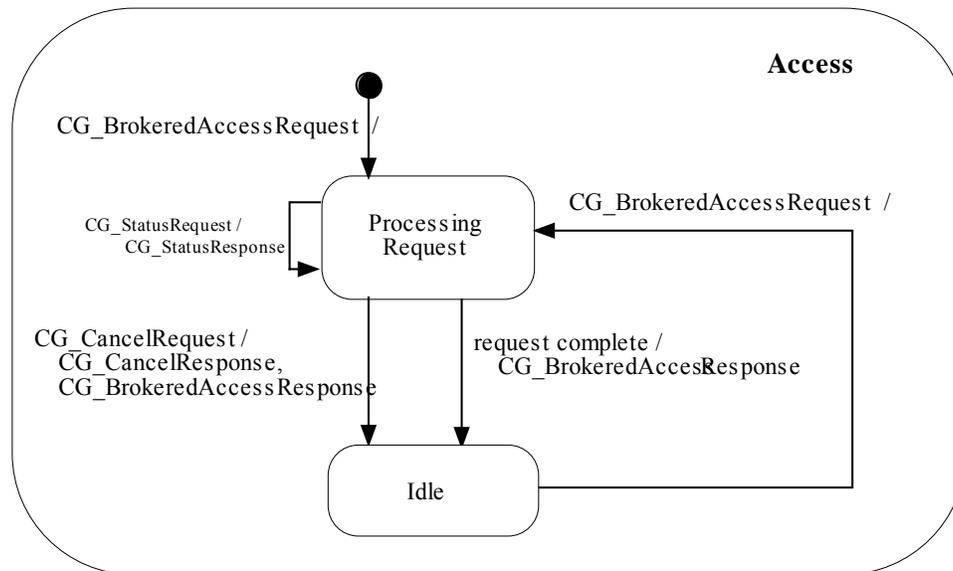


Figure 21 — Access state diagram

7.4.6 Management state

The Management State Diagram is shown in Figure 24. A session can be in the Management substate, once a successful initialization has occurred at which time the Management substate will be in the initial state. The requests are independent and paired, i.e., the response upon leaving the Processing Request state is determined by the request that caused the transition into the Processing Request State.

Once the processing of a request has completed a response is sent and the state transitions to Idle. Transition out of the Idle state occurs upon the client sending a subsequent management request in which case the state transitions to Processing Request. When a TerminateRequest is received, the Catalogue Interface state, which includes the Management substate, transitions from Main to the end state also closing the Management state.

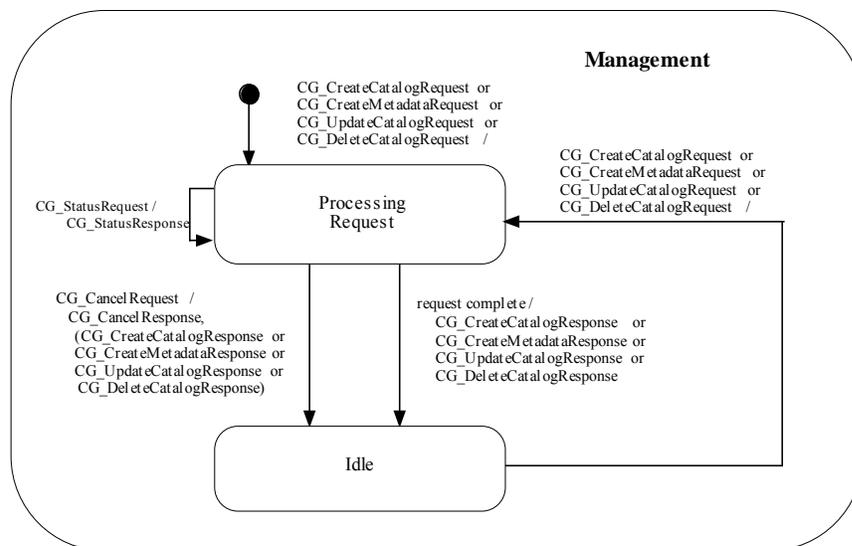


Figure 24 — Management state diagram

7.4.7 Explain state diagram

The Explain State Diagram is shown in Figure 25. A session can be in the Explain substate, once a successful initialization has occurred at which time the Explain substate will be in the initial state. The requests are independent and paired, i.e., the response upon leaving the Processing Request state is determined by the request that caused the transition into the Processing Request State.

Once the processing of a request has completed a response is sent and the state transitions to Idle. Transition out of the Idle state occurs upon the client sending a subsequent explain request in which case the state transitions to Processing Request. When a TerminateRequest is received, the Catalogue Interface state, which includes the Explain substate, transitions from Main to the end state also closing the Explain state.

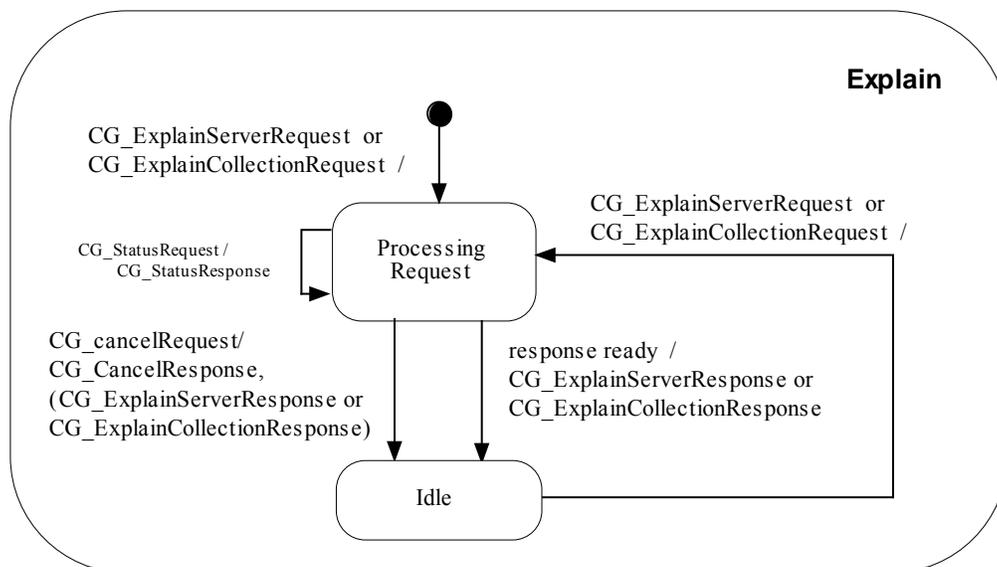


Figure 25 — Explain state diagram

8 Z39.50 protocol binding

8.1 Architecture

8.1.1 Introduction

The Z39.50 Protocol binding uses a message-based client server architecture implemented using the ANSI/NISO Z39.50 Application Service Definition and Protocol Specification [ISO 23950]. This protocol binding maps each of the general model operations to a corresponding service specified in the ANSI/NISO/ISO standard <http://lcweb.loc.gov/z3950/agency/document.html>. For compliance, clients and servers shall support Z39.50 Version 3.

At a minimum, Catalogue Services implemented using the Z39.50 protocol binding shall support the Discovery and Session operation groupings.

The Z39.50 Protocol binding offers the choice of the following transport mechanisms:

- a) Directly over TCP where services are encoded using the Basic Encoding Rules (BER) [ISO 8825].
- b) Z39.50 – Next Generation Search/Retrieve – Web Service (SRW) or its URL Access Mechanism (SRU) Version 1.0 [<http://www.loc.gov/standards/sru/>].

BER over TCP/IP is the historical implementation of Z39.50 transport, whereas SRW/SRU provides a transition strategy from legacy Z39.50/BER to a Web Services interface.

8.1.2 Supported services

Each operation specified in this protocol binding corresponds to a Z39.50 Service, and each consists of paired client requests and server responses in a session-based environment. It is possible to implement “piggy-back present” wherein session facilities are not required. The Z39.50 Services used in this protocol binding include the Init, Search, Present, Resource Control, Trigger Resource Control, Sort, Extended Services and Close.

8.1.3 Core queryable elements

The OGC Catalogue Services specification requires abstract query of a small number of metadata elements for cross-collection, cross-discipline search. These elements are described using the following Use Attribute, Structure, and Relation mappings.

Namespace prefixes denote Z39.50 profiles

[<http://lcweb.loc.gov/z3950/agency/profiles/profiles.html>].

Table 45 — Correspondence of Z39.50 Attributes to general model equivalents

General model queryable name	Use attribute	Data type	Valid relations
BoundingBox	BoundingCoordinates (geo:2060)	Coordinate String (geo:201)	Overlaps (geo:7)
CRS	Any (bib1:1016)	Word (bib-1:2), Word List (bib-1:6)	Equal (bib-1:3)
Subject	ThemeKeyword (geo:2002)	Word List (bib-1:6)	Equal (bib-1:3)
Title	Title (bib-1:4)	Word List (bib-1:6)	Equal (bib-1:3)
Abstract	Abstract (bib-1:62)	Word List (bib-1:6)	Equal (bib-1:3)
Type	Geospatial Data Presentation Form (geo:3805)	Word List (bib-1:6)	Equal (bib-1:3)
AnyText	Any (bib1:1016)	Word (bib-1:2), Word List (bib-1:6)	Equal (bib-1:3)
Identifier	CatalogIdentifier catid (bib-1:12)	Word (bib-1:2)	Equal (bib-1:3)
Modified	Metadata Date metd (bib 1:1012)	Date String (bib-1:210)	Equal (bib-1:3), Before (bib-1:14), Before-or-During (bib-1:15), During (bib-1:16), During-or-After (bib-1-17, After (bib-1:18)

8.2 General model to Z39.50 protocol binding operations mapping

Table 46 provides a mapping between general model operations and the Z39.50 Protocol binding services. The messages listed under the Z39.50 Protocol binding Service column are representative operations from the ISO 23950 standard that provide appropriate functionality. Further interpretation is provided through details in the footnotes. This table is provided to orient the programmer in correspondence with the general model but

does not provide parameter-level mapping. This table also only depicts the mandatory (Discovery) catalogue services operations and does not declare equivalence for the optional management and access operations in this version.

Table 46 — General Model to Z39.50 protocol binding operations mapping and obligation

General model operation	Z39.50 BER equivalent	SRW equivalent	Obligation
Session.initialise	initRequest ¹		Conditional, if session used
Session.close	close ²		Conditional, if session used
OGC_Service.getCapabilities	searchRequest ^{3,4}	explainRequest	Conditional, if service supports Explain
Session.status	triggerResourceControlRequest		Optional
Session.cancel	triggerResourceControlRequest		Optional
Discovery.query	searchRequest ³ and sortRequest	SearchRetrieveRequest	Mandatory
Discovery.present	presentRequest	SearchRetrieveRequest	Mandatory
Discovery.describeRecordType	SearchRequest ⁶	Inferred by XML namespace reference	Conditional, if service supports Explain
BrokeredAccess.order	ExtendedServicesRequest ⁷		Optional

¹ The following init Options are used in this protocol binding: search, present, sort, extended-services, trigger-resource-control, named result sets, and resource-control.

² Although Z39.50 permits both the client and server to initiate a Close request, for conformance with the general model, only the client is permitted to initiate a Close request. In practice, a server may terminate a session after a reasonable amount of idle client activity.

³ Note that the ResultType values of results and hits are supported in this protocol binding. The ResultType values of result set ID and validate are unsupported.

⁴ The OGC_Service.getCapabilities is implemented using a searchRequest on the Explain Database with ExplainCategory = TargetInfo and DatabaseInfo.

⁵ The CatalogEntryType and QueryScope parameters in the QueryRequest are implemented in the Z39.50 Protocol binding as external elements of the SearchRequest.

⁶ The Discovery.describeRecordType is implemented using a searchRequest on the Explain Database with ExplainCategory = TargetInfo and RetrievalRecordDetails.

⁷ Brokered Access is implemented in the Z39.50 Protocol binding using the Order Extended Service defined in Subclause 8.3.4. The Order Extended Service uses the Z39.50 Extended Service mechanism.

8.3 Z39.50 BER implementation notes

8.3.1 Introduction

Z39.50 using Basic Encoding Rules over TCP is implemented using registered profiles listed by the maintenance agency: [<http://lcweb.loc.gov/z3950/agency/profiles/profiles.html>]. These profiles indicate required operations, sets of registered Use Attributes (search fields), Relations (operators), Z39.50 datatypes, Element Sets (named sets of returned fields/elements), Preferred Syntaxes (encoding format), and information sets (metadata standards and/or schemas returned). The most relevant Community Profiles for the discovery of geospatial data resources are the Geospatial Metadata Profile, GEO, Version 2.2, and the Catalogue Interoperability Profile, CIP, Version 2.4.

In the context of a managed session, the client transmits request messages to the server and the server returns response messages to the client directly over TCP as specified in *IETF RFC 1729: Using the Z39.50 Information Retrieval Protocol in the Internet Environment* [<ftp://ftp.ietf.org/rfc/rfc1729.txt>], where all request and response messages are encoded using BER.

Figure 26 illustrates a typical set of transactions that may occur between a client and server, and between the server and its interface to an external catalogue. The client sends an `initRequest` message to the server, the external system processes the `initRequest` message by initializing a session with the client and the server returns an `initResponse` message to the client. This interaction establishes a session in which all subsequent interactions occur.

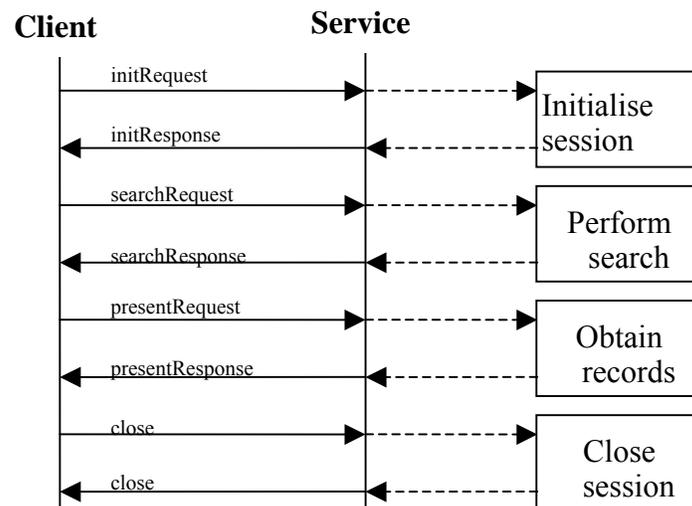


Figure 26 — Z39.50 Protocol binding Sequence Diagram

Next the client constructs a query and sends the query in the `searchRequest` message to the server. The server runs the search on the external catalogue system, and returns the requested results in the `searchResponse` message. If the search was successful, a virtual

result set is created and the client may request records from the result set using the presentRequest message. In the presentRequest, the client may request any contiguous set of records from the result set (e.g., records 10 through 20). The server returns the records to the client in the presentResponse message. The client may continue to perform additional searches and record retrievals, or may close the session with the server by sending a close message. Optionally, the server may respond with a close message.

8.3.2 Message encoding

For Z39.50 over TCP, messages are encoded using the Basic Encoding Rules (BER) from the ASN.1 specification of Z39.50 available from <http://lcweb.loc.gov/z39.50/agency/document.html>.

8.3.3 Additional search info

This subclause contains the parameters used in the "otherInfo" part of a Z39.50 searchRequest in order to implement the CatalogEntryType and QueryScope parameters in the QueryRequest of the General Model.

"otherInfo" in a SearchRequest may be used by the origin to specify the scope of a search, i.e., whether the search domain is local to a server or distributed to many servers. This is achieved using the SearchControl EXTERNAL in otherInfo. SearchControl is defined below using ASN.1 notation. If otherInfo is not provided, the type of item descriptors to be searched shall be derived from the query definition and/ or the content of the collection and the default scope of a local search shall be assumed.

The Search Control structure contains two items: itemDescriptorType which maps to CatalogEntryType and searchScope which maps to QueryScope. The CIP-Release-B-APDU {Z39.50-CIP-B-APDU 1} defines the following items:

```
SearchControl ::= SEQUENCE
{
    itemDescriptorType [1] IMPLICIT INTEGER
    {
        collectionDescriptorSearch (1),
        productDescriptorSearch (3),
        serviceDescriptorSearch (4),
        catalogDescriptorSearch (5)
    }
    searchScope [2] IMPLICIT INTEGER
    {
        localSearch (1),
        wideSearch (2)
    }
}
```

For further information, see Subclause 3.5.2.5 and Appendix E. 6.1. of Catalogue Interoperability Protocol (CIP) Specification - Release B, CEOS/WGISS/PTT/CIP-B, June 1998, Issue 2.4, Committee on Earth Observation Satellites (CEOS) (ftp://harp.gsfc.nasa.gov/incoming/fed/cip_spec24.pdf).

8.3.4 Order extended service

The Order Extended Service, which is a custom Z39.50 Extended Service, allows an origin to order products previously queried. The Order ES is presented in Table 47.

Further information describing the Order Extended Service can be found in Catalogue Interoperability Protocol (CIP) Specification - Release B, CEOS/WGISS/PTT/CIP-B, June 1998, Issue 2.4, Committee on Earth Observation Satellites (CEOS) (ftp://harp.gsfc.nasa.gov/incoming/fed/cip_spec24.pdf).

Table 47 — Order extended service

ASN.1 definition	Meaning
<pre> {Z39.50-CIP-Order-ES} DEFINITIONS ::= BEGIN IMPORTS OtherInformation, InternationalString, IntUnit FROM Z39.50-APDU-1995; CIPOrder ::= CHOICE { esRequest [1] IMPLICIT SEQUENCE { toKeep [1] OriginPartToKeep, notToKeep [2] OriginPartNotToKeep}, taskPackage [2] IMPLICIT SEQUENCE { originPart [1] OriginPartToKeep, targetPart [2] TargetPart} } </pre>	<p>The Order Extended Service uses the Z39.50 Extended Service Facility.</p>
<pre> OriginPartToKeep ::= SEQUENCE { action [1] IMPLICIT INTEGER { orderEstimate (1), orderQuoteAndSubmit (2), orderMonitor (3), orderCancel (4)}, orderId [2] InternationalString OPTIONAL, orderSpecification [3] OrderSpecification OPTIONAL, statusUpdateOption [4] StatusUpdateOption OPTIONAL, userInformation [5] UserInformation OPTIONAL, otherInfo [6] OtherInformation OPTIONAL } </pre>	<p>The OriginPartToKeep contains the following:</p> <p>action, which indicates the type of operation that is requested to be performed for the order request. The supported operations are the following:</p> <p>orderEstimate, which is used to validate and obtain the estimate of an order specification.</p> <p>orderQuoteAndSubmit, which is used to quote 1 and submit an order specification.</p> <p>orderMonitor, which is used to monitor the progress of the processing of an order request.</p> <p>orderCancel, which is used to cancel an order request.</p>

¹ The estimate for an order is approximate and non-binding, whereas the quote for an order is precise and binding.

ASN.1 definition	Meaning
	<p>orderId, which is the identifier of the order request as provided as input by the origin.</p> <p>orderSpecification, which is the specification of the order request as provided as input by the origin. Note that, in principle, the order request specified by the origin is unstructured, i.e. it contains a list of item descriptor identifiers and the order options related to them, but does not attempt to group them into packages and delivery units.</p> <p>statusUpdateOption, which indicates how the origin wishes to be kept up to date as to the status of the order processing.</p> <p>userInformation, which contains the personal user information as provided as input by the origin.</p> <p>otherInformation, which contains additional information not specified by the CIP.</p>
<pre>OriginPartNotToKeep ::= SEQUENCE { orderId [1] InternationalString OPTIONAL, orderSpecification [2] OrderSpecification OPTIONAL, userInformation [3] UserInformation OPTIONAL, otherInfo [4] OtherInformation OPTIONAL }</pre>	<p>The OriginPartNotToKeep2 contains the following:</p> <p>orderId, which is the identifier of the order request.</p> <p>orderSpecification, which is the specification of the order request.</p> <p>userInformation, which contains the personal user information.</p> <p>otherInformation, which contains additional information not specified by the CIP.</p>

² The definitions used in OriginPartNotToKeep are strictly identical to the ones provided in OriginPartToKeep. The former is used as input by the target (which may overwrite some values as appropriate) for the definition of TargetPart, whereas the latter remains unmodified and is stored in the task package. This duplication therefore allows the comparison of the order as specified by the origin (OriginPartToKeep) with the order as returned by the target (TargetPart).

ASN.1 definition				Meaning
TargetPart	::=	SEQUENCE		<p>The TargetPart contains the following:</p> <p>orderId, which is the identifier of the order request as provided as output by the target.</p> <p>orderSpecification, which is the specification of the order request as provided as output by the target. This order specification provided by the target overrides the specification provided as input by the origin in originPartNotToKeep. It contains the item descriptors and order options supplied as input, with any necessary modifications or additions, in a structured manner, i.e. the item descriptors are grouped into packages and delivery units.</p> <p>orderStatusInfo, which indicates the status of the order request being performed³.</p> <p>userInformation, which contains the personal user information.</p> <p>otherInfo, which contains additional information not specified by the CIP</p>
{				
orderId	[1]	InternationalString,	OPTIONAL,	
orderSpecification	[2]	OrderSpecification	OPTIONAL,	
orderStatusInfo	[3]	OrderStatusInfo	OPTIONAL,	
userInformation	[4]	UserInformation	OPTIONAL,	
otherInfo	[5]	OtherInformation	OPTIONAL	
}				
StatusUpdateOption	::=	CHOICE		<p>The StatusUpdateOption provides options for how the user will receive updates on the status of an extended service request. The parameters are:</p> <p>manual the user performs the status request.</p> <p>automatic where the OHS filing the order provides status updates for the user via email⁴.</p>
{				
manual	[1]	NULL,		
automatic	[2]	IMPLICIT INTEGER {	(1)}	
eMail				
}				

³ Note the difference between the operationStatus, which is provided in the ES Response, and the orderStatusInfo, which is included in the task package. operationStatus provides status information for the ES operation as a whole and indicates whether the ES operation has been performed successfully or not by the target. orderStatusInfo provides status information for the order specified in the task package and indicates the state of the order or the process being performed for an order at the LOHS.

⁴ This could be expanded in the future to include, for example, automatic update via the origin.

ASN.1 definition	Meaning
<pre> UserInformation ::= SEQUENCE { userId [1] InternationalString, userName [2] InternationalString OPTIONAL, userAddress [3] PostalAddress OPTIONAL, telNumber [4] InternationalString OPTIONAL, faxNumber [5] InternationalString OPTIONAL, emailAddress [6] InternationalString OPTIONAL, networkAddress [7] InternationalString OPTIONAL, billing [8] Billing OPTIONAL } </pre>	<p>The Userinformation structure is presented by the origin part of a request to a target. The information provided contains mandatory fields (the user identifier) and optional fields. The target will allow the Userinformation structure contents to be used as an input to the delivery specification for elements which can be altered by the user. The target will refer to the local database contents for the user and will use the contents of the database, or the Userinformation structure depending on the privilege of the user to offer alternative information. The UserInformation structure consists of the following attributes:</p> <ul style="list-style-type: none"> userId the user identifier, the identifier which the user provides as part of an InitialiseRequest. userName the full name of the user. userAddress a structure to hold the users address. telNumber the users telephone number. faxNumber the fax number for the user. emailAddress the electronic mail address for the user. networkAddress the network address to send files to electronically. For Internet addresses, the address is written in URL format to allow directories as well as domains to be specified. billing the method of payment (and hence of billing) available for the user.

ASN.1 definition		Meaning
OrderSpecification ::= SEQUENCE		
{		
orderingCentreId	[1] InternationalString,	
orderPrice	[2] PriceInfo	OPTIONAL,
orderDeliveryDate	[3] InternationalString	
		OPTIONAL,
orderCancellationDate	[4] InternationalString	
deliveryUnits	[5] SEQUENCE OF DeliveryUnitSpec,	
otherInfo	[6] OtherInformation	OPTIONAL
}		
<p>The OrderSpecification is the specification of the order request and contains the following:</p> <p>orderingCentreId, which identifies the ordering centre at which the order will be performed.</p> <p>orderPrice, which is the price for the whole order.</p> <p>orderDeliveryDate, which is the latest date at which the order can be expected to be delivered to the user.</p> <p>orderCancellationDate, which is the latest date at which the user can cancel the order.</p> <p>deliveryUnits, which contains the definition of the delivery units which compose the order.</p> <p>otherInfo, which may be used to provide additional information not specified by the CIP.</p>		
DeliveryUnitSpec ::= SEQUENCE		
{		
deliveryUnitId	[1] InternationalString	OPTIONAL,
deliveryUnitPrice	[2] PriceInfo	
		OPTIONAL,
deliveryMethod	[3] DeliveryMethod	
billing	[4] Billing	OPTIONAL,
packages	[5] SEQUENCE OF PackageSpec,	
otherInfo	[6] OtherInformation	OPTIONAL
}		
<p>The DeliveryUnitSpec contains the specification of a single delivery unit (i.e. part of an order that is delivered as a unit):</p> <p>deliveryUnitId, which is the identifier of the delivery unit.</p> <p>deliveryUnitPrice, which is the price of the delivery unit.</p> <p>deliveryMethod, which is the method with which the delivery unit is delivered to the user.</p> <p>billing, which is the method with which the user is going to be billed.</p> <p>packages, which contains the definition of the packages which compose the delivery unit.</p> <p>otherInfo, which may be used to provide additional information not specified by the CIP.</p>		

ASN.1 definition	Meaning
<pre> DeliveryMethod ::= CHOICE { eMail [1] InternationalString, ftp [2] FTPDelivery, mail [3] PostalAddress, otherInfo [4] OtherInformation } </pre>	<p>The DeliveryMethod defines the method with which a delivery unit is delivered to the user and is one of the following:</p> <p>eMail, which specifies the email address that the order will be delivered to</p> <p>ftp, which specifies that the order will be delivered via ftp, the type of transfer and the ftp address</p> <p>mail, which specifies that the order will be delivered via mail and provides the postal address</p> <p>otherInfo, which may be used to provide additional information (such as an alternative delivery method) not specified by the CIP.</p>
<pre> FTPDelivery ::= SEQUENCE { transferDirection [1] IMPLICIT INTEGER { push (0), pull (1) }, ftpAddress [2] InternationalString } </pre>	<p>The FTPMethod defines the method with which a delivery unit is delivered to the user and is one of the following:</p> <p>transferDirection, which specifies that the order will be delivered via e-mail.</p> <p>ftpAddress, which specifies that the order will be delivered via ftp.</p>
<pre> Billing ::= SEQUENCE { paymentMethod [1] PaymentMethod, customerReference [2] IMPLICIT CustomerReference, customerPONumber [3] IMPLICIT InternationalString OPTIONAL } </pre>	<p>The Billing structure⁵ contains attributes which describe the method by which a user will pay for a service, together with supporting information regarding the payment. The attributes are:</p> <p>paymentMethod indicates the method of payment used.</p> <p>customerReference is the customer provided reference for the order.</p> <p>customerPONumber is the purchase order provided by the customer for the order.</p>

⁵ The Billing structure used by the Order Extended Service is derived from the addBilling structure defined in the Item Order ES.

ASN.1 definition		Meaning
PaymentMethod ::= CHOICE { billInvoice [0] IMPLICIT NULL, prepay [1] IMPLICIT NULL, depositAccount [2] IMPLICIT NULL, privateKnown [3] IMPLICIT NULL, privateNotKnown [4] IMPLICIT EXTERNAL}, }		The PaymentMethod structure contains attributes which describe the method by which a user will pay for a service. The attributes are: billInvoice indicates that an invoice is to be sent to the user (or payee). prepay indicates that payment has already been agreed/performed. depositAccount indicates that there is a deposit account for the payment. privateKnown indicates that the payment method is private and known. privateNotKnown contain private unknown payment method information.
CustomerReference ::= SEQUENCE { customerId [1] InternationalString, accounts [2] SEQUENCE OF InternationalString }		The CustomerReference structure contains attributes which provide a customer reference for the order. The attributes are: customerId indicates the customer identifier at the LOHS. accounts is the name of the account(s) available to apply charges to on behalf of the user.
PostalAddress ::= SEQUENCE { streetAddress [1] InternationalString, city [2] InternationalString, state [3] InternationalString, postalCode [4] InternationalString, country [5] InternationalString }		PostalAddress contains the postal address for a user and consists of: streetAddress , which is the street name and number. city , which is the name of the city (or nearest city). state , which is the name of the state or county. postalCode , which is the country specific postal code. country , which is the name of the country.

ASN.1 definition		Meaning
<pre> PackageSpec ::= SEQUENCE { packageId [1] InternationalString OPTIONAL, packagePrice [2] PriceInfo OPTIONAL, package [3] CHOICE { predefinedPackage [1] PredefinedPackage, adHocPackage [2] AdHocPackage }, packageMedium [4] InternationalString, packageKByteSise [5] INTEGER, otherInfo [6] OtherInformation OPTIONAL } </pre>	<p>The PackageSpec contains the specification of a single package (i.e. part of an order that is delivered on a single medium):</p> <ul style="list-style-type: none"> packageId, which is the identifier of the package. packagePrice, which is the price of the package. package, which contains the specification of the package. The package is one of the following: <ul style="list-style-type: none"> predefinedPackage, which is a package pre-defined by the data provider. adHocPackage, which is a package constructed ad-hoc by the data provider to fulfil the order request. packageMedium, which is the medium on which the package will be delivered to the user. packageKByteSise, which contains the size of the package in kilobytes. otherInfo, which may be used to provide additional information not specified by the CIP. 	
<pre> PredefinedPackage ::= SEQUENCE { collectionId [1] InternationalString, orderItems [2] SEQUENCE OF OrderItem, otherInfo [3] OtherInformation OPTIONAL } </pre>	<p>A PredefinedPackage contains the definition of a package that is pre-defined by the data provider. A PredefinedPackage is a collection that is stored in advance (i.e. not to fulfil a specific order) on a medium and is defined as follows:</p> <ul style="list-style-type: none"> collectionId, which is the identifier of the pre-packaged collection. Must be formatted according to the naming convention for collection identifiers specified in Appendix E. orderItems, which contains the list of the order items contained in the package. otherInfo, which may be used to provide additional information not specified by the CIP. 	

ASN.1 definition		Meaning
AdHocPackage ::=	SEQUENCE OF OrderItem	An AdHocPackage is a package that is defined ad-hoc by a data provider to fulfil a specific order. An AdHocPackage contains the list of the order items contained in the package.
OrderItem ::=	SEQUENCE { productId [1] InternationalString, productPrice [2] PriceInfo OPTIONAL, productDeliveryOptions [3] ProductDeliveryOptions OPTIONAL, processingOptions [5] ProcessingOptions OPTIONAL, sceneSelectionOptions [6] SceneSelectionOptions OPTIONAL, orderStatusInfo [7] OrderStatusInfo OPTIONAL, otherInfo [8] OtherInformation OPTIONAL }	The OrderItem contains the specification of a single order item (i.e. the product that is ordered and that is to be delivered): productId , which is the identifier of the ordered product. productPrice , which is the price of the product. productDeliveryOptions , which contains delivery options for the product. processingOptions , which specifies the processing options that are to be applied on the product before delivery. sceneSelectionOptions , which specifies the selection of the scene from the whole product that is to be delivered. orderStatusInfo , which indicates the status of the order item ⁶ . otherInfo , which may be used to provide additional information not specified by the CIP.
ProductDeliveryOptions ::=	SEQUENCE { productByteSise [1] INTEGER OPTIONAL, productFormat [2] InternationalString OPTIONAL, productCompression [3] InternationalString OPTIONAL, otherInfo [4] OtherInformation OPTIONAL }	The ProductDeliveryOptions contains the specification of the options regarding the delivery of a product: productByteSise , which contains the size of the product in bytes. productFormat , which specifies the format of the product. productCompression , which specifies the compression mechanism applied to the product. otherInfo , which may be used to provide additional information not specified by the CIP.

⁶ Note the difference between the orderStatusInfo in TargetPart, which indicates the state, or the process being performed for, an order as a whole at the LOHS, and the orderStatusInfo in OrderItem, which indicates the state, or the process being performed for, a specific order item within an order at the LOHS.

ASN.1 definition	Meaning
<pre> ProcessingOptions ::= CHOICE { formattedProcessingOptions [1] EXTERNAL, unformattedProcessingOptions [2] InternationalString } </pre>	<p>The ProcessingOptions specifies the processing options that are to be applied on the product before delivery and is one of the following:</p> <ul style="list-style-type: none"> formattedProcessingOptions, which specifies the processing options according to the format specified in [ORD]. unformattedProcessingOptions, which specifies the processing options in a free-text form.
<pre> SceneSelectionOptions ::= CHOICE { formattedSceneSelectionOptions [1] EXTERNAL, unformattedSceneSelectionOptions [2] InternationalString } </pre>	<p>The SceneSelectionOptions specifies the selection of the scene from the whole product that is to be delivered and is one of the following:</p> <ul style="list-style-type: none"> formattedSceneSelectionOptions, which specifies the scene selection options according to the format specified in [ORD]. unformattedSceneSelectionOptions, which specifies the scene selection options in a free-text form.
<pre> PriceInfo ::= SEQUENCE { price [1] IntUnit, priceExpirationDate [2] InternationalString, additionalPriceInfo [3] InternationalString OPTIONAL } </pre>	<p>The PriceInfo contains the information related to the price of an item:</p> <ul style="list-style-type: none"> price, which contains the price of the item. priceExpirationDate, which specifies the latest date at which the price provided is valid (i.e. until the expiration date the origin is guaranteed that the price will not vary. However, after the expiration date the price may change). additionalPriceInfo, which may be used to provide a textual explanation when the price of a item differs from the sum of the elements which compose this item (e.g. it can be used to explain why the price of a delivery unit differs from the sum of the prices of the packages which compose the delivery unit).
<pre> OrderStatusInfo ::= SEQUENCE { orderState [1] CHOICE { staticState [1] StaticState, dynamicState [2] DynamicState }, additionalStatusInfo [2] InternationalString OPTIONAL } </pre>	<p>OrderStatusInfo describes the status of an extended service order request. The different status values are:</p> <ul style="list-style-type: none"> orderState indicates the state of the order request or the processing being performed for the order: staticState indicates the state of the order when no order request is being performed. dynamicState indicates the processing that is currently performed for an order request. additionalStatusInfo contains additional status information provided by the LOHS (e.g. to clarify the meaning of the orderState).

ASN.1 definition	Meaning
<pre> StaticState ::= [1] IMPLICIT INTEGER { orderNotValid (1), orderEstimated (2), orderCompleted (3) } </pre>	<p>StaticState describes the state of an order when no order request is active. The possible states are:</p> <p>orderNotValid indicates that the order has not been successfully validated.</p> <p>orderEstimated indicates that the order has been successfully validated and that an estimate is provided.</p> <p>orderCompleted indicates that the order has been completed.</p>
<pre> DynamicState ::= [2] IMPLICIT INTEGER { orderBeingEstimated (4), orderBeingQuoted (5), orderBeingProcessed (6), orderBeingCancelled (7), orderBeingDeleted (8) } END </pre>	<p>DynamicState describes the state of an order when an order request is active and thus being process. The possible states are:</p> <p>orderBeingEstimated the order is currently being estimated by the target order handling system.</p> <p>orderBeingQuoted the order is currently being quoted by the target order handling system.</p> <p>orderBeingProcessed the order is currently being processed by the target order handling system.</p> <p>orderBeingCancelled the order request which was previously sent to the target is being cancelled.</p> <p>orderBeingDeleted the order is being deleted.</p>

8.4 Search/Retrieve Web Service (SRW/SRU) implementation notes

SRW is the “Search/Retrieve Web Service” variant of Z39.50 that implements a simplified, stateless approach to catalogue services that preserves core functions of the Information retrieval protocol but offers them over HTTP. SRW provides Simple Object Access Protocol (SOAP) access to post and receive messages as formatted XML; SRU is a URL-based access method that employs keyword-value pairs (KVP) using GET interfaces. SRU supports a simplified predicate language that can be mapped to the OGC Common Query Language. SRW permits the use of the KVP query or Xpath expressions for search over SOAP.

SRW defines a web service combining several Z39.50 features, most notably, the Search, Present, and Sort Services. Additional features/services may be added later or defined later as new web services. The Z39.50 concepts retained in SRW include result sets, abstract access points, abstract record schemas, explain, and diagnostics. SRW features which differ from Z39.50 include the result set is named by the server rather than the client, lack of connections or sessions, and the fact that a service/server is synonymous with a database or target (services infer a single database). All SRW records are retrieved according to a single record syntax (XML) and therefore the Z39.50 concept of record syntax is not meaningful in SRW. The Z39.50 concepts of element set/specification and schema are represented by XML schemas. Explain information identifies supported access points and record schemas. Finally, XML is used in place of ASN.1 and BER.

Table 48 represents the request arguments for simplified search using SRW.

Table 48 — SearchRetrieve request parameters

Name	Type	Obligation	Description
query	xsi:string	Only one of xQuery or query shall be present	Contains a query expressed in CQL to be processed by the server. This parameter may only be present if 'xQuery' is not present.
xQuery	srw:xcqlType	Only one of xQuery or query shall be present	Contains a query expressed in XCQL to be processed by the server. This parameter may only be present if 'query' is not present. This parameter is not valid for SRU.
sortKeys	xsi:string	Optional	Contains a sequence of sort keys to be applied to the results, if any. The keys are expressed in the simple string format for sort in SRW.
xSortKeys	srw:xsortType	Optional	Contains a sequence of sort keys to be applied to the results, if any. The keys are expressed in the XML format for sort in SRW. This parameter is not valid for SRU.
startRecord	xsi:integer	Optional	The position within the sequence of matched records of the first record to be returned. The first position in the sequence is 1. The value supplied shall be greater than 0. Default value if not supplied is 1.
maximumRecords	xsi:integer	Optional	The number of records requested to be returned. The value shall be 0 or greater. Default value if not supplied is determined by the server.
recordSchema	xsi:string	Optional	The schema in which any records should be returned. The value is the URI identifier for the schema. The default value if not supplied is determined by the server.

Table 49 — SearchRetrieve response parameters

Name	Type	Obligation	Description
numberOfRecords	<i>xsi:integer</i>	Mandatory	The number of records matched by the query. If the query fails this will be 0.
resultSetId	<i>xsi:string</i>	Optional	The identifier for a result set that was created through the execution of the query.
resultSetIdleTime	<i>xsi:integer</i>	Optional	The number of seconds in which the created result set will be destroyed. The result set may be destroyed before this by the server.
records	<i>array of records</i>	Optional	A sequence of records matched by the query, or surrogate diagnostics.
diagnostics	<i>array of diagnostics</i>	Optional	A sequence of non surrogate diagnostics generated during execution.
nextRecordPosition	<i>xsi:integer</i>	Optional	The next position within the result set after the final record returned. If there are no remaining records, this value will be 0.
echoedRequest	<i>xsi:string</i>	Optional	The request parameters echoed back to the client in a simple XML form.

The Explain functionality is further described in the SRW document.

9 CORBA/IIOP protocol binding

9.1 Architecture

This clause describes the CORBA protocol binding. The intention of the CORBA protocol binding is to follow the General Model closely. This enables the building of lightweight bridges between the CORBA protocol binding and the Z39.50 Protocol binding or the HTTP protocol binding. The CORBA protocol binding is described in IDL (interface definition language) of OMG (the Object Management Group). The interfaces follow the General Model as closely as possible. Table 50 provides a mapping between general model operations and the CORBA Protocol binding services.

9.1.1 Supported services

The core of the CORBA protocol binding consists of only one interface: CatalogServices. The separate services of the General Model (Discovery, BrokeredAccess, Manager and Session) are defined in separate interfaces to reflect the General Model. They are all inherited by the central interface CatalogServices. At a minimum, Catalogue Services implemented using the CORBA protocol binding shall support the Discovery and Session interfaces as described in Table 50. The Manager- and BrokeredAccess-interfaces are optional.

The operations of CatalogServices, without exception, take a request message as an input parameter and return a response parameter. All messages are filled with standard or compound CORBA structures. Name value pairs, an optional way to transfer meta information, are borrowed from the OMG CORBA 2.3 *Dynamic Any* specification.

9.1.2 Core queryable elements

The OGC Catalogue Services specification requires an abstract query of a small number of metadata elements for cross-collection, cross-discipline search (see).

9.2 Content types (Catalogue entry types)

The content types define the type of resources a catalogue can contain (parameter "ContentType" in the minimal OGC model). In the former Catalogue Service version the allowed values for this parameter were restricted to: 'product, collection, catalog, service'. Now we have 'product' (maps to 'dataset' in the general model), 'collection' (maps to 'datasetcollection' in the general model), 'service' (as in the general model)

9.3 Supported query languages

The CORBA protocol binding supports the following query languages:

- a) CQL (Common Query Language, see Subclause 6.2.2) (mandatory)

- b) OGC Filter Encoding, Version 1.1.0 [OGC 04-095] (optional)
- c) Z39.50 Type-1 (optional)
- d) SQL3_SimpleFeature (optional)

9.4 Result set encodings

9.4.1 XML

The default encoding for returning results is XML. The General Model clause describes the common Summary set fields and XML syntax to be supported by all Protocol Bindings (see Subclause 6.3).

9.4.2 Name-Value pairs

Additional, the CORBA profile adds a Name-Value (**NV**) entry to the message format enumeration (see Subclause 9.6). Specifying NV lets the server return results as name-value pairs. Name-value pairs are specified in the OMG CORBA 2.3 DynamicAny specification, but to be complete, the definition is repeated below and in the IDL. Usage of NameValuePair specification from OMG CORBA 2.3 DynamicAny aligns Catalogue Services CORBA Profile with revision 1.1 (draft 3) of Simple Feature Access for CORBA.

```
enum MessageFormat {XML, HTML, TXT, NV};
module DynamicAny
{
    struct NameValuePair
    {
        string name;
        any value;
    };

    typedef sequence<NameValuePair> NameValuePairSeq;
};
```

So if the server gives the results back as XML in the next example:

```
<?xml version="1.0"?>
<!DOCTYPE Metadata SYSTEM "min.dtd" >
<Metadata>
  <Title>Countries of Europe</Title>
  <Abstract>This dataset contains the countries of Europe</Abstract>
  <GeographicBoundingBox>
    <westBoundLongitude>-24.17</westBoundLongitude>
    <eastBoundLongitude>40.71</eastBoundLongitude>
    <northBoundLatitude>71.26</northBoundLatitude>
    <southBoundLatitude>27.63</southBoundLatitude>
  </GeographicBoundingBox>
</Metadata>
```

Name-value pair results are as follows:

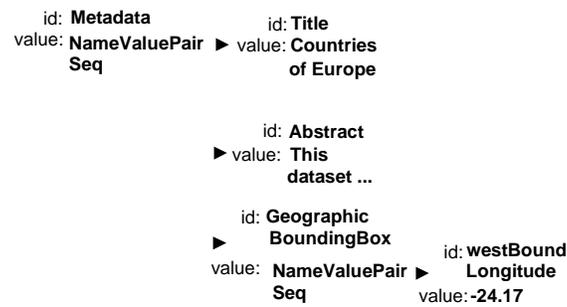


Figure 27 — Name-value pair results

The advantage is that pure CORBA environments do not have to parse the XML to get the results. They receive them in a suitable general structure. If the CORBA server is combined with another type of client, e.g. a Web client, then probably XML (the default) will be preferred.

The **any value** member can contain any type: standard types as long, double, string, types as NameValuePair or NameValuePairSeq (this gives the possibility to create recursive structures) or user-defined types.

9.5 General model to CORBA protocol binding operations mapping

Table 50 provides a mapping between general model operations and the CORBA Protocol binding services. The CORBA Protocol binding messages are defined in Subclause 9.6. The messages listed under the CORBA Protocol binding Equivalent column are the operations that provide appropriate functionality used in the CORBA interfaces. Further interpretation is provided through details in the footnotes. This table is provided to orient the programmer in correspondence with the general model but does not provide parameter-level mapping.

Table 50 — General Model to CORBA protocol binding operations mapping and obligation

General model operation	CORBA protocol binding equivalent	Obligation
Session.initialize	OGC_StatefulService.initSession	Conditional, if session used
Session.close	OGC_StatefulService.terminateSession ²	Conditional, if session used
OGC_Service.getCapabilities	OGC_Service.explainServer	Mandatory
Session.status	OGC_StatefulService.status	Optional
Session.cancel	OGC_StatefulService.cancel	Optional
Discovery.query	CG_Discovery.query ³	Mandatory
Discovery.present	CG_Discovery .present	Mandatory
Discovery.describeRecordType	CG_Discovery.explainCollection	Conditional, if service supports Explain
Discovery.getDomain	CG_Discovery.getDomain	Conditional, if service supports getDomain
Manager.transaction	CG_Manager.transaction	Optional
Manager.harvestRecords	--	--
BrokeredAccess.order	CG_Access.brokeredAccess	Optional

² Although the CORBA protocol binding permits both the client and server to initiate a terminateSession request, for conformance with the general model, only the client is permitted to initiate a terminateSession request. In practice, a server may terminate a session after a reasonable amount of idle client activity.

³ Note that the ResultType values “results”, “hits”, “resultSetID” and “validate” are supported in this protocol binding.

9.6 Interface definition - IDL

9.6.1 Introduction

This subclause describes the CORBA IDL. It first describes enumerations and then structures, unions, and messages, respectively. It concludes with a description of the CatalogServices interface, the core of the profile, and other interfaces.

All enumerations, structures, unions, messages and interfaces are part of the OGC_CatalogService module. Module names have to be harmonized across all OGC CORBA specifications and have to be prefixed by opengis.org.

```
#pragma prefix "opengis.org"
module OGC_CatalogService
{
...
};
```

Throughout the module OGC_CatalogService the IDL types `wstring` and `wchar` are used instead of `string` and `char` to allow usage of different character codesets (other than Unicode) for internationalization (i18n).

In CORBA IDL type definitions for sequences containing different element datatypes are used to avoid anonymous sequences in IDL mappings for some programming languages.

9.6.2 Enumerations

Enumerations can be modelled by a direct translation of all code-lists of the General Model. The following enumerations are borrowed literally:

```
enum AttributeCategory {queriable, presentable, both};
enum CatalogEntryType {product, collection, service};
enum CharacterSet {ASCII, UniCode, ShiftJIS};
enum PredefinedPresentationType {full, summary, brief};
enum QueryLanguage {OGC_Common, OGC_Filter, Z3950_TypeOne
SQL3_SimpleFeature };
enum QueryScope {distributed, locale};
enum ResultType {validate, resultSetID, hits, results};
enum SortOrder {ascending, descending};
enum Status {success, successResultsAvailable, processingNormal,
processingQueued, processingPausedOrSuspended, failure,
failureAccessDenied};
```

Additional, the CORBA profile adds an `NV` entry to the message format enumeration (see Subclause 9.4.2):

```
enum MessageFormat {XML, HTML, TXT, NV};
```

9.6.3 Structures and unions

Most of the structures and unions from the General Model can be translated directly into CORBA structs and unions. Here the `collectionID` of the general model is translated as follows:

```
union CollectionName
  switch(long)
  {
    case 1 : wstring collectionID;
    case 2 : wstring collectionName;
  };
```

A special capability is present in `QueryExpression` in the CORBA Profile that allows passing of parameters that can't be converted to strings but shall be bound to variables in string *theQuery* (e.g. "?" in JDBC). For example, references or handles for metadata retrieved from related collections in previous queries. *queryParameters* might contain a `NameValuePairSeq` or non ASCII XML Data. The additional member aligns Catalogue Services query facilities with respective Simple Feature Access for CORBA query facilities.

```

struct QueryExpression
{
    wstring theQuery;
    wstring theNamespace;
    QueryLanguage theLanguage;
    any queryParameters;
};

```

To allow for globally unique *sessionID* a long long (Long) is used as datatype instead of long (Integer).

```

struct RequestID
{
    long long sessionID;
    long counter;
};
struct SortField
{
    wstring attributeName;
    SortOrder sortOrder;
};

```

The CORBA protocol binding specifies an **any** structure member for the retrievedData in a way that strings (e.g. xml) or name-value pairs or sequences can be stored.

```

struct ReturnData
{
    MessageFormat encoding;
    any payload;
};

```

The responseElements in the general model specify a set name or a list of metadata elements to be returned in the context of a specific metadata structure. In CORBA IDL this is specified by a union which can represent a sequence of attribute names or alternative a PredefinedPresentationType.

```

typedef sequence<wstring> StringSeq;
union PresentationDescription
    switch(long)
    {
        case 1 : StringSeq attributes; // TupleType in GM
        case 2 : PredefinedPresentationType presentationType;
    };

```

The SchemeID structure uses a structure member Schema. This is in the CORBA profile defined as a sequence of name-value pairs from the OMG CORBA 2.3 DynamicAny module. All names, types, and used sequences can be specified in name-value pairs. A schema, tuple-type or a dictionary is not needed here.

```

typedef DynamicAny::NameValuePairSeq Schema;
struct SchemaID
{
    wstring schemeName;
    Schema schema;
};

```

```
};
```

The `getDomain` operation in the general model specifies descriptions of domains of one or more requested metadata properties or requested parameters in their responseElements which are implemented as follows:

```

    // Enumeration for the type of the metadata property or request
parameter
enum DomainType
    { domainTypeString, domainTypeDate, domainTypeInteger,
domainTypeReal,
        domainTypeSpatial };

    // Enumeration for the type how valid values can be described
enum DomainValueType
    { valueRef, valueRange };

// valid range, composed of lower and upper boundary value
struct ValueRange
{
    wstring lowerValue;
    wstring upperValue;
};
typedef sequence<ValueRange> ValueRangeSeq;

// Definition of the real type by totalDigits and fractionDigits
struct RealTypeDef
{
    long totalDigits;
    long fractionDigits;
};

// type definition of the metadata property or request parameter
union DomainDef
    switch(DomainType)
    {
        case domainTypeString : long length;
        case domainTypeDate : wstring dateFormat;
        case domainTypeInteger : long totalDigits;
        case domainTypeReal : CG_RealTypeDef realTypeDef;
        case domainTypeSpatial : long dimensions;
    };

// definition of valid values (range, value list) of the type
union DomainValueDef
    switch(DomainValueType)
    {
        case valueRef : StringSeq valueList;
        case valueRange : ValueRangeSeq rangeList;
    };

// definition of domain values (typical value, valid values, value
name, value description) of the type
struct DomainValue
{

```

```

    wstring value;
    DomainValueDef valuesDef;
    wstring title;
    wstring description;
    wstring metadataURL;
};
typedef sequence<DomainValue> DomainValueSeq;

// Descriptions of domain of one requested metadata properties or
request parameter
struct Domain
{
    wstring attributeName;
    DomainDef domainDef;
    DomainValueSeq domainValue;
};
typedef sequence<Domain> DomainSeq;

```

9.6.4 Definitions for brokered access

The General Model defines some code-lists and structures for brokered access. These definitions are directly translated into their CORBA counterparts:

```

enum BrokeredAccessType {orderEstimate, orderQuoteAndSubmit,
    orderMonitor, orderCancel};
struct OrderItem
{
    // Note: datatypes not provided by GM
    any productID;
    any productPrice;
    any productDeliveryOptions;
    any processingOptions;
    any sceneSelectionOptions;
};
struct OrderSpecification
{
    // Note: datatypes not provided by GM
    any orderCentreID;
    any orderPrice;
    any orderDeliveryDate;
    any orderCancellationDate;
    any deliveryMethod;
    any package;
};
enum OrderStatus {orderBeingEstimated, orderEstimated,
    orderBeingQuoted, orderBeingProcessed,
    orderCompleted, orderNotValid, orderCancelled};
enum PackagingType {predefinedPackage, adhocPackage};
struct PackageSpecification
{
    // Note: datatypes not provided by GM
    any packageId;
    any packagePrice;
    PackagingType package;
    any packageMedium;
};

```

```

    long packageSize;
};
enum PaymentMethod {credit, cash, purchaseOrder};
enum StatusUpdateType {manual, automatic};
struct UserInformation
{
    wstring userName;
    wstring userAddress;
    wstring phoneNumber;
    wstring faxNumber;
    wstring emailAddress;
    wstring netAddress;
    PaymentMethod paymentMethod;
};

```

9.6.5 Capabilities

The capabilities in the General Model are designed with inheritance. In CORBA designing capabilities as interfaces can reflect this, but this is not useful. Capabilities like messages (see below) have to be transferred over the network. Therefore, they are defined as either type definitions or structures.

```

typedef boolean AllSupportedRequest;
typedef boolean Defaults;
struct DefaultTimeOut
{
    unsigned long long timeOut;
    // used to be OGC_Basic::UomTime, but OGC_Basic is no longer
    maintained as normative part of the Catalogue Services Specification
};
typedef boolean Explain;
struct Messaging
{
    CharacterSet characterSet;
    MessageFormat messageFormat;
};
struct Query
{
    wstring version;
    CharacterSet characterSet;
    QueryLanguage queryLanguage;
};
struct Session
{
    wstring language;
    wstring catalogSpecificationVersion;
    CharacterSet characterSet;
};
struct SoftwareInformation
{
    wstring vendor;
    wstring SVversionNumber;
    wstring IFversionNumber;
};
typedef sequence<CollectionName> SupportedCollections;

```

To be able to make a sequence of different capabilities, a union Capability is created, encompassing all derived capabilities.

A union normally has a discriminator. This can be a **long** value, but this is generally not preferred because you have to remember the value indicating the intended capability. Therefore, an enumeration of capabilities is included in the CORBA profile.

```
enum CapabilityType
{ ctAllSupportedRequest, ctDefaults, ctDefaultTimeOut, ctExplain,
ctMessaging, ctQuery, ctSession, ctSoftwareInformation,
ctSupportedCollections
};
union Capability
switch(CapabilityType)
{
case ctAllSupportedRequest : AllSupportedRequest
allSupportedRequest;
case ctDefaults : Defaults defaults;
case ctDefaultTimeOut : DefaultTimeOut timeOut;
case ctExplain : Explain explain;
case ctMessaging : Messaging messaging;
case ctQuery : Query query;
case ctSession : Session session;
case ctSoftwareInformation : SoftwareInformation
softwareInformation;
case ctSupportedCollections : SupportedCollections
supportedCollections;
};
```

9.6.6 General messages

The General Model is a message-based model, where messages are designed in the form of a class hierarchy. In CORBA IDL, the messages are translated as structs. Writing them in the form of interfaces is not useful. In CORBA, the objects (instances of interfaces) stay on a remote server machine and are referred to by a client machine. They are not transferred over the network. This is definitely not the intention for messages.

All messages have the same form as the messages described in the General Model. However, messages in the form of structs cannot inherit from each other in CORBA. Therefore the Message class is also included in the CORBA profile and a member of all other messages, called 'base'.

```
struct Message
{
long long sessionID;
wstring destinationID;
RequestID requestID;
wstring additionalInfo;
};
```

All other messages, which in the General Model inherit from Message, have in the CORBA profile the Message as a structure member. The next messages do not add extra

structure members. Alternatively, they might have been modelled by a typedef. But to be consistent with the rest of the messages these message have **base** as a structure member.

Note that the response in the General Model also contains a string structure member **diagnostic**. This parameter is not specified in the CORBA profile. Error handling will be handled by exceptions, the standard CORBA facility. Exceptions are described below. WWW/CORBA bridges can catch these exceptions and convert them into diagnostic info if necessary.

```
struct InitSessionRequest
{
    Message base;
};
struct InitSessionResponse
{
    Message base;
};
struct TerminateRequest
{
    Message base;
};
struct TerminateResponse
{
    Message base;
    Status status;
};
```

The status and cancel messages add a few structure members in addition to the **base** structure member.

```
struct StatusRequest
{
    Message base;
    RequestID requestIDtoStatus;
};
struct StatusResponse
{
    Message base;
    RequestID requestIDtoStatus;
    Status status;
};
struct CancelRequest
{
    Message base;
    RequestID requestIDtoCancel;
    boolean freeResources;
};
struct CancelResponse
{
    Message base;
    Status status;
    RequestID canceledRequest;
};
```

The explain server messages add sequence of capabilities to the base message. The capability-type sequence can be filled with capability-types to specify which capabilities are requested from the server. The server responds with reporting each capability in a sequence of capabilities.

```
typedef sequence<CapabilityType> CapabilityTypeSeq;
struct ExplainServerRequest
{
    Message base;
    CapabilityTypeSeq capabilities;
};
struct ExplainServerResponse
{
    Message base;
    CapabilityTypeSeq capabilities;
};
```

9.6.7 Discovery messages

There are four request/response message pairs in the discovery service. To enhance distributed searching, an additional structure member for the query message is provided. This member is not included in the General Model. This structure member **asynchronous** can be set to **true** to force asynchronous searching. The query method will return immediately, setting structure member **hits** in the response to zero. Query results can be retrieved later on, when the query is ready. The progress of the query can be examined with the status messages. The query can be cancelled with the cancel messages.

NOTE 1 This asynchronous behaviour is only specified for the query request message. All other operations (e.g. init, terminate, status, cancel, explain, present) are not considered as time-consuming and return immediately after processing.

Another structure member, **maxLevel**, is added to have more control in the range of the distribution. If one catalogue contains another one, that other one contains a third one, and so on, you will possibly specify that only two levels of sub-catalogs will be searched. Setting the **maxLevel** member to two will force this. Setting **maxLevel** to -1 forces searching all sub-catalogs.

NOTE 2 If the **queryScope** is **Local** there is no distributed search at all.

```
typedef sequence<SortField> SortFieldSeq;
struct QueryRequest
{
    Message base;
    QueryExpression queryExpression;
    ResultType resultType;
    long iteratorSize;
    long cursor;
    MessageFormat returnFormat;
    PresentationDescription presentation;
    SortFieldSeq sortField;
    QueryScope queryScope;
    CollectionName collectionID;
    CatalogEntryType resourceType;
```

```

    boolean asynchronous;
    long maxLevel;
};
struct QueryResponse
{
    Message base;
    ReturnData retrievedData;
    CollectionName resultSetID;
    Status status;
    long hits;
    long cursor;
};
struct PresentRequest
{
    Message base;
    CollectionName resultSetID;
    PresentationDescription presentation;
    SortFieldSeq sortField;
    MessageFormat returnFormat;
    long iteratorSize;
    long cursor;
};
struct PresentResponse
{
    Message base;
    ReturnData retrievedData;
    long cursor;
    long hits;
    Status status;
};
struct ExplainCollectionRequest
{
    Message base;
    AttributeCategory attributeCategory;
    CollectionName collectionID;
    MessageFormat returnFormat;
};
struct ExplainCollectionResponse
{
    Message base;
    CollectionName collectionID;
    SchemaID dataModel;
    Status status;
};
struct GetDomainRequest
{
    Message base;
    StringSeq attributes;
};
struct GetDomainResponse
{
    Message base;
    DomainSeq attributeDomains;
};

```

9.6.8 Management messages

The General Model defines messages for managing catalogs. These messages are translated to the CORBA profile literally.

```
typedef sequence<DynamicAny::NameValuePairSeq> NameValuePairSeqSeq;

// Messages for managing functions

// Insert-Structure for inserting metadata (payload) to a collection
// of a defined catalogType in a predefined format (encoding)
struct InsertMetadata
{
    CollectionName collectionID;
    CatalogEntryType catalogType;

    MessageFormat encoding;

    // the data to insert, e.g. a list of NV-Pairs or ISO19139-
    // XML
    any payload;
};

// Update-Structure for updating a certain set (defined by the
// queryExpression) of metadata (payload) of a
// collection/catalogType in a predefined format (encoding)
struct UpdateMetadata
{
    CollectionName collectionID;
    CatalogEntryType catalogType;

    MessageFormat encoding;

    QueryExpression queryExpression;

    // the data to insert, e.g. a list of NV-Pairs or ISO19139-
    // XML
    any payload;
};

// Delete-Structure for deleting a certain set (defined by the
// queryExpression) of metadata of a
// collection/catalogType in a predefined format (encoding)
struct DeleteMetadata
{
    CollectionName collectionID;
    CatalogEntryType catalogType;

    QueryExpression queryExpression;
};

// Enumeration for the type of a single manipulation in a
// transaction
enum MetadataManipulationType
{ ctInsert, ctUpdate, ctDelete };

```

```

// Structure for the content of a single manipulation in a
// transaction
union MetadataManipulation
    switch(MetadataManipulationType)
    {
        case ctInsert : InsertMetadata insertData;
        case ctUpdate : UpdateMetadata updateData;
        case ctDelete : DeleteMetadata deleteData;
    };

typedef sequence<MetadataManipulation> MetadataManipulationSeq;

// Struct of a transaction, as a sequence of single manipulations
struct Transaction
{
    MetadataManipulationSeq manipulations;
};

struct TransactionRequest
{
    Message base;
    Transaction transactionData;
};

struct TransactionResponse
{
    Message base;
    Status status;
    // value of -1 means not calculated
    long totalInserted;
    long totalUpdated;
    long totalDeleted;
    // List of newly generated catalogue entry identifiers
    // assigned to the new catalogue entry instances
    NameValuePairSeqSeq newKeyList;
};

```

9.6.9 Access messages

The General Model specifies direct access and brokered access. Direct access is provided by interfaces such as the OGC Simple Features and Coverage interfaces for CORBA. If a catalogue entry denotes an OGC Feature, a Feature Collection or a Coverage, the meta-information of this entry can be populated with an **ior** (interoperable object reference). This meta-information entity is called **ior** and is filled with the standard representation of an **ior**, specified by the OMG (Object Management Group), the creators of CORBA. In XML this looks like the following (abbreviated) example:

```
<ior>IOR:010631002800000049444c3a6f6d672e6f...</ior>
```

Brokered access is specified by a request and a response message, conform all operations of the General Model. The messages are listed below.

```

struct BrokeredAccessRequest
{
    Message base;
};

```

```

wstring productHandle;
OrderSpecification orderInformation;
wstring orderID;
BrokeredAccessRequestType requestType;
UserInformation userInformation;
StatusUpdateType statusOrderUpdateType;
PackageSpecification packageSpecification;
};
typedef sequence<long> LongSeq;
struct BrokeredAccessResponse
{
    Message base;
    OrderStatus orderStatus;
    LongSeq resourceEstimate;
    CollectionName order;
    wstring orderID;
    Status status;
    BrokeredAccessRequestType requestType;
};

```

9.6.10 Exceptions

Exceptions are not specified in the General Model because they are profile specific. In CORBA exceptions are considered as an appropriate way to notify error situations to clients. The CORBA profile specifies exceptions. The **diagnostic** structure member of the response messages are not used in the CORBA profile, their role is taken over by the exceptions. Some exceptions specify the **diagnostic** (w) string as an exception parameter. By other exceptions this is not necessary, as the exceptions are self-explaining.

```

exception InvalidRequest{};
exception InvalidSession{};
exception InvalidCollection{ wstring diagnostic; };

```

The exception InvalidQuery is thrown if the client specifies an invalid query.

NOTE The exception is not thrown if the **resultType** field is set to **validate**.

```

exception InvalidQuery{ wstring diagnostic; };

```

The exception NotImplemented is defined in cases where the client asks for not-implemented behaviour. This might occur by requesting the optional access or management services.

```

exception NotImplemented{ wstring diagnostic; };

```

The NotSupported exception is thrown if the client specifies something in a request parameter that is not implemented by the server. For example the client can specify its query in Z3950_TypeOne, but the server can only interpret OGC_Common queries.

```

exception NotSupported{ wstring diagnostic; };

```

The last exception, CatalogError, indicates an error when none of the above exceptions is appropriate.

```
exception CatalogError{ wstring diagnostic; };
```

9.6.11 Catalogue Service interfaces

The interface Discovery implements methods for discovery: **query**, **present**, **explainCollection** and **getDomain**. These methods take a request message as input parameter and return a response message as output parameter.

```
interface Discovery
{
    QueryResponse query(in QueryRequest request)
        raises(InvalidSession, InvalidQuery, InvalidCollection,
NotSupported, CatalogError);
    PresentResponse present(in PresentRequest request)
        raises(InvalidSession, InvalidCollection, NotSupported,
CatalogError);
    ExplainCollectionResponse explainCollection(in
ExplainCollectionRequest request)
        raises(CatalogError);
    GetDomainResponse getDomain(in GetDomainRequest request)
        raises(CatalogError);
};
```

The next interface describes the Manager interface, which defines catalogue management functions. The operation transaction is taken literally from the General Model. This operation can create, update, or delete catalogue entries. The appropriate meta information will be provided in the request messages.

```
interface Manager
{
    // This operation is used to by a client that has the appropriate
    // user privileges to execute a whole set of insert, update and
    // delete operation of metadata to a catalog.
    TransactionResponse
        transaction(in TransactionRequest request)
            raises(NotImplemented, CatalogError);
};
```

The interface Access is the interface for access messages. It describes only one operation: the brokeredAccess function which has the request as input and which returns the response. Direct access is provided by interfaces as the Simple Feature interface and the Coverage interface. These interfaces are not described here. The client can get a reference to these interfaces by examining the **ior** field in the meta-information.

```
interface Access
{
    BrokeredAccessResponse
        brokeredAccess(in BrokeredAccessRequest request)
            raises(NotImplemented, CatalogError);
};
```

The OGC_StatefulService interface provides four operations for interactive sessions between a server and a client. All operations have a comparable form of the operations specified in the General Model.

```

interface OGC_StatefulService : OGC_Service
{
InitSessionResponse initSession(in InitSessionRequest
request)
    raises(CatalogError);
TerminateResponse terminateSession(in TerminateRequest
request)
    raises(InvalidSession, CatalogError);
StatusResponse status(in StatusRequest request)
    raises(InvalidSession, InvalidRequest,
CatalogError);
CancelResponse cancel(in CancelRequest request)
    raises(InvalidSession, InvalidRequest,
CatalogError);
};

```

The CatalogServices interface is the core of the CORBA profile.

The CatalogServices inherits from the interfaces Discovery, Access and Manager. In this way these services are realized.

NOTE Access and manager services are optional. If a server does not implement these services it throws the exception **NotImplemented**.

The CatalogServices also inherits from OGC_StatefulService that is described below.

```

interface CatalogServices : OGC_StatefulService, Discovery, Access,
Manager
{
};

```

9.6.12 Basic interfaces

Because of the asynchronous behaviour of the query operation, a callback notifying the termination of the query might be useful. The Observer Design Pattern [GAMMA97] describes a standard mechanism for notifications to one or more clients. We envision that such a mechanism will be useful for many operations in the OpenGIS world. Therefore the OGC_Observer and the OGC_Subject interfaces are modelled separately. These interfaces might be moved to an OGC general module, in the same or a similar form. The next interfaces describe the mechanism.

NOTE They are not mentioned in the General Model, as this is a CORBA specific behaviour.

```

interface OGC_Observer;
interface OGC_Subject
{
    void attachObserver(in OGC_Observer Observer);
    void detachObserver(in OGC_Observer Observer);
    void notifyObserver();
};
interface OGC_Observer
{
    void updateSubject(in OGC_Subject ChangedSubject);
};

```

The CatalogServices interface inherits from **OGC_Service**. This is envisioned as the basic interface for all OpenGIS services. As it does not exist yet, the content of this interface is not clear.

```
interface OGC_Service : OGC_Subject
{
};

interface OGC_StatefulService : OGC_Service
{
    InitSessionResponse initSession(in InitSessionRequest request)
        raises(CatalogError);
    TerminateResponse terminateSession(in TerminateRequest request)
        raises(InvalidSession, CatalogError);
    StatusResponse status(in StatusRequest request)
        raises(InvalidSession, InvalidRequest, CatalogError);
    CancelResponse cancel(in CancelRequest request)
        raises(InvalidSession, InvalidRequest, CatalogError);
};\
```

9.6.13 Complete IDL

```
//-----
// Module      : CORBA protocol binding of the OpenGIS Catalogue
Services Specification 2.0
//      described in IDL (interface definition language) of
the OMG (the Object
//      Management Group).
//-----
// Purpose      : The intention of this CORBA protocol binding is to
follow the General
//      Model closely.
//-----
// Authors      :
//      Uwe Voges, con terra GmbH, Germany
//      Barend Gehrels, Geodan IT b.v., the Netherlands
//      Joined Catalogue Response Team
// Date         : july 13, 1999
//      july 26, 1999: errata based upon minor GM changes
//      july 30, 2000: Juergen Ebbinghaus (SICAD)
//      and Barend Gehrels:
//      changes based on SICAD Review
//      - string -> wstring
//      - long SessionID -> long long
//      - e.g. sequence<type> TypeSeq
//      april 2, 2003: Uwe Voges (con terra)
//      added management-/transaction interface
//      feb 17, 2004: Uwe Voges (con terra)
//      adapted to 2.0 general model: new types for new
operations like
//      getDomain,...
```

```

//      april 29, 2004: Uwe Voges (con terra)
//      minor changes for CS 2.0 r2
//-----
-----

#pragma prefix "opengis.org"

module DynamicAny
{
    struct NameValuePair
    {
        string name;
        any value;
    };

    typedef sequence<NameValuePair> NameValuePairSeq;
};

module OGC_CatalogService
{
    //-----
    // Parameter type definitions
    //-----
    enum CG_AttributeCategory {queriable, presentable, both};
    enum CG_BrokeredAccessRequestType {orderEstimate,
orderQuoteAndSubmit,
        orderMonitor, orderCancel};
    enum CG_CatalogEntryType {product, collection, service};
    enum CG_CharacterSet {ASCII, UniCode, ShiftJIS};
    union CG_CollectionName
        switch(long)
        {
            case 1 : wstring collectionID;
            case 2 : wstring collectionName;
        };
    enum CG_MessageFormat {XML, HTML, TXT, NV};

    struct CG_OrderItem
    {
        // Note: datatypes not provided by GM
        any productID;
        any productPrice;
        any productDeliveryOptions;
        any processingOptions;
        any sceneSelectionOptions;
    };

    struct CG_OrderSpecification
    {
        // Note: datatypes not provided by GM
        any orderCentreID;
        any orderPrice;
        any orderDeliveryDate;
        any orderCancellationDate;
        any deliveryMethod;
        any package;
    };
};

```

```

};

enum CG_OrderStatus {orderBeingEstimated, orderEstimated,
    orderBeingQuoted, orderBeingProcessed,
    orderCompleted, orderNotValid, orderCancelled};

enum CG_PackagingType {predefinedPackage, adhocPackage};
struct CG_PackageSpecification
{
    // Note: datatypes not provided by GM
    any packageId;
    any packagePrice;
    CG_PackagingType package;
    any packageMedium;
    long packageSize;
};

enum CG_PaymentMethod {credit, cash, purchaseOrder};

enum CG_PredefinedPresentationType {full, summary, brief};

typedef sequence<wstring> StringSeq;

union CG_PresentationDescription
    switch(long)
    {
        case 1 : StringSeq attributes; // CG_TupleType in GM
        case 2 : CG_PredefinedPresentationType presentationType; //
name in GM
    };

enum CG_QueryLanguage {OGC_Common, OGC_Filter, Z3950_TypeOne,
SQL3_SimpleFeature};
struct CG_QueryExpression
{
    wstring theQuery;
    wstring theNamespace;
    CG_QueryLanguage theLanguage;
    any queryParameters;
};

enum CG_QueryScope {distributed, locale};

struct CG_RequestID
{
    long long sessionID;
    long counter;
};

enum CG_ResultType {validate, resultSetID, hits, results};

struct CG_ReturnData
{
    CG_MessageFormat encoding;
    any payload;
    // XML,HTML,TEXT will return a string

```

```

        // NV will return a DynamicAny::NameValuePairSeq (from CORBA 2.3
Dynamic Any)
    };

typedef DynamicAny::NameValuePairSeq CG_Schema;
struct CG_SchemaID
{
    wstring schemeName;
    CG_Schema schema;
};

enum CG_SortOrder {ascending, descending};
struct CG_SortField
{
    wstring attributeName;
    CG_SortOrder sortOrder;
};

enum CG_Status {success, successResultsAvailable, processingNormal,
processingQueued, processingPausedOrSuspended, failure,
failureAccessDenied};

enum CG_StatusUpdateType {manual, automatic};

struct CG_UserInformation
{
    wstring userName;
    wstring userAddress;
    wstring phoneNumber;
    wstring faxNumber;
    wstring emailAddress;
    wstring netAddress;
    CG_PaymentMethod paymentMethod;
};
//-----
// Capabilities, 3.2.7.3
//-----
enum CG_CapabilityType
{ ctAllSupportedRequest, ctDefaults, ctDefaultTimeOut,
    ctExplain, ctMessaging, ctQuery, ctSession,
    ctSoftwareInformation, ctSupportedCollections };

typedef boolean CG_AllSupportedRequest;

typedef boolean CG_Defaults;

struct CG_DefaultTimeOut
{
    unsigned long long timeOut;
};
typedef boolean CG_Explain;
struct CG_Messaging
{
    CG_CharacterSet characterSet;
    CG_MessageFormat messageFormat;
};

```

```

struct CG_Query
{
    wstring version;
    CG_CharacterSet characterSet;
    CG_QueryLanguage queryLanguage;
};

struct CG_Session
{
    wstring language;
    wstring catalogSpecificationVersion;
    CG_CharacterSet characterSet;
};

struct CG_SoftwareInformation
{
    wstring vendor;
    wstring SVversionNumber;
    wstring IFversionNumber;
};

typedef sequence<CG_CollectionName> CG_SupportedCollections;

union CG_Capability
    switch(CG_CapabilityType)
    {
        case ctAllSupportedRequest : CG_AllSupportedRequest
allSupportedRequest;
        case ctDefaults : CG_Defaults defaults;
        case ctDefaultTimeOut : CG_DefaultTimeOut timeOut;
        case ctExplain : CG_Explain explain;
        case ctMessaging : CG_Messaging messaging;
        case ctQuery : CG_Query query;
        case ctSession : CG_Session session;
        case ctSoftwareInformation : CG_SoftwareInformation
softwareInformation;
        case ctSupportedCollections : CG_SupportedCollections
supportedCollections;
    };

//-----
// DomainType
//-----

// Enumeration for the type of the metadata property or request
parameter
enum CG_DomainType
    { domainTypeString, domainTypeDate, domainTypeInteger,
domainTypeReal,
    domainTypeSpatial };

// Enumeration for the type how possible values can be described
enum CG_DomainValuesType
    { valueRef, valueRange };

```

```

// valid range, composed of lower and upper boundary value
struct CG_ValueRange
{
    wstring lowerValue;
    wstring upperValue;
};
typedef sequence<CG_ValueRange> CG_ValueRangeSeq;

// Definition of the real type by totalDigits and fractionDigits
struct CG_RealTypeDef
{
    long totalDigits;
    long fractionDigits;
};

// type definition of the metadata property or request parameter
union CG_DomainDef
    switch(CG_DomainType)
    {
        case domainTypeString : long length;
        case domainTypeDate : wstring dateFormat;
        case domainTypeInteger : long totalDigits;
        case domainTypeReal : CG_RealTypeDef realTypeDef;
        case domainTypeSpatial : long dimensions;
    };

// definition of valid values (range, value list) of the type
union CG_DomainValueDef
    switch(CG_DomainValueType)
    {
        case valueRef : StringSeq valueList;
        case valueRange : CG_ValueRangeSeq rangeList;
    };

// definition of domain values (typical value, valid values, value
name, value description)
// of the type
struct CG_DomainValue
{
    wstring value;
    CG_DomainValueDef valuesDef;
    wstring title;
    wstring description;
    wstring metadataURL;
};
typedef sequence<CG_DomainValue> CG_DomainValueSeq;

// Description of domain of one requested metadata property or
request parameter
struct CG_Domain
{
    wstring attributeName;
    CG_DomainDef domainDef;
    CG_DomainValueSeq domainValue;
};

```

```

typedef sequence<CG_Domain> CG_DomainSeq;

//-----
// Messages
//-----
struct CG_Message
{
    long long sessionID;
    wstring destinationID;
    CG_RequestID requestID;
    wstring additionalInfo;
};
struct CG_InitSessionRequest
{
    CG_Message base;
};
struct CG_InitSessionResponse
{
    CG_Message base;
};
struct CG_TerminateRequest
{
    CG_Message base;
};
struct CG_TerminateResponse
{
    CG_Message base;
    CG_Status status;
};
typedef sequence<CG_CapabilityType> CG_CapabilityTypeSeq;
struct CG_ExplainServerRequest
{
    CG_Message base;
    CG_CapabilityTypeSeq capabilities;
};
struct CG_ExplainServerResponse
{
    CG_Message base;
    CG_CapabilityTypeSeq capabilities;
};
struct CG_StatusRequest
{
    CG_Message base;
    CG_RequestID requestIDtoStatus;
};
struct CG_StatusResponse
{
    CG_Message base;
    CG_RequestID requestIDtoStatus;
    CG_Status status;
};
struct CG_CancelRequest
{
    CG_Message base;
    CG_RequestID requestIDtoCancel;
    boolean freeResources;
};

```

```

};
struct CG_CancelResponse
{
    CG_Message base;
    CG_Status status;
    CG_RequestID canceledRequest;
};
typedef sequence<CG_SortField> CG_SortFieldSeq;
struct CG_QueryRequest
{
    CG_Message base;
    CG_QueryExpression queryExpression;
    CG_ResultType resultType;
    long iteratorSize;
    long cursor;
    CG_MessageFormat returnFormat;
    CG_PresentationDescription presentation;
    CG_SortFieldSeq sortField;
    CG_QueryScope queryScope;
    CG_CollectionName collectionID;
    CG_CatalogEntryType catalogType;
    boolean asynchronous;
    long maxLevel;
};
struct CG_QueryResponse
{
    CG_Message base;
    CG_ReturnData retrievedData;
    CG_CollectionName resultSetID;
    CG_Status status;
    long hits;
    long cursor;
};
struct CG_PresentRequest
{
    CG_Message base;
    CG_CollectionName resultSetID;
    CG_PresentationDescription presentation;
    CG_SortFieldSeq sortField;
    CG_MessageFormat returnFormat;
    long iteratorSize;
    long cursor;
};
struct CG_PresentResponse
{
    CG_Message base;
    CG_ReturnData retrievedData;
    long cursor;
    long hits;
    CG_Status status;
};
struct CG_ExplainCollectionRequest
{
    CG_Message base;
    CG_AttributeCategory attributeCategory;
    CG_CollectionName collectionID;
};

```

```

    CG_MessageFormat returnFormat;
};
struct CG_ExplainCollectionResponse
{
    CG_Message base;
    CG_CollectionName collectionID;
    CG_SchemaID dataModel;
    CG_Status status;
};

struct CG_GetDomainRequest
{
    CG_Message base;
    StringSeq attributes;
};

struct CG_GetDomainResponse
{
    CG_Message base;
    CG_DomainSeq attributeDomains;
};

// Messages for access
struct CG_BrokeredAccessRequest
{
    CG_Message base;
    wstring productHandle;
    CG_OrderSpecification orderInformation;
    wstring orderID;
    CG_BrokeredAccessRequestType requestType;
    CG_UserInformation userInformation;
    CG_StatusUpdateType statusOrderUpdateType;
    CG_PackageSpecification packageSpecification;
};

typedef sequence<long> LongSeq;
struct CG_BrokeredAccessResponse
{
    CG_Message base;
    CG_OrderStatus orderStatus;
    LongSeq resourceEstimate;
    CG_CollectionName order;
    wstring orderID;
    CG_Status status;
    CG_BrokeredAccessRequestType requestType;
};

typedef sequence<DynamicAny::NameValuePairSeq> NameValuePairSeqSeq;

// Messages for managing functions

// Insert-Structure for inserting metadata (payload) to a collection
// of a defined catalogType in a predefined format (encoding)
struct CG_InsertMetadata
{
    CG_CollectionName collectionID;

```

```

    CG_CatalogEntryType catalogType;

    CG_MessageFormat encoding;

    // the data to insert, e.g. a list of NV-Pairs or ISO19139-XML
    any payload;
};

// Update-Structure for updating a certain set (defined by the
queryExpression) of metadata
// (payload) of a collection/catalogType in a predefined format
(encoding)
struct CG_UpdateMetadata
{
    CG_CollectionName collectionID;
    CG_CatalogEntryType catalogType;

    CG_MessageFormat encoding;

    CG_QueryExpression queryExpression;

    // the data to replace, e.g. a list of NV-Pairs or ISO19139-XML
    any payload;
};

// Delete-Structure for deleting a certain set (defined by the
queryExpression) of metadata
// of a collection/catalogType
struct CG_DeleteMetadata
{
    CG_CollectionName collectionID;
    CG_CatalogEntryType catalogType;

    CG_QueryExpression queryExpression;
};

// Enumeration for the type of a single manipulation in a
transaction
enum CG_MetadataManipulationType
{ ctInsert, ctUpdate, ctDelete };

// Structure for the content of a single manipulation in a
transaction
union CG_MetadataManipulation
switch(CG_MetadataManipulationType)
{
    case ctInsert : CG_InsertMetadata insertData;
    case ctUpdate : CG_UpdateMetadata updateData;
    case ctDelete : CG_DeleteMetadata deleteData;
};
typedef sequence<CG_MetadataManipulation> MetadataManipulationSeq;

// Struct of a transaction, as a sequence of single manipulations
struct CG_Transaction
{
    MetadataManipulationSeq manipulations;
};

```

```

};

struct CG_TransactionRequest
{
    CG_Message base;
    CG_Transaction transactionData;
};
struct CG_TransactionResponse
{
    CG_Message base;
    CG_Status status;

    // value of -1 means not calculated
    long totalInserted;
    long totalUpdated;
    long totalDeleted;

    // List of newly generated catalogue entry identifiers
    // assigned to the new catalogue entry instances
    NameValuePairSeqSeq newKeyList;
};

//-----
// Exceptions
//-----
exception InvalidSession{};
exception InvalidRequest{};
exception InvalidCollection{ wstring diagnostic; };
exception InvalidQuery{ wstring diagnostic; };
exception NotImplemented{ wstring diagnostic; };
exception NotSupported{ wstring diagnostic; };
exception CatalogError{ wstring diagnostic; };

//-----
// Interfaces
//-----
interface OGC_Observer;
interface OGC_Subject
{
    oneway void attachObserver(in OGC_Observer Observer);
    oneway void detachObserver(in OGC_Observer Observer);
    oneway void notifyObserver();
};
interface OGC_Observer
{
    void updateSubject(in OGC_Subject ChangedSubject);
};

interface OGC_Service : OGC_Subject
{
    CG_ExplainServerResponse explainServer(in CG_ExplainServerRequest
request)
        raises(CatalogError);
};

interface OGC_StatefulService : OGC_Service

```

```

    {
        CG_InitSessionResponse initSession(in CG_InitSessionRequest
request)
            raises(CatalogError);
        CG_TerminateResponse terminateSession(in CG_TerminateRequest
request)
            raises(InvalidSession, CatalogError);
        CG_StatusResponse status(in CG_StatusRequest request)
            raises(InvalidSession, InvalidRequest, CatalogError);
        CG_CancelResponse cancel(in CG_CancelRequest request)
            raises(InvalidSession, InvalidRequest, CatalogError);
    };

interface CG_Discovery
{
    CG_QueryResponse query(in CG_QueryRequest request)
        raises(InvalidSession, InvalidQuery, InvalidCollection,
NotSupported, CatalogError);
    CG_PresentResponse present(in CG_PresentRequest request)
        raises(InvalidSession, InvalidCollection, NotSupported,
CatalogError);
    CG_ExplainCollectionResponse explainCollection(in
CG_ExplainCollectionRequest request)
        raises(CatalogError);
    CG_GetDomainResponse getDomain(in CG_GetDomainRequest request)
        raises(CatalogError);
};

interface CG_Access
{
    // Direct access is provided by the IOR fields in the meta-
information
    // itself
    // Brokered access
    CG_BrokeredAccessResponse
        brokeredAccess(in CG_BrokeredAccessRequest request)
            raises(NotImplemented, CatalogError);
};

interface CG_Manager
{
    // This operation is used to by a client that has the appropriate
user
    // privileges to execute a whole set of insert, update and delete
operation
    // of metadata to a catalog.
    CG_TransactionResponse
        transaction(in CG_TransactionRequest request)
            raises(NotImplemented, CatalogError);
};

interface CG_CatalogServices : OGC_StatefulService, CG_Discovery,
CG_Access, CG_Manager
{
};
};

```

10 HTTP protocol binding (Catalogue Services for the Web, CSW)

10.1 Architectural principles

The purpose of this clause is to describe the request and response messages that are common to all web-based catalogue services. The basic message exchange pattern is illustrated in .

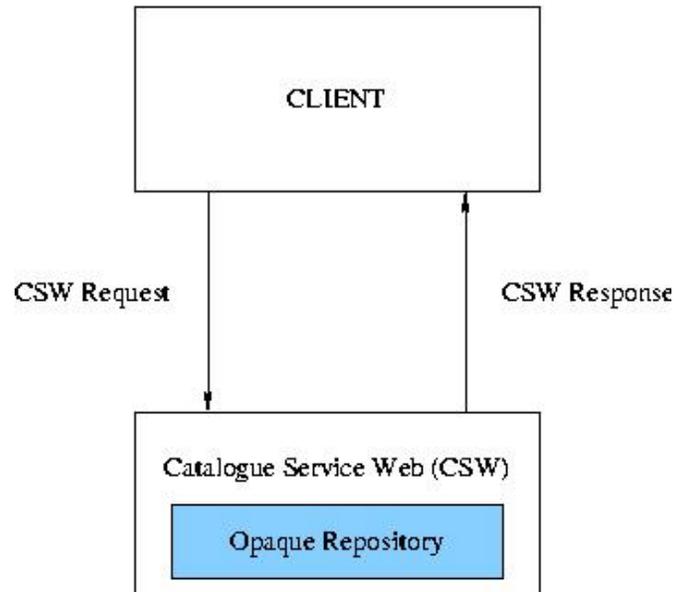


Figure 28 — Catalogue service web

The interaction between a client and a server is accomplished using a standard request-response model of the HTTP protocol. That is, a client sends a request to a server using HTTP, and expects to receive a response to the request or an exception message.

Request and response messages are encoded as keyword-value pairs within a request URI or using an XML entity-body. Requests may also be embedded in a messaging framework such as SOAP.

10.2 The HTTP protocol

10.2.1 Overview

The Hypertext Transfer Protocol (HTTP) is a generic, stateless, application-level protocol that is widely used to exchange information on the web. The HTTP/1.1 specification is published by the Internet Engineering Task Force (IETF) as RFC 2616:

<http://www.ietf.org/rfc/rfc2616>. The "http" URI scheme is used to locate network resources using the HTTP protocol; consult Section 3.2 of RFC 2616 and RFC 2396 for details. The general syntax of the scheme is summarised below for convenience:

```
http_URL = "http:" "://" host [ ":" port ] [ abs_path [ "?" query ] ]
```

HTTP messages have a simple line-oriented structure. The three basic parts of a message are summarised in Table 51 (RFC 2616, 4.1):

Table 51 — HTTP message elements

Start line	Indicates what to do for a request or what happened for a response.
Header fields	Zero or more header fields, each consisting of a name and a value separated by a colon (:). Four categories of headers provide metainformation about the message: general, request, response, and entity. The headers part ends with a blank line.
Body	An optional body containing the message content that conforms to some Internet media type.

Note that URIs are generally case-sensitive except for the scheme and host names; furthermore, if the port number is not specified it is equivalent to the default TCP port number (80). As noted in Section 3 of RFC 2396, certain special characters are reserved within various URI components; if the data within a URI component conflicts with the reserved purpose it shall be escaped before forming the URI.

This subclause clarifies some aspects of HTTP usage for catalogue application profiles that employ the protocol to exchange request and response messages. In many cases this means turning a SHOULD level requirement into a SHALL level requirement; what results is effectively an HTTP profile for OGC catalogue services.

10.2.2 Message headers

The standard headers are defined in Section 14 of RFC 2616. Some of these are of particular significance to catalogue operations.

Any HTTP/1.1 message containing an entity-body shall include a Content-Type header field defining the media type of that body (RFC 2616, 7.2.1); the charset parameter shall also be specified for text.

EXAMPLES 1

```
Content-Type: application/xml; charset=utf-8
Content-Type: application/octet-stream
Content-Type: multipart/related; boundary="part-boundary";
  start="<urn:uuid:e3fec7a9-cc5d-45ba-87a5-8a2a27f6fb5b>"
  type="application/xml"
```

A user agent may use the Accept request header to declare a set of preferred Internet media types for the response. The IANA registry of media types is available online: <http://www.iana.org/assignments/media-types/>.

EXAMPLES 2

```
Accept: application/xml
Accept: application/xhtml+xml, text/html; q=0.5
```

Implementers should be mindful that the function of some request parameters (e.g. outputFormat) overlaps the functionality of some of the HTTP message headers (e.g., Accept). When a request to a CSW service includes both a request parameter (e.g.,

outputFormat) and a message header that perform similar functions (e.g., Accept), the server shall use the value of the request parameter (e.g., outputFormat) to process the request. When the request parameter is not specified, the value of the HTTP message header shall be used. If the value of the request parameter and the HTTP message header disagree, server shall raise an exception. If neither is specified, the default value of the request parameter shall be used if one is specified; otherwise the server shall raise an exception as specified in Subclause 10.3.7.

10.2.3 Content encoding

The Content-Encoding entity-header may be used to indicate any additional content encodings that have been applied to the entity body, usually for the purpose of data compression or encryption. This header shall be included if a non-identity encoding has been applied.

EXAMPLE 1

```
Content-Encoding: gzip
```

A user agent may specify a preferred content encoding using the Accept-Encoding header. If no such request-header is included, the server shall use the "identity" encoding.

EXAMPLE 2

```
Accept-Encoding: gzip;q=1.0, identity; q=0.5
```

10.2.4 Request methods

The HTTP/1.1 specification defines eight methods for manipulating and retrieving representations of resources. Within an application profile abstract catalogue operations shall be mapped to one or more of these methods; these mappings should be consistent with HTTP/1.1 semantics. The methods that are most relevant to catalogue services are summarised in below:

Table 52 — Selected HTTP Request Methods

Method name	Semantics
GET	Used to retrieve whatever information (in the form of an entity) is identified by the Request-URI
POST	Used to request that the origin server accept the entity enclosed in the request as data to be processed by the resource identified by the Request-URI in the Request-Line

The Request-URI is a Uniform Resource Identifier that identifies the target resource to which the request is applied; it has the following syntax:

```
Request-URI = "*" | absoluteURI | abs_path ["?" query] | authority
```

10.2.5 Message payload

10.2.5.1 Introduction

All implementations of the HTTP protocol, and application profiles derived from the CSW protocol binding, shall support the response schema described in this subclause. This response schema is an XML realization of the core metadata properties described abstractly in Subclause 6.3.

An application profile may specify other allowable payloads that constitute the body of a request or response message (if applicable). Entities shall conform to a registered Internet media type, but there is otherwise no restriction on the content; the actual payload is dependent upon the information model supported by the profile. In all cases, elements of the underlying information model shall be mapped to the core metadata properties described in Subclause 6.3 which are concretely expressed in the HTTP protocol binding using the schema described in this subclause.

10.2.5.2 Abstract Record

An abstract element, called **csw:AbstractRecord**, is declared which is the head of a substitution group representing all the views (i.e. brief, summary and full) of the core metadata properties.

The following XML-Schema fragment declares the **csw:AbstractRecord** element:

```
<xsd:element name="AbstractRecord"
             type="csw:AbstractRecordType" abstract="true"/>
<xsd:complexType name="AbstractRecordType" abstract="true"/>
```

10.2.5.3 Core queryable and returnable realization

10.2.5.3.1 Introduction

The common CSW record syntax is an XML-based encoding of Dublin Core metadata terms; it represents a concrete realization of the core metadata properties abstractly specified in subclause 6.3.2.

The full set of core properties is concretely materialized by the **csw:Record** element. Two addition elements, **csw:BriefRecord** and **csw:SummaryRecord** materialize the brief and summary views of the full set of core properties.

Table 53 maps the Dublin Core element names of the OGC query and response elements specified in Subclause 6.3 to the concrete XML elements specified in this subclause.

Table 53 — Mapping of Dublin Core names to XML element names

Dublin Core element name	OGC queryable term	XML element name
title	Title	dc:title
creator		dc:creator
subject	Subject	dc:subject
description	Abstract	dct:abstract
publisher		dc:publisher
contributor		dc:contributor
date	Modified	dct:modified
type	Type	dc:type
format	Format	dc:format
identifier	Identifier	dc:identifier
source	Source	dc:source
language		dc:language
relation	Association	dc:relation
coverage	BoundingBox	ows:BoundingBox
rights		dc:rights

NOTE Other XML elements names may be substitutable for the XML elements listed in Table 53. The schema rec-dcterms.xsd contains a complete list of substitutable XML elements.

10.2.5.3.2 Full record

The XML encoding of the element representing the full catalogue record, **csw:Record**, is defined by the following XML-Schema fragment:

```
<xsd:element name="Record"
  type="csw:RecordType"
  substitutionGroup="csw:AbstractRecord"/>
<xsd:complexType name="RecordType">
  <xsd:complexContent>
    <xsd:extension base="csw:DCMIRecordType">
      <xsd:sequence>
        <xsd:element name="AnyText" type="csw:EmptyType"
          minOccurs="0"/>
        <xsd:element ref="ows:BoundingBox"
          minOccurs="0" maxOccurs="unbounded"/>
      </xsd:sequence>
    </xsd:extension>
  </xsd:complexContent>
</xsd:complexType>
```

The full record is based on the XML schemas developed for the Dublin Core Metadata Initiative with the addition of the **csw:AnyText** and **ows:BoundingBox** elements.

All elements contained in the **csw:Record** element shall be available as queryables. In the event that a catalogue implementation does not have a value to map to an element

contained in **csw:Record**, its value shall be considered to be NULL for the purpose of querying.

In the Dublin Core schema, the elements **dc:identifier** and **dc:title**, or any element that can substitute for them, are optional. However for the purposes of this specification, these elements shall be considered mandatory queryables and presentables.

The **csw:AnyText** element shall only be available as a queryable, and is intended as a query target for a full text query of the catalogue's records. This element is not a presentable and shall never appear in a response message. Even though the content model for the **csw:AnyText** element is empty, the catalogue shall interpret its value to be the full text of all text fields in the catalogue record.

The **ows:BoundingBox** element is used to express the spatial extent of the record.

It is anticipated that the number of optional presentables will vary from one record instance to another inside the catalogue; and so the remaining elements, or any element that can substitute for them, may appear as presentables if the catalogue has a corresponding value to present. Otherwise, elements for which no value is available may be omitted from the response.

EXAMPLE

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<Record
  xmlns="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:dc="http://purl.org/dc/elements/1.1/"
  xmlns:dct="http://purl.org/dc/terms/"
  xmlns:ows="http://www.opengis.net/ows"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.opengis.net/cat/csw/2.0.2
    ../../../../csw/2.0.2/record.xsd">
  <dc:identifier>00180e67-b7cf-40a3-861d-b3a09337b195</dc:identifier>
  <dc:title>Image2000 Product 1 (at1) Multispectral</dc:title>
  <dct:modified>2004-10-04 00:00:00</dct:modified>
  <dct:abstract>IMAGE2000 product 1 individual orthorectified scenes.
  IMAGE2000 was produced from ETM+ Landsat 7 satellite data and provides
  a consistent European coverage of individual orthorectified scenes in
  national map projection systems.</dct:abstract>
  <dc:type>dataset</dc:type>
  <dc:subject>imagery</dc:subject>
  <dc:subject>baseMaps</dc:subject>
  <dc:subject>earthCover</dc:subject>
  <dc:format>BIL</dc:format>
  <dc:creator>Vanda Lima</dc:creator>
  <dc:language>en</dc:language>
  <ows:WGS84BoundingBox>
    <ows:LowerCorner>14.05 46.46</ows:LowerCorner>
    <ows:UpperCorner>17.24 48.42</ows:UpperCorner>
  </ows:WGS84BoundingBox>
</Record>
```

NOTE The XML examples and schema fragments in Clause 10 often include the new Catalogue 2.0.2 XML namespace identifier URI, namely <http://www.opengis.net/cat/csw/2.0.2>. The Catalogue 2.0.2 XML Schema

Documents use this new namespace identifier URI instead of the previous namespace identifier URI, namely <http://www.opengis.net/cat/csw>. This change was made based on the TC decision described in Clause 11 in OGC Best Practices Paper 06-135r1. That decision was based on the discussion in document 05-065r3 “Change Request for Namespaces for versions and profiles of XML Schemas”. A new namespace identifier URI is used for this Catalogue version 2.0.2 because these reasons apply to this version 2.0.2. However, using this new namespace identifier URI could produce problems for the GetCapabilities operation request and response, which have not yet been resolved.

10.2.5.3.3 Summary record

The XML encoding of the element representing the summary view, **csw:SummaryRecord**, is defined by the following XML-Schema fragment:

```
<xsd:element name="SummaryRecord"
  type="csw:SummaryRecordType"
  substitutionGroup="csw:AbstractRecord" />
<xsd:complexType name="SummaryRecordType" final="#all">
  <xsd:complexContent>
    <xsd:extension base="csw:AbstractRecordType">
      <xsd:sequence>
        <xsd:element ref="dc:identifier"
          minOccurs="1" maxOccurs="unbounded" />
        <xsd:element ref="dc:title"
          minOccurs="1" maxOccurs="unbounded" />

        <xsd:element ref="dc:type"
          minOccurs="0" />
        <xsd:element ref="dc:subject"
          minOccurs="0" maxOccurs="unbounded" />
        <xsd:element ref="dc:format"
          minOccurs="0" maxOccurs="unbounded" />
        <xsd:element ref="dc:relation"
          minOccurs="0" maxOccurs="unbounded" />
        <xsd:element ref="dct:modified"
          minOccurs="0" maxOccurs="unbounded" />
        <xsd:element ref="dct:abstract"
          minOccurs="0" maxOccurs="unbounded" />
        <xsd:element ref="dct:spatial"
          minOccurs="0" maxOccurs="unbounded" />
        <xsd:element ref="ows:BoundingBox"
          minOccurs="0" maxOccurs="unbounded" />
      </xsd:sequence>
    </xsd:extension>
  </xsd:complexContent>
</xsd:complexType>
<xsd:complexType name="AnyText" />
```

The **dc:identifier** and **dc:title** elements, or any element that can be substituted for them, are mandatory presentables and shall always appear in the response. The remaining elements, or any element that can substitute for them, can appear as a presentable if the catalogue has a corresponding value to present. The **csw:SummaryRecord** element cannot be the target of a query operation.

EXAMPLE

```
<?xml version="1.0" encoding="ISO-8859-1"?>
```

```

<SummaryRecord
  xmlns="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:dc="http://purl.org/dc/elements/1.1/"
  xmlns:dct="http://purl.org/dc/terms/"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.opengis.net/cat/csw/2.0.2
    ../../../../csw/2.0.2/record.xsd">
  <dc:identifier>00180e67-b7cf-40a3-861d-b3a09337b195</dc:identifier>
  <dc:title>Image2000 Product 1 (at1) Multispectral</dc:title>
  <dc:type>dataset</dc:type>
  <dc:subject>imagery</dc:subject>
  <dc:subject>baseMaps</dc:subject>
  <dc:subject>earthCover</dc:subject>
  <dc:format>BIL</dc:format>
  <dct:modified>2004-10-04 00:00:00</dct:modified>
  <dct:abstract>IMAGE2000 product 1 individual orthorectified scenes.
  IMAGE2000 was produced from ETM+ Landsat 7 satellite data and provides
  a consistent European coverage of individual orthorectified scenes in
  national map projection systems.</dct:abstract>
</SummaryRecord>

```

10.2.5.3.4 Brief record

The XML encoding of the element representing the brief view, **csw:BriefRecord**, is defined by the following XML-Schema fragment:

```

<xsd:element name="BriefRecord"
  type="csw:BriefRecordType"
  substitutionGroup="csw:AbstractRecord" />
<xsd:complexType name="BriefRecordType" final="#all">
<xsd:complexContent>
  <xsd:extension base="csw:AbstractRecordType">
    <xsd:sequence>
      <xsd:element ref="dc:identifier"
        minOccurs="1" maxOccurs="unbounded"/>
      <xsd:element ref="dc:title"
        minOccurs="1" maxOccurs="unbounded"/>
      <xsd:element ref="dc:type"
        minOccurs="0"/>
      <xsd:element ref="ows:BoundingBox"
        minOccurs="0" maxOccurs="unbounded"
      />
    </xsd:sequence>
  </xsd:extension>
</xsd:complexContent>
</xsd:complexType>

```

The **dc:identifier** and **dc:title** elements, or any element that can be substituted for them, are mandatory presentables and shall always appear in the response. The remaining elements, or any element that can be substituted for them, can appear as a presentable if the catalogue has a corresponding value to present. The **csw:BriefRecord** element cannot be the target of a query operation.

EXAMPLE

```
<?xml version="1.0" encoding="ISO-8859-1"?>
```

```

<BriefRecord
  xmlns="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:dc="http://purl.org/dc/elements/1.1/"
  xmlns:dct="http://purl.org/dc/terms/"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.opengis.net/cat/csw/2.0.2
    ../../../../csw/2.0.2/record.xsd">
  <dc:identifier>00180e67-b7cf-40a3-861d-b3a09337b195</dc:identifier>
  <dc:title>Image2000 Product 1 (at1) Multispectral</dc:title>
  <dc:type>dataset</dc:type>
</BriefRecord>

```

10.3 Operation request and response encoding

10.3.1 Introduction

Only the GET and POST methods are employed in the HTTP binding. Table 54 summarizes the allowed HTTP method bindings and request data encodings for all CSW requests; optional method bindings and data encodings are enclosed in parentheses.

Table 54 — HTTP method bindings

Request	HTTP method binding(s)	Data encoding(s) ^{a, b}
GetCapabilities	GET (POST)	KVP (XML)
DescribeRecord	POST (GET)	XML (KVP)
GetDomain	POST (GET)	XML (KVP)
GetRecords	POST (GET)	XML (KVP)
GetRecordById	GET (POST)	KVP (XML)
Harvest	POST (GET)	XML (KVP)
Transaction	POST	XML
^a XML = application/xml using POST (with a charset parameter if necessary—UTF-8 is strongly recommended) ^b KVP = URL-encoded key/value pairs using GET or application/x-www-form-urlencoded using POST		

10.3.2 SOAP

This subclause specifies the use of SOAP (formerly Simple Object Access Protocol) messages for communication between a catalogue client and a CSW.

SOAP is a communication protocol for communication between applications. It defines a format for sending messages between communicating applications via the Internet and specifically using HTTP. Soap is platform independent, language independent and SOAP messages are encoded using XML. This means that SOAP provides a way to communicate between applications running on different operating systems, with different technologies and programming languages.

A SOAP message is an ordinary XML document containing the following elements:

- a) A required Envelope element that identifies the XML document as a SOAP message
- b) An optional Header element that contains header information
- c) A required Body element that contains call and response information
- d) An optional Fault element that provides information about errors that occurred while processing the message

All the elements above are declared in the default namespace for the SOAP envelope:

<http://www.w3.org/2003/05/soap-envelope>

and the default namespace for SOAP encoding and data types is:

<http://www.w3.org/2003/05/soap-encoding>

The SOAP specification defines a number of syntax rules. Among the most important are:

- a) A SOAP message shall be encoded using XML
- b) A SOAP message shall use the SOAP Envelope namespace
- c) A SOAP message shall use the SOAP Encoding namespace
- d) A SOAP message shall not contain a DTD reference
- e) A SOAP message shall not contain XML Processing Instructions

The following XML fragment illustrates a skeleton SOAP message:

```
<?xml version="1.0"?>
<soap:Envelope
  xmlns:soap=http://www.w3.org/2003/05/soap-envelope
  soap:encodingStyle="http://www.w3.org/2003/05/soap-encoding">

  <soap:Header>
    ...
    ...
  </soap:Header>

  <soap:Body>
    ...
    ...
    <soap:Fault>
      ...
      ...
    </soap:Fault>
  </soap:Body>
```

```
</soap:Envelope>
```

A client may send CSW requests to a compatible catalogue using the body of a SOAP envelope. The client simply encodes the CSW request as the content of the `<soap:Body>` element in the request message. The CSW shall then response by generating a SOAP message where the response to the client's request is the content of the **<soap:Body>** element.

If an exception is encountered while processing a CSW request encoded in a SOAP envelope, the CSW server shall generate a SOAP response message where the content of the **<soap:Body>** element is a **<soap:Fault>** element. The following skeleton XML fragment shall be used when generating the **<soap:Body>** element in the event that the CSW server encounters an exception:

```
<soap:Body>
  <soap:Fault>
    <soap:Code>
      <soap:Value>soap:Server</soap:Value>
    </soap:Code>
    <soap:Reason>
      <soap:Text>A server exception was encountered.</soap:Text>
    </soap:Reason>

    <soap:Detail>
      <ows:ExceptionReport>
        ...
      </ows:ExceptionReport>
    </soap:Detail>
  </soap:Fault>
</soap:Body>
```

The **<soap:Value>** element in the **<soap:Code>** element shall have the content *soap:Server* indicating that this is a server exception. The **<soap:Text>** element in the **<soap:Reason>** element shall have the content "Server exception was encountered.". This fixed string is used since the details of the exception will be specified in the **<soap:Detail>** element using an **<ows:ExceptionReport>** element as defined in document [OGC 05-008c1].

The **<soap:Detail>** element shall contain an **<ows:ExceptionReport>** element detailing the specific exception that the server encountered.

The use of the **<soap:Header>** element is not discussed in this version of this specification.

10.3.3 Namespaces

Namespaces [W3C Recommendation January 1999] are used to discriminate XML vocabularies from one another. For the CSW there are two normative namespace definitions, namely:

(<http://www.opengis.net/cat/csw/2.0.2>)- for CSW interface vocabulary
 (<http://www.opengis.net/ogc>) - for OGC Filter vocabulary

A given CSW implementation will make use of one or more XML Schemas describing the metadata that is being manipulated and these schemas will, in turn, use one or more namespaces (e.g. <http://www.someserver.com/myns><http://www.someserver.com/myns>
<urn:oasis:names:tc:ebxml-regrep:xsd:rim:3.0>).

10.3.4 Predicate languages

The general model allows catalogue clients to specify the predicate language used to constrain operations. The HTTP protocol binding schemas define two predicate languages, based on the BNF in Subclause 6.2.2, which may be used. The two predicate languages are:

- a) CQL_TEXT is a text encoding of the BNF.
- b) FILTER is an XML encoding of the BNF grammar and is normatively defined in the Filter Encoding Implementation Specification, version 1.1.0 [OGC 04-095]. All CSW implementations are required to support this filter syntax.

Table 55 defines the parameters required to specify a predicate in keyword-value pair encoded CSW operation requests.

NOTE In Clause 10, tables are used to describe the parameters of each CSW request. The name of each parameter is taken from the XML encoding of the request, unless the parameter is not applicable in the XML encoding in which case the name used in the KVP encoding is used. XML element names are case sensitive and are presented in the tables using the correct case. KVP parameter names are case insensitive and are presented in the tables in UPPER case. For example, in Table 55, the parameter "Constraint" is specified using its XML encoding, while the parameter CONSTRAINT_LANGUAGE (which is implied in the XML encoding) is specified using the KVP encoding and is in upper case.

Table 55 — KVP encoding for constraints

Keyword ^b	Description	Data type and value	Optionality
CONSTRAINT_LANGUAGE	Identifies the predicate language used for the value of the Constraint	Code List with allowed values: <i>CQL_TEXT</i> , used to indicate CQL. <i>FILTER</i> , used to indicate OGC Filter.	Zero or one (Optional) ^a Must be specified with the Constraint
CONSTRAINT_LANGUAGE_VERSION	Identifies the version of the predicate language used.	Character String	Zero or one There is no default as the parameter is specified if required to indicate which version of a specification the value of the constraint parameter conforms to.
Constraint	Text of query constraint in the predicate language identified by the CONSTRAINT_LANGUAGE	Character String	Zero or one (Optional) Must be specified with the CONSTRAINT_LANGUAGE
<p>a The CONSTRAINT_LANGUAGE parameter contains the same information as the contents of the <Constraint> element in XML encoding.</p> <p>b Parameter keywords, for KVP encoding, are case insensitive.</p>			

The following XML schema fragments define how the predicate language may be XML encoded in CSW operations that allow constraints to be defined (Query, Update and Delete):

```
<xsd:complexType name="QueryConstraintType" id="QueryConstraintType">
  <xsd:choice>
    <xsd:element ref="ogc:Filter"/>
    <xsd:element name="CqlText" type="xsd:string"/>
  </xsd:choice>
  <xsd:attribute name="version" type="xsd:string" use="required">
  </xsd:attribute>
</xsd:complexType>
```

The **version** parameter may be used to specify a version number indicating which version of a specification the constraint conforms to. For example, in the XML encoding, if the <ogc:Filter> element is being used, the **version** parameter could be set to “1.1.0” indicating that the filter conforms to version 1.1.0 of the *Filter Encoding Implementation Specification* [OGC 04-095].

10.3.5 General model message mapping

Table 56 maps the general model operations, defined in Clause 7, to the Catalogue Service for the Web (CSW) operations. This table does not list the general model operations that are not mapped to CSW operations.

Table 56 — General model to CSW mapping

General Model Operation	CSW Operation
OGC_Service.getCapabilities	OGC_Service.GetCapabilities
Discovery.query	CSW-Discovery.GetRecords
Discovery.present	CSW-Discovery.GetRecordById
Discovery.describeRecordType	CSW-Discovery.DescribeRecord
Discovery.getDomain	CSW-Discovery.GetDomain
Manager.transaction	CSW-Publication.Transaction
Manager.harvestRecords	CSW-Publication.Harvest

10.3.6 Common request parameters

All CSW operation requests except for GetCapabilities shall include the three parameters specified in Table 20 of [05-008c1]. Only one of these parameters is included in the general catalogue model, the others are specific to the HTTP protocol binding.

In KVP encoding, these common parameters in CSW operation requests are encoded as shown in Table 57. Note that the parameter names in all KVP encodings shall be handled in a **case insensitive** manner while parameter values shall be handled in a **case sensitive** manner.

Table 57 — KVP encoding of common operation request parameters

Keyword	Datatype and value	Optionality	Parameter in general model
REQUEST	Character String type Value is operation name (e.g., "DescribeRecord")	One (Mandatory)	(none)
service	Character String type Fixed values of "CSW"	One (Mandatory)	serviceId
version	Character String type Fixed value of "2.0.2"	One (Mandatory)	(none)

In XML encoding, all operation request elements, except for GetCapabilities, are extended from the following XML Schema fragment:

```
<xsd:complexType name="RequestBaseType">
  <xsd:attribute name="service" type="xsd:string"
    use="required" fixed="CSW"/>
  <xsd:attribute name="version" type="ows:ServiceType"
    use="required" fixed="2.0.2"/>
</xsd:complexType>
```

The "service" parameter is used to indicate that the request is a CSW request. This parameter shall be specified for all CSW requests. The "version" parameter is used to indicate that the associated CSW request conforms to this specification. This is indicated by setting the value of the **version** parameter to 2.0.2. This XML Schema fragment does

not include a “request” attribute, since the name of the operation requested is always the name of the XML element encoding the request.

10.3.7 Exception reporting

In the event that a catalogue service encounters an error while processing a request or receives an unrecognised request, it shall generate an XML document indicating that an error has occurred. The format of the XML error response is specified by, and shall validate against, the exception response schema defined in clause 8 of the OWS Common Implementation Specification [OGC 05-008c1].

An **<ows:ExceptionReport>** element may contain one or more catalogue processing exceptions specified using the **<ows:Exception>** element. The mandatory **version** attribute is used to indicate the version of the service exception report schema. For this version of the specification, this value is fixed at 1.2.0. The optional **language** attribute may be used to indicate the language used. The code list for the language parameter is defined in IETF RFC 1766.

Individual exception messages are contained within the **<ows:ExceptionText>** element. The mandatory **code** attribute shall be used to associate an exception code with the accompanying message. The optional **locator** attribute may be used to indicate where an exception was encountered in the request that generated the error. A number of elements defined in this document include a **handle** attribute and/or a **requestId** attribute that may be used to associate a mnemonic name with the element. If such a **handle** or **requestId** attribute exists, its value may be reported using the **locator** attribute of the **<ows:ExceptionText>** element in order to correlate the element with the exception text. If the **handle** or **requestId** attribute is not specified, then a catalogue is not required to specify a value for the **locator** attribute or may, optionally, attempt to locate the error using other means such as line numbers.

EXAMPLE The following is an example of an exception message

```
<?xml version="1.0" ?>
<ExceptionReport
  version="1.2.0"
  xmlns="http://www.opengis.net/ows"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.opengis.net/ows
    ../../ows/1.0.0/owsExceptionReport.xsd">
  <Exception exceptionCode="ParsingError" locator="INSERT STMT 01">
  <ExceptionText>parse error: missing closing tag for element</ExceptionText>
  </Exception>
</ExceptionReport>
```

NOTE This example uses the **exceptionCode** value "ParsingError", which has not yet been specified, meaning that an error was detected in validating a XML encoded operation request.

10.4 Operations overview

Figure 29 shows the request/response message pairs for all the operations defined for the web catalogue service (CSW). There are three classes of operations: *service operations*

which are operations a client may use to interrogate the service to determine its capabilities; *discovery operations* which a client may use to determine the information model of the catalogue and query catalogue records; and *management operations* which are used to create or change records in the catalogue.

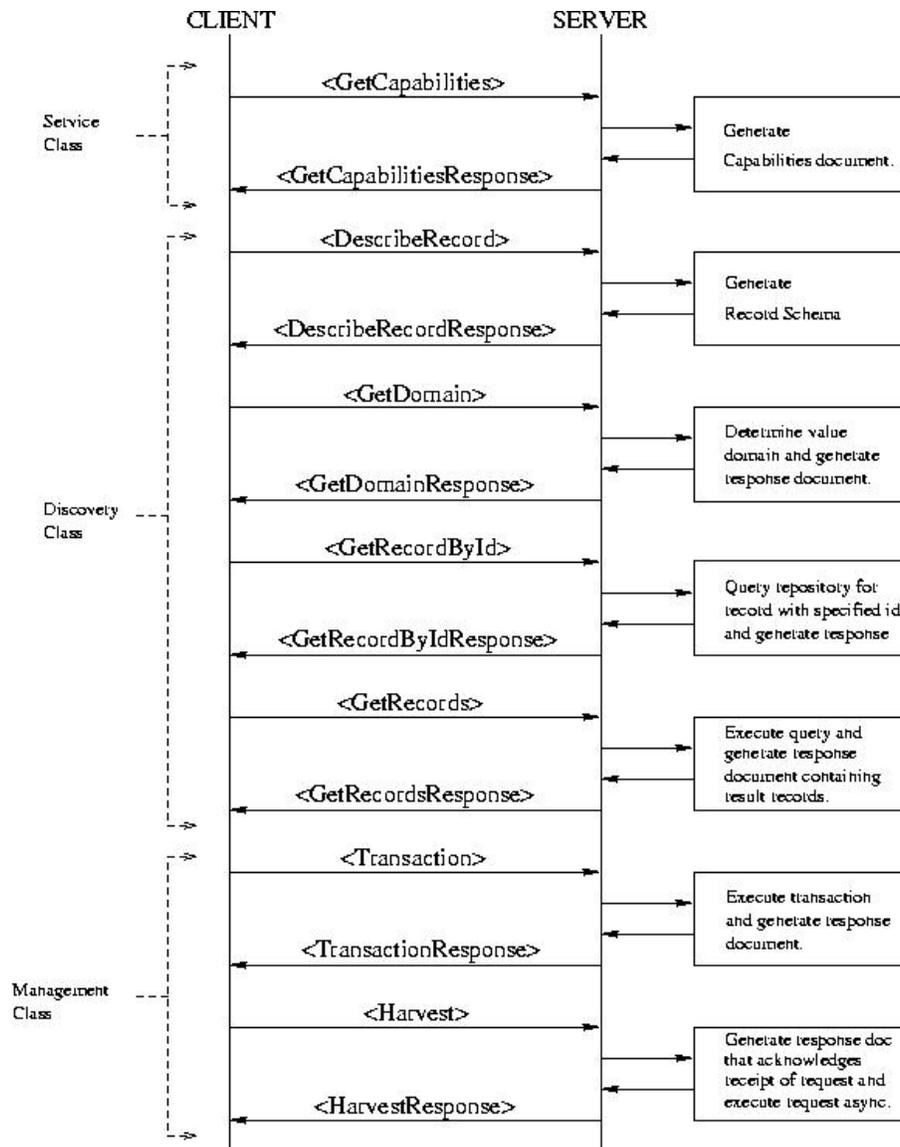


Figure 29 — Protocol sequence diagram

Figure 30 depicts a conceptual architecture to illustrate the relationship of these interfaces to service consumers and producers. The arrows represent the CSW requests that producers and consumers of metadata may generate. For example, to create metadata, a metadata producer may invoke the **Transaction** request or the **Harvest** request. Similarly, a consumer of metadata may invoke the **GetRecords** request to query the catalogue.

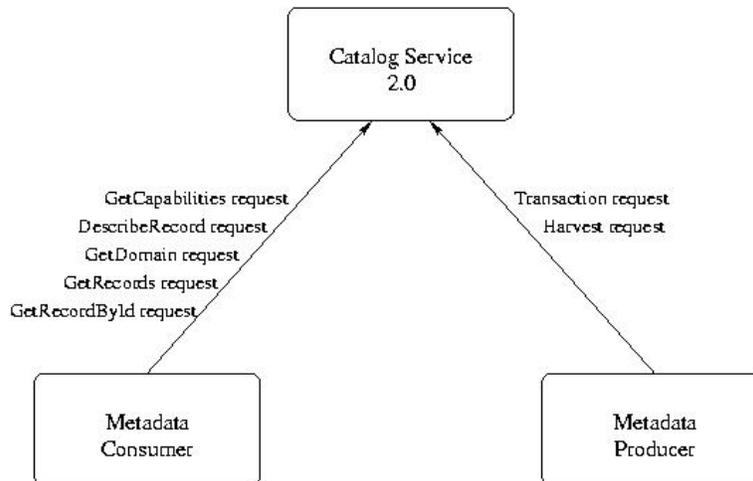


Figure 30 — Conceptual architecture

10.5 GetCapabilities operation

10.5.1 Introduction

The mandatory GetCapabilities operation allows CSW clients to retrieve service metadata from a server. The response to a GetCapabilities request shall be an XML document containing service metadata about the server. This subclause specifies the XML document that a CSW server shall return to describe its capabilities.

10.5.2 Operation request

The GetCapabilities operation request is defined in Subclause 7.2 of the *OGC Web Services Common Specification 1.0* [OGC 05-008]. CSW servers shall implement the *request*, *service*, *Sections*, *AcceptVersions* and *AcceptFormats* operation request parameters, and may implement the *updateSequence* parameter. If the **updateSequence** parameter is not implemented, a catalogue server shall always return the most up-to-date version of the capabilities document. All CSW servers shall implement the HTTP GET transfer using the keyword-value pair encoding of the GetCapabilities operation. Servers may optionally implement the HTTP POST transfer using XML encoding only.

The value of the *service* parameter shall be “CSW” or the identifier of a specific CSW Application Profile as specified in that profile. The common service metadata elements that may be included in a Capabilities document are specified in Subclause 7.4 of OGC 05-008; a catalogue service that implements the CSW binding may also include the elements listed in Table 58. An application profile may introduce additional service information items as needed by extending the `csw:CapabilitiesType` definition.

Table 58 — Additional section name value and meaning

Section name	Meaning
Filter_Capabilities	<p>A Filter_Capabilities section shall be included in the service metadata to describe which elements of the predicate language are supported. All CSW implementations shall support at least the following filter operators:</p> <ul style="list-style-type: none"> • logical operators: And, Or, Not • comparison operators: PropertyIsEqualTo, PropertyIsNotEqualTo, PropertyIsLessThan, PropertyIsGreaterThan, PropertyIsLessThanOrEqualTo, PropertyIsGreaterThanOrEqualTo, PropertyIsLike • spatial operators: BBOX.

10.5.3 Operation response

The service metadata document shall contain the sections specified in Table 59. Depending on the values in the *Sections* parameter of the GetCapabilities operation request, any combination of these sections can be requested and returned. If the Sections parameter is not specified, then all sections shall be returned.

Table 59 — Section names and contents

Section name	Contents
ServiceIdentification	Metadata about a specified CSW implementation. The contents and schema of this section shall be as specified in Subclause 7.4.3 and owsServiceIdentification.xsd of [OGC 05-008c1].
ServiceProvider	Metadata about the organization offering the CSW service. The contents and schema of this section shall be as specified in Subclauses 7.4.4 and owsServiceProvider.xsd of [OGC 05-008c1].
OperationsMetadata	Metadata about the CSW operations offered by a specific CSW implementation, including the URLs for operation requests. The contents and schema of this section shall be as specified in Subclauses 7.4.5 and owsOperationsMetadata.xsd of [OGC 05-008c1]. The specific operations that may be listed in the OperationsMetadata section are specified in Subclause 10.5.4 or this document.
Filter_Capabilities	Metadata about the filter capabilities of the server if the server implements the Filter predicate encoding as defined in [OGC 04-095]..

10.5.4 OperationsMetadata section standard contents

The OperationsMetadata element shall list all operations implemented by the service, as described in Subclause 7.4.5 of OGC 05-008. An application profile may restrict the **<ExtendedCapabilities>** element to provide additional computational metadata (e.g., WSDL service descriptions, OWL-S resource definitions) compliant with this specification. Table 61 lists the optional values of OperationsMetadata section attributes for additional operations that a CSW may offer. If a specified server implementation offers one or more of these operations, they shall be listed in the OperationsMetadata section of the capabilities document. In both tables, the “Attribute name” column uses dot-separator notation to specify parts of a parent item. The “Attribute value” column references an operation parameter, and the meaning of including that value is listed in the right column.

Table 60 specifies the required values of OperationsMetadata section attributes for operations that a CSW server shall implement to be minimally compliant with this specification. Table 61 lists the optional values of OperationsMetadata section attributes for additional operations that a CSW may offer. If a specified server implementation offers one or more of these operations, they shall be listed in the OperationsMetadata section of the capabilities document. In both tables, the “Attribute name” column uses dot-separator notation to specify parts of a parent item. The “Attribute value” column references an operation parameter, and the meaning of including that value is listed in the right column.

Table 60 — Required values of the OperationsMetadata section attributes

Attribute name	Attribute value	Meaning of attribute value
OperationsMetadata.Operation.name	GetCapabilities	The GetCapabilities operation is implemented by this server.
OperationsMetadata.Operation.name	DescribeRecord	The DescribeRecord operation is implemented by this server.
OperationsMetadata.Operation.name	GetRecords	The GetRecords operation is implemented by this server.

Table 61 — Optional values of the OperationsMetadata section attributes

Attribute name	Attribute value	Meaning of attribute value
OperationsMetadata.Operation.name	GetRecordById	The GetRecordById operation is implemented by this server.
OperationsMetadata.Operation.name	GetDomain	The GetDomain operation is implemented by this server.
OperationsMetadata.Operation.name	Harvest	The Harvest operation is implemented by this server.
OperationsMetadata.Operation.name	Transaction	The Transaction operation is implemented by this server.

In addition to the items listed in Table 61, there are many optional values of “name” attributes and “value” elements in the OperationsMetadata section, primarily for recording the domain of various parameters and quantities. For example, the domain of the exceptionCode parameter could record all the codes implemented for each operation by that specific server. Similarly, each of the GetCapabilities operation request parameters might have its domain recorded. For example, the domain of the Sections parameter could record all the sections implemented by that specific server.

CSW servers shall indicate whether or not SOAP encodings of operation requests can be transferred, using HTTP POST as specified in Subclause 10.3.2. This indication shall use an **ows:Constraint** element named **PostEncoding**, within the **ows:OperationsMetadata** section of the Capabilities document. This **PostEncoding** constraint shall specify the formats that can be used with HTTP POST transfer of operation requests. The value **SOAP** for **PostEncoding** shall indicate that SOAP encoded operation requests can be handled, as described in Subclause 6.7. The value **XML** shall indicate that (bare) XML encoded operation requests can be handled. If the connect point URL is the same for all

SOAP-encoded and base-XML operation requests, the **ows:Constraint** element shall be included in the **ows:OperationsMetadata** element. If the connect point URL is different for SOAP-encoded and base-XML operation requests, this **ows:Constraint** element shall be included in each **ows:Post** element.

10.5.5 Examples⁷

KVP Encoding:

<http://www.someserver.com/wrs.cgi?REQUEST=GetCapabilities&SERVICE=CSW&ACCEPTVERSION=2.0.2,2.0.0,0.7.2&outputFormat=application/xml>

XML Encoding:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<GetCapabilities
  xmlns="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:ows="http://www.opengis.net/ows"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.opengis.net/cat/csw/2.0.2
    ../../../../csw/2.0.2/CSW-discovery.xsd"
  service="CSW">
  <ows:AcceptVersions>
    <ows:Version>2.0.2</ows:Version>
    <ows:Version>2.0.0</ows:Version>
    <ows:Version>0.7.2</ows:Version>
  </ows:AcceptVersions>
  <ows:AcceptFormats>
    <ows:OutputFormat>application/xml</ows:OutputFormat>
  </ows:AcceptFormats>
</GetCapabilities>
```

Annex D includes a sample capabilities document for a CSW server that supports only the default message payload. As described in clause 10.2.5, the default message payload is the core queryable elements defined by the **csw:Record** element; this default may be overridden in an application profile. The **csw:Record** element is the root element of the XML encoding of the core queryable element and is define in the schema file record.xsd.

10.6 DescribeRecord operation

10.6.1 Introduction

The mandatory **DescribeRecord** operation allows a client to discover elements of the information model supported by the target catalogue service. The operation allows some or all of the information model to be described.

⁷ All examples in clause 10 are informative. In addition, the examples do not include all the XML syntax required to validate. This is done intentionally so as not to obfuscate the examples with XML syntax.

10.6.2 KVP encoding

Table 62 specifies the keyword-value pair (KVP) encoding for the DescribeRecord operation request. This encoding is suitable for the HTTP GET binding.

NOTE To reduce the need for readers to refer to other parts of this document, the first three parameters listed below are copied from Table 57 in Subclause 10.3.5 of this document.

Table 62 — KVP encoding for DescribeRecord operation request

Keyword ^c	Data type and value	Optionality and use	Parameter in general model
REQUEST	Character String Fixed value of <i>DescribeRecord</i> , case insensitive	One (Mandatory) ^a	(none)
service	Character String Fixed value of <i>CSW</i>	One (Mandatory)	serviceId
version	Character String Fixed value of <i>2.0.2</i>	One (Mandatory)	(none)
NAMESPACE	List of Character String, comma separated Used to specify namespace(s) and their prefix(es) Format is <code>xmlns([prefix]=)namespace-url</code> . If prefix is not specified, then this is the default namespace.	One (Optional) ^b Include declarations for each namespace used in a TypeName	(none)
TypeName	List of Character String, comma separated One or more qualified type names to be described	Zero or one (Optional) Default action is to describe all types known to server	typeName
outputFormat	Character String A MIME type indicating the format that the output document should have	Zero or one (Optional) Default value is application/xml	returnFormat
schemaLanguage	Character String	Zero or one (Optional) Default value is XMLSCHEMA	schemaLanguage
<p>^a The REQUEST parameter contains the same information as the name of the < DescribeRecord> element in XML encoding.</p> <p>^b The NAMESPACE parameter contains the same information as the xmlns attributes which may be used to define and bind namespaces in XML encoding.</p> <p>^c Parameter keywords are case insensitive for KVP encoding. Parameters values are case sensitive..</p>			

10.6.3 XML encoding

The following XML-Schema fragment defines the XML encoding for the **DescribeRecord** operation request:

```
<xsd:element name="DescribeRecord" type="csw:DescribeRecordType"/>
<xsd:complexType name="DescribeRecordType">
  <xsd:complexContent>
```

```

<xsd:extension base="csw:RequestBaseType">
  <xsd:sequence>
    <xsd:element name="TypeName" type="xsd:QName"
      minOccurs="0" maxOccurs="unbounded"/>
  </xsd:sequence>
  <xsd:attribute name="outputFormat" type="xsd:string"
    use="optional"
    default="application/xml"/>
  <xsd:attribute name="schemaLanguage" type="xsd:anyURI"
    use="optional"
    default="http://www.w3.org/XML/Schema"/>
</xsd:extension>
</xsd:complexContent>
</xsd:complexType>

```

10.6.4 Parameter descriptions

10.6.4.1 NAMESPACE parameter

The **DescribeRecord** operation depends on namespace declarations in order to know exactly which types to describe.

For XML encoded **DescribeRecord** requests, the namespace declarations are specified using standard XML conventions (**xmlns** attributes) and described in the document "Namespaces in XML" [<http://www.w3.org/TR/1999/REC-xml-names-19990114>].

For the KVP encoding, namespace declarations are specified using the **NAMESPACE** parameter. The **NAMESPACE** parameter is a comma separated list of namespace declarations of the form `xmlns([prefix=]namespace-url)`.

The following is an example declaration:

```

...NAMESPACE=xmlns(gml=http://www.opengis.org/gml),xmlns(wfs=http://
www.opengis.org/wfs)...

```

The value of the **NAMESPACE** parameter shall be properly escaped for url encoding, which is not shown in this example for the sake of clarity.

10.6.4.2 TypeName parameter

The **TypeName** parameter specifies a list of type names that are to be described by the catalogue. A type name is the name of a queryable entity from the information model of the catalogue (e.g., `ebrim:ExtrinsicObject`).

Every type name shall be fully qualified in order to indicate the target namespace for the type definition. If no type names are provided, then entire schemas from the target namespace are returned. For XML-encoded **DescribeRecord** requests, the namespace declarations are specified using the `targetNamespace` attribute of the `TypeName` element.

If the **DescribeRecord** request is XML encoded, then namespaces shall be declared according to the conventions of XML. If the **DescribeRecord** request is KVP encoded, then the namespaces referenced shall be declared using the **NAMESPACE** parameter.

10.6.4.3 outputFormat parameter

The **outputFormat** parameter specifies the MIME type of the response document. The default output format attribute is the MIME type *application/xml*. All supported output formats shall be declared in the Capabilities document.

10.6.4.4 schemaLanguage parameter

The **schemaLanguage** parameter is used to specify the schema language that should be used to describe the specified types. The default value is *XMLSCHEMA*, which indicates that the XML-Schema, schema description language shall be used. Other schemas languages are possible as long as the required **schemaLanguage** values are declared in the Capabilities document.

10.6.5 Response

The following XML Schema fragment defines the response to a **DescribeRecord** operation when the **schemaLanguage** parameter is set to the value *XMLSCHEMA*.

```
<xsd:element name="DescribeRecordResponse"
  type="csw:DescribeRecordResponseType" />
<xsd:complexType name="DescribeRecordResponseType">
  <xsd:sequence>
    <xsd:element name="SchemaComponent"
      type="csw:SchemaComponentType"
      minOccurs="0" maxOccurs="unbounded" />
  </xsd:sequence>
</xsd:complexType>
<xsd:complexType name="SchemaComponentType" mixed="true">
  <xsd:sequence>
    <xsd:any namespace="##any" processContents="lax" />
  </xsd:sequence>
  <xsd:attribute name="targetNamespace" type="xsd:anyURI"
    use="required" />
  <xsd:attribute name="parentSchema" type="xsd:anyURI"
    use="optional" />
  <xsd:attribute name="schemaLanguage" type="xsd:anyURI"
    use="required" />
</xsd:complexType>
```

The **<DescribeRecordResponse>** element is a container for zero or more **<SchemaComponent>** elements, each of which contains the description of one or more type names in the requested schema language. The **<SchemaComponent>** element may contain any content so that schema language descriptions other than XML Schema may be accommodated. For example, the content could be SQL DDL.

10.6.6 Examples

KVP encoded example

```
http://www.someserver.com/csw/csw.cgi?request=DescribeRecord&version=2.0
.2&outputFormat=application/xml&schemaLanguage=XMLSCHEMA&typeName=csw:Re
cord&namespace=xmlns(csw=http://www.opengis.org/cat/csw)
```

XML encoded example

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<DescribeRecord
  service="CSW"
  version="2.0.2"
  outputFormat="application/xml"
  schemaLanguage="http://www.w3.org/2001/XMLSchema"
  xmlns="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:csw="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.opengis.net/cat/csw/2.0.2
  ../../../../csw/2.0.2/CSW-discovery.xsd">
  <TypeName>csw:Record</TypeName>
</DescribeRecord>
```

10.7 GetDomain operation

10.7.1 Introduction

The optional **GetDomain** operation is used to obtain runtime information about the range of values of a metadata record element or request parameter. The runtime range of values for a property or request parameter is typically much smaller than the value space for that property or parameter based on its static type definition. For example, a property or request parameter defined as a 16bit positive integer in a database may have a value space of 65535 distinct integers but the actual number of distinct values existing in the database may be much smaller.

This type of runtime information about the range of values of a property or request parameter is useful for generating user interfaces with meaningful pick lists or for generating query predicates that have a higher chance of actually identifying a result set.

It should be noted that the **GetDomain** operation is a “best-effort” operation. That is to say that a catalogue tries to generate useful information about the specified request parameter or property if it can. It is entirely possible that a catalogue may not be able to determine anything about the values of a property or request parameter beyond the basic type; in this case only a type reference or a type description will be returned.

10.7.2 KVP encoding

Table 64 specifies the keyword-value pair encoding for the **GetDomain** operation request.

NOTE To reduce the need for readers to refer to other parts of this document, the first three parameters listed below are copied from Table 57 in Subclause 10.3.5 of this document.

Table 63 — KVP encoding for GetDomain operation request

Keyword ^b	Data type and value	Optionality and use	Parameter in general model
REQUEST	Character String Fixed value of "GetDomain", case insensitive	One (Mandatory) ^a	(none)
service	Character String Fixed values of "CSW"	One (Mandatory)	serviceId
version	Character String Fixed value of "2.0.2"	One (Mandatory)	(none)
ParameterName	List of Character String, comma separated Unordered list of names of requested parameters, of the form <i>OperationName.ParameterName</i>	Zero or one (Conditional) Include when PropertyName not included	parameterName
PropertyName	List of Character String, comma separated Unordered list of names of requested properties, from the information model that the catalogue is using	Zero or one (Conditional) Include when ParameterName not included	parameterName
<p>a The REQUEST parameter contains the same information as the name of the <GetDomain> element in XML encoding.</p> <p>b Parameter keywords are case insensitive for KVP encoding. Parameters values are case sensitive..</p>			

10.7.3 XML encoding

The following XML-Schema fragment defines that XML encoding for the **GetDomain** operation request:

```
<xsd:element name="GetDomain" type="csw:GetDomainType" />
<xsd:complexType name="GetDomainType">
  <xsd:complexContent>
    <xsd:extension base="csw:RequestBaseType">
      <xsd:sequence>
        <xsd:choice>
          <xsd:element name="PropertyName" type="xsd:anyURI" />
          <xsd:element name="ParameterName" type="xsd:anyURI" />
        </xsd:choice>
      </xsd:sequence>
    </xsd:extension>
  </xsd:complexContent>
</xsd:complexType>
```

10.7.4 Parameter descriptions

10.7.4.1 PropertyName parameter

The **PropertyName** parameter is used to specify the name of a property that is defined in the information model for which value domain information is desired. An example of a property name might be the name of one of the core queryable properties described in Subclause 6.3.2. The **PropertyName** value may be specified using an absolute or a relative URI; the precise syntax or permissible values are defined in an application profile.

10.7.4.2 ParameterName parameter

The **ParameterName** parameter is used to specify the name of an interface parameter for which value domain information is desired. Table 64 defines the list of interface parameters that may be interrogated using the **GetDomain** operation.

Table 64 — Interface parameters that may be interrogated using GetDomain operation

Parameter Name
GetRecords.resultType
GetRecords.outputFormat
GetRecords.outputRecType
GetRecords.typeName
GerRecords.ElementSetName
GetRecords.ElementName
GetRecords.CONSTRAINTLANGUAGE
GetRecordById.ElementSetName
DescribeRecord.typeName
DescribeRecord.schemaLanguage

10.7.5 Response

The following XML-Schema fragment defines the response to a **GetDomain** operation.

```
<xsd:element name="GetDomainResponse"
type="csw:GetDomainResponseType">
  <xsd:complexType name="GetDomainResponseType">
    <xsd:sequence>
      <xsd:element name="DomainValues"
        type="csw:DomainValuesType"
        maxOccurs="unbounded"/>
    </xsd:sequence>
  </xsd:complexType>
  <xsd:complexType name="DomainValuesType">
    <xsd:sequence>
      <xsd:choice>
        <xsd:element name="PropertyName" type="xsd:QName"/>
      </xsd:choice>
    </xsd:sequence>
  </xsd:complexType>
</xsd:element>
```

```

        <xsd:element name="ParameterName" type="xsd:QName"/>
    </xsd:choice>
    <xsd:choice minOccurs="0">
        <xsd:element name="ListOfValues"
            type="csw:ListOfValuesType"/>
        <xsd:element name="ConceptualScheme"
            type="csw:ConceptualSchemeType" />
        <xsd:element name="RangeOfValues"
            type="csw:RangeOfValuesType" />
    </xsd:choice>
</xsd:sequence>
<xsd:attribute name="type" type="xsd:QName" use="required" />
<xsd:attribute name="uom" type="xsd:anyURI" use="optional" />
</xsd:complexType>
<xsd:complexType name="ListOfValuesType">
    <xsd:sequence>
        <xsd:element name="Value" type="xsd:anyType"
            minOccurs="1" maxOccurs="unbounded"/>
    </xsd:sequence>
</xsd:complexType>
<xsd:complexType name="ConceptualSchemeType">
    <xsd:sequence>
        <xsd:element name="Name" type="xsd:string"/>
        <xsd:element name="Document" type="xsd:anyURI"/>
        <xsd:element name="Authority" type="xsd:anyURI"/>
    </xsd:sequence>
</xsd:complexType>
<xsd:complexType name="RangeOfValuesType">
    <xsd:sequence>
        <xsd:element name="MinValue" type="xsd:anyType"/>
        <xsd:element name="MaxValue" type="xsd:anyType"/>
    </xsd:sequence>
</xsd:complexType>

```

The response is composed of one or more **<DomainValues>** elements. The domain values may be a list of enumerated values (i.e. **<ListOfValues>**), one or more ranges of values (i.e. **<RangeOfValues>**), or a reference to some authoritative vocabulary (i.e. **<ConceptualScheme>**). An example of an authoritative vocabulary might be a standard list of animal and plant species names.

If the only child element of the **<DomainValue>** element is the **<PropertyName>** or **<ParameterName>** element, this shall be taken to mean that the catalogue was unable to determine anything about the specified property or parameter.

10.7.6 Examples

KVP encoded example:

```
http://www.someserver.com/csw/csw.cgi?request=GetDomain&version=2.0.2&parameterName=GetRecords.outputFormat
```

XML encoded example:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
```

```

<GetDomain
  service="CSW"
  version="2.0.2"
  xmlns="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:csw="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.opengis.net/cat/csw/2.0.2
    ../../../../csw/2.0.2/CSW-discovery.xsd">
  <ParameterName>GetRecords.outputFormat</ParameterName>
</GetDomain>

```

10.8 GetRecords operation

10.8.1 Introduction

The primary means of resource discovery in the general model are the two operations **search** and **present**. In the HTTP protocol binding these are combined in the form of the mandatory **GetRecords** operation, which does a search and a piggybacked present.

The search portion of the **GetRecords** operation is encoded using the **Query** element. The **Query** element includes the parameters **typeName** and **Constraint**. The **typeName** parameter is used to specify which entities, from the information model of the catalogue, shall be queried. The **Constraint** parameter is used to specify which query constraints shall be applied to identify the request set.

The present portion of the **GetRecords** operation is encoded using the **outputSchema** parameter and the **ElementName/ElementSetName** parameter(s). The **outputSchema** parameter indicates which schema shall be used to generate the response to the **GetRecords** operation. The **ElementName** or **ElementSetName** parameter is used to specify which properties of the **outputSchema** to include in each record in the **GetRecords** response.

10.8.2 KVP encoding

Table 65 specifies the keyword-value pair encoding for the **GetRecords** operation request. This encoding is suitable for the HTTP GET binding.

NOTE To reduce the need for readers to refer to other parts of this document, the first three parameters listed below are copied from Table 57 in Subclause 10.3.5 of this document.

Table 65 — KVP encoding for GetRecords operation request

Keyword ^d	Data type and value	Optionality and use	Parameter in general model
REQUEST	Character String Fixed value of <i>GetRecords</i> , case insensitive	One (Mandatory) ^a	(none)
service	Character String Fixed values of “CSW”	One (Mandatory)	serviceId
version	Character String Fixed value of 2.0.2	One (Mandatory)	(none)
NAMESPACE	List of Character String, comma separated Used to specify a namespace and its prefix Format shall be xmlns([prefix=]namespace-url). If the prefix is not specified then this is the default namespace.	Zero or one (Optional) ^b Include value for each distinct namespace used by all qualified names in the request. If not included, all qualified names are in default namespace	(none)
resultType	CodeList with allowed values: ”hits”, “results” or ”validate”	Zero or one (Optional) Default value is ”hits”	resultType
requestId	URI	Zero or one (Optional) Include when client chooses to assign requestId	requestId
outputFormat	Character String Value is Mime type The only value that is required to be supported is <i>application/xml</i> . Other supported values may include <i>text/html</i> and <i>text/plain</i>	Zero or one (Optional) Default value is <i>application/xml</i>	returnFormat
outputSchema	Any URI..	Zero or one (Optional) Default value is <i>http://www.opengis.net/cat/csw/2.0.2</i> .	responseSchema
startPosition	Non-Zero Positive Integer	Zero or one (Optional) Default value is 1	cursorPosition
maxRecords	Positive Integer	Zero or one (Optional) Default values is 10	iteratorSize

Keyword^d	Data type and value	Optionality and use	Parameter in general model
typeName	List of Character String, comma separated Unordered List of object types implicated in the query	One (Mandatory)	collectionID
ElementSetName	List of Character String	Zero or one (Mutually exclusive with ElementName)	responseElements
ElementName	List of Character String	Zero or more (Mutually exclusive with ElementSetName)	responseElements
CONSTRAINTLANGUAGE	CodeList with allowed values: <i>CQL_TEXT</i> or <i>FILTER</i>	Zero or one (Optional) ^e Include when Constraint included	queryLanguage
Constraint	Character String Predicate expression specified in the language indicated by the CONSTRAINTLANGUAGE parameter	Zero or one (Optional) Default action is to execute an unconstrained query	predicate
SortBy	List of Character String, comma separated Ordered list of names of metadata elements to use for sorting the response Format of each list item is <i>metadata_element_name:A</i> indicating an ascending sort or <i>metadata_element_name:D</i> indicating descending sort	Zero or one (Optional) Default action is to present the records in the order in which they are retrieved	sortField and sortOrder
DistributedSearch	Boolean	Zero or one (Optional) Default value is FALSE	queryScope
hopCount	Integer	Zero or one (Optional) Include only if DistributedSearch parameter is included Default value is 2	queryScope
ResponseHandler	Any URI	Zero or one (Optional) If not included, process request synchronously	responseHandler

Keyword ^d	Data type and value	Optionality and use	Parameter in general model
a	The REQUEST parameter contains the same information as the name of the <GetRecords> element in XML encoding.		
b	The NAMESPACE parameter contains the same information as the xmlns attributes which may be used for encoding namespace information in XML encoding.		
c	The CONSTRAINTLANGUAGE parameter contains the same information as the root element of the content of the <Constraint> element indicates the predicate language being used in XML encoding.		
d	Parameter keywords are case insensitive for KVP encoding. Parameters values are case sensitive..		

10.8.3 XML encoding

The following XML-Schema fragment defines the XML encoding of the **GetRecords** operation request:

```

<xsd:element name="GetRecords" type="csw:GetRecordsType"
  id="GetRecords"/>
<xsd:complexType name="GetRecordsType" id="GetRecordsType">
  <xsd:complexContent>
    <xsd:extension base="csw:RequestBaseType">
      <xsd:sequence>
        <xsd:element name="DistributedSearch"
          type="csw:DistributedSearchType"
          minOccurs="0"/>
        <xsd:element name="ResponseHandler"
          type="xsd:anyURI"
          minOccurs="0" maxOccurs="unbounded"/>
        <xsd:choice>
          <xsd:element ref="csw:AbstractQuery"/>
          <xsd:any processContents="strict"
            namespace="##other" />
        </xsd:choice>
      </xsd:sequence>
      <xsd:attribute name="requestId" type="xsd:anyURI"
        use="optional" />
      <xsd:attribute name="resultType" type="csw:ResultType"
        use="optional" default="hits"/>
      <xsd:attributeGroup ref="csw:BasicRetrievalOptions"/>
    </xsd:extension>
  </xsd:complexContent>
</xsd:complexType>

<xsd:complexType name="DistributedSearchType">
  <xsd:attribute name="hopCount" type="xsd:positiveInteger"
    use="optional" default="2"/>
</xsd:complexType>

<xsd:attributeGroup name="BasicRetrievalOptions">
  <xsd:attribute name="outputFormat" type="xsd:string"
    use="optional" default="application/xml"/>
  <xsd:attribute name="outputSchema" type="xsd:anyURI"
    use="optional"/>
  <xsd:attribute name="startPosition" type="xsd:positiveInteger"
    use="optional" default="1"/>
  <xsd:attribute name="maxRecords" type="xsd:nonNegativeInteger"

```

```

        use="optional" default="10"/>
</xsd:attributeGroup>

<xsd:simpleType name="ResultType" id="ResultType">
  <xsd:restriction base="xsd:string">
    <xsd:enumeration value="results"/>
    <xsd:enumeration value="hits"/>
    <xsd:enumeration value="validate"/>
  </xsd:restriction>
</xsd:simpleType>

<xsd:element name="Query" type="csw:QueryType"
  substitutionGroup="csw:AbstractQuery"/>
<xsd:complexType name="QueryType">
  <xsd:complexContent>
    <xsd:extension base="csw:AbstractQueryType">
      <xsd:sequence>
        <xsd:choice>
          <xsd:element ref="csw:ElementSetName"/>
          <xsd:element name="ElementName"
            type="xsd:QName"
            minOccurs="1" maxOccurs="unbounded"/>
        </xsd:choice>
        <xsd:element ref="csw:Constraint"
          minOccurs="0" maxOccurs="1"/>
        <xsd:element ref="ogc:SortBy"
          minOccurs="0" maxOccurs="1"/>
      </xsd:sequence>
      <xsd:attribute name="typeNameNames"
        type="csw:TypeNameListType" use="required"/>
    </xsd:extension>
  </xsd:complexContent>
</xsd:complexType>

<xsd:simpleType name="TypeNameListType">
  <xsd:list itemType="xsd:QName"/>
</xsd:simpleType>

<xsd:element name="Constraint" type="csw:QueryConstraintType"/>
<xsd:complexType name="QueryConstraintType">
  <xsd:choice>
    <xsd:element ref="ogc:Filter"/>
    <xsd:element name="CqlText" type="xsd:string"/>
  </xsd:choice>
  <xsd:attribute name="version" type="xsd:string" use="required"/>
</xsd:complexType>

<xsd:element name="ElementSetName" type="csw:ElementSetNameType"
  default="summary"/>
<xsd:complexType name="ElementSetNameType" id="ElementSetNameType">
  <xsd:simpleContent>
    <xsd:extension base="csw:ElementSetType">
      <xsd:attribute name="typeNameNames"
        type="csw:TypeNameListType" use="optional"/>
    </xsd:extension>
  </xsd:simpleContent>

```

```

</xsd:complexType>

<xsd:simpleType name="ElementSetType" id="ElementSetType">
  <xsd:restriction base="xsd:string">
    <xsd:enumeration value="brief"/>
    <xsd:enumeration value="summary"/>
    <xsd:enumeration value="full"/>
  </xsd:restriction>
</xsd:simpleType>

```

10.8.4 Parameter descriptions

10.8.4.1 NAMESPACE parameter

The **NAMESPACE** parameter is included in the KVP encoding to allow clients to bind any namespace prefixes that might be used for qualified names specified in other parameters. For example, the **typeName** parameter may include qualified names of the form *namespace prefix:name*.

The value of the **NAMESPACE** parameter is a comma separated list of character strings of the form *xmlns([prefix=]namespace url)*. Not including the name namespace prefix binds the specified URL to the default namespace. As in XML, only one default namespace may be bound.

This parameter is not required for the XML encoding since XML includes a mechanism for binding namespace prefixes.

10.8.4.2 requestId parameter

The **requestId** parameter may be included to uniquely identify this request message. If the client specifies a **requestId** in its request, the server shall use it in its responses. If the client does not specify a **requestId** but includes a **ResponseHandler** in its request and if the service supports asynchronous processing, the server shall generate a **requestId** and put it in the acknowledgement message and in the response. If the client specifies neither a **requestId** nor a **ResponseHandler** in its request, the server shall not put a **requestId** in its response. Whether generated by the client or by the server, the **requestId** specified shall include a UUID (Universal Unique Identifier), generated using the mechanism described in the X.667 specification. The **ResponseHandler** should use **requestId** as part of its value.

10.8.4.3 resultType parameter

The **resultType** parameter may have the value “hits”, “results”, or “validate”; the value determines whether the catalogue service returns just a summary of the result set, includes one or more records from the result set, or validates the request message and processes it asynchronously.

If the **resultType** parameter is set to “hits”, the catalogue service shall return a **<GetRecordsResponse>** element containing an empty **<SearchResults>** element that

indicates the estimated size of the result set. Optional attributes may or may not be set accordingly.

If the **resultType** parameter is set to “results”, the catalogue service shall include any matching records within the <**SearchResults**> element, up to the maximum number of records specified in the request.

If the **resultType** parameter is set to “validate”, the catalogue service shall validate the request and return an <**Acknowledgement**> message if validation succeeds. An <**ows:ExceptionReport**> element, as described in Subclause 10.3.7, is returned if validation fails. If the catalogue supports asynchronous query processing, the acknowledgment response shall include a **RequestId** element that may be subsequently used to retrieve the result set when processing is complete.

10.8.4.4 outputFormat parameter

The **outputFormat** parameter is used to control the format of the output that is generated in response to a **GetRecords** request. Its value shall be a MIME type. The default value, “**application/xml**”, means that the output shall be an XML document. All catalogs shall at least support XML as an output format. Other output formats may be supported and may include output formats such as TEXT (MIME type *text/plain*), HTML (MIME type *text/html*) or XML without an associated schema (MIME type *text/xml*). The list of output formats that a CSW instance provides shall be advertised in the Capabilities document.

In the case where the output format is *application/xml*, the CSW shall generate an XML document that validates against a schema document that is specified in the output document via the **xsi:schemaLocation** attribute defined in XML.

10.8.4.5 outputSchema parameter

The **outputSchema** parameter is used to indicate the schema of the output that is generated in response to a **GetRecords** request. The default value for this parameter shall be:

<http://www.opengis.net/cat/csw/2.0.2>

indicating that the schema for the core returnable properties (as defined in Subclause 10.2.5) shall be used. Application profiles may define additional values for **outputSchema**, but all profiles shall support the value

<http://www.opengis.net/cat/csw/2.0.2>.

Although the value of this parameter can be any URI, any additional values defined for the **outputSchema** parameter should be the target namespace of the additionally supported output schemas and should include a version number. For example, the value *urn:oasis:names:tc:ebxml-regrep:rim:xsd:2.5* might be used to indicate an ebRIM v2.5 output schema, while the value *urn:oasis:names:tc:ebxml-regrep:rim:xsd:3.0* might be used to indicate an ebRIM v3.0 output schema.

The list of supported output schemas shall be advertised in the capabilities document of the service using the Parameter element as outlined in OGC 05-008.

10.8.4.6 startPosition parameter

The **startPosition** parameter is used to indicate at which record position the catalogue should start generating output. The default value is *1* meaning it starts at the first record in the result set.

10.8.4.7 maxRecords attribute

The **maxRecords** parameter is used to define the maximum number of records that should be returned from the result set of a query. If it is not specified, then 10 records shall be returned. If its value is set to zero, then the behaviour is identical to setting “*resultType=hits*” as described in subclause 10.8.4.2.

10.8.4.8 typeNames parameter

The **typeNames** parameter is a list of one or more names of queryable entities in the catalogue's information model that may be constrained in the predicate of the query. In the case of XML realization of the OGC core metadata properties (Subclause 10.2.5), the element *csw:Record* is the only queryable entity. Other information models may include more than one queryable component. For example, queryable components for the XML realization of the ebRIM include *rim:Service*, *rim:ExtrinsicObject* and *rim:Association*. In such cases the application profile shall describe how multiple **typeNames** values should be processed.

In addition, all or some of the these queryable entity names may be specified in the query to define which metadata record elements the query should present in the response to the **GetRecords** operation.

NOTE The **typeNames** parameter is different from the Type core queryable metadata property defined in Subclause 6.3. The **typeNames** parameter is composed of one or more names of queryable entities in the information model of the catalogue. The core queryable Type is used to indicate the type or class of a resource being described by the catalogue. Typically the value of the Type property is taken from some controlled vocabulary.

10.8.4.9 ElementName or ElementSetName parameter

The **ElementName** parameter is used to specify one or more metadata record elements, from the output schema specified using the **outputSchema** parameter, that the query shall present in the response to the a **GetRecords** operation. Since clause 10.2.5 realizes the core metadata properties using XML schema, the value of the **ElementName** parameter would be an XPath expression perhaps using qualified names. In the general case, a complete XPath expression may be required to correctly reference an element in the information model of the catalog.

EXAMPLE 1 Using an XPath expression to reference a property in the information model of a catalogue.

```
... <csw:Query typeNames="rim:ExtrinsicObject rim:Association">
    <csw:ElementName>/rim:ExtrinsicObject/@status</csw:ElementName>
    <csw:ElementName>/rim:ExtrinsicObject/@home</csw:ElementName>
```

```
...
</csw:Query> ...
```

However, in the case where the `typeNames` attribute on the `Query` element contains a single value, the catalogue can infer the first step in the path expression and it can be omitted. This is usually the case when querying the core metadata properties since the only queryable target is `csw:Record`.

EXAMPLE 2 A query example where the first step of the path expression is inferred.

```
... <csw:Query typeNames="csw:Record">
  <csw:ElementName>dc:identifier</csw:ElementName>
  <csw:ElementName>dct:modified</csw:ElementName>
  ...
</csw:Query> ...
```

If the metadata record element names are not from the schema specified using the **outputSchema** parameter, then the service shall raise an exception as described in Subclause 10.3.7.

As mentioned in Subclause 10.8.4.4, if the **outputFormat** parameter is set to *application/xml*, then the response to the **GetRecords** operation shall validate against a schema document that is referenced in the response using the **xmlns:schemaLocation** attribute. If the set of metadata record elements that the client specifies in the query is insufficient to generate a valid XML response document, a CSW shall augment the list of elements presented to the client in order to be able to generate a valid XML document. Thus a client application should expect to receive more than the requested elements if the output format is set to generate an XML document.

Well known sets of elements may be named, in which case the **ElementSetName** parameter can be used (e.g., *brief*, *summary* or *full*) to indicate which named set the service shall present to the client. In the case where the query includes more than one entity name as the value of the **typeName** attribute on the **Query** element, the **typeName** attribute on the **ElementSetName** element can be used to discriminate which element set or sets should be presented. The names specified for the **typeName** attribute on the **ElementSetName** element shall be a proper subset of the names specified as the value of the **typeName** attribute on the **Query** element. If the **typeName** attribute is not included on the **ElementSetName** element, then the named element sets for all entities specified as the value of the **typeName** attribute on the **Query** element shall be presented. If any entity name specified as a value of the **typeName** attribute on the **ElementSetName** element does not match a name specified as a value of the **typeName** attribute on the **Query** element, then the server shall raise an exception as described in Subclause 10.3.7.

EXAMPLE 3 The following XML fragment shows how the `typeName` attribute on the `ElementSet` element can be used to indicate which element set should be presented when more than one entity in the information model of the catalogue is being queried. In this case, the query is performing a join between the `rim:Service`, `rim:Classification` and `rim:Association` entities but presenting the `Brief` element set of the `rim:Service` entity.

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<GetRecords
  service="CSW"
  version="2.0.2"
```

```

outputFormat="application/xml "
outputSchema="urn:oasis:names:tc:ebxml-regrep:xsd:rim:3.0 "
xmlns="http://www.opengis.net/cat/csw/2.0.2"
xmlns:csw="http://www.opengis.net/cat/csw/2.0.2"
xmlns:rim="urn:oasis:names:tc:ebxml-regrep:xsd:rim:3.0"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://www.opengis.net/cat/csw/2.0.2
                    ../../csw/2.0.2/CSW-discovery.xsd
                    urn:oasis:names:tc:ebxml-regrep:xsd:rim:3.0
                    ../../ebrim/3.0/schema/rim.xsd">
<Query typeNameNames="rim:Service rim:Classification rim:Association">
  <ElementSetName typeNameNames="rim:Service">brief</ElementSetName>
  <!-- ... -->
</Query>
</GetRecords>

```

The ElementName and ElementSetName parameters are mutually exclusive. Either an ElementSetName parameter OR one or more ElementSetName parameters shall be specified in a query.

10.8.4.10 Predicate languages

Each request encoding (XML and KVP) has a specific mechanism for specifying the predicate language that will be used to constrain a query.

In the XML encoding, the element **<csw:Constraint>** is used to define the query predicate. The root element of the content of the **<csw:Constraint>** element defines the predicate language that is being used. Two possible root elements are **<ogc:Filter>** for the OGC XML filter encoding, and **<csw:CqlText>** for a common query language string. An example predicate specification in the XML encoding is:

```

<Constraint>
  <CqlText>prop1!=10</CqlText>
</Constraint>

```

or, using the a Filter expression:

```

<Constraint>
  <ogc:Filter>
    <ogc:Not>
      <ogc:PropertyIsEqualTo>
        <ogc:PropertyName>prop1</ogc:PropertyName>
        <ogc:Literal>10</ogc:Literal>
      </ogc:PropertyIsEqualTo>
    </ogc:Not>
  </ogc:Filter>
</Constraint>

```

In the KVP encoding, the parameter **CONSTRAINTLANGUAGE** is used to specify the predicate language being used.

The **Constraint** parameter is used to specify the actual predicate. For example, to specify a CQL predicate, the following parameters would be set in the KVP encoding:

```
...CONSTRAINTLANGUAGE=CQL_TEXT&CONSTRAINT="prop1!=10"...
```

The value of the CONSTRAINT parameter in the KVP encoding may also include an XML encoded Filter expression although some clients may impose length limitations.

10.8.4.11 OGC filter syntax

10.8.4.11.1 Introduction

The XML implementation of the BNF in Subclause 6.3.2 may be found in OGC document 04-095, Filter Encoding Implementation Specification. The intent of the XML encoding of the OGC common query language is that it be easily parsable using readily available XML parsers and be easily translatable into a target predicate language such as a SQL where clause or an XQuery predicate.

10.8.4.11.2 Provide functional extensibility

One feature of the OGC common query language and the XML implementation is that the predicate language is functionally extensible. This means that functions may be added to the filter predicate language without having to change the underlying schema. The relevant schema fragment from OGC 04-095 is:

```
<xsd:element name="Function"
  type="ogc:FunctionType"
  substitutionGroup="ogc:expression"/>
<xsd:complexType name="FunctionType">
  <xsd:complexContent>
    <xsd:extension base="ogc:ExpressionType">
      <xsd:sequence>
        <xsd:element ref="ogc:expression"
          minOccurs="0" maxOccurs="unbounded"/>
      </xsd:sequence>
      <xsd:attribute name="name" type="xsd:string" use="required"/>
    </xsd:extension>
  </xsd:complexContent>
</xsd:complexType>
```

According to the schema fragment, any function may added to the filter predicate language simply by specifying its name and including zero or more ogc expressions as content of the **<Function>** element which represent the arguments of the function. The following example shows how a function may be called using the filter syntax:

```
<Function name="MAX">
  <PropertyName>DEPTH</PropertyName>
</Function>
```

In this example, the MAX() function is invoked to find the maximum value of the property DEPTH.

Any function may be called using the filter syntax as long as the function is advertised in the filter capabilities section (Clause 16 of OGC 04-095) of an OGC capabilities document.

10.8.4.11.3 Precedence

The XML notation does not provide parentheses to indicate operator precedence as specified in the BNF. The Filter Specification uses the nested structure of the XML notation to indicate this relationship.

In this example, a more complex scalar predicate is encoded using the logical operators AND and OR. The example is equivalent to the expression ((FIELD1=10 OR FIELD1=20) AND (STATUS="VALID")):

```
<Filter xmlns="http://http://www.opengis.net/ogc"
  xmlns:foo="http://www.someverser.com/foo/">
  <And>
    <Or>
      <PropertyIsEqualTo>
        <PropertyName>foo:FIELD1</PropertyName>
        <Literal>10</Literal>
      </PropertyIsEqualTo>
      <PropertyIsEqualTo>
        <PropertyName>foo:FIELD1</PropertyName>
        <Literal>20</Literal>
      </PropertyIsEqualTo>
    </Or>
    <PropertyIsEqualTo>
      <PropertyName>foo:STATUS</PropertyName>
      <Literal>VALID</Literal>
    </PropertyIsEqualTo>
  </And>
</Filter>
```

10.8.4.11.4 Tight and loose queries

The following examples show the implementation of tight and loose queries using the Filter Grammar.⁸ In both cases the query is directed to a federation of image catalogs that are compliant to an application profile where an optional searchable attribute “cloudcover” is defined as the percentage of the target obscured by clouds.

In Case 1 the investigator wants to get only images he is sure he can use so he requests only images where the cloud cover is less than 5 percent. This is the normal case for querying known schema.

```
<ogc:Filter xmlns:ogc="http://http://www.opengis.net/ogc">
  <ogc:PropertyIsLessThan>
    <ogc:PropertyName>cloudcover</ogc:PropertyName>
    <ogc:Literal>5</ogc:Literal>
  </ogc:PropertyIsLessThan>
</ogc:Filter>
```

⁸ This capability is included in future version 1.2 of the OGC Filter Encoding Specification.

In Case 2 the investigator is aware that 5% cloud cover is very rare over the target area but he requires only images with less than 5% cloud cover. In this case he wants any images he might be able to use so he requests images which meets his criteria and images where cloud cover is unknown because the catalogue has chosen not to include cloudcover in its searchable attribute set.

```
<ogc:Filter xmlns:ogc="http://www.opengis.net/ogc">
  <ogc:Or>
    <ogc:PropertyIsLessThan>
      <ogc:PropertyName>cloudcover</ogc:PropertyName>
      <ogc:Literal>5</ogc:Literal>
    </ogc:PropertyIsLessThan>
    <ogc:PropertyValueDoesNotExist>
      <ogc:PropertyName>cloudcover</ogc:PropertyName>
    </ogc:PropertyValueDoesNotExist>
  </ogc:Or>
</ogc:Filter>
```

10.8.4.11.5 Property references

The **<ogc:PropertyName>** element is used in Filter expressions to reference properties from the information model of a catalogue. If the catalogue's information model is being realized as an XML document with an accompanying XML Schema, then XPath expressions shall be used to reference the various metadata properties.

This is the case when the core metadata properties that are realized as an XML document, using XML Schema, as described in Subclause 10.2.5. This implies the XPath expressions shall be used to reference the various core metadata properties in Filter expressions. This is also the case for predicates encoded using CQL text if the predicates reference the XML realization of the core metadata properties.

EXAMPLE 1 The following examples shows an XPath expression being used in a Filter to reference one of the core queryable properties described in Subclause 10.2.5.

```
... <csw:Query typeName="csw:Record">
  <csw:Constraint>
    <ogc:Filter>
      <ogc:PropertyIsBetween>
        <ogc:PropertyName>/csw:Record/dct:modified</ogc:PropertyName>
        <ogc:LowerBoundary>
          <ogc:Literal>2006-01-01</ogc:Literal>
        </ogc:LowerBoundary>
        <ogc:UpperBoundary>
          <ogc:Literal>2006-05-30</ogc:Literal>
        </ogc:UpperBoundary>
      </ogc:PropertyIsBetween>
    </ogc:Filter>
  </csw:Constraint>
</csw:Query> ...
```

This following is the same expression represented using CQL text but referencing the XML realization of the core metadata properties.

```

... <csw:Query typeName="csw:Record">
  <csw:Constraint>
    <CqlText>/csw:Record/dct:modified BETWEEN 2006-01-01 AND 2006-05-
30</CqlText>
  </csw:Constraint>
</csw:Query> ...

```

In certain cases, such as when the **typeName** attribute on the **Query** element only contains the name of a single entity, the root path step may be omitted since the catalogue is able to infer what the first step in the path would be.

EXAMPLE 2 This is same example as above but using an inferred initial path step.

```

... <csw:Query typeName="csw:Record">
  <csw:Constraint>
    <ogc:Filter>
      <ogc:PropertyIsBetween>
        <ogc:PropertyName>dct:modified</ogc:PropertyName>
        <ogc:LowerBoundary>
          <ogc:Literal>2006-01-01</ogc:Literal>
        </ogc:LowerBoundary>
        <ogc:UpperBoundary>
          <ogc:Literal>2006-05-30</ogc:Literal>
        </ogc:UpperBoundary>
      </ogc:PropertyIsBetween>
    </ogc:Filter>
  </csw:Constraint>
</csw:Query> ...

```

The CQL text encoding of the same query would be:

```

...<csw:Query typeName="csw:Record">
  <csw:Constraint>
    <CqlText>dct:modified BETWEEN 2006-01-01 AND 2006-05-30</CqlText>
  </csw:Constraint>
</csw:Query> ...

```

Annex F describes the minimum subset of XPath that catalogues implementing the HTTP protocol binding shall support.

10.8.4.12 SortBy parameter

The result set may be sorted by specifying one or more metadata record elements upon which to sort.

In KVP encoding, the **SORTBY** parameter is used to specify the list of sort elements. The value for the **SORTBY** parameter is a comma-separated list of metadata record element names upon which to sort the result set. The format for each element in the list shall be either *element name:A* indicating that the element values should be sorted in ascending order or *element name:D* indicating that the element values should be sorted in descending order.

For XML encoded requests, the `<ogc:SortBy>` element is used to specify a list of sort metadata record elements. The attribute `sortOrder` is used to specify the sort order for each element. Valid values for the `sortOrder` attribute are *ASC* indicating an ascending sort and *DESC* indicating a descending sort.

10.8.4.13 DistributedSearch parameter

The **DistributedSearch** parameter may be used to indicate that the query should be distributed. The default query behaviour, if the **DistributedSearch** parameter is set to *FALSE* (or is not specified at all), is to execute the query on the local server. In the XML encoding, if the `<DistributedSearch>` element is not specified then the query is executed on the local server.

The **hopCount** parameter controls the distributed query behaviour by limiting the maximum number of message hops before the search is terminated. Each catalogue decrements this value by one when the request is received and does not propagate the request if the `hopCount=0`.

Catalogues may advertise, in the capabilities document, to which other catalogues a query is distributed using an operation constraint called “FederatedCatalogues”. Operation constraints are described in Subclause 7.4.5 of OGC 05-008c1.

EXAMPLE The following XML fragment shows how the FederatedCatalogues constraint can be used to list the catalogues on which distributed queries are executed.

```
<ows:OperationsMetadata>
  .
  .
  .
  <ows:Constraint name="FederatedCatalogues">
    <ows:Value>http://www.mycatalogue.com</ows:Value>
    <ows:Value>http://www.yourcatalogue.com</ows:Value>
    <ows:Value>http://www.theotherguyscatalogue.com</ows:Value>
  </ows:Constraint>
</ows:OperationsMetadata>
```

10.8.4.14 ResponseHandler parameter

The **ResponseHandler** parameter is a flag that indicates how the **GetRecords** operation should be processed by a CSW.

If the parameter is not present, then the **GetRecords** operation is processed synchronously meaning that the client sends the **GetRecords** request to a CSW and waits to receive a valid response or exception message as described in Subclause 10.3.7. The CSW immediately processes the **GetRecords** request while the client waits for a response. The problem with this mode of operation is that the client may timeout waiting for the CSW to process the request.

If the **ResponseHandler** parameter is present, the **GetRecords** operation is processed asynchronously. In this case, the CSW responds immediately to a client's request with an acknowledgment message that tells the client that the request has been received and validated, and notification of completion will be sent to the URI specified as the value of

the **ResponseHandler** parameter. The following XML Schema fragment defines the acknowledgement message:

```
<xsd:element name="Acknowledgement" type="csw:AcknowledgementType"/>
<xsd:complexType name="AcknowledgementType">
  <xsd:sequence>
    <xsd:element name="EchoedRequest"
      type="csw:EchoedRequestType" />
    <xsd:element name="RequestId"
      type="xsd:anyURI" minOccurs="0"/>
  </xsd:sequence>
  <xsd:attribute name="timeStamp" type="xsd:dateTime"
use="required"/>
</xsd:complexType>
<xsd:complexType name="EchoedRequestType">
  <xsd:sequence>
    <xsd:any namespace="##any" processContents="lax" />
  </xsd:sequence>
</xsd:complexType>
```

The acknowledgment message shall echo the exact XML text of the client's request, using the **<EchoedRequest>** element, and may include an optionally generated request identifier using the **<RequestId>** element. The echoed request is used to correlate the acknowledgement message with the originating request.

The **GetRecords** request may then be processed at some later time, taking as much time as is required to complete the operation. When the operation is completed, a **<csw:GetRecordsResponse>** message or an exception message per Subclause 10.3.7 (if a problem was encountered) shall be sent to the URI specified as the value of the **ResponseHandler** parameter using the protocol encoded therein. Common protocols include *ftp* and *mailto*.

10.8.5 Response

The following XML-Schema fragment defines the XML format response to a **GetRecords** operation:

```
<xsd:element name="GetRecordsResponse"
  type="csw:GetRecordsResponseType"
  id="GetRecordsResponse"/>
<xsd:complexType name="GetRecordsResponseType">
  <xsd:sequence>
    <xsd:element name="RequestId"
      type="xsd:anyURI" minOccurs="0"/>
    <xsd:element name="SearchStatus"
      type="csw:RequestStatusType"/>
    <xsd:element name="SearchResults"
      type="csw:SearchResultsType"/>
  </xsd:sequence>
  <xsd:attribute name="version" type="xsd:string" use="optional"/>
</xsd:complexType>
<xsd:complexType name="RequestStatusType" id="RequestStatusType">
```

```

    <xsd:attribute name="timestamp"
                  type="xsd:dateTime" use="optional"/>
</xsd:complexType>

<xsd:complexType name="SearchResultsType" id="SearchResultsType">
  <xsd:sequence>

    <xsd:choice>
      <xsd:element ref="csw:AbstractRecord"
                  minOccurs="0" maxOccurs="unbounded" />
      <xsd:any processContents="strict" namespace="##other"
              minOccurs="0" maxOccurs="unbounded" />
    </xsd:choice>
  </xsd:sequence>

  <xsd:attribute name="resultSetId"
                 type="xsd:anyURI"
                 use="optional"/>
  <xsd:attribute name="elementSet"
                 type="csw:ElementSetType"
                 use="optional"/>
  <xsd:attribute name="recordSchema"
                 type="xsd:anyURI"
                 use="optional"/>
  <xsd:attribute name="numberOfRecordsMatched"
                 type="xsd:nonNegativeInteger" use="required"/>
  <xsd:attribute name="numberOfRecordsReturned"
                 type="xsd:nonNegativeInteger" use="required"/>
  <xsd:attribute name="nextRecord"
                 type="xsd:nonNegativeInteger" use="required"/>
  <xsd:attribute name="expires"
                 type="xsd:dateTime"
                 use="optional"/>
</xsd:complexType>

```

The **<GetRecordsResponse>** element is a container for the response of the **GetRecords** operation. Three levels of detail may be contained in the response document.

The **<RequestId>** element may be used to correlate the response to a **GetRecords** request for which a value was defined for the **requestId** attribute.

The **<SearchStatus>** element shall be present and indicates the status of the response. The search status consists of a timestamp attribute indicating when the result set was created.

The **<SearchResults>** element is a generic XML container for the actual response to a **GetRecords** request. The content of the **<SearchResults>** element is the set of records returned by the **GetRecords** operation. Table 66 describes the attributes that can appear on the **<SearchResults>** element.

Table 66 — <SearchResults> Parameters

Parameters	Description
resultSetId	A server-generated identifier for the result set. May be used in subsequent GetRecords operations to further refine the result set. If the server does not implement this capability then the attribute should be omitted.
elementSet	The element set returned (brief, summary or full).
recordSchema	A reference to the type or schema of the records returned.
numberOfRecordsMatched	Number of records found by the GetRecords operation.
numberOfRecordsReturned	Number of records actually returned to client. This may not be the entire result set since some servers may limit the number of records returned to limit the size of the response package transmitted to the client. Subsequent queries may be executed to see more of the result set. The <i>nextRecord</i> attribute will indicate to the client where to begin the next query.
nextRecord	Start position of next record. A value of 0 means all records have been returned.
expires	An ISO 8601 format date indicating when the result set will expire. If this value is not specified then the result set expires immediately.

10.8.6 Examples

KVP encoded example

```
http://www.someserver.com/csw/csw.cgi?request=GetRecords&version=2.0.2&
outputFormat=application/xml&outputSchema=http://www.opengis.net/cat/csw/2.0.2&namespace=csw:http://www.opengis.org/cat/csw&ResponseHandler="mailto:pvretano@cubewerx.com"&typeName=csw:Record&elementSetName=brief&constraintlanguage=QLTEXT&constraint="csw:AnyText Like '%pollution%'"
```

XML encoded request

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<GetRecords
  service="CSW"
  version="2.0.2"
  maxRecords="5"
  startPosition="1"
  resultType="results"
  outputFormat="application/xml"
  outputSchema="http://www.opengis.net/cat/csw/2.0.2"
  xmlns="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:ogc="http://www.opengis.net/ogc"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.opengis.net/cat/csw/2.0.2
    ../../../../csw/2.0.2/CSW-discovery.xsd">
  <ResponseHandler>ftp://www.myserver.com/pub/MyQuery_Resp.xml</ResponseHandler>
  <Query typeName="Record">
    <ElementSetName>brief</ElementSetName>
    <Constraint version="1.1.0">
      <ogc:Filter>
```

```

        <ogc:PropertyIsLike wildCard="% " singleChar="_"
escapeChar="\ ">
            <ogc:PropertyName>AnyText</ogc:PropertyName>
            <ogc:Literal>%pollution%</ogc:Literal>
        </ogc:PropertyIsLike>
    </ogc:Filter>
</Constraint>
</Query>
</GetRecords>

```

10.9 GetRecordById operation

10.9.1 Introduction

The mandatory **GetRecordById** request retrieves the default representation of catalogue records using their identifier. The **GetRecordById** operation is an implementation of the **Present** operation from the general model. This operation presumes that a previous query has been performed in order to obtain the identifiers that may be used with this operation. For example, records returned by a **GetRecords** operation may contain references to other records in the catalogue that may be retrieved using the **GetRecordById** operation. This operation is also a subset of the **GetRecords** operation, and is included as a convenient short form for retrieving and linking to records in a catalogue.

10.9.2 KVP encoding

Table 67 specifies the keyword value pair encoding for the **GetRecordById** operation request.

NOTE To reduce the need for readers to refer to other parts of this document, the first three parameters listed below are copied from Table 57 in Subclause 10.3.5 of this document.

Table 67 — KVP encoding for GetRecordById operation request

Keyword ^b	Data type and value	Optionality and use	Parameter in general model
REQUEST	Character String Fixed value of “GetRecordById”, case insensitive.	One (Mandatory) ^a	(none)
ElementSetName	CodeList with allowed values: “brief”, “summary” or “full”	Zero or one (Optional) Default value is “summary”	responseElements
outputFormat	Character String Value is a MIME type The only value that is required to be supported is <i>application/xml</i> . Other supported values, that may include <i>text/html</i> and <i>text/plain</i> , shall be advertised in the capabilities document.	Zero or one (Optional) Default value is <i>application/xml</i>	returnFormat
outputSchema	URI Reference to the preferred schema of the response	Zero or one (Optional) Default value depends on schema of catalogue's information model	responseSchema
Id	Comma separated list of anyURI	One (Mandatory)	
<p>a The REQUEST parameter contains the same information as the name of the < GetRecordById> element in XML encoding.</p> <p>b Parameter keywords are case insensitive for KVP encoding. Parameters values are case sensitive.</p>			

10.9.3 XML encoding

The following XML-Schema fragment defines the XML encoding for the **GetRecordById** operation request:

```

<xsd:element name="GetRecordById" type="csw:GetRecordByIdType"/>
<xsd:complexType name="GetRecordByIdType">
  <xsd:complexContent>
    <xsd:extension base="csw:RequestBaseType">
      <xsd:sequence>
        <xsd:element name="Id" type="xsd:anyURI"
maxOccurs="unbounded"/>
        <xsd:element ref="csw:ElementSetName" minOccurs="0"/>
      </xsd:sequence>
      <xsd:attribute name="outputFormat"
type="xsd:string"
use="optional" default="application/xml"/>
      <xsd:attribute name="outputSchema"
type="xsd:anyURI"
use="optional"/>
    </xsd:extension>
  </xsd:complexContent>
</xsd:complexType>

```

10.9.4 Parameter descriptions

10.9.4.1 ElementSetName parameter

The **ElementSetName** parameter is used to specify a named, predefined set of metadata record elements from each source record that should be presented in the response to the operation. The predefined set names of *brief*, *summary* and *full* are meant to represent different level of detail of the source record with *brief* representing the least amount of detail and *full* representing all the metadata record elements. The sets of metadata record element names that correspond to *brief*, *summary* and *full* shall be defined in an Application Profile. If the ElementSetName parameter is not set in a request the service shall respond with one or more *summary* records.

10.9.4.2 Id parameter

The **Id** parameter is a comma-separated list of record identifiers for the records that a CSW shall return to the client. In the XML encoding, one or more **<Id>** elements may be used to specify the record identifier to retrieve.

If any of the identifiers specified in the operation is invalid, then the operation should fail and an exception message should be returned as described in subclause 10.3.7.

10.9.4.3 outputFormat parameter

The outputFormat parameter is used to control the format of the output that is generated in response to a GetRecordById request. Its value shall be a MIME type. The default value, “*application/xml*”, means that the output shall be an XML document. All catalogues shall at least support XML as an output format. Other output formats may be supported and may include output formats such as TEXT (MIME type *text/plain*), or HTML (MIME type *text/html*). The list of output formats that a CSW instance provides shall be advertised in the capabilities document.

In the case where the output format is *application/xml*, the CSW shall generate an XML document that validates against a schema document that is specified in the output document via the **xsi:schemaLocation** attribute defined in XML.

10.9.4.4 outputSchema parameter

The outputSchema parameter is included in the GetRecordById operation to allow clients to request a representation of a catalogue record other than the default representation.

By default, the GetRecordById operation generates a response that validates against the schema of the information model that the catalogue implements. This is the default representation of a catalogue record. Thus, a catalogue that implements the ebRIM would generate a response that validates against *rim.xsd*, the ebRIM schema. Similarly, a catalogue that uses ISO19115/ISO19119 as its underlying information model would generate a response that validates against those schemas.

A CSW server shall, in its Capabilities document, advertise which schemas it can use to represent a record in response to a GetRecordById request. This shall be done using the Parameter element inside the Operation element in the Capabilities document. The list of values shall include the value “<http://www.opengis.net/cat/csw/2.0.2>”, representing the schema of the common queryable and returnable elements that all CSW implementation must support. This list may include a value that represents the schema of the underlying information model, if that schema is different than the OGC common schema. The default value shall be the first value in the list.

EXAMPLE 1 The following XML fragment lists the possible values for the outputSchema parameter for a catalogue that only support the OGC common schema:

```
<Operation name="GetRecordById">
  ...
  <Parameter name="outputSchema">
    <Value>http://www.opengis.net/cat/csw/2.0.2</Value>
  </Parameter>
  ...
</Operation>
```

EXAMPLE 2 The following XML fragment lists the possible values for the outputSchema parameter for a catalogue that uses eBRIM v3.0 as its underlying information model:

```
<Operation name="GetRecordById">
  ...
  <Parameter name="outputSchema">
    <Value>urn:oasis:names:tc:ebxml-regrep:xsd:rim:3.0</Value>
    <Value>http://www.opengis.net/cat/csw/2.0.2</Value>
  </Parameter>
  ...
</Operation>
```

EXAMPLE 3 The following XML fragment lists the possible values for the outputSchema parameter for a catalogue that uses eBRIM v3.0 as its underlying information model but also supports mappings to other catalogue information models:

```
<Operation name="GetRecordById">
  ...
  <Parameter name="outputSchema">
    <Value>urn:oasis:names:tc:ebxml-regrep:xsd:rim:3.0</Value>
    <Value>http://www.opengis.net/cat/csw/2.0.2</Value>
    <Value>urn:oasis:names:tc:ebxml-regrep:xsd:rim:2.5</Value>
  </Parameter>
  ...
</Operation>
```

10.9.5 Response

The following XML Schema fragment defines the response to a GetRecordById request:

```
<xsd:element name="GetRecordByIdResponse"
  type="csw:GetRecordByIdResponseType"/>
<xsd:complexType name="GetRecordByIdResponseType"/>

  <xsd:sequence>
    <xsd:choice>
```

```

        <xsd:element ref="csw:AbstractRecord"
            minOccurs="0" maxOccurs="unbounded">
            <xsd:any processContents="strict" namespace="##other"
                minOccurs="0" maxOccurs="unbounded"/>
        </xsd:element>
    </xsd:choice>
</xsd:sequence>
</xsd:complexType>

```

This is simply the list of requested records. The response records should substitute for the element **<csw:AbstractRecord>**.

10.9.6 Examples

KVP encoded request:

<http://www.someserver.com/csw/csw.cgi?request=GetRecordById&version=http://www.someserver.com/csw/csw.cgi?request=GetRecordById&version=2.0.0&id=REC-10,REC-11,REC-122.0.2&id=REC-10,REC-11,REC-12>

XML encoded request:

```

<?xml version="1.0" encoding="ISO-8859-1"?>
<GetRecordById
  service="CSW"
  version="2.0.2"
  outputFormat="application/xml"
  outputSchema="http://www.opengis.net/cat/csw/2.0.2"
  xmlns="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.opengis.net/cat/csw/2.0.2
    ../../../../csw/2.0.2/CSW-discovery.xsd">
  <Id>REC-10</Id>
  <Id>REC-11</Id>
  <Id>REC-12</Id>
</GetRecordById>

```

Another example request:

```

<?xml version="1.0" encoding="ISO-8859-1"?>
<GetRecordById
  service="CSW"
  version="2.0.2"
  outputFormat="application/xml"
  outputSchema="http://www.opengis.net/cat/csw/2.0.2"
  xmlns="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.opengis.net/cat/csw/2.0.2
    ../../../../csw/2.0.2/CSW-discovery.xsd">
  <Id>efc40467-284d-4fee-af2a-522c717e7165</Id>
  <ElementSetName>summary</ElementSetName>
</GetRecordById>

```

Without the declaration of an Element Set or Output Schema, the server will respond with a default encoding of the results as *summary* records.

XML encoded response:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<SummaryRecord
  xmlns="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:dc="http://purl.org/dc/elements/1.1/"
  xmlns:dct="http://purl.org/dc/terms/"
  xmlns:ows="http://www.opengeospatial.net/ows"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.opengis.net/cat/csw/2.0.2
    ../../../../csw/2.0.2/record.xsd">
  <dc:identifier>00180e67-b7cf-40a3-861d-b3a09337b195</dc:identifier>
  <dc:title>title: 1141987690552</dc:title>
  <dc:type>service</dc:type>
  <dc:subject>City</dc:subject>
  <dc:format>digital</dc:format>
  <dc:relation>5a389ad2-22dd-11d1-aa77-002035b29093</dc:relation>
  <dc:relation>5a389ad2-22dd-11d1-aa77-002035b29014</dc:relation>
  <dct:modified>2004-10-04 00:00:00</dct:modified>
  <dct:abstract>Orthophotos 1:5 000 (OP5) einer Stadt</dct:abstract>
</SummaryRecord>
```

10.10 Record locking

This specification does not define a locking interface, instead relying on the underlying repository to mediate concurrent access to catalogue records. A profile of this specification may define a locking interface if it is found to be needed.

10.11 Transaction operation

10.11.1 Introduction

The optional **Transaction** operation defines an interface for creating, modifying and deleting catalogue records. The specific payload being manipulated shall be defined in a profile.

10.11.2 KVP encoding

There is no KVP encoding for transaction operation request, because there is no convenient way to encode the transaction payloads using keyword-value pairs. In addition, the actual text of a transaction message may be very long, again making it inconvenient to use KVP encoding.

10.11.3 XML encoding

10.11.3.1 Overview

The following XML schema fragment defines the XML encoding of the Transaction operation request:

```
<xsd:element name="Transaction" type="csw:TransactionType"/>
<xsd:complexType name="TransactionType">
  <xsd:complexContent>
    <xsd:extension base="csw:RequestBaseType">
      <xsd:sequence>
        <xsd:choice minOccurs="1" maxOccurs="unbounded">
          <xsd:element name="Insert" type="csw:InsertType"/>
          <xsd:element name="Update" type="csw:UpdateType"/>
          <xsd:element name="Delete" type="csw>DeleteType"/>
        </xsd:choice>
      </xsd:sequence>
      <xsd:attribute name="requestId" type="xsd:anyURI"
        use="optional"/>
      <xsd:attribute name="verboseResponse"
        type="xsd:boolean"
        use="optional" default="false"/>
    </xsd:extension>
  </xsd:complexContent>
</xsd:complexType>
```

The **<Transaction>** element defines an atomic unit of work and is a container for one or more insert, update and/or delete actions.

The **requestId** attribute may be used by a client application to associate a user-defined identifier with the operation.

The **verboseResponse** attribute is a boolean that may be used by a client to indicate to a server the amount of detail to generate in the response. A value of FALSE means that a CSW should generate a terse or brief transaction response. A value of TRUE, or the absence of the attribute, means that the normal detailed transaction response should be generated. The schema of transaction response is defined in Subclause 10.11.4.

10.11.3.2 Insert action

The following XML-Schema fragment defines an insert action:

```
<xsd:complexType name="InsertType" id="InsertType">
  <xsd:sequence>
    <xsd:any processContents="strict" namespace="##other"
      maxOccurs="unbounded" />
  </xsd:sequence>
  <xsd:attribute name="handle" type="xsd:ID" use="optional"/>
</xsd:complexType>
```

The **<Insert>** element is a container for one or more records that are to be inserted into the catalogue. The schema of the record(s) shall conform to the schema of the

information model that the catalogue supports as described using the **DescribeRecord** operation.

The **handle** attribute is an additional parameter not defined in the general model. It is used in the XML encoding to associate a mnemonic name with each action contained in a **<Transaction>** element for the purpose of error handling. If a CSW encounters an error processing a transaction request and the **handle** attribute is defined, the CSW can localize the source of the problem for the client by specifying the handle in the exception response as described in Subclause 10.3.7.

10.11.3.3 Update action

The following XML Schema fragment defines an update action:

```
<xsd:complexType name="UpdateType" id="UpdateType">
  <xsd:sequence>
    <xsd:choice>
      <xsd:any processContents="strict" namespace="##other" />
      <xsd:sequence>
        <xsd:element ref="csw:RecordProperty"
maxOccurs="unbounded"/>
        <xsd:element ref="csw:Constraint"/>
      </xsd:sequence>
    </xsd:choice>
  </xsd:sequence>
  <xsd:attribute name="typeName" type="xsd:anyURI" use="optional"/>
  <xsd:attribute name="handle" type="xsd:ID" use="optional"/>
</xsd:complexType>
<xsd:complexType name="RecordPropertyType">
  <xsd:sequence>
    <xsd:element name="Name" type="xsd:string"/>
    <xsd:element name="Value" type="xsd:anyType" minOccurs="0"/>
  </xsd:sequence>
</xsd:complexType>
```

The **<Update>** element is used to specify values to be used to change existing information in the catalogue. If the **<Update>** element contains a child element declared in some other namespace (i.e., **not** “<http://www.opengis.net/cat/csw>”), then the corresponding record in the catalogue shall be replaced. The record type shall be defined in an application profile. If individual record property values are specified in the **<Update>** element, using the **<RecordProperty>** element, then those individual property values of the catalogue record shall be updated.

The **<RecordProperty>** element contains a **<Name>** element and a **<Value>** element. The **<Name>** element is used to specify the name of the record property to be updated. The value of the **<Name>** element may be a path expression to identify complex properties. The **<Value>** element contains the value that will be used to update the record in the catalogue.

The number of records affected by an **<Update>** action is determined by the contents of the **<Constraint>** element. The **<Constraint>** element is defined in Subclause 10.3.7 and is used to define the set of catalogue records that the update operation will affect. In

order to prevent all records in a catalogue from inadvertently being updated, the **<Constraint>** element shall be specified.

The optional **typeName** attribute may be used to specify the collection name from which records will be updated.

The **handle** attribute is described in subclause 10.11.3.2.

10.11.3.4 Delete action

The following XML Schema fragment defines a delete action:

```
<xsd:complexType name="DeleteType" id="DeleteType">
  <xsd:sequence>
    <xsd:element ref="csw:Constraint" minOccurs="1" maxOccurs="1"/>
  </xsd:sequence>
  <xsd:attribute name="typeName" type="xsd:anyURI" use="optional"/>
  <xsd:attribute name="handle" type="xsd:ID" use="optional"/>
</xsd:complexType>
```

The **<Delete>** element contains a **<csw:Constraint>** element (see Subclause 10.3.7) that identifies a set of records that are to be deleted from the catalogue. The **<csw:Constraint>** element shall be specified in order to prevent every record in the catalogue from inadvertently being deleted.

The **typeName** attribute is used to specify the collection name from which records will be deleted.

The **handle** attribute is described in subclause 10.11.3.2.

10.11.4 Response

The following XML Schema fragment defines the response that shall be generated after a CSW server executes a transaction request:

```
<xsd:element name="TransactionResponse"
  type="csw:TransactionResponseType"/>
<xsd:complexType name="TransactionResponseType">
  <xsd:sequence>
    <xsd:element name="TransactionSummary"
      type="csw:TransactionSummaryType"/>
    <xsd:element name="InsertResult"
      type="csw:InsertResultType"
      minOccurs="0" maxOccurs="unbounded"/>
  </xsd:sequence>
  <xsd:attribute name="version" type="xsd:string" use="optional"/>
</xsd:complexType>

<xsd:complexType name="TransactionSummaryType">
  <xsd:sequence>
    <xsd:element name="totalInserted"
      type="xsd:nonNegativeInteger"
      minOccurs="0"/>
  </xsd:sequence>
</xsd:complexType>
```

```

        <xsd:element name="totalUpdated"
            type="xsd:nonNegativeInteger"
            minOccurs="0"/>
        <xsd:element name="totalDeleted"
            type="xsd:nonNegativeInteger"
            minOccurs="0"/>
    </xsd:sequence>
    <xsd:attribute name="requestId" type="xsd:anyURI"
use="optional"/>
</xsd:complexType>

<xsd:complexType name="InsertResultType">
    <xsd:sequence>
        <xsd:element ref="csw:BriefRecord" maxOccurs="unbounded"/>
    </xsd:sequence>
    <xsd:attribute name="handleRef" type="xsd:anyURI"
use="optional"/>
</xsd:complexType>

```

The transaction response message conveys two pieces of information. First, it reports a summary of the transaction by indicating the number records created, updated or deleted by the transaction. Second, the transaction response message indicates the results of each insert operation found in the transaction in the form of the **<InsertResult>** element.

The **<InsertResult>** element may appear zero or more times in the transaction response. It is used to report to the client a brief representation of each new record, including the record identifier, created in the catalogue. The records shall be reported in the same order in which the **<Insert>** elements appear in a transaction request and shall map 1-to-1. Optionally, the handle attribute may be used to correlate a particular **<Insert>** element in the Transaction request with an **<InsertResult>** element found in the transaction response.

10.12 Harvest operation

10.12.1 Introduction

The general model defines two operations in the Manager class that may be used to create or update records in the catalogue. They are the **transaction** operation and the **harvestRecords** operation. The **transaction** operation may be used to "push" data into the catalogue and is defined in Subclause 10.11. This subclause defines the optional **Harvest** operation, which is an operation that "pulls" data into the catalogue. That is, this operation only references the data to be inserted or updated in the catalogue, and it is the job of the catalogue service to resolve the reference, fetch that data, and process it into the catalogue.

The **Harvest** operation had two modes of operation, controlled by a flag in the request. The first mode of operation is a synchronous mode in which the CSW receives a **Harvest** request from the client, processes it immediately, and sends the results to the client while the client waits. The second mode of operation is asynchronous in that the server receives a **Harvest** request from the client, and sends the client an immediate acknowledgement that the request has been successfully received. The server can then process the **Harvest**

request whenever it likes, taking as much time as is required and then send the results of the processing to a URI specified in the original **Harvest** request. This latter mode of operation is included to support **Harvest** requests that could run for a period of time longer than most HTTP timeout's will allow.

Processing a **Harvest** request means that the CSW resolves the URI pointing to the metadata resource, parses the resource, and then creates or modifies metadata records in the catalogue in order to register the resource. This operation may be performed only once or periodically depending on how the client invokes the operation.

10.12.2 KVP encoding

Table 68 specifies the keyword-value pair encoding for the Harvest operation request.

NOTE To reduce the need for readers to refer to other parts of this document, the first three parameters listed below are copied from Table 57 in Subclause 10.3.5 of this document.

Table 68 — KVP encoding for Harvest operation request

Keyword ^c	Data type and value	Optionality and use	Parameter in general model
REQUEST	Character String Fixed value of <i>Harvest</i> , case insensitive	One (Mandatory) ^a	(none)
service	Character String Fixed values of “CSW”	One (Mandatory)	serviceId
version	Character String Fixed value of 2.0.2	One (Mandatory)	(none)
NAMESPACE	List of Character String, comma separated Used to specify a namespace and its prefix Format shall be <i>xmlns([prefix]=Jurl)</i> . If the <i>prefix</i> is not specified then this is the default namespace.	Zero or one (Optional) ^b Include value for each distinct namespace used by all qualified names in the request. If not included, all qualified names are in default namespace	(none)
Source	URI Reference to the source from which the resource is to be harvested	One (Mandatory)	Source
ResourceType	Character String Reference to the type of resource being harvested, see Subclause 10.12.4.2	One (Mandatory)	Type
ResourceFormat	Character String MIME type indicating format of the resource being harvested	Zero or one (Optional) Default value is <i>application/xml</i>	resourceFormat
ResponseHandler	URL A reference to a person or entity that the CSW should respond to when it has completed processing Harvest request asynchronously	Zero or one (Optional) If not included, process request synchronously	responseHandler
HarvestInterval	Period Must conform to ISO8601 Period syntax.	Zero or one (Optional) If not specified, then harvest only once in response to the request.	harvestInterval
<p>a The REQUEST parameter contains the same information as the name of the <Harvest> element in XML encoding.</p> <p>b The NAMESPACE parameter contains the same information as the xmlns attributes which may be used to convey namespace information in XML encoding.</p> <p>c Parameter keywords are case insensitive for KVP encoding. Parameters values are case sensitive..</p>			

10.12.3 XML encoding

The following XML-Schema fragment defines the XML encoding for a **Harvest** operation request:

```
<xsd:element name="Harvest" type="csw:HarvestType"/>
<xsd:complexType name="HarvestType">
  <xsd:complexContent>
    <xsd:extension base="csw:RequestBaseType">
      <xsd:sequence>
        <xsd:element name="Source" type="xsd:anyURI"/>
        <xsd:element name="ResourceType" type="xsd:anyURI"/>
        <xsd:element name="ResourceFormat" type="xsd:string"
          minOccurs="0" default="application/xml"/>
        <xsd:element name="HarvestInterval" type="xsd:duration"
          minOccurs="0"/>
        <xsd:element name="ResponseHandler" type="xsd:anyURI"
          minOccurs="0" maxOccurs="unbounded"/>
      </xsd:sequence>
    </xsd:extension>
  </xsd:complexContent>
</xsd:complexType>
```

10.12.4 Parameter descriptions

10.12.4.1 Source parameter

The **Source** parameter is used to specify a URI reference to the metadata resource to be harvested.

10.12.4.2 ResourceType parameter

The **ResourceType** parameter references a document that defines the structure of the resource being harvested. For high interoperability, this resource should be an XML document, and the ResourceType parameter string value should be a URI that references the structure of that XML document (i.e., its XML Schema namespace identifier URI). If a server can harvest resources in the schema of an information model it supports, the ResourceType URI should be the same as the outputSchema parameter URI defined for the GetRecords operation.

Table 69 defines a set of URIs that may be used to identify well know metadata formats. CSW implementations and CSW application profiles may support other values of the ResourceType parameter for harvesting additional artifact types. If creating new ResourceType URNs is needed, the format of the values may be *urn:ogc:def:resourceType:CSW:token*, where *token* is a placeholder for a value that denotes the specific resource type.

Table 69 — URIs for well known metadata standards

URI	Description
http://www.opengis.net/wms	WMS capability document, all current versions
http://www.opengis.net/wfs	WFS capability document, versions 1.0 and 1.1
http://www.opengis.net/wfs/1.2.0	WFS capability document, version 1.2
http://www.opengis.net/wcs	WCS capability document, version 1.0
http://www.opengis.net/wcs/1.1	WCS capability document, version 1.1
http://www.opengis.net/cat/csw	CSW capability document, versions 2.0.0 and 2.0.1
http://www.opengis.net/cat/csw/2.0.2	CSW capability document, version 2.0.2
http://www.fgdc.gov/metadata/csdgm	Content Standard for Digital Geospatial Metadata (CSDGM), Vers. 2 (FGDC-STD-001-1998)
http://www.auslig.gov.au/dtd/anzmeta-1.3.dtd	Australian Spatial Data Infrastructure Standard
http://www.isotc211.org/schemas/2005/gmd/	ISO19139 document
http://metadata.dod.mil/mdr/ns/DDMS/1.3/	DEPARTMENT OF DEFENSE DISCOVERY METADATA STANDARD (DDMS)

A CSW may accept a MIME type as the string value of the Resource Type parameter. If MIME types are accepted, the value of the MIME type shall match the value of the 'Content Type:' MIME header when the resource is accessed by the catalogue. If these two values do not agree, the catalogue shall raise an exception.

A compliant server shall, in its Capabilities document, advertise the resource type URIs it recognizes using the Parameter element within the Operation element.

EXAMPLE The following XML fragment illustrates how a catalogue server could advertise which resource types it can harvest:

```
<ows:Operation name="Harvest">
  ...
  <ows:Parameter name="ResourceType">
    <ows:Value>http://www.opengis.net/wms</ows:Value>
    <ows:Value>http://www.opengis.net/wfs</ows:Value>
    <ows:Value>http://www.opengis.net/wms</ows:Value>
    <ows:Value>http://www.opengis.net/wfs</ows:Value>
    <ows:Value>http://www.opengis.net/cat/csw/2.0.2</ows:Value>
    <ows:Value>urn:ogc:def:resourceType:CSW:FGDC</ows:Value>
    <ows:Value>http://www.auslig.gov.au/dtd/anzmeta-1.3.dtd</ows:Value>
    <ows:Value>http://metadata.dod.mil/mdr/ns/DDMS/1.3/</ows:Value>
    <ows:Value>audio/mpeg</ows:Value>
    <ows:Value>application/pdf</ows:Value>
  </ows:Parameter>
  ...
</ows:Operation>
```

10.12.4.3 ResourceFormat parameter

The **ResourceFormat** parameter is used to indicate the encoding used for the resource being harvested. This parameter is included to support the harvesting of metadata resources available in various formats such as plain text, XML or HTML. The values of

this parameter shall be a MIME type. If the parameter is not specified then the default value of *application/xml* shall be assumed.

10.12.4.4 ResponseHandler parameter

The **ResponseHandler** parameter is a flag that indicates how the **Harvest** operation should be processed by a CSW server.

If the parameter is not present, then the **Harvest** operation is processed synchronously meaning that the client sends the **Harvest** request to a CSW and then waits to receive a **HarvestResponse** or exception message as described in Subclause 10.3.7. The CSW immediately processes the **Harvest** request, while the client waits for a response. The problem with this mode of operation is that the client may timeout waiting for the server to process the request.

If the parameter is present, the **Harvest** operation is processed asynchronously. In this case, the server responds immediately to a client's request with an acknowledgement message as defined in Subclause 10.8.4.13. The acknowledgment message echoes the client's request, using the **<EchoedRequest>** element, and may include an optionally generated request identifier using the **<RequestId>** element. The acknowledgement message tells the client that the request has been received and notification of completion will be sent to the URL specified as the value of the **ResponseHandler** parameter. The **Harvest** request may then be processed at some later time taking as much time as is required to complete the operation. When the operation is completed, a **HarvestResponse** message or exception message per Subclause 10.3.7 (if a problem was encountered) is sent to the URL specified as the value of the **ResponseHandler** parameter.

10.12.4.5 HarvestInterval Parameter

The **HarvestInterval** parameter is used to specify the period of time, in ISO 8601 period format, that should elapse before a CSW attempts to re-harvest the specified resource thus refreshing its copy of a resource.

If no **HarvestInterval** parameter is specified then the resource is harvested only once in response to the **Harvest** request.

10.12.5 Response

The following XML-Schema fragment defines the **HarvestResponse** message:

```
<xsd:element name="HarvestResponse" type="csw:HarvestResponseType"/>
<xsd:complexType name="HarvestResponseType"
id="HarvestResponseType">
  <xsd:choice>
    <xsd:element ref="csw:Acknowledgement"/>
    <xsd:element ref="csw:TransactionResponse"/>
  </xsd:choice>
</xsd:complexType>
```

The **Harvest** operation can respond in one of two ways depending on the presence or absence of the **ResponseHandler** parameter.

If the **ResponseHandler** parameter is present, then the CSW server should verify the request syntax and immediately respond to the client with an acknowledgment message as defined in Subclause 10.12.4.4. Later, after the server has processed the request, it shall generate a **HarvestResponse** message and send it to the URI specified by the **ResponseHandler** parameter using the protocol encoded therein. Common protocols are *ftp* for sending the response to a ftp server and *mailto* which may be used to send the response to an email address.

If the **ResponseHandler** parameter is not present, the CSW server shall process the **Harvest** request immediately and respond to the waiting client with a **HarvestResponse** message.

If the **Harvest** attempt is successful, the response may include summary representations of the newly created or modified catalogue object(s). The response is the same as the **TransactionResponse**.

10.12.6 Examples

KVP encoded example:

```
http://www.myserver.com/csw/csw.cgi?request=Harvest&version="2.0.2"&source=http://www.yourserver.com/metadata.xml&resourcetype=http://www.fgdc.gov/metadata/csdgm&resourceformat=application/xml&responsehandler=mailto:pvretano@cubewerx.com&harvestinterval=P2W
```

XML encoded example:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<Harvest
  service="CSW"
  version="2.0.2"
  xmlns="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.opengis.net/cat/csw/2.0.2
    ../../../../csw/2.0.2/CSW-publication.xsd">
  <Source>http://www.yourserver.com/metadata.xml</Source>
  <ResourceType>http://www.fgdc.gov/metadata/csdgm</ResourceType>
  <ResourceFormat>application/xml</ResourceFormat>
  <HarvestInterval>P14D</HarvestInterval>
  <ResponseHandler>
ftp://ftp.myserver.com/HarvestResponses</ResponseHandler>
</Harvest>
```

10.13 XML Schemas

The CSW abilities specified in this Clause directly and indirectly use several newly specified XML Schema Documents, included in the zip file with this document. These new XML Schema files are named:

- a) CSW-discovery.xsd
- b) CSW-publication.xsd
- c) records.xsd
- d) rec-dcmes.xsd
- e) rec-dcterms.xsd

NOTE These XML Schema Documents use a new namespace identifier URI instead of the previous namespace identifier URI, namely <http://www.opengis.net/cat/csw/2.0.2> instead of <http://www.opengis.net/cat/csw>. This change was made based on the TC decision described in Clause 11 in OGC Best Practices Paper 06-135r1. That decision was based on the discussion in document 05-065r3 “Change Request for Namespaces for versions and profiles of XML Schemas”. A new namespace identifier URI is used for this Catalogue version 2.0.2 because these reasons apply to this version 2.0.2. However, using this new namespace identifier URI could produce problems for the GetCapabilities operation request and response, which have not yet been resolved.

After this document is approved, these new XML Schema files will be posted at the URL <http://schemas.opengis.net/cat/csw/2.0.2/> for electronic access. In the event of a discrepancy between the attached and online versions of the XML Schema files, the online files shall be considered authoritative.

These new XML Schemas build on the XML Schemas defined in the Filter Encoding Implementation Specification [OGC 04-095], and the OWS Common Implementation Specification [OGC 05-008c1], and described in those documents.

11 Specializing general model through protocol bindings and profiles

11.1 Introduction

This subclause provides an overview of the core elements of the General Catalogue Model and how these may be used in protocol bindings and application profiles.

The General Catalogue Model consists of an abstract model and a General Interface Model. The abstract query model specifies a BNF grammar for a minimal query syntax and a set of core search attributes (names, definitions, conceptual datatypes). The General Interface Model specifies a set of service interfaces that support the discovery, access, maintenance and organization of catalogues of geospatial information and related resources; these interfaces may be bound to multiple application protocols, including the HTTP protocol that underlies the World Wide Web. This specification also specifies bindings to CORBA/IIOP and Z39.50. Furthermore, all behaviour requiring sessions is expressed by a dynamic model of conversation state and state transitions. The model expresses the states and messages that trigger the changes in state.

Implementations are constrained by the protocol binding clauses of this specification. Each protocol binding includes a mapping from the general interfaces, operations, and parameters specified in this clause to the constructs available in a chosen protocol. Application profiles are intended to further document implementation choices.

An Application Profile is based on one of the protocol bindings in the base specification. In the case of the Catalogue Services Specification, a profile could reference

CORBA/IIOP, Z39.50, or the HTTP/1.1 protocol bindings. In most, but not all, protocol bindings, there may be restrictions or refinements on implementation agreed within an implementation community. A graphic model of the relationships is shown in Figure 31.

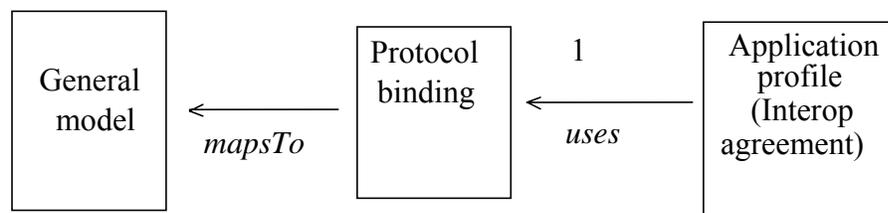


Figure 31 — Relationship of general model, protocol binding, and application profile

11.2 Interface definitions

The various elements of the General Catalogue Interface Model provide functional behaviours and capabilities to address particular areas of concern. A protocol binding may realise specific configurations of these components to serve different purposes (e.g. a read-only catalogue for discovery, a transactional catalogue for discovery and publication, or a ‘stateful’ catalogue that also supports session management).

A compliant protocol binding of the catalogue service is required to implement the OGC_Service, Catalogue Service, and Discovery classes. A protocol binding may also include any of the optional classes associated with the Catalogue Service class. A compliant implementation of a protocol binding shall recognise all operations defined within each class included in the protocol binding, and shall generate a service exception report indicating when a particular operation is not implemented (in such cases the operation is abstract—an implementation is not required).

The protocol binding clauses of this specification provides more detail on the implementation of the general interfaces. In effect, each binding maps these interfaces to a particular application protocol. For example, the names of the classes and operations in this general UML model are changed in some of the protocol bindings. The names of some operation parameters are also changed in some protocol bindings. However, the interfaces and operations specified in all Protocol Bindings shall be consistent with the semantics and granularity of interaction specified in the General Interface Model.

Application profiles, which will appear as separate documents may further specialise the implementation of these interfaces and their operations, including adding classes and parameters. However, the application profile is a specialization of the parent protocol binding, in that the names of the operations and the parameters cannot be changed.

11.3 Query model components

11.3.1 Query language/model

Many OGC service operations have the requirement to pass and process a query as a structure to perform a search. There are several query languages and messaging mechanisms identified within OGC specifications. Binding protocols and application profiles should be explicit about the selected query languages and any features peculiar to a scope of application. The following items should be addressed in specialization of a Protocol Binding or an Application profile with respect to query language support:

- a) Support for “abstract” queries, against well-known access points (e.g. core search properties). Some specifications promote or require the exposure of well-known field-like objects as common search targets (queryables), allowing interrogation of a service without prior negotiation on information content. The mandatory queryable attributes which shall be recognised by all OGC Catalogue Services is discussed in Subclause 6.3.
- b) Selection of a query language. Identify the name and version of required query language(s) anticipated by this Protocol Binding or Application Profile for use.
- c) Supported data types (e.g., character, integer, coordinate, date, polygon) and operator types (e.g., inequality, proximity, partial string, spatial, temporal). Query languages may be restricted in their implementation or extended with functions not described in the base specification. This would need to be done if the base query language did not support a data type required by the OGC Common Query Language discussed in Clause 6 such as envelope.

In addition, an application profile may extend the OGC CQL or Filter syntax with functions not described in the base specification through use of the “function “construct in CQL and the “Filter “ language. Use of this construct is discussed in Subclause TBD of this document. If an application protocol uses this extension method, the profile documentation should include an updated BNF grammar in addition to lists or reference documents with the enumerated data types and operator types required by this Application Profile. In addition, any description of special techniques (e.g. supporting joins or associations) that are expected by an Application Profile should be described.

11.3.2 Common search and retrieval elements

The abstract information model is discussed in Clause 6; this model consists of a small set of abstract search elements and the specification of a common “summary” element set to allow queries across protocol bindings and even from outside the OGC domain. Each Protocol Binding should specialize this model by:

- a) Specify the syntax of the globally unique Identifiers including any registration authority information
- b) Map the core search (queryable) elements into a concrete syntax based on the chosen record format(s)

- c) Define a “summary” element set that corresponds to the “summary” element set in the Catalogue general model

An application profile is expected to fully specify the conceptual information model adopted by the user community. This process and resulting artefacts are further discussed in Subclause 6.2.5 and the remainder of this clause.

11.4 Catalogue Application Profiles

ISO TR 10000-1:1998 describes a general framework for functional standardization and defines the concept of a profile. A profile identifies the use of particular options available in one or more base standards and it also provides a basis for developing conformance tests; a compliant profile shall not contradict the base specifications or otherwise give rise to non-conforming conditions. An *application profile* specifies the use of an application-layer protocol (e.g., Z39.50, HTTP/1.1, CORBA/IIOP) in order to provide for the structured transfer of information between systems (ISO/IEC TR 10000-2:1998).

A catalogue application profile binds a set of functional components (with interfaces specified as part of a protocol binding) to an abstract information model—expressed using UML—that has one or more concrete representations of catalogue content. Each representation is an Internet media type that conforms to a schema defined using some schema language (e.g., ASN.1, XML Schema, RDF Schema). An application profile specifies a set of functional components that are provided by a conforming implementation (Figure 32).

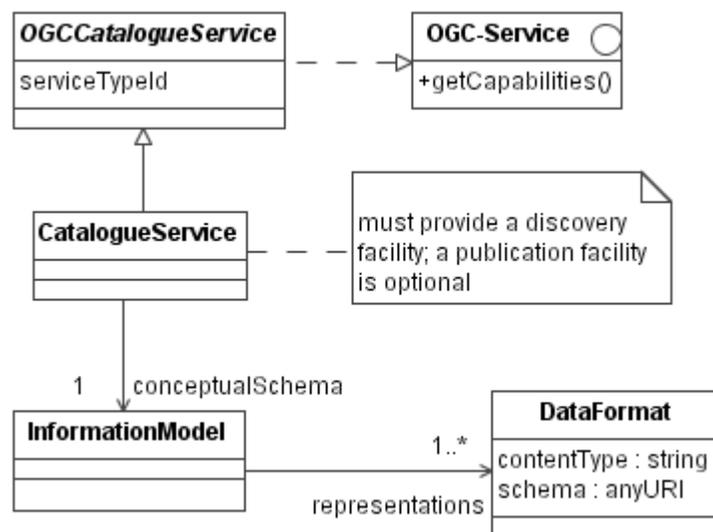


Figure 32 — Application profiles specify concrete catalogue services

An application profile is derived from one or more base specifications in order to address particular needs or requirements. The general OGC catalogue model defines common behaviours and interfaces that have general utility, but in practice there is no single solution that fits everyone’s needs. Catalogue application profiles specify refinements or extensions that are targeted toward specific implementation communities; for these

communities it is the application profile that represents the standard for conformance. Following ISO 19106, a **Level 1** profile is defined as a pure subset of one or more ISO standards; a **Level 2** profile includes allowable extensions and may also depend on non-ISO specifications.

Clause 10 in the ISO 19119 standard distinguishes *platform-neutral* from *platform-specific* specifications and assumes that one of the former will constitute the basis for one or more of the latter. That is, a single platform-neutral specification will give rise to multiple platform-specific specifications, each of which is bound to a particular distributed computing protocol. The OGC catalogue framework upholds this basic distinction: the general interface model is a platform-neutral description of catalogue operations; each application profile is platform-specific—it makes use of one of the protocol bindings defined in the catalogue specification.

The relationships between base specifications, application profiles, and catalogue service implementations are illustrated in Figure 33. The platform-neutral specification is one of the base specifications with which the application profile complies. A given catalogue implementation in turn conforms to one or more application profiles. The relationships shown in Figure 33 are consistent with the standard terms defined in ISO 10746 (*Information technology – Open Distributed Processing*) and ISO 9646-1 (*Information technology – Open Systems Interconnection – Conformance testing methodology and framework – Part 1: General concepts*)⁹.

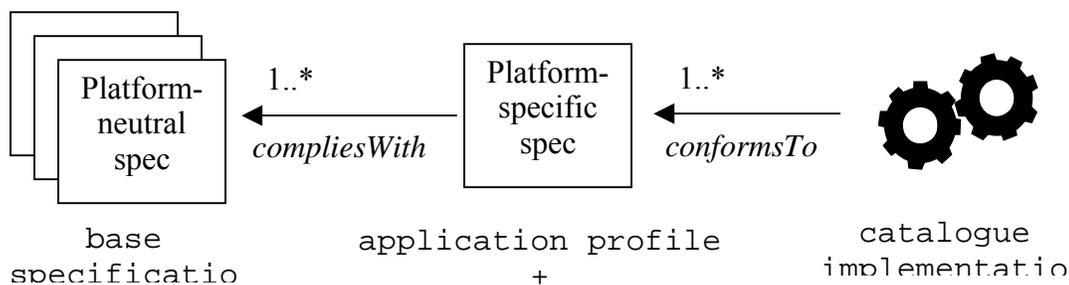


Figure 33 — Relationships between base specs, profiles, and implementations

Note that in the figure application profiles will reflect differing degrees of “thickness”. For example, if a profile employs a very simple conceptual model that embodies a limited set of simple properties then its ‘native’ representation may include little more than the common search and retrieval elements. Profiles that utilize more sophisticated models will define a native representation that provides more information; in this case the common search and retrieval elements shall be mapped to the catalogue information model.

⁹ This usage is also closely aligned with the notion of profiles expressed in Appendix 16 of the Z39.50 information retrieval standard: “The overall goal of profiles is to improve interoperability between systems conforming to a specific profile. The implication is that an implementor does not “implement the standard” but rather, configures a Z39.50 client and/or Z39.50 server to conform to one or more profiles.”

11.5 Structure and format

11.5.1 Introduction

All application profiles shall be structured as shown in Table 71. This organization complies with clause 12.3 of ISO 19106 (*Geographic information – Profiles*). A profile may introduce additional (sub)clauses as required.

Table 70 — Structure of an application profile

Clause	Title
(front matter)	Preface
	Submitting organizations
	Revision history
	Changes to the OpenGIS Catalogue Services Specification
	Future work
	Forward
	Introduction
1	Scope
2	Conformance
3	Normative references
4	Terms and definitions
5	Symbols and abbreviations
6	System context <ul style="list-style-type: none"> 6.1 Application domain 6.2 Essential use cases
7	Information models <ul style="list-style-type: none"> 7.1 Capability classes 7.2 Catalogue information model 7.3 Supported data bindings 7.4 Service information model 7.5 Native language support
8	External interfaces <ul style="list-style-type: none"> 8.1 Imported protocol bindings 8.2 Interface A 8.3 Interface B ... 8.i Query facilities 8.j General implementation guidance 8.k Security considerations
Annex A	Abstract test suite (normative)
Annex B	Design rationale (informative)

Clauses 6 through 8 convey the particulars of the application profile in terms of three ‘views’ (these correspond to the following standard ODP viewpoints: Enterprise, Information, and Computational). The three views describe various aspects of the catalogue service with respect to the base specifications; taken together they constitute the basic application architecture. The essential content of these views is summarized in the following subclauses; additional guidance can be found in the annotated profile template (OGC Document 03-101).

11.5.2 System context

This view focuses on the purpose, scope, and policies of the catalogue service (i.e., what is the system used for). It documents special requirements¹⁰ and describes the context of use as suggested in Table 71.

Table 71 — System context: required subclauses

Subclause	Topical content
Application domain	The subject domain being addressed—identify whether this profile has a specific disciplinary focus (e.g. oceanography), or is of interest to a broader community (e.g. research, public access, or libraries) The prospective stakeholders or community of practice
Essential use cases	What the system should be able to do, what it will be used for, who will use it Typical scenarios that encompass a series of interactions between users and the catalogue system being described in order to fulfill the needs of stakeholders. The inclusion of narrative use cases with accompanying interaction and/or sequence diagrams is recommended

11.5.3 Information models

This view primarily focuses on the information structures and the semantics of information processing (i.e., what the system is about); it describes the public information model that is employed by the catalogue service and the interfaces through which it is accessed. The syntax for all supported representations of the catalogued resources shall also be defined (Table 72).

¹⁰ Clause 7 of ISO 19106 stipulates that a profile must clearly identify the specific user requirements that are satisfied by the profile.

Table 72 — Information models: required subclauses

Subclause	Topical content
Capability classes	Capabilities provided by the application profile, including a high-level summary of the interfaces provided (and conformance classes/levels if these are distinguished)
Catalogue information model	Kinds of information objects managed by the catalogue using UML notation—a catalogue may offer discovery and publication support for many different types of information resources (services, data sets, schemas, style sheets, reference documents, software components, ontologies, thesauri, etc.) Mappings to the common XML Record format
Supported data formats	Supported representations of the information objects using an appropriate syntax, one of which shall be designated as the default representation Supported element sets (schemas) for each format
Service information model	Content model and syntax for service information
Native language support	How the catalogue service supports multiple languages and character encodings (i.e. internationalization and localization issues)

11.5.4 External interfaces

This view primarily focuses on documenting the externally visible behaviour of the system, including the interfaces provided by its components and the supported protocol binding(s). This view shall define the request and response message structures as part of the operation signatures; it also documents supported query facilities and any relevant security considerations (Table 73). Most of the request and response message elements are imported with the protocol binding, but a (Level 2) profile may introduce extensions to meet more specialized requirements.

Table 73 — Public interfaces: required subclauses

Subclause	Topical content
Imported protocol binding	How the interfaces or functions specified for the profile are related to the imported protocol binding.
Interface specifications	Syntax and semantics of the operations provided by each interface, including relevant preconditions, postconditions, and other usage constraints Formal, language-independent interface specifications that admit multiple programming language bindings (e.g. W3C WSDL, OMG IDL) Error conditions that can be raised and how they're handled Any restrictions or variations on the use of the supported protocol binding (e.g. CORBA/IIOP, Z39.50, HTTP/1.1)
Query facilities	Supported query languages (e.g. OGC CQL/Filter, SQL-92, XPath, XQuery, etc.) extensions or restrictions to any of the above languages
Implementation guidance	Any additional information (typically non-normative) that may be helpful to implementers
Security considerations	Information regarding the provision of security functions: authentication, access control, message integrity, confidentiality, non-repudiation, audit trails

The inclusion of a UML diagram is recommended to provide an overview of the interfaces provided by a given service type, where each type provides a different—perhaps overlapping—set of interfaces (e.g. a read-only catalogue, a catalogue that allows a ‘push’ style of publication).

11.6 Compliance

A compliant application profile shall:

- a) Include the (sub)clauses indicated in Table 73 (additional clauses MAY also be included);
- b) Define the supported catalogue information model using UML as the conceptual schema language;
- c) Define a set of mappings for the common XML record format data format;
- d) Specify the ‘native’ representation of information model elements (additional representations MAY also be specified);
- e) Define any extensions to the imported protocol binding.
- f) Indicate how the elements of the general model are related to the corresponding elements of the profile-specific interfaces;
- g) Include a conformance test suite (web-based services can do so using the OGC CITE notation).

Annex A
(normative)

Abstract conformance test suite

EDITOR'S NOTE In a future version, this annex will specify conformance checking requirements for each protocol binding.

Annex B (informative)

Description of Distributed Search

To enable Distributed Searching, the following items are needed:

- a) A multi-tier Reference Architecture as provided by this specification (as defined in Subclause 7.1)
- b) A data model to define how searches are to be distributed as defined by an information community
- c) Messages with elements applicable to Distributed Searching as provided by this specification

To support distributed searching, a community develops a data model that determines how a search will be distributed to coordinated data servers. The OGC Catalogue General Model allows data model neutrality with respect to distributed searching.

Several of the Discovery messages defined in Subclause 10.8 contain elements that pertain to distributed searching. The query message contains elements that allow the client to request certain search behaviour with respect to distribution. The request and response messages define elements that allow for the retrieval and comprehension of a distributed result set. The request and response messages contain elements that allow for understanding the status of distributed searches.

Distributed searches can cause specific problems that should be addressed in detail by an application profile. These problems result from the possibility that within one distributed, or cascading, query a catalogue service node may be approached multiple times, resulting in duplication results or, even worse, in loops causing the whole distributed system to potentially fail. Other problems are caused by duplicate metadata entries that are served by different catalogue services.

Figure B.1 displays a case resulting in duplicates due to the same catalogue service node being queried twice.

In this case the loop causes infinite recursion – the same query is sent again and again resulting in system failure and/or timeout.

It is important to notice that both of these problems can be solved by restricting the search hierarchy to two levels – one client queries a number of catalogue services that are allowed to cascade once. Generally speaking one method to avoid these problems is to control the network topology manually. Before a query is issued, the query topology is checked for duplicates or loops.

To allow an automatic solution to this problem, mechanisms should be specified in Application Profiles of the Catalogue Specification. To make sure that one catalogue node is not approached more than once within one query, whether this happens in a loop or not, one may wish to track the nodes already accessed. A cascading catalogue service would make sure that the list of already accesses nodes of the query gets added its own identifier – most likely as a URI. By the same mechanism, a counter could be implemented, reflecting the number of catalogue services the query already passed. This could help restricting the search depth of one query. An application profile could state the maximum number of cascading a query would traverse.

A possible solution to the problem of duplicates of metadata entries is to define unique identifiers for metadata entries that are applicable to a whole network, e.g. based on a combination of the server URL and the file identifier specific to the catalogue. Node-specific File identifiers alone are not the solution to the problem, as different catalogues can use the same ID without referring to the same dataset.

Annex C (informative)

Details of order operation

C.1 Introduction

This annex provides the details of the Order operation messages from OGC Catalogue Specification 1.1.1. It is included here as background for profiles or protocol bindings that have maintained backwards compatibility with these more detailed specifications. Note the message and parameter names are the names from Catalogue Version 1.1.1.

C.2 BrokeredAccessRequest

C.2.1 Message

The BrokeredAccessRequest is a service requesting data that cannot be made available directly.

```
BrokeredAccessRequest ::= sessionID destinationID requestID additionalInfo
                           productHandle orderInformation orderID requestType
                           userInformation statusOrderUpdateType
```

```
sessionID ::= Integer
```

```
destinationID ::= CharacterString
```

```
requestID ::= RequestID
```

```
additionalInfo ::= CharacterString
```

```
productHandle ::= CharacterString
```

```
orderInformation ::= OrderSpecification
```

```
orderID ::= CharacterString
```

```
requestType ::= BrokeredAccessRequestType
```

```
userInformation ::= UserInformation
```

```
statusOrderUpdateType ::= StatusUpdateType
```

C.2.2 Message Parameters:

productHandle: Type = CharacterString

The product handle is the identifier for a specific product taken from the catalogue metadata for that product.

orderInformation: Type = OrderSpecification

For BrokeredAccessRequestType = orderEstimate or OrderQuoteAndSubmit, the specification of the current order request as provided as by the client or modified by the server during the estimation process.

For BrokeredAccessRequestType = orderMonitor or orderCancel, OrderSpecification is ignored and may not be supplied.

orderID: type = CharacterString

The orderID parameter provides a unique identifier for an order in progress. This ID can be used to inquire about the status of the order as it is being processed. For BrokeredAccessRequestType = orderMonitor or orderCancel, orderID shall be supplied. For requestType = orderEstimate or OrderQuoteAndSubmit, orderID shall be empty.

requestType: Type = BrokeredAccessRequestType

The request type parameter identifies the type of service the client needs from the server. Valid values are estimate, submit, monitor and cancel. Estimate is used to check if the order is valid and to request an estimate of resources required to fill the order. Submit is a request to order and deliver the products(s). Monitor provides the current status of the order. Cancel requests that the order be cancelled. The server shall grant cancellation of the order. BrokeredAccessRequestType is formally defined in Subclause C.4.2.

userInformation: Type = UserInformation

To receive products it is necessary to provide requester identification, billing and delivery data as part of the order. This parameter is used to provide that data.

statusOrderUpdateType : Type = StatusUpdateType

How a given client likes to be kept informed about the status of a given order.

C.2.3 Message Operations: None

C.3 BrokeredAccessResponse

C.3.1 Response

The server generates the BrokeredAccessResponse message in response to a BrokeredAccessRequest.

BrokeredAccessResponse ::= sessionID destinationID requestID additionalInfo
diagnostic format orderStatus resourceEstimate order orderID status requestType

sessionID ::= Integer

destinationID ::= CharacterString

requestID ::= RequestID

additionalInfo ::= CharacterString

diagnostic ::= CharacterString

format ::= MessageFormat

orderStatus ::= OrderStatus

resourceEstimate ::= CharacterString

order ::= CollectionName

orderID ::= CharacterString

status ::= Status

requestType ::= BrokeredAccessRequestType

orderInformation ::= OrderSpecification

C.3.2 Message Parameters:

orderStatus Type ::= OrderStatus

This parameter indicates the status of the order. The status of the order is different than the status of an Access message. The status of the message is reported in the response in the status parameter. The OrderStatus type is formally defined in Subclause C.4.6 of this specification.

resourceEstimate: Type = CharacterString

This parameter reports back on the resources needed to process and/or deliver the requested resource. Examples of these resources are time until delivery and cost.

order: Type = CollectionName

The order parameter returns a name or id of the requested product object online. This parameter can be used for direct access (such as through simple features) to the online product. The CollectionName type is formally defined in Subclause C.4.3 of this specification.

orderID: type = CharacterString

The orderID parameter provides a unique identifier for an order in progress. This ID can be used to inquire about the status of the order as it is being processed. This number is generated by the server in response to a BrokeredAccessRequest where requestType = orderEstimate or OrderQuoteAndSubmit

status: Type = Status

The Status parameter conveys the status of the requested product. The Status type is formally defined in Subclause C.4.11.

requestType: Type = BrokeredAccessRequestType

The request type parameter identifies the type of service the client needs from the server. BrokeredAccessRequestType is formally defined in Subclause C.4.2.

orderInformation: Type ::= OrderSpecification

For BrokeredAccessRequestType = orderEstimate or OrderQuoteAndSubmit, the specification of the current order request as provided as by the client or modified by the server during the estimation process. .

For BrokeredAccessRequestType = orderMonitor or orderCancel, OrderSpecification is ignored and may not be supplied.

C.3.3 Message Operations: None

C.4 Parameter Type Definitions

C.4.1 Introduction

This annex provides definitions for all of the parameter data types used in Request-Response Message Pairs. These definitions assume the use of the OGC well known data types where applicable.

C.4.2 BrokeredAccessRequestType

Recommended Implementation Type: Code_List

Used By: BrokeredAccessRequest

BrokeredAccessRequestType is a code list for identifying the nature of a brokered access request. Valid values for this type are shown in Table C.1.

Table C.1 — Brokered Access Request Types

Value	Explanation
orderEstimate	Validate and obtain the estimate of an order specification
orderQuoteAndSubmit	Obtain a quote and subsequently submit an order specification
orderMonitor	Monitor the progress of an order request
orderCancel	Cancel an order request

C.4.3 CollectionName

Recommended Implementation Type: Union data

Used By: BrokeredAccessResponse

Collection Name is a type that identifies a catalogue data resource. It can point to a catalog, catalogue entry, named catalogue subspace, named catalogue superspace or a result set. This type is a “C” union of two base types:

- collection ID (character string)
- collection Name (character string).

C.4.4 OrderItem

Recommended Implementation Type: Data Structure

Used by: GC_BrokeredAccessRequestType

This data structure contains the specification of a single order item (i.e. e. the product that is ordered and that is to be delivered):

- a) productId, which is the identifier of the ordered product.
- b) productPrice, which is the price of the product.
- c) productDeliveryOptions, which contains delivery options for the product.
- d) processingOptions, which specifies the processing options that are to be applied on the product before delivery.
- e) sceneSelectionOptions, which specifies the selection of the scene from the whole product that is to be delivered.

C.4.5 OrderSpecification

Recommended Implementation Type: Data Structure

Used By: BrokeredAccessRequest

The specification of the order request as provided as input by the client if BrokeredAccessRequestType = orderEstimate or OrderQuoteAndSubmit.

The structure contains the following information about the product specification:

- a) orderCentreID – identifies the order centre at which the order will be performed
- b) orderPrice –the price for the whole order
- c) orderDeliveryDate - the latest date at which the order can be expected to be delivered to the user.
- d) orderCancellationDate – the latest date at which the user can cancel the order.
- e) deliveryMethod – how the order will be delivered to the user: e-mail, ftp or mail.
- f) package – contains the definition of how the packages which compose the order

C.4.6 OrderStatus

Recommended Implementation Type: Code_List

Used By: BrokeredAccessResponse

OrderStatus is a code list for identifying the status of an order. Valid values for this type are:

Table C.2 — Order Status Codes

Value	Explanation
orderBeingEstimated	the order is currently being estimated by the target order handling system. An Estimate is an approximation only.
orderEstimated	indicates that the order has been successfully validated and that an estimate is provided.
orderBeingQuoted	the order is currently being quoted by the target order handling system. A Quote shall be considered contractually binding.
orderBeingProcessed	the order is currently being processed by the target order handling system.
orderCompleted	processing of order has been completed.
orderNotValid	the order has not been successfully validated.
orderCancelled	the order has been cancelled

C.4.7 PackageSpecification

Recommended Implementation Type: Data Structure

Used By: BrokeredAccessRequest, OrderSpecification, PackagingType

The specification of a single package or multiple packages.

The structure contains the following information about the packaging order:

- a) `packageId` – the identifier of the ordered package
- b) `packagePrice` –the price for the package
- c) `package` – the detailed information concerning the specification of package. (See `packagingType`)
- d) `packageMedium` –the medium on which the package will be delivered to a user.
- e) `packageSize` – the size of the package in kilobytes.

C.4.8 PackagingType

Recommended Implementation Type: Code List

Used By: `PackageSpecification`, `BrokeredAccessRequest`

The specification of the packaging method used to deliver an order to a user.

- a) `predefinedPackage`: A package predefined by the given catalogue service
- b) `adhocPackage`: A package constructed of `OrderItems` to fulfill a particular order

C.4.9 PaymentMethod

Recommended Implementation Type: Code_List

Used By: `UserInformation`

This code list contains the payment methods for an order secured through using a `Access` operation. The supported methods are the following:

- a) `credit`
- b) `cash`
- c) `purchaseOrder`

C.4.10 QueryScope

Recommended Implementation Type: Code_List

Used By: `QueryRequest`

`QueryScope` is a code list describing the size of the search space for a query. Current valid values for this type are:

distributed

local

The Reference Architecture for the OGC Catalogue allows for catalogue requests to be distributed to multiple catalogs. The architecture allows for a Catalogue to accept a request from a client and distribute the request to other Catalogs. For the OGC Catalogue Service, Distributed Catalogue Searching is defined as a service that involves services of multiple Catalogue Servers, in addition to the primary client-server interaction. A catalogue server may be able to perform Distributed Searching by propagating secondary catalogue service requests to other catalogue servers. See Annex D for more explanation.

C.4.11 Status

Recommended Implementation Type: Code_List

Used By: BrokeredAccessResponse

Status is a code list for representing the current status of a resource or request. The valid values for this type are the following:

- a) success: the request has been processed without error.
- b) successResultsAvailable: the request has been processed without error and outputs of the processing can be retrieved.
- c) processingNormal: the requested operations have begun but are not completed. No errors have been identified.
- d) processingQueued: the requested operations have begun but are not completed. No errors have been identified. The processing has been temporally suspended and will resume when other processing has been completed.
- e) processingPausedOrSuspended: the requested operations have begun but are not completed. No errors have been identified. The processing has been temporally suspended and will resume when triggered by an external event.
- f) failure: the request could not be completed due to errors being encountered. On a best effort basis the server has returned to the state prior to the request.
- g) failureAccessDenied : the request could not be completed because the privileges of the client did not permit the operation. On a best effort basis the server has returned to the state prior to the request.

C.4.12 StatusUpdateType

Recommended Implementation Type: Code List

Used By: OrderStatusUpdateType

This parameter defines how the user requesting the order desires to be kept informed about the order processing.

- a) manual: The user performs the status request using the Catalogue Interface
- b) automatic: The OHS filling the order provides status updates for the user via email

C.4.13 UserInformation

Recommended Implementation Type: Data Structure

Used By: BrokeredAccessRequest

This parameter type is a data structure used to provide information about the user.

- a) userName: (type = Character String) – name of the user
- b) userAddress: (type = CharacterString) – billing, home or delivery address of user
- c) phoneNumber: (type = CharacterString) – home or office phone number for user
- d) faxNumber: (type = CharacterString) – home or office fax number for user
- e) emailAddress: (type = CharacterString) – e-mail address for the user
- f) NetAddress: (type = CharacterString) – Address of the users' primary computer.
- g) PaymentMethod: (type = PaymentMethod) – defines the payment method

Annex D (informative)

Sample CSW capabilities document

```

<csw:Capabilities
  version="2.0.2"
  updateSequence="0"
  xmlns:ows="http://www.opengis.net/ows"
  xmlns:ogc="http://www.opengis.net/ogc"
  xmlns:csw="http://www.opengis.net/cat/csw/2.0.2"
  xmlns:xlink="http://www.w3.org/1999/xlink"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation=
    "http://www.opengis.net/ows
    http://www.pvretano.com/schemas/ows/1.0.0/owsCommon.xsd
    http://www.opengis.net/ogc

http://www.pvretano.com/schemas/filter/1.1.0/filterCapabilities.xsd
  http://www.opengis.net/cat/csw/2.0.2
  http://www.pvretano.com/schemas/csw/2.0.2/CSW-discovery.xsd
  http://www.w3.org/1999/xlink
  http://www.pvretano.com/schemas/xlink/1.0.0/xlinks.xsd">
<ows:ServiceIdentification>
  <ows:ServiceType>CSW</ows:ServiceType>
  <ows:ServiceTypeVersion>0.7.8</ows:ServiceTypeVersion>
  <ows:ServiceTypeVersion>2.0.0</ows:ServiceTypeVersion>
  <ows:ServiceTypeVersion>2.0.2</ows:ServiceTypeVersion>
  <ows:Title>CubeWerx CSW</ows:Title>
  <ows:Abstract>
    A catalogue service that conforms to the HTTP protocol
    binding of the OpenGIS Catalogue Service specification
    version 2.0.2.
  </ows:Abstract>
  <ows:Keywords>
    <ows:Keyword>CSW</ows:Keyword>
    <ows:Keyword>CubeWerx</ows:Keyword>
    <ows:Keyword>geospatial</ows:Keyword>
    <ows:Keyword>catalogue</ows:Keyword>
  </ows:Keywords>
  <ows:Fees>NONE</ows:Fees>
  <ows:AccessConstraints>NONE</ows:AccessConstraints>
</ows:ServiceIdentification>
<ows:ServiceProvider>
  <ows:ProviderName>CubeWerx Inc.</ows:ProviderName>
  <ows:ProviderSite xlink:href="http://www.cubewerx.com"/>
  <ows:ServiceContact>
    <ows:IndividualName>Panagiotis A.
Vretanos</ows:IndividualName>
    <ows:PositionName>Senior Developer</ows:PositionName>
    <ows:ContactInfo>
      <ows:Phone>

```

```

        <ows:Voice>+1-819-771-8303</ows:Voice>
        <ows:Facsimile>+1-819-771-8303</ows:Facsimile>
    </ows:Phone>
    <ows:Address>
        <ows:DeliveryPoint>200 rue Montcalm</ows:DeliveryPoint>
        <ows:DeliveryPoint>Suite R-13</ows:DeliveryPoint>
        <ows:City>Gatineau</ows:City>
        <ows:AdministrativeArea>Quebec</ows:AdministrativeArea>
        <ows:PostalCode>J8Y 3B5</ows:PostalCode>
        <ows:Country>CANADA</ows:Country>
        <ows:ElectronicMailAddress>
            pvretano@cubwerx.com
        </ows:ElectronicMailAddress>
    </ows:Address>
    <ows:OnlineResource
xlink:href="mailto:pvretano@cubwerx.com"/>
    </ows:ContactInfo>
</ows:ServiceContact>
</ows:ServiceProvider>
<ows:OperationsMetadata>
    <ows:Operation name="GetCapabilities">
        <ows:DCP>
            <ows:HTTP>
                <ows:Get
xlink:href="http://www.cubwerx.com/cwsw.cgi?"/>
                <ows:Post
xlink:href="http://www.cubwerx.com/cwsw.cgi"/>
            </ows:HTTP>
        </ows:DCP>
    </ows:Operation>
    <ows:Operation name="DescribeRecord">
        <ows:DCP>
            <ows:HTTP>
                <ows:Get
xlink:href="http://www.cubwerx.com/cwsw.cgi?"/>
                <ows:Post
xlink:href="http://www.cubwerx.com/cwsw.cgi"/>
            </ows:HTTP>
        </ows:DCP>
        <ows:Parameter name="typeName">
            <ows:Value>csw:Record</ows:Value>
        </ows:Parameter>
        <ows:Parameter name="outputFormat">
            <ows:Value>application/xml</ows:Value>
        </ows:Parameter>
        <ows:Parameter name="schemaLanguage">
            <ows:Value>XMLSCHEMA</ows:Value>
        </ows:Parameter>
    </ows:Operation>
    <ows:Operation name="GetRecords">
        <ows:DCP>
            <ows:HTTP>
                <ows:Get
xlink:href="http://www.cubwerx.com/cwsw.cgi?"/>
                <ows:Post
xlink:href="http://www.cubwerx.com/cwsw.cgi"/>
            </ows:HTTP>
        </ows:DCP>
    </ows:Operation>

```

```

        </ows:HTTP>
    </ows:DCP>
    <ows:Parameter name="TypeName">
        <ows:Value>csw:Record</ows:Value>
    </ows:Parameter>
    <ows:Parameter name="outputFormat">
        <ows:Value>application/xml</ows:Value>
        <ows:Value>text/html</ows:Value>
        <ows:Value>text/plain</ows:Value>
    </ows:Parameter>
    <ows:Parameter name="outputSchema">
        <ows:Value>http://www.opengis.net/cat/csw/2.0.2</ows:DefaultValue>
    </ows:Parameter>
    <ows:Parameter name="resultType">
        <ows:Value>hits</ows:DefaultValue>
        <ows:Value>results</ows:Value>
        <ows:Value>validate</ows:Value>
    </ows:Parameter>
    <ows:Parameter name="ElementSetName">
        <ows:Value>brief</ows:Value>
        <ows:Value>summary</ows:Value>
        <ows:Value>full</ows:Value>
    </ows:Parameter>
    <ows:Parameter name="CONSTRAINTLANGUAGE">
        <ows:Value>Filter</ows:Value>
    </ows:Parameter>
</ows:Operation>
<ows:Operation name="GetRecordById">
    <ows:DCP>
        <ows:HTTP>
            <ows:Get
xlink:href="http://www.cubewerx.com/cwcsq.cgi?"/>
            <ows:Post
xlink:href="http://www.cubewerx.com/cwcsq.cgi?"/>
            </ows:HTTP>
        </ows:DCP>
        <ows:Parameter name="ElementSetName">
            <ows:Value>brief</ows:Value>
            <ows:Value>summary</ows:Value>
            <ows:Value>full</ows:Value>
        </ows:Parameter>
    </ows:Operation>
<ows:Operation name="GetDomain">
    <ows:DCP>
        <ows:HTTP>
            <ows:Get
xlink:href="http://www.cubewerx.com/cwcsq.cgi?"/>
            <ows:Post
xlink:href="http://www.cubewerx.com/cwcsq.cgi?"/>
            </ows:HTTP>
        </ows:DCP>
        <ows:Parameter name="ParameterName">
            <ows:Value>GetRecords.resultType</ows:Value>
            <ows:Value>GetRecords.outputFormat</ows:Value>
            <ows:Value>GetRecords.outputRecType</ows:Value>

```

```

        <ows:Value>GetRecords.typeNames</ows:Value>
        <ows:Value>GetRecords.ElementSetName</ows:Value>
        <ows:Value>GetRecords.ElementName</ows:Value>
        <ows:Value>GetRecords.CONSTRAINTLANGUAGE</ows:Value>
        <ows:Value>GetRecordById.ElementSetName</ows:Value>
        <ows:Value>DescribeRecord.typeName</ows:Value>
        <ows:Value>DescribeRecord.schemaLanguage</ows:Value>
    </ows:Parameter>
</ows:Operation>
<ows:Operation name="Harvest">
    <ows:DCP>
        <ows:HTTP>
            <ows:Get
xlink:href="http://www.cubewerx.com/cwcsq.cgi?"/>
            <ows:Post
xlink:href="http://www.cubewerx.com/cwcsq.cgi"/>
            </ows:HTTP>
        </ows:DCP>
        <ows:Parameter name="ResourceType">
            <ows:Value>csw:Record</ows:DefaultValue>
            <ows:Value>http://www.fgdc.gov/metadata/csdgm</ows:Value>
            <ows:Value>http://www.opengis.net/wfs</ows:Value>
            <ows:Value>http://www.opengis.net/wms</ows:Value>
        </ows:Parameter>
        <ows:Parameter name="ResourceFormat">
            <ows:Value>application/xml</ows:DefaultValue>
            <ows:Value>text/plain</ows:Value>
        </ows:Parameter>
    </ows:Operation>
<ows:Operation name="Transaction">
    <ows:DCP>
        <ows:HTTP>
            <ows:Get
xlink:href="http://www.cubewerx.com/cwcsq.cgi?"/>
            <ows:Post
xlink:href="http://www.cubewerx.com/cwcsq.cgi"/>
            </ows:HTTP>
        </ows:DCP>
    </ows:Operation>
<ows:Parameter name="service">
    <ows:Value>CSW</ows:DefaultValue>
</ows:Parameter>
<ows:Parameter name="version">
    <ows:Value>2.0.1</ows:Value>
    <ows:Value>2.0.0</ows:Value>
    <ows:Value>2.0.2</ows:Value>
</ows:Parameter>
<ows:ExtendedCapabilities></ows:ExtendedCapabilities>
</ows:OperationsMetadata>
<ogc:Filter_Capabilities>
<ogc:Spatial_Capabilities>
    <ogc:Spatial_Operators>
        <ogc:BBOX/>
        <ogc:Equals/>
        <ogc:Intersect/>
        <ogc:Touches/>

```

```

        <ogc:Crosses/>
        <ogc:Contains/>
        <ogc:Overlaps/>
    </ogc:Spatial_Operators>
</ogc:Spatial_Capabilities>
<ogc:Scalar_Capabilities>
    <ogc:Logical_Operators/>
    <ogc:Comparison_Operators>
        <ogc:Simple_Comparisons/>
        <ogc:Like/>
        <ogc:Between/>
        <ogc:NullCheck/>
    </ogc:Comparison_Operators>
    <ogc:Arithmetic_Operators>
        <ogc:Simple_Arithmetic/>
        <ogc:Functions>
            <ogc:Function_Names>
                <ogc:Function_Name
nArgs="1">COUNT</ogc:Function_Name>
                <ogc:Function_Name
nArgs="1">DISTINCT</ogc:Function_Name>
                <ogc:Function_Name nArgs="1">MIN</ogc:Function_Name>
                <ogc:Function_Name nArgs="1">MAX</ogc:Function_Name>
                <ogc:Function_Name
nArgs="1">UPPER</ogc:Function_Name>
            </ogc:Function_Names>
        </ogc:Functions>
    </ogc:Arithmetic_Operators>
</ogc:Scalar_Capabilities>
</ogc:Filter_Capabilities>
</csw:Capabilities>

```

Annex E (Normative) Minimum XPath subset

The XML Path Language [XPath] specification is a language for addressing parts of a XML document. In this standard, XPath is used as a selection mechanism for metadata properties in the information model of the catalogue and their values. This ANNEX does not use the full XPath language but only a well-defined subset:

- a) Abbreviated relative location paths are used to reference properties from the information model of a catalogue. Each step of a relative location path is separated by a '/' separator as specified in the XPath grammar [XPath]. The first step shall select a queryable type within the information model of the catalog by specifying its qualified name (e.g. csw:Record) which shall be one of the names that appears as list items in the value of the **typeNames** attribute of the **GetRecords** element. This first step allows the selection of the type of the entities that the XPath value is referring to.
- b) Subsequent steps in the XPath expression reference metadata property names from the information model of the catalogue (e.g. dct:modified) to select the value the XPath expression points to. Each step uses the abbreviate *child::axis* specifier of XPath.
- c) Finally, the last part of the XPath expression is used to select a specific part of a value if the value is complex. In the case of a collection of values, an indexed XPath expression may be used to select a sub-element of a value.
 - 1) The selection of the sub-element of a value can use the indexed XPath expression (i.e. '[]') to access a specific element of a collection.
 - 2) The *attribute::axis* specifier '@' may be used to access a sub-element encoded as an XML attribute (i.e. '@attribute_name').
 - 3) In the case where the type name or any object name is enclosed in the XPath function *schema-element()*, as defined in [XPath], the catalogue shall consider not only the specified object type but also all the types that are derived.

Bibliography

- [1] ISO/IEC 8825:1990, Information technology -- Open Systems Interconnection -- Specification of Basic Encoding Rules for Abstract Syntax Notation One (ASN.1)
- [2] ISO/IEC TR 10000-1:1998, Information technology – Framework and taxonomy of International Standardised Profiles – Part 1: General principles and documentation framework.
- [3] ISO/IEC TR 10000-2:1998, Information technology – Framework and taxonomy of International Standardised Profiles – Part 2: Principles and Taxonomy for OSI Profiles
- [4] ISO 19101:2002, Geographic information -- Reference model
- [5] ISO 19103 (DTS), Geographic information - Conceptual schema language, (Draft Technical Specification)
- [6] ISO 19106:2002 (DIS), Geographic information - Profiles
- [7] ISO 19108:2002, Geographic information - Temporal schema
- [8] ISO 19109:2002 (DIS), Geographic information - Rules for application schema
- [9] ISO 19110:2001 (DIS), Geographic information - Methodology for feature cataloguing
- [10] ISO 19113:2002, Geographic information - Quality principles
- [11] ISO 19114:2001, (DIS) Geographic information - Quality evaluation procedures
- [12] ISO 19118:2002, (DIS) Geographic information - Encoding
- [13] ISO 23950:1998, Information and documentation -- Information retrieval (Z39.50) -- Application service definition and protocol specification