



1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34

**SDMX REGISTRY SPECIFICATION:
LOGICAL INTERFACES**

(VERSION 2.0)

November 2005



- 35
- 36
- 37
- 38
- 39
- 40
- 41
- 42
- 43
- 44
- 45
- 46
- 47
- 48
- 49
- 50
- 51
- 52
- 53
- 54
- 55
- 56
- 57
- 58
- 59
- 60
- 61
- 62
- 63
- 64
- 65
- 66
- 67
- 68
- 69
- 70
- 71
- 72
- 73
- 74
- 75
- 76
- 77
- 78
- 79

© SDMX 2005
<http://www.sdmx.org/>



80 Contents

81	1	INTRODUCTION	6
82	2	SCOPE AND NORMATIVE STATUS.....	7
83	3	OBJECTIVES OF THE SDMX REGISTRY/ REPOSITORY	7
84	4	SDMX REGISTRY/REPOSITORY ARCHITECTURE.....	8
85	4.1	Architectural Schematic.....	8
86	4.2	Structural Metadata Repository	9
87	4.3	Provisioning Metadata Repository.....	9
88	4.4	Registry Services.....	9
89	4.4.1	Data and Metadata Registration Service.....	9
90	4.4.2	Data and Metadata Discovery	10
91	4.4.3	Subscription and Notification	11
92	5	SDMX INFORMATION MODEL (SDMX-IM)	11
93	5.1	Identification, Versioning, and Maintenance.....	11
94	5.1.1	Identification, Versioning and Maintenance Model.....	12
95	5.2	Unique identification of SDMX objects	13
96	5.2.1	Universal Resource Name (URN)	13
97	5.2.2	Identification components of SDMX objects.....	15
98	6	STRUCTURAL DEFINITION METADATA	18
99	6.1	Introduction	18
100	6.2	Item Scheme, Structure.....	19
101	6.2.1	Structure and Item Scheme: Basic Concepts.....	19
102	6.2.2	Structure and Item Scheme Model.....	20
103	6.2.3	Structure and Item Scheme: Functions and Behaviour.....	20
104	6.3	Structure Usage.....	26
105	6.3.1	Structure Usage: Basic Concepts.....	26
106	6.3.2	Structure Usage Model.....	26
107	6.3.3	Structure Usage: Functions and Behaviour.....	27



108	7	DATA AND METADATA PROVISIONING	29
109	7.1	Provision Agreement, Data and Metadata Sources	29
110	7.1.1	Provisioning Agreement: Basic concepts	29
111	7.1.2	Provisioning Agreement Model.....	29
112	7.1.3	Provisioning Agreement: Functions and Behaviour	32
113	7.2	Data and Metadata Constraints.....	33
114	7.2.1	Data and Metadata Constraints: Basic Concepts.....	33
115	7.2.2	Data and Metadata Constraints: Model.....	34
116	7.2.3	Data and Metadata Constraints: Functions and Behaviour.....	37
117	8	DATA AND METADATA REGISTRATION.....	39
118	8.1	Basic Concepts.....	39
119	8.2	The Registration Request.....	39
120	8.2.1	Registration Request Model	39
121	8.2.2	Registration Constraints	40
122	8.2.3	Registration Request: Functions and Behaviour	42
123	8.2.4	Registration Registry Service	43
124	8.3	Registration Response	43
125	8.3.1	Registration Response Model	43
126	8.3.2	Registration Response: Functions and Behaviour	45
127	9	DATASET AND METADATA SET DISCOVERY	46
128	9.1	Finding the Data and Metadata We Want	46
129	9.2	Building a Query	47
130	9.2.1	The Data/Metadata Query Model	48
131	9.2.2	Data/Metadata Query”Functions and Behaviour	49
132	9.3	Processing the Query	50
133	9.3.1	No Content (Key Set or Cube Region) Constraints.....	50
134	9.3.2	Cube Region Content Constraint.....	50
135	9.3.3	Key Set Content Constraint.....	50



136	9.4	Handling the Results	50
137	9.4.1	The Results Model.....	50
138	9.4.2	Query Response: Functions and Behaviour.....	52
139	10	SUBSCRIPTION AND NOTIFICATION SERVICE.....	53
140	10.1	Subscription	53
141	10.1.1	Common Information.....	53
142	10.1.2	Structural Metadata Events.....	54
143	10.1.3	Provisioning Metadata Events.....	54
144	10.1.4	Registration Events	54
145	10.2	Notification	55
146	10.2.1	Structural Metadata Notification	55
147	10.2.2	Provisioning Metadata Notification.....	55
148	10.2.3	Registration Notification	55
149	10.2.4	Subscription and Notification: Functions and Behaviour.....	56
150	11	SDMX-RR LOGICAL MODEL AND SDMX-ML	58
151			



1 INTRODUCTION

152

153 The business vision documents for SDMX envisage the promotion of a “data sharing”
154 model to facilitate low-cost, high-quality statistical data and metadata exchange.
155 Data sharing reduces the reporting burden of organisations by allowing them to
156 publish data once, and let their counterparties “pull” data and related metadata as
157 required. The scenario is based on:

158

- 159 • the availability of an abstract information model capable of supporting any
160 time-series and cross-sectional data, structural metadata and reference
161 metadata (SDMX-IM)
- 162 • standardised XML schemas derived from the model (SDMX-ML)
- 163 • the use of web-services technology (XML, XSD, WSDL, ebXML Registries)

164

165 Such an architecture needs to be well organised, and the SDMX Registry/Repository
166 (SDMX-RR) is tasked with providing structure and organisation for most of the SDMX
167 components required to support the data-sharing vision.

168

169 Standard formats for the exchange of aggregated statistical data and metadata as
170 prescribed in SDMX v1.0 (ISO/TS 17369:2005 SDMX) are envisaged to bring benefits
171 to the statistical community because data reporting and dissemination processes can
172 be more efficient.

173

174 As organisations migrate to SDMX enabled systems, many XML (and conventional)
175 artefacts will be produced [e.g. Key Family, Metadata Structures, Code List and
176 Concept definitions (often called structural metadata), XML schemas generated from
177 key families, XSLT style-sheets for transformation and display of data and metadata,
178 terminology references, etc.]. The data sharing model is based on interoperability,
179 and the discovery and sharing of these artefacts between parties in a controlled and
180 organized way.

181

182 With these fundamental standards in place, a set of architectural standards are
183 needed to address some of the processes involved in statistical data and metadata
184 exchange, with an emphasis on promoting the data-sharing vision. Architectural
185 standards address the ‘how’ rather than the ‘what’, and are aimed at enabling
186 existing SDMX standards to achieve their mission. The architectural standards
187 address registry services which initially comprise: structural metadata repository,
188 data provisioning repository, data and metadata registration and query. The registry
189 services outlined in this specification are designed to help the SDMX community
190 manage the proliferation of SDMX assets and to support data sharing for reporting
191 and dissemination.

192 2 SCOPE AND NORMATIVE STATUS

193 The scope of this document is to specify the logical interfaces for the SDMX registry
194 in terms of the functions required and the data that may be present in the function
195 call. This interface specification is syntax independent.

196
197 In this document, functions and behaviours of the registry interfaces are described in
198 three ways:

- 199
- 200 • In text
- 201 • With tables
- 202 • With UML diagrams excerpted from the SDMX Information Model
- 203

204 For those interested in seeing a syntax-bound implementation, the registry interfaces
205 are implemented as part of the SDMX-ML schemas.

206
207 Whilst the introductory section contains some information on the role of the registry, it
208 is assumed that the reader is familiar with the uses of a registry in providing shared
209 metadata across a community of counterparties. The SDMX Implementers' Guide
210 contains more details on the specific roles the registry plays.

211
212 Note that section 5.2 contains normative rules regarding the identification of registry
213 objects. Further, the minimum standard for access to the registry is via a REST
214 interface (HTTP or HTTPS), as described in the appropriate sections. The notification
215 mechanism must support e-mail and HTTP/HTTPS protocols as described.
216 Normative registry interfaces are specified in the SDMX-ML implementation. All other
217 sections of this document are explanatory.

218 3 OBJECTIVES OF THE SDMX REGISTRY/ 219 REPOSITORY

220 The objective of the SDMX Registry/Repository is, in broad terms, to allow
221 organisations to publish statistical data and metadata in known formats such that
222 interested third parties can discover these data and interpret them accurately and
223 correctly and within the shortest possible timescale. The mechanism for doing this is
224 set out below in high-level terms:

225
226 Setting up structural metadata and the exchange context (referred to as “data
227 provisioning”) involves the following steps for maintenance agencies:

- 228
- 229 • Agreeing and creating a specification of the structure of the data (called a key
230 family) which defines the dimensions, measures and attributes of a dataset
231 and their valid value set
- 232 • Defining a subset or view of a key family which allows some restriction of
233 content called a “dataflow definition”
- 234 • Agreeing and creating a specification of the structure of metadata (metadata
235 structure definition) which defines the attributes and presentational
236 arrangement of a metadataset and their valid values and content
- 237 • Defining a subset or view of a metadata structure definition which allows
238 some restriction of content called a “metadataflow definition”



- 239 • Defining which subject matter domains are related to the dataflow and
240 metadataflow definitions to enable browsing
- 241 • Defining one or more lists of data providers (which includes metadata
242 providers)
- 243 • Defining which data providers have agreed to publish a given dataflow and/or
244 metadataflow definition - this is called a provision agreement

245

246 Publishing the data and metadata involves the following steps for a data provider:

247

- 248 • Making the metadata and data available in SDMX-ML conformant data files or
249 databases (which respond to an SDMX-ML query with SDMX-ML data) - the
250 data and metadata files or databases must be web-accessible, and must
251 conform to an agreed dataflow or metadataflow definition (key family or
252 metadata structure definition subset)
- 253 • Registering the published metadata and data files or databases with one or
254 more SDMX Registries

255

256 Notifying interested parties of newly published or re-published data, metadata or
257 changes in structural metadata:

258

- 259 • The Registry can optionally support a subscription-based notification service
260 which sends an email announcing all published data that meets the criteria
261 contained in the subscription request

262

263 Discovering published data and metadata involves the following steps:

264

- 265 • Optionally browsing a subject matter domain category scheme to find
266 dataflow definitions (and hence key families) and metadataflows which
267 structure the type of data and/or metadata being sought
- 268 • Build a query, in terms of the selected key family or metadata structure
269 definition, which specifies what data are required
- 270 • Submit the query to an SDMX Registry which will return a list of (URLs of)
271 data and metadata files and databases which satisfy the query
- 272 • Processing the query result set and retrieving data and/or metadata from the
273 supplied URLs

274 **4 SDMX REGISTRY/REPOSITORY ARCHITECTURE**

275 **4.1 Architectural Schematic**

276 The architecture of the SDMX Registry/Repository is derived from the objectives
277 stated above. It is a layered architecture that is founded by a structural metadata
278 repository which supports a provisioning metadata repository which supports the
279 registry services.

280

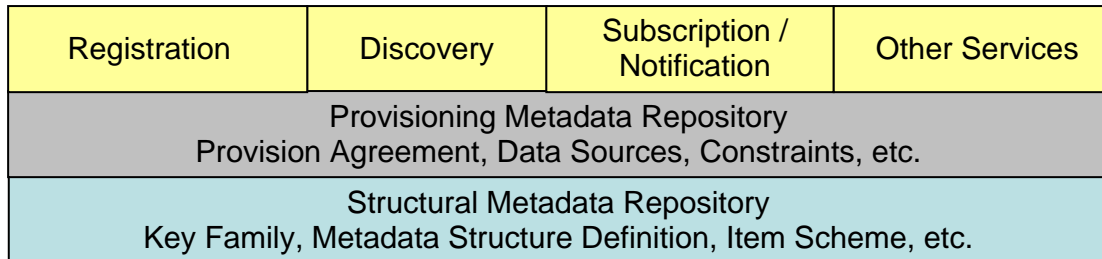

 281
282

Figure 1: Schematic Architecture of SDMX Registry/Repository

283 **4.2 Structural Metadata Repository**

284 The basic layer is that of a structural metadata service which supports the lifecycle of
 285 SDMX structural metadata artefacts such as Maintenance Agencies, Key Families,
 286 Reference Metadata Structure Definitions etc. This layer must allow structural
 287 definitions to be created, modified, and removed in a controlled fashion. It must also
 288 allow the structural metadata artefacts to be queried and retrieved either in part or as
 289 a whole. In order for the architecture to be scalable, the finest-grained piece of
 290 structural metadata that can be processed by the SDMX-RR is a
 291 MaintainableArtefact (see next section on the SDMX Information Model).

292 **4.3 Provisioning Metadata Repository**

293 Supported by the structural metadata service (effectively a repository) the
 294 provisioning service allows the definition of the data and metadata supply chain by
 295 the introduction of provision agreements. This is analogous to a “service level
 296 agreement” whereby a data provider commits to publishing a dataflow or
 297 metadataflow according to an agreed schedule. This repository also supports the
 298 definition of various types of data-store which model SDMX-conformant databases or
 299 files and can be specified for a data provider, or specifically for a data or metadata
 300 flow. In addition, the provisioning repository supports the definition of constraints
 301 which can define a valid sub-set of a key family or metadata structure definition.
 302 Constraints can be used for validation and also for attaching reference metadata to
 303 specific points within a dataflow or dataset. This layer must allow provisioning
 304 definitions to be created, modified, and removed in a controlled fashion. It must also
 305 allow the provisioning metadata artefacts to be queried and retrieved either in part or
 306 as a whole.

307 **4.4 Registry Services**

308 **4.4.1 Data and Metadata Registration Service**

309 The registration service allows SDMX conformant XML files and web-accessible
 310 databases containing published data and reference metadata to be registered in the
 311 SDMX Registry. The registration process involves validating the content of the data-
 312 sets or metadata-sets, extracting a concise representation of the contents in terms of
 313 concept values and storing this as a record in the registry to enable discovery of the
 314 original data-set or metadata-set.

315

316 The registration of data-sets and metadata-sets (that is, indexing of values found in
 317 the data and metadata sets) is supported by a service which depends upon the
 318 provisioning metadata repository for the following information:

- 319
- Is the data provider allowed to register the data-set or metadata-set?
- 320
- Does the content of the data-set or metadata-set meet the validation
- 321
- constraints?
- 322
- Has the data-set or metadata-set been published in a timely fashion?
- 323
- Has a database already been set up as a data-source for this data provider
- 324
- and data flow?

325

326 This service also depends upon the structural metadata repository for the following

327 information:

- 328
- What key family or metadata structure definition is used by the registered
- 329
- data?
- 330
- What are the components (dimensions, attributes, measures, identifier
- 331
- components etc.) of the key family or metadata structure definition?
- 332
- What are the valid representations of the concepts to which these
- 333
- components correspond?

334 **4.4.2 Data and Metadata Discovery**

335 The discovery service allows:

- 336
- dataflows (or metadataflows) to be found within a statistical subject-matter
- 337
- category scheme
- 338
- data-sets and metadata-sets to be located based on the dataflow (or
- 339
- metadataflow), key family (or metadata structure definition), or provider
- 340
- agreement and a specification of content via concept value definitions.

341

342 The discovery of data-sets (or metadata-sets) is supported by a service which

343 depends upon the registration service for the following information:

- 344
- Has an SDMX conformant XML file or database been registered for the
- 345
- dataflow (or metadataflow), key family (or metadata structure definition) or
- 346
- provision agreement in question?
- 347
- Does the content of the registered data-set (or metadata-set) meet the
- 348
- discovery query constraints (including both concept values and reference
- 349
- periods)?

350

351 The discovery service also depends upon the provisioning metadata repository for

352 the following information:

- 353
- Has an SDMX-conformant database been registered for the data provider or
- 354
- provision agreement in question?
- 355
- Does the content of a selected SDMX-conformant database meet the
- 356
- discovery query constraints (including both concept values and reference
- 357
- periods)?

358

359 This service also depends upon the structural metadata repository for the following

360 information:

- 361
- What dataflow (or metadataflow) is related to a selected statistical subject-
- 362
- matter category?
- 363
- What key family or metadata structure definition is used by the dataflow (or
- 364
- metadataflow)?
- 365
- What are the components (dimensions, attributes, measures, identifier
- 366
- components etc.) of the key family or metadata structure definition?
- 367
- What are the valid representations of the concepts to which these
- 368
- components correspond?

369

370 **4.4.3 Subscription and Notification**

371 The data sharing paradigm relies upon the consumers of data and metadata being
372 able to pull information from data providers' dissemination systems. For this to work
373 efficiently, a data consumer needs to know when to pull data, i.e. when something
374 has changed in the registry (e.g. a dataset has been updated and re-registered).
375 Additionally, many SDMX stakeholders will also want to know if a new key family,
376 code-list or metadata structure definition has been added. The subscription and
377 notification service comprises two parts: subscription management, and notification.

378

379 Subscription management involves an authenticated user submitting a subscription
380 request which contains:

- 381 • A query or constraint expression which defines the events for which the user
382 is interested(e.g. new data for a specific dataflow, or for a domain category, or
383 changes to a key family).
- 384 • A list of URIs or end-points to which an XML notification message can be
385 sent. Supported end-point types will be email (mailto:) and HTTP POST (a
386 normal http:// address).

387 Subscription management must also provide a way to list subscriptions and to delete
388 subscriptions.

389

390 Notification requires the structural metadata repository, the provisioning metadata
391 repository and the data and metadata registration service to monitor any event which
392 is of interest to a consumer (the object of a subscription request query) and to issue
393 an SDMX-ML notification document to the end-points specified in the relevant
394 subscriptions.

395 **5 SDMX INFORMATION MODEL (SDMX-IM)**

396 The data requirements presented here are consistent with, and have been derived
397 from, the SDMX-IM version 2.0. In some cases, to aid clarity and to avoid the
398 necessity of the reader to have concurrent access to the SDMX-IM, parts of the
399 model are reproduced here.

400

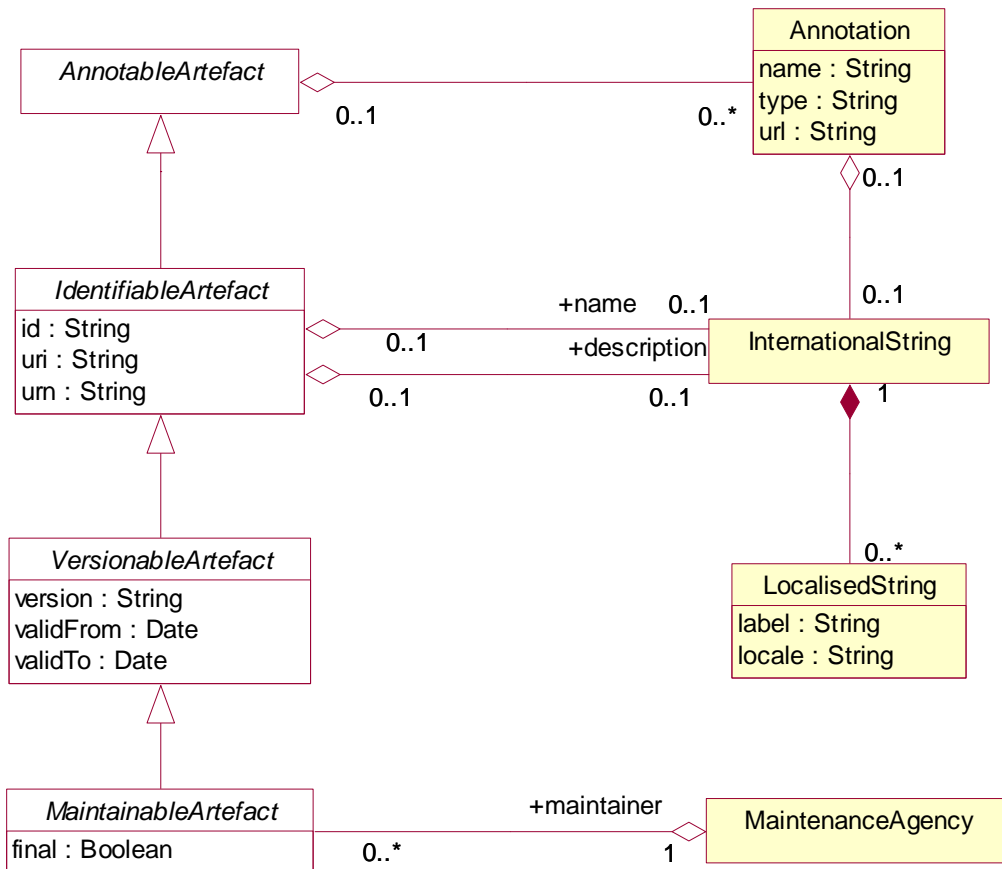
401 In some cases the model presented in this document is a logical model of the data
402 requirements of the interface function and as such is not a part of the SDMX-IM but is
403 drawn here to aid understanding. These specialised class diagrams are annotated as
404 such in this document.

405 **5.1 Identification, Versioning, and Maintenance**

406 All major classes of the SDMX Information model inherit from one of:

- 407 • IdentifiableArtefact - gives an object the ability to be uniquely identified (see
408 following section on identification), to carry a multi-lingual name and
409 description, to have a user-defined URI, and to carry multi-lingual
410 annotations.
- 411 • VersionableArtefact - has all of the above features plus a version number and
412 a validity period.
- 413 • MaintainableArtefact - has all of the above features plus an association to
414 the maintenance agency of the object.

415

416 **5.1.1 Identification, Versioning and Maintenance Model**


417

418

Figure 2: Class diagram of fundamental artefacts in the SDMX-IM

 419 The table below shows the identification and related data attributes to be stored in a
 420 registry for objects that are one of:

421

422

- Annotable

423

- Identifiable

424

- Versionable

425

- Maintainable

Object Type	Data Attributes	Status	Data type	Notes
Annotable	name	M	string	
	type	C	string	
	uri	C	string	
	annotation text	M		This can have language-specific variants.
Identifiable	id	M	string	
	uri	C	string	

Object Type	Data Attributes	Status	Data type	Notes
	urn	M	string	Although the urn is computable and therefore may not be stored physically, it is logically mandatory as applications must return the urn for each object, and must be able to service a query on an object referenced solely by its urn.
	name	M	string	This can have language-specific variants.
	description	C	string	This can have language-specific variants.
Versionable	All data as for Identifiable plus			
	version	M	string	This is the version number
	dateFrom	C	Date/time	
	dateTo	C	Date/time	
Maintainable	All data as for Versionable plus			
	final		boolean	Value of "true" indicates that this is a final specification and it cannot be changed unless as a new version.
	Maintenance Agency Id	M	string	The object must be linked to a maintenance agency

426

Table 1: Common Attributes of Object Types

427 5.2 Unique identification of SDMX objects

428 5.2.1 Universal Resource Name (URN)

429 To provide interoperability between SDMX Registry/Repositories in a distributed
 430 network environment, it is important to have a scheme for uniquely identifying (and
 431 thus accessing) all first-class (Identifiable) SDMX-IM objects. Most of these unique
 432 identifiers are composite (containing maintenance agency, or parent object
 433 identifiers), and to build references in XML would be very cumbersome (a reference
 434 to a maintenance agency has one component, but to a dimension has five). To
 435 simplify this, all SDMX objects also have a globally unique identifier called a universal
 436 resource name (URN) which is generated from the actual key components by SDMX-
 437 RR APIs. Note that this technique is used in ISO 15000 (ebXML registry
 438 specification), thus the adoption of a similar technique for the SDMX-RR makes it
 439 easier to map SDMX artifacts into an ISO 15000 compliant registry.

440

441 The structure of the URN is as follows:

442

443 `Prefix.SDMX-IM package name.classname=agency id:object id`



444 Or, where the object is in a maintained list of objects (called Item Scheme in the
445 SDMX-IM)

446

447 Prefix.SDMX-IM package name.classname=agency id:Item Scheme

448 id.object id

449 Or, where the object is in a maintained list of objects that can be hierarchic and
450 where the object id may not be unique in itself but only within the context of the
451 hierarchy

452

453 Prefix.SDMX-IM package name.classname=agency id:Item Scheme

454 id.object id.object id

455 (Maintenance) agency identifiers always precede major SDMX artefact identifiers and
456 act as a namespace or definition of ownership. Agency identifiers are always
457 separated from the maintainable artefact identifier by a colon ':'. All other identifiers
458 are separated by a period(.)

459

460 Where:

461

462 **Prefix:** urn:sdmx:

463

464 **SDMX-IM package name:** sdmx.infomodel.package.

465

466 The packages are:

467

base

468

codelist

469

conceptscheme

470

keyfamily

471

categoryscheme

472

registry

473

metadatastructure

474

transformation

475

process

476

mapping

477

478 Thus: urn:sdmx:org.sdmx.infomodel.registry.ProvisionAgreement

479

480 Examples

481

482 The key family EXT_DEBT maintained by the BIS would have the URN:

483

484 urn:sdmx:org.sdmx.infomodel.keyfamily.KeyFamily=BIS:EXT_DEBT

485 The URN for a code for Argentina could be:

486

487 urn:sdmx:org.sdmx.infomodel.codelist.Code=ISO:CL_3166A2.AR

488 The URN for a category in a hierarchic category scheme could be:

489

490 urn:sdmx:org.sdmx.infomodel.categoryscheme.Category=IMF:SDDS_C

491 ATEGORIES.REAL_SECTOR.BOP

492 **5.2.2 Identification components of SDMX objects**

 493 The table below describes the identification components for all SDMX object types
 494 that have identification. Note the the actual attributes are all Id, but have been
 495 prefixed by their class name or multiple class names to show navigation, e.g.
 496 conceptSchemeAgencyId is really the Id attribute of the Agency class that is
 497 associated to the ConceptScheme. * indicates that the object is maintainable.
 498

SDMX Class	Key attribute(s)	Example
Agency	Note. On the model this is an organisation role and so is theoretically identified the same way as data provider. However, all agencies are in the one scheme, maintained by a fixed and known agency, which does not need to be identified. agencyId	IMF
Concept Scheme*	conceptSchemeAgencyId:conceptSchemeId	SDMX:CORE
Concept	conceptSchemeAgencyId:conceptSchemeId.conceptId	SDMX:CORE.FREQ
CodeList*	codeListAgencyId:codeListId	SDMX:CL_FREQ
Code	codeListAgencyId:codeListId.codeId.codeId	SDMX:CL_FREQ.Q For hierarchical list it is IMF:CL_HIERARCHY.level_1_code.level_2_code
KeyFamily*	keyFamilyAgencyId:keyFamilyId	TFFS:EXT_DEBT
KeyDescriptor Measure Descriptor Attribute Descriptor	keyFamilyAgencyId:keyFamilyId. componentListId	TFFS:EXT_DEBT.KeyDescriptor TFFS:EXT_DEBT.MeasureDescriptor TFFS:EXT_DEBT.AttributeDescriptor
GroupKey Descriptor	keyFamilyAgencyId:keyFamilyId. groupKeyDescriptorId	TFFS:EXT_DEBT.SIBLING
Dimension Measure Attribute	keyFamilyAgencyId:keyFamilyId. ConceptSchemeAgencyId:conceptSchemeId.conceptId	TFFS:EXT_DEBT.SDMX:CORE.FREQ
Category Scheme*	categorySchemeAgencyId:categorySchemeId	IMF:SDDS
Reporting Taxonomy*	reportingTaxonomyAgencyId:reportingTaxonomyId	BIS:BASELII

SDMX Class	Key attribute(s)	Example
Category	categorySchemeAgencyId, categorySchemeId. categoryId.categoryId	IMF:SDDS: level_1_category.level_2_category ...
Organisation Scheme*	organisationSchemeAgencyId, organisationSchemeId,	SDMX:DATA_PROVIDER
Data Provider	organisationSchemeAgencyId: organisationSchemeId. dataProviderId	SDMX:DATA_PROVIDER.1F
Metadata Structure Definition*	MSDAgencyId:MSDId	IMF:SDDS_MSD
FullTarget Identifier	MSDAgencyId: MSDId.fullTargetId	IMF:SDDS_MSD.FULL_IDENTIFIER
PartialTarget Identifier	MSDAgencyId: MSDId.partialTargetId	IMF:SDDS_MSD.CATEGORY
Metadata Attribute	MSDAgencyId: MSDId. ConceptschemeAgencyId: conceptSchemeId. conceptId.conceptId	IMF:SDDS_MSD.SDMX:CORE.COMPI LATION
Identifiable Object Type	objectTypeSchemeAgency id: objectTypeSchemeId. objectType	SDMX:IM_OBJECT_TYPES.DATA_PR OVIDER
Dataflow*	dataflowAgencyId: dataflowId	TFFS:CRED_EXT_DEBT
Provision Agreement*	organisationSchemeAgencyId: organisationSchemeId. dataProviderId. dataflowAgencyId: dataflowId	SDMX:DATA_PROVIDER.1F.TFFS:CR ED_EXT_DEBT
Content Constraint	ConstrainableId.ContentConstr aintId	TFFS:CREDITOR_DATA.CONTENT_C ONSTRAINT
Attachment Constraint	ConstrainableId. AttachmentConstraintId	TFFS:CREDITOR_DATA.ATTACHMEN T_CONSTRAINT_ONE
Metadadataflow*	metadadataflowAgencyId: metadadataflowId	IMF:SDDS_CONTACTS
Dataset	organisationSchemeAgencyId: organisationSchemeId. dataProviderId.datasetId	Note that the dataset id is a unique id for the data provider that identifies this data set SDMX:DATA_PROVIDER.1F.DATA123
XSDataset	organisationSchemeAgencyId: organisationSchemeId. dataProviderId. xsdatasetId	SDMX:DATA_PROVIDER.1F.XSDATA 789
Metadata Set	organisationSchemeAgencyId: organisationSchemeId.	SDMX:DATA_PROVIDER.1F.METADA TA456

SDMX Class	Key attribute(s)	Example
	dataProviderId. metadadatasetId	
Component	StructureAgencyId: StructureId. ConceptSchemeAgencyId: ConceptSchemeId. ConceptId	This is an abstract class, so does not have a concrete example. Refer to Dimension, Attribute or Measure which are the concrete manifestations.
Hierarchical CodeList*	hierarchicalCodeListAgencyId: hierarchicalCodeListId	SDMX:CL_REF_AREA
Hierarchy	hierarchicalCodeListAgencyId: hierarchicalCodeListId: hierarchyId	SDMX:CL_REF_AREA.MX For hierarchical list it is IMF:CL_REF_AREA.level_1_code.level_2_code
StructureSet*	StructureSetAgencyId: StructureSetId	SDMX:BOP_STRUCTURES
StructureMap	StructureSetAgencyId: StructureSetId. StructureMapId	SDMX:BOP_STRUCTURES.TABLE1_ TABLE2
Component Map	StructureSetAgencyId: StructureSetId. StructureMapId. ComponentMapId	SDMX:BOP_STRUCTURES.TABLE1_ TABLE2. REFAREA_REPCOUNTRY
CodeListMap	StructureSetAgencyId: StructureSetId. ComponentMapId. CodeListMapId	SDMX:BOP_STRUCTURES.TABLE1_ TABLE2.REFAREA_REPCOUNTRY.C LREFAREA_CLREPCOUNTRY
CodeMap	StructureSetAgencyId: StructureSetId. ComponentMapId. CodeListMapId. CodeMapId	SDMX:BOP_STRUCTURES.TABLE1_ TABLE2.REFAREA_REPCOUNTRY.C LREFAREA_CLREPCOUNTRY.DE_ G E R
Category SchemeMap	StructureSetAgencyId: StructureSetId. CategorySchemeMapId	SDMX:BOP_STRUCTURES.SDMX_EU ROSTAT
CategoryMap	StructureSetAgencyId: StructureSetId. CategorySchemeMapId. CategoryMapId	SDMX:BOP_STRUCTURES.SDMX_EU ROSTAT.2.3.1_3.4.2.1
Organisation SchemeMap	StructureSetAgencyId: StructureSetId. OrganisationSchemeMapId	SDMX:BOP_STRUCTURES.DATA_PR OVIDER_MAP
Organisation RoleMap	StructureSetAgencyId: StructureSetId. OrganisationSchemeMapId. OrganisationRoleMapId	SDMX:BOP_STRUCTURES. DATA_PROVIDER_MAP.IMF_1C0
Concept SchemeMap	StructureSetAgencyId: StructureSetId. ConceptSchemeMapId	SDMX:BOP_STRUCTURES.SDMX_O ECD

SDMX Class	Key attribute(s)	Example
ConceptMap	StructureSetAgencyId: StructureSetId. ConceptSchemeMapId. ConceptMapId	SDMX:BOP_STRUCTURES.SDMX_O ECD.COVERAGE_AVAILABILITY
Process*	ProcessAgencyId: ProcessId	BIS:PROCESS1
ProcessStep	ProcessAgencyId: ProcessId. ProcessStepId	BIS:PROCESS1.STEP1
Transition	ProcessAgencyId: ProcessId. TransitionId	BIS:PROCESS1.STEP1_STEP3
Transformation Scheme*	TransformationSchemeAgency Id: TransformationSchemeId	BIS:TRANSFORMS
Expression Node	TransformationSchemeAgency Id: TransformationSchemeId. ExpressionNodeId	BIS:TRANSFORMS.EXPRESSION1

499

Table 2: Table of identification components for SDMX Identifiable Artefacts

500

6 STRUCTURAL DEFINITION METADATA

501

6.1 Introduction

502

The SDMX Registry must have the ability to support maintenance agencies in their role of defining and disseminating structural metadata artefacts. These include key families, code lists, concepts etc. and are fully explained in the SDMX Implementer's Guide. An authenticated maintenance agency may submit valid structural metadata definitions which must be stored in the registry. Note that the term "structural metadata" refers as a general term to all structural metadata (key families/data structure definitions, metadata structure definitions, codelists, concept schemas, etc.)

503

504

505

506

507

508

509

510

At a minimum, structural metadata definitions stored in the registry must be

511

accessible via an HTTP/HTTPS GET in the form of an SDMX-ML structure message,

512

or may be submitted as part of the registry interface. The use of SOAP is also

513

recommended, as described in the SDMX Web Services Guidelines. Other protocols

514

may be supported. The message may contain all structural metadata items for the

515

whole registry, structural metadata items for one maintenance agency, or individual

516

structural metadata items.

517

518

Structural metadata items may only be modified by the maintenance agency which

519

created them. They may only be deleted by the agency which created them, and then

520

only if they are not referred to by other items. The level of granularity for the

521

maintenance of SDMX Structural Metadata objects in the registry is the Maintainable

522

Artefact. In other words, any function such as add, modify, delete is at the level of the

523

Maintainable Artefact. For instance, if a Code is added to a Code List, or the Name of

524

a Code is changed, the registry function is "Modify Codelist" and whole of the Code

525

List will be presented in the interface.

526

527

The following table lists the Maintainable Artefacts.

528

Maintainable Artefacts		Content
Abstract Class	Concrete Class	
Item Scheme	Code List	Code
	Hierarchical Codelist	Hierarchy
	Concept Scheme	Concept
	Category Scheme	Category
	Organisation Scheme	Data Provider, Data Consumer, Maintenance Agency
	Reporting Taxonomy	Category
	Object Type Scheme	Identifiable Object Type
Structure	Key Family	Key Descriptor, Group Key Descriptor, Dimension, Attribute Descriptor, Data Attribute, Measure Descriptor, Measure
	Metadata Structure Definition	Full Target Identifier, Partial Target Identifier, Identifier Component, Report Structure, Metadata Attribute
Structure Usage	Dataflow Definition	
	Metadataflow Definition	
Process	Process	Process Step
Structure Set	Structure Set	Maps of various types
Provision Agreement	Provision Agreement	

529

Table 3: Table of Maintainable Artefacts for Structural Definition metadata

530

6.2 Item Scheme, Structure

531

6.2.1 Structure and Item Scheme: Basic Concepts

532

The artefacts included in the structural definitions are:

533

534

- All types of Item Scheme (Code List, Concept Scheme, Category Scheme, Organisation Scheme, Object Type Scheme)

535

536

- All types of Structure (Key Family, Metadata Structure Definition)

537

- All types of Structure Usage (Dataflow Definition, Metadataflow Definition)

538

539

The SDMX Registry must have the ability to support maintenance agencies in their role of defining and disseminating structural metadata artefacts. These include key families, code lists, concepts etc. An authenticated maintenance agency may submit valid structural metadata definitions which must be stored in the registry.

540

541

542

543

544

Structural metadata definitions stored in the registry must be accessible via an HTTP/HTTPS GET and optionally via SOAP (as explained in the SDMX Web

545

546 Services Guidelines) or via other protocol. The message may contain all structural
547 metadata items for the whole registry, structural metadata items for one maintenance
548 agency, or individual structural metadata items.

549

550 Structural metadata items may only be modified by the maintenance agency which
551 created them. They may only be deleted by the agency which created them if they
552 are not referred to by other items.

553

554 The SDMX-RR must support a globally unique identification scheme such that all
555 definitions of SDMX artefacts are globally unique and resolvable. The SDMX-RR
556 architecture must be able to resolve the unique identifier of an SDMX artefact to
557 produce an XML definition of that artefact.

558 **6.2.2 Structure and Item Scheme Model**

559 See the SDMX Information Model packages:

560

561 • Base – Item Scheme pattern

562 • Base – Organisations

563 • Base – Structure pattern

564 • Code List

565 • Concept Scheme

566 • Category Scheme

567 • Key Family

568 • Metadata Structure Definition

569 **6.2.3 Structure and Item Scheme: Functions and Behaviour**

570 The table below defines the data that could be present for the various functions. The
571 data required for the reference to “Identifiable”, “Versionable”, and “Maintainable” (and
572 which of the attributes are mandatory) are specified in Table 1 above.

573

574 An (s) after the Function indicates that multiple objects of the specified type may be
575 presented in the one function call.

576

577
 578

Function	Comprises object types	Data attributes	Status	Notes
<i>Reference to Item scheme and Item in this table infers functions generic to any type of Item Scheme i.e. Code List, Category Scheme, Organisation Scheme, Object Type Scheme</i>				
Add Item Scheme(s)		Item Scheme type	M	
		Maintainable	M	
		- final	C	Set to "true" if the specification is final and cannot be modified
	Item	Parent Item Id	C	Required if the scheme is hierarchic
		Versionable	M	
Modify Item Scheme(s)		Same data as for Add Item Scheme as the whole scheme is replaced	M	Cannot be modified if the "final" attribute is "true"
Delete Item Scheme(s)		Item Scheme urn or Maintainable or both		Deletion is not allowed if the Item Scheme or one of the Items in the scheme is referenced in the registry
<i>Item Scheme – functions specific to a specific type of Item Scheme</i>				
Concept Scheme	Concept	Representation - Date Range - Numeric Range - Pattern - Sequence		Required if a core representation is specified.



Function	Comprises object types	Data attributes	Status	Notes
		Representation is Item Scheme – reference to one of: - Code List - Category Scheme - Concept Scheme - Organisation Scheme		The Id of the relevant Item Scheme (see Table 2). Required if the core representation is an Item Scheme.
Add/Modify/Delete Data Provider Scheme	Data Provider			This is an Organisation Scheme which contains data providers. Therefore, the functions are as for Item Scheme.
Add/Modify/Delete Object Type Scheme	Identifiable Object Type			This is an Object Type Scheme which contains Identifiable Object Types. Therefore, the functions are as for Item Scheme.
Add/Modify/Delete Maintenance Agency Scheme	Maintenance Agency			This is an Organisation Scheme which contains maintenance agencies. Therefore, the functions are as for Item Scheme. For SDMX there is only one such scheme and it is maintained by SDMX
Add Key Family		Maintainable	M	
		- final	C	Set to “true” if the specification is final and cannot be modified.
	Key Descriptor	Identifiable – no data required as the Id is implied	M	
	- Dimension(s)	Identifiable – no data required as the Id is implied	M	The Dimensions must be present in the sequence required.
		Concept reference	M	The Id of the Concept – see Table 2

Function	Comprises object types	Data attributes	Status	Notes
		Representation - Date Range - Numeric Range - Pattern - Sequence	C	Required if a local representation is specified.
		Representation - Code List reference		The Id of the Code List (see Table 2). Required if the local representation is a Code List.
	Group Key Descriptor(s)	Identifiable – Id only	C	Only present if groups are required.
	- Dimension(s)	Dimension reference	M	Must reference one or more of the Dimensions of the Key Descriptor.
	Attribute Descriptor	Identifiable – no data required as the Id is implied	C	
	- Data Attribute(s)	Identifiable – no data required as the Id is implied	M	
		Concept reference	M	The Id of the Concept – see Table 2
		Representation - Date Range - Numeric Range - Pattern - Sequence	C	Required if a local representation is specified.
		Representation - Code List reference		The Id of the Code List (see Table 2). Required if the local representation is a Code List.
	Measure Descriptor	Identifiable – no data required as the Id is implied	C	
	- Measure(s)	Identifiable – no data	M	

Function	Comprises object types	Data attributes	Status	Notes
		required as the Id is implied		
		Concept reference	M	The Id of the Concept – see Table 2
		Representation - Date Range - Numeric Range - Pattern - Sequence	C	Required if a local representation is specified.
		Representation - Code List reference	C	The Id of the Code List (see Table 2). Required if the local representation is a Code List.
Modify Key Family		Same data as for Add Key Family as the whole definition is replaced		Cannot be modified if the “final” attribute is “true”.
Delete key Family		Key Family urn		Cannot be deleted if the Key Family is referenced from a Dataflow Definition.
Add Metadata Structure Definition		Maintainable	M	Set to “true” if the specification is final and cannot be changed.
		- final	C	
	Full Target Identifier	Identifiable – Id only	M	
	- Identifier Component(s)	Identifiable – no data required	M	
		Object type reference	M	The id of the Identifiable Object Type – see Table 2.
		Item Scheme – reference to one of: - Code List - Category Scheme - Concept Scheme - Organisation Scheme	C	The Id of the relevant Item Scheme (see Table 2).

Function	Comprises object types	Data attributes	Status	Notes
	Partial Target Identifier(s)	Identifiable – Id only	C	
	- Identifier Component(s)	Identifier Component reference	M	Must reference one or more of the Identifier Components of the Full Target Identifier.
	Report Structure	Identifiable – Id only	M	
	- Metadata Attributes	Identifiable – no data required	M	
		Concept reference	M	The Id of the Concept – see Table 2
		Representation - Date Range - Numeric Range - Pattern - Sequence	C	Required if a local representation is specified.
		Representation - Code List reference	C	The Id of the Code List (see Table 2). Required if the local representation is a Code List.

Table 4: Table of functions and data for Structure and Item Scheme

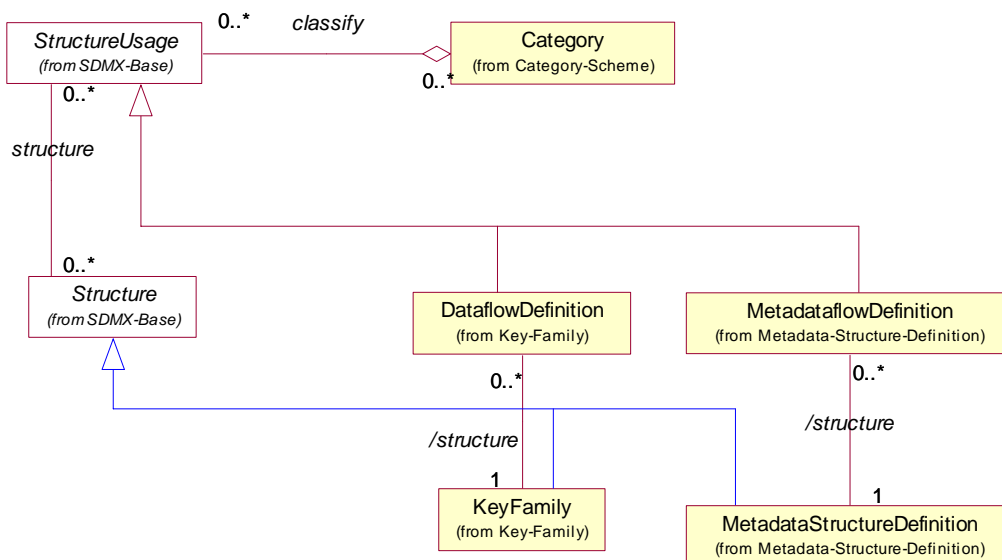
579

580

581 **6.3 Structure Usage**

 582 **6.3.1 Structure Usage: Basic Concepts**

583 The Structure Usage defines, in its concrete classes of Dataflow Definition and
 584 Metadataflow Definition, which flows of data and metadata use which specific
 585 Structure, and importantly for the support of data and metadata discovery, the
 586 Structure Usage can be linked to one or Category in one or more Category Scheme.
 587 This gives the ability for an application to discover data and metadata by “drilling
 588 down” the Category Schemes.

 589 **6.3.2 Structure Usage Model**


590

591

Figure 3: Class diagram of links from Structure Usage in the SDMX-IM

592 In addition to the maintenance of the Dataflow Definition and the Metadataflow
 593 Definition the following links require to be maintained in the registry:

594

595

- Dataflow Definition to Key Family

596

- Category to Dataflow Definition

597

- Metadataflow Definition to Metadata Structure Definition

598

- Category to Metadataflow Definition

599

600

601



602 **6.3.3 Structure Usage: Functions and Behaviour**

603

Function	Linked object types	Data attributes	Status	Notes
Add Dataflow Definition(s)	Dataflow Definition	Maintainable	M	
Modify Dataflow Definition(s)	This is a complete substitution of the Dataflow Definition by the revised specification	Maintainable	M	
Delete Dataflow Definition		urn of the Dataflow Definition		The Dataflow Definition cannot be deleted if there are any Data Sets registered which reference the Dataflow Definition
Add Metadataflow Definition(s)	Metadataflow Definition	Maintainable	M	
Modify Metadataflow Definition(s)	This is a complete substitution of the Metadataflow Definition by the revised specification	Maintainable	M	
Delete Metadataflow Definition		urn of the Metadataflow Definition		The Metadataflow Definition cannot be deleted if there are any Metadata Sets registered which reference the Metadataflow Definition
Link Objects	Category to Key Family	urn of the Category and urn of the Key Family	M	
	Category to Dataflow Definition	urn of the Category and urn of the Dataflow Definition	M	
	Category to Metadata Structure Definition	urn of the Category and urn of the Metadata Structure Definition	M	



604

Function	Linked object types	Data attributes	Status	Notes
	Category to Metadataflow Definition	urn of the Category and urn of the Metadataflow Definition	M	

605

Table 5: Table of functions and data for data and metadata flows

606

607 **7 DATA AND METADATA PROVISIONING**

608 **7.1 Provision Agreement, Data and Metadata Sources**

609 **7.1.1 Provisioning Agreement: Basic concepts**

610 Data provisioning defines a framework in which the provision of different types of
611 statistical data and metadata by various data providers can be specified and
612 controlled. This framework is the basis on which the existence of data can be made
613 known to the SDMX-enabled community and hence on which data can subsequently
614 be discovered. Such a framework can be used to regulate the data content to
615 facilitate the building of intelligent applications. It can also be used to impose service
616 level agreements, or other provisioning agreements in those scenarios that are
617 based on legal directives. Additionally, quality and timeliness information can be
618 supported by this framework which makes it practical to implement information
619 supply chain monitoring.

620

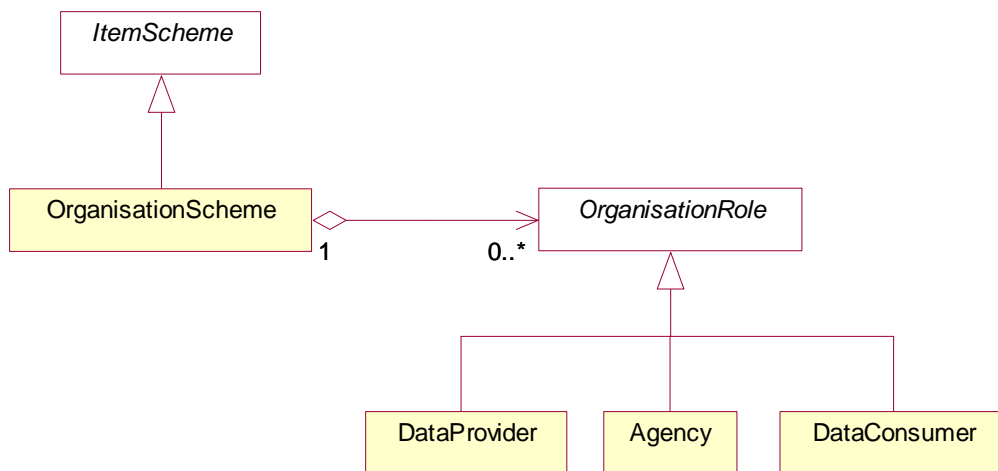
621 Note that in the SDMX-IM the class “Data Provider” encompasses both data and
622 metadata and the term “data provisioning” here includes both the provisioning of data
623 and metadata.

624

625 Although the SDMX registry work has been conducted with the data-sharing “pull”
626 model in mind, the information model does have limited support for the role of data
627 consumer, and so the specification could be extended to better support “push”
628 exchanges (bilateral and gateway scenarios). It should be noted, too, that in any
629 exchange scenario, the registry functions as a repository of structural metadata.

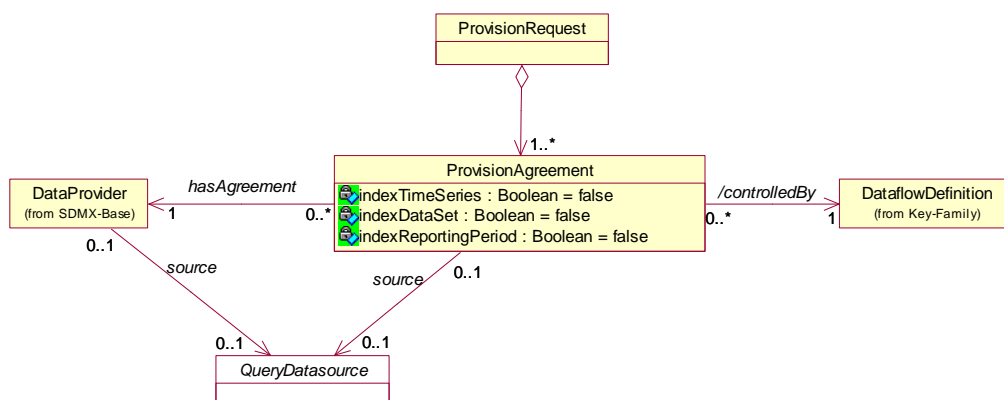
630 **7.1.2 Provisioning Agreement Model**

631 An organisation (or sub-unit) which publishes statistical data and wishes to make it
632 available to an SDMX enabled community is called a Data Provider. In terms of the
633 SDMX Information model, an Organisation can act in many roles (OrganisationRole)
634 and the diagram below depicts three cases of these: MaintenanceAgency
635 (abbreviated to Agency), DataProvider and DataConsumer. For those readers
636 unfamiliar with the term maintenance agency, it is part of an organisation which
637 defines structural definitions such as code lists and key families (maintainable
638 artefacts). The OrganisationScheme is a special kind of ItemScheme which can hold
639 zero or more DataProviders, Agencies or DataConsumers (since these are types of
640 OrganisationalRole).



641
642
643
644

The diagram below is a logical representation of the data required in order to maintain Provision Agreements.



645
646
647

Figure 4: Logical class diagram of the use of the Provision Request for submitting a Provision Agreement

648 A Provision Request contains a number of Provision Agreements, for each of which
649 there must be a reference to a Data Provider and a Dataflow Definition. The Data
650 Provider and the Dataflow Definition must exist already in order to set up a Provision
651 Agreement.

652

653 The Provision Agreement has three Boolean attributes which may be present to
654 determine how an SDMX compliant Dataset indexing application must index any
655 Datasets which are registered against it. The indexing application behaviour is as
656 follows:

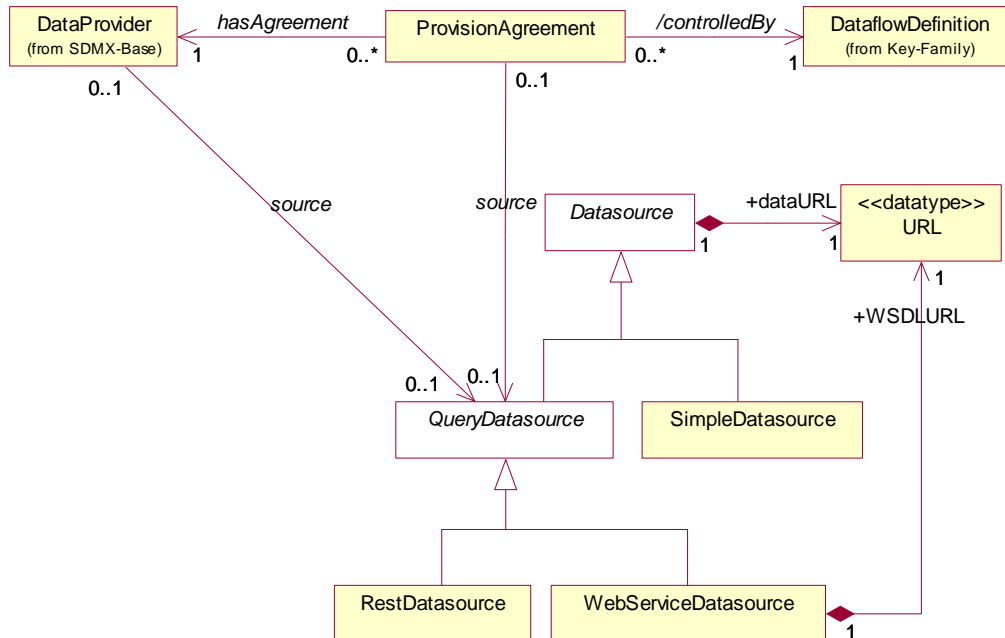
657

658 If `indexTimeSeries` is true, a compliant indexing application must index all the time
659 series keys when a Dataset is registered.

660

661 If `indexDataSet` is true, a compliant indexing application must index the range of
662 actual (present) values for each dimension of the Dataset. Note that this requires
663 much less storage than full key indexing, but this method cannot guarantee that a
664 specific combination of Dimension values (the Key) is actually present in the Dataset.
665

666 If indexReportingPeriod is true then a compliant indexing application must index the
 667 time period range(s) for which data are present in the Dataset.
 668



669
 670

Figure 5: Logical class diagram of data sources

671

672 The Query Datasource is an abstract class that represents a data source which can
 673 understand an SDMX-ML query and respond appropriately. There are two types of
 674 query data source, the Rest Datasource which is accessed via the REST protocol
 675 (HTTP/HTTPS) and the Web Service Datasource which supports SOAP based
 676 queries. Each of these different data sources inherit the dataURL from Datasource,
 677 and the Web Service Datasource has an additional URL to locate a WSDL document
 678 to describe how to access it. All other supported protocols are assumed to use the
 679 Simple Datasource URL.

680

681 A Simple Datasource is used to describe physical SDMX-ML files that are available
 682 at a URL. If the URL changes with successive releases of data then the Simple
 683 Datasource must be passed to the SDMX registration service for each release (see 8
 684 below).

685

686 A Data Provider may use one database system to supply all its SDMX datasets. In
 687 this case, it must be a QueryDatasource, and the query passed to it must specify (as
 688 a minimum) which data flow must be returned. (Since QueryDatasource is abstract,
 689 the Data Provider must specify one of RestDatasource or WebServiceDatasource).

690

691 A Data Provider may opt to specify a Datasource per data flow. This would be done
 692 by attaching the Datasource to the Provision Agreement.

693

694

695 **7.1.3 Provisioning Agreement: Functions and Behaviour**

Function	Comprises object types	Data attributes	Status	Notes
Add Provision Request	Provision Agreement(s)	DataflowDefinition Reference	M	URN or Maintainable reference to Dataflow Definition
		Data Provider Reference	M	URN. The Data Provider is maintained in an Organisation Scheme
		indexTimeSeries	C	Enables fine-grained indexing of data
		indexDataSet	C	Enables coarse-grained indexing of data
		indexReportingPeriod	C	Enables time-based indexing of data
	QueryDatasource		C	Specifies the data source
		dataUrl	M	Data access URL for data source
		WSDLUrl	C	WSDL URL for data source
Modify Provision Request	Provision Agreement(s)	URN	M	This is a complete replacement and the content is the same as for Add Provision Agreement
Delete Provision Request	Provision Agreement(s)	URN	M	Removes all ProvisionAgreements specified in the Request
Set Datasource	Constrainable Reference	URN	M	Identifies a Data Provider or Provision Agreement
	Datasource	dataURL	M	Data access URL for data source
		WSDLUrl	C	WSDL URL for data source
Unset Datasource	Constrainable Reference	URN	M	Identifies a Data Provider or Provision Agreement

Table 6: Table of functions and data for Provision Agreement

696

697

698 **7.2 Data and Metadata Constraints**

699 **7.2.1 Data and Metadata Constraints: Basic Concepts**

700 Constraints are metadata about data provisioning artefacts defined above (data
701 providers, data flows and/or provision agreements) which restrict and define these
702 artefacts. Specification of constraints gives enhanced semantics to data provisioning
703 artefacts, enabling more automated processing of the “information supply chain”.

704
705 Constraints comprise the specification of subsets of key values or attribute values
706 that are contained in a Datasource or is to be provided for a Dataflow or
707 Metadataflow Definition. This is important metadata because the full range of
708 possibilities which is implied by the Key Family Definition (e.g. the complete set of
709 valid keys is the Cartesian product of all the values in the code lists for each of the
710 Dimensions) is often more than is actually present in any specific Datasource, or
711 more than is intended to be supplied according to a specific Dataflow Definition.

712
713 Often a Data Provider will not be able to provide data for all key combinations, either
714 because the combination itself is not meaningful, or simply that the provider does not
715 have the data for that combination. In this case the Data Provider would constrain the
716 Datasource (at the level of the Provision Agreement or the Data Provider) by
717 supplying metadata that defined the key combinations or cube regions that are
718 available.

719
720 Furthermore, it is often useful to define subsets or views of the Key Family which
721 restrict values in some code lists, especially where many such subsets restrict the
722 same Key Family Definition. Such a view is called a Dataflow Definition, and there
723 can be one or more defined for any Key Family.

724
725 Whenever data is published or made available by a Data Provider, it must conform to
726 a Dataflow Definition (and hence to a Key Family). The Dataflow Definition is thus a
727 means of enabling content based processing. A Dataflow Definition may carry
728 confidential data, and this needs to be known to control access to potential
729 consumers.

730

731 7.2.2 Data and Metadata Constraints: Model

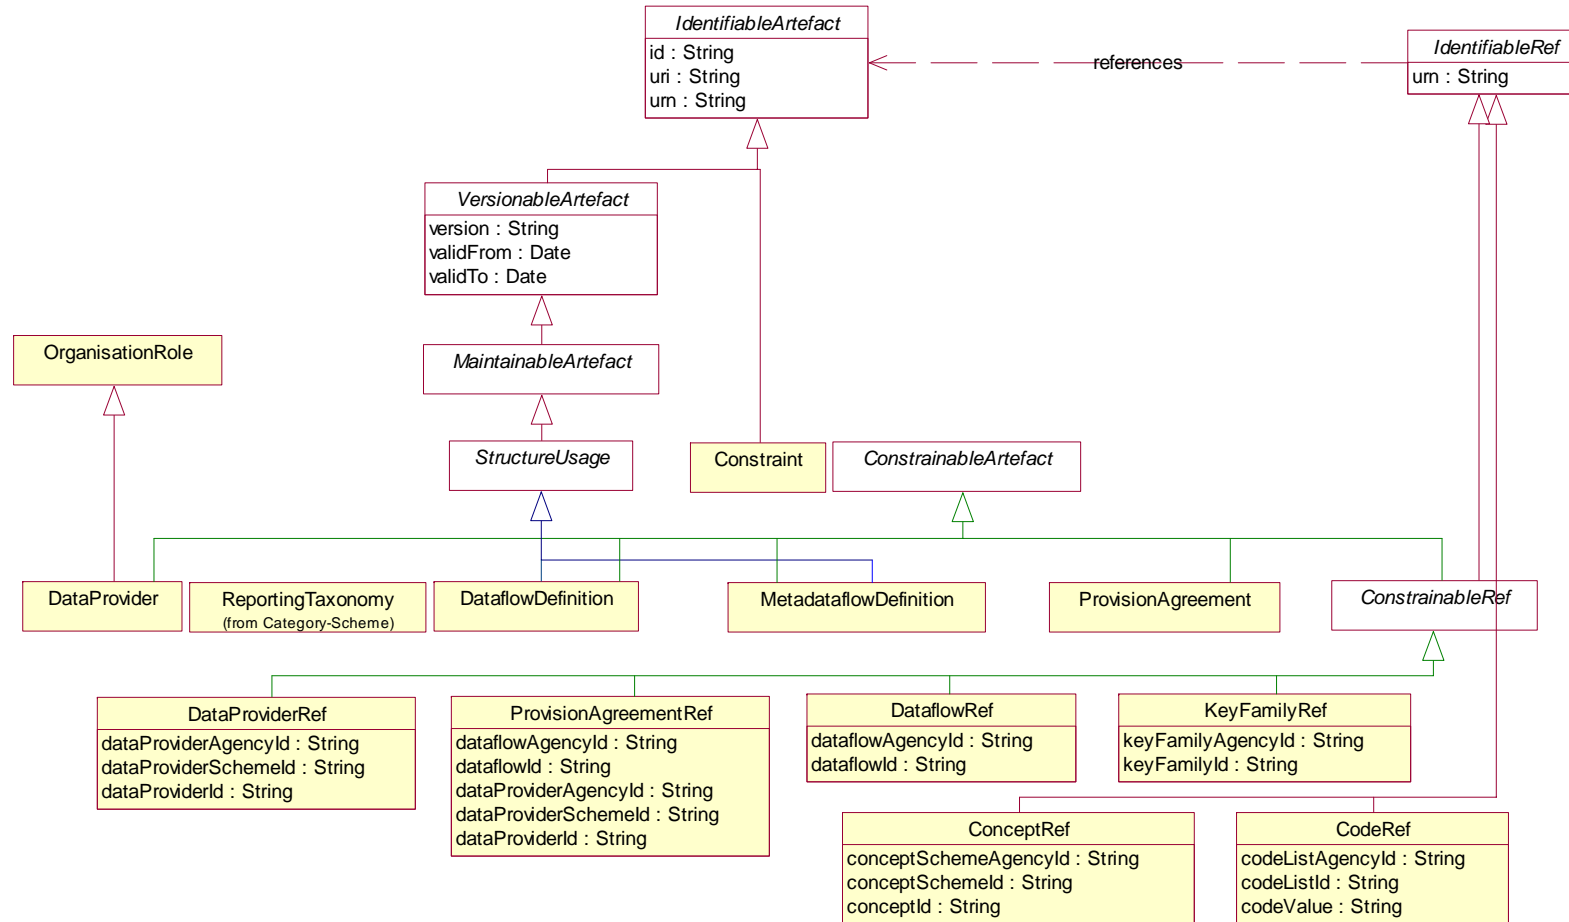


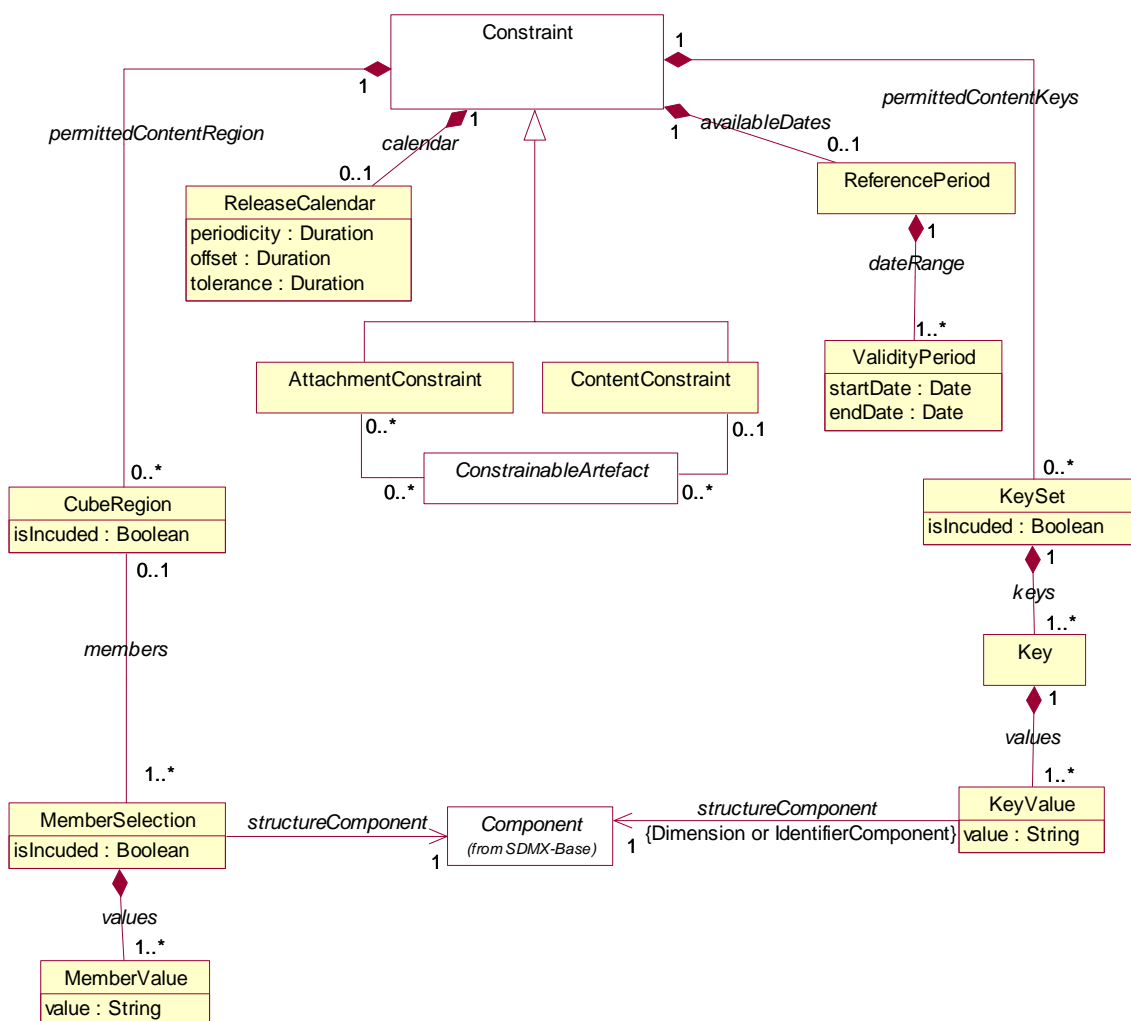
Figure 6: Class diagram inheritance between constrainable artifacts

732
733

734 The class diagram above shows that DataProvider, DataflowDefinition,
 735 MetadataflowDefinition, and ProvisionAgreement are all concrete sub-classes of
 736 ConstrainingArtefact. Note also the reference classes which inherit from
 737 ConstrainingArtefact: these are used as proxies to refer to the provisioning classes.

738

739 All sub-classes of ConstrainingArtefact may have one ContentConstraint and any
 740 number of AttachmentConstraints. The content constraint is used for defining legal
 741 values that dimension, attributes and measures may take. The content constraint
 742 may apply to a DataflowDefinition, thus defining a sub-set of the KeyFamily. If it is
 743 applied to a DataProvider or ProvisionAgreement then it specifies the content of the
 744 respective QueryDatasource. Constraints are also used in dataset registration and
 745 querying.



746

747

Figure 7: Class diagram showing Constraints

748 Any number of attachment constraints may be defined. They are used to sub-set the
 749 key family for metadata attribute attachment, rather than for constraining content.

750 Attachment constraints are not of direct interest to SDMX registry services.

751 The content constraint comprises four parts, one or more of which must be specified

752 Cube Region constraints limit the values that a dimension or coded attribute
 753 (Component) may take in a dataset which corresponds to the constrained data flow
 754 definition. This is done by specifying the legal values (MemberValue) for any
 755 dimension, attribute or measure. During the provisioning process, a Cube Region
 756 constraint can be used to define a sub-region of the whole key family. A simple
 757 example is that for an external debt dataflow, a cube region constraint could define
 758 that the “reporting role” dimension is “creditor”, thus no valid dataset would be
 759 allowed to contain “debtor” data. This also can be applied to attributes and measures:
 760 it allows the definition of, say, a single currency for EXT_DEBT statistics, even
 761 though at a national level a different dataflow could be used that published in national
 762 currency.

763
 764 The Key Set constraint performs a similar function, except that it is used to define
 765 legal combinations of values that may exist within a dataset. If KeyValues are
 766 specified for all Dimensions, then a time-series is defined. This mechanism can be
 767 used for building dependencies between dimensions.

768
 769 It is important to know which time periods a Provision Agreement covers. There is
 770 little point directing data consumers interested in 10 year trends to a data source
 771 which has data from 2001 onwards. The ReferencePeriod class must provide a
 772 standard way of defining time periods for which data is held. There may be many dis-
 773 contiguous periods for a data flow agreement.

774 If data publication falls under a legal directive or commercial contract, it is important
 775 to know when each data provider commits to publishing new data. The
 776 ReleaseCalendar must support enough calendaring information to support this.
 777 Making the data release process transparent amongst all data providers can provide
 778 great encouragement to meet targets.

779 Three attributes of type Duration (as per XML schema specification) are defined:

780

attribute	Description	example
periodicity	the period between releases of a dataset	P3M
offset	The period between the first publication date in the year and January 1 st .	P7D
tolerance	the point at which the dataset is deemed late	P3D

781
 782 The above example indicates that a datasource is published every quarter, it is
 783 published one week after the beginning of a quarter, and it must be released within
 784 three days of the expected publish date.

785

786 .

787 7.2.3 Data and Metadata Constraints: Functions and Behaviour

788

Function	Comprises object types	Data attributes	Status	Notes
Set Constraint		ConstrainableArtefactURN	M	URN of constrainable artefact, one of: Data Provider, Provision Agreement, or Dataflow Definition
	Constraint	Id	M	Id has fixed value "Content" for content constraint, otherwise user defined
		Type	M	Content or Attachment
	Cube Regions(s)		C	Defines a subset of the key family
		isIncluded	M	Is this cube region included or excluded from the whole cube?
	Member Selection(s)	isIncluded	M	Are these concept values included or to be excluded?
		ComponentRef	M	Reference to the Component (Dimension, Attribute or Measure) whose values are specified
	MemberValue(s)	Value	M	The value selected for a component
	Key Set(s)		C	Defines a specific combination of single-valued key family components
		isIncluded	M	Is this key set included or excluded from the whole cube?
	Key(s)		M	
	Key Value(s)	ComponentRef	M	Reference to the Component (Dimension, Attribute or Measure) whose value is specified
		Value	M	The value selected for a component



Function	Comprises object types	Data attributes	Status	Notes
	Release Calendar		C	Defines when data is to be published
		Periodicity	M	the period between releases of a dataset
		Offset	M	The period between the first publication date in the year and January 1 st .
		Tolerance	M	the point at which the dataset is deemed late

Table 7: Table of functions and data for Data and Metadata Constraints

789

790

791 8 DATA AND METADATA REGISTRATION

792 8.1 Basic Concepts

793 A data provider has published a new dataset conforming to an existing dataflow
794 definition (and hence key family). This is implemented as either a web-accessible
795 SDMX-ML file, or in a database which has a web-services (SOAP or REST) interface
796 capable of responding to and SDMX-ML Query with an SDMX-ML data stream.

797
798 The data provider wishes to make this new data available to a wide audience without
799 having to physically distribute the data. To do this, the data provider registers the
800 new dataset with one or more SDMX conformant registries that have been configured
801 with structural and provisioning metadata. In other words, the registry “knows” the
802 data provider and “knows” what data-flows the data provider has agreed to make
803 available.

804
805 SDMX-RR supports dataset and metadata set registration via the Registration
806 Request, which can be created by the data provider (giving the data provider
807 maximum control) or generated by a registry service on behalf of the data provider.
808 The registry responds to the registration request with a registration response which
809 indicates if the registration was successful. In the event of an error, the error
810 messages are returned as a Registry Exception within the response.

811

812 8.2 The Registration Request

813 8.2.1 Registration Request Model

814 The following UML diagram shows the composition of the registration request. Each
815 request is made up of one or more Registrations, one per dataset to be registered.
816 The Registration has a validity period (valid from/to) which instructs the registry
817 whether the registration is active or not. This allows a data provider to retire datasets
818 after they are no longer current. The last updated date is needed during the
819 discovery process to make sure the client knows which data is freshest.

820

821 The Registration has an action attribute which takes one of the following values:

urn:sdmx:registry:register:action:append,	Add this registration to the registry
urn:sdmx:registry:register:action:replace	As above, but replace any registrations for same Provision Agreement Reference
urn:sdmx:registry:register:action:delete	Delete registrations for Provision Agreement Reference

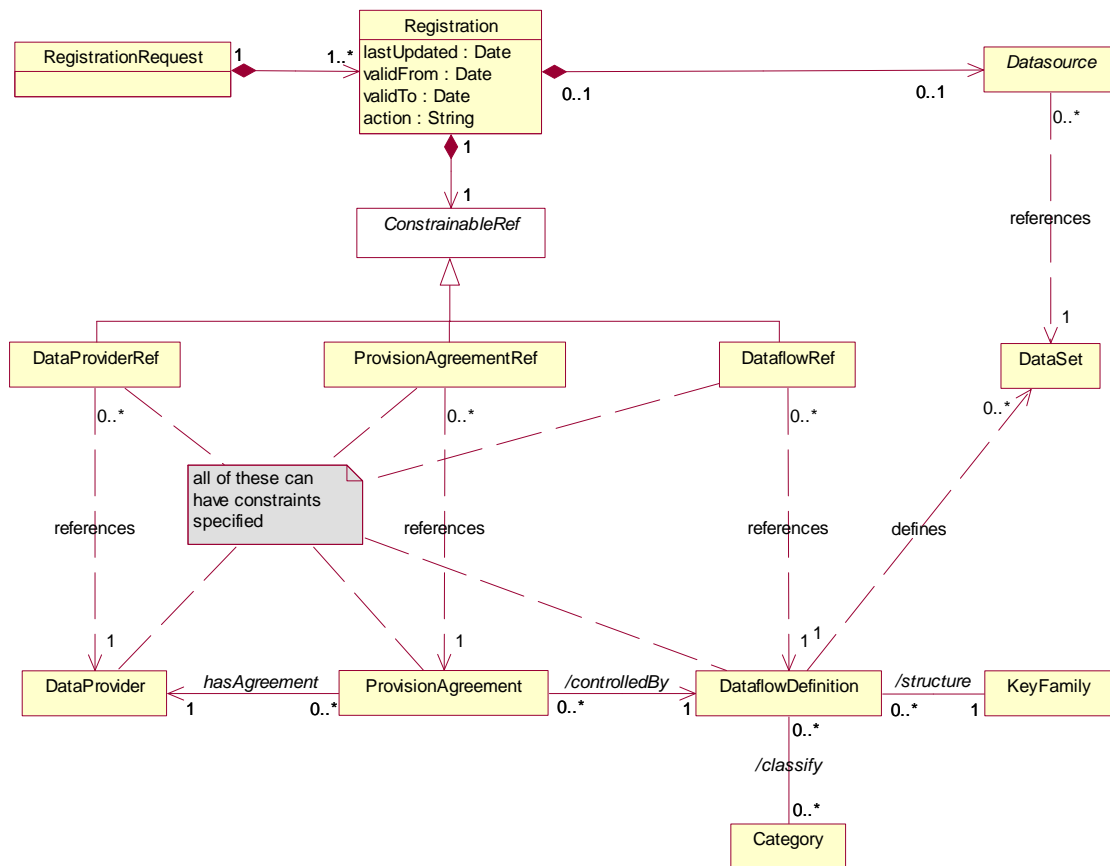
822

823 Each Registration may optionally specify a Datasource. If a QueryDatasource has
824 already been specified at the provision agreement (or data provider) level, then this
825 may also be used. When a dataset is published as an SDMX-ML file accessible
826 through a web server, then a SimpleDatasource would be specified with the
827 Registration.

828

829 The UML diagram also shows that the Registration must specify one
830 ConstrainableRef object. The concrete classes of ConstrainableRef are
831 DataProviderRef, ProvisionAgreementRef and DataflowRef (or MetadataflowRef).
832 Normally, when a dataset is registered, both the data provider and the

833 dataflow/metadataflow are known, so the Registration would specify a
 834 ProvisionAgreementRef. A DataProviderRef would indicate that the contents of the
 835 dataset are unknown and a DataflowRef (or MetadataflowRef) would indicate that the
 836 data provider is unknown: there may be use cases for these two scenarios, but this is
 837 not the focus of the current implementation. The reference objects are just simple
 838 references to existing objects in the registry. They contain the URN of the object to
 839 which they refer and also have the necessary attributes so that the URN can be
 840 calculated. For example, a DataflowRef has attributes or dataflowAgencyId and
 841 dataflowId from which a valid URN can be constructed automatically.
 842



843
 844

Figure 8: Class diagram showing the registration request

845 8.2.2 Registration Constraints

846 All ConstrainableRef objects can carry constraints (as discussed earlier for
 847 provisioning) although this is not shown on the preceding UML diagram. During the
 848 registration process, content constraints defined in the Registration Request are used
 849 to define the contents of the dataset to the registry. This is use case dependent: if the
 850 data provider would like clients to be able to locate an individual time-series via the
 851 registry, then each time-series key must be represented in a Key Set constraint. If the
 852 data provider is registering a database as a QueryDatasource, it is probable that a
 853 Cube Region constraint will be used to outline the probability of finding a particular
 854 time series. In scenarios where the registry is used for feeding a database, then
 855 content constraints are less important, as the database most likely would be fed
 856 using all data available.



857

858 ReferencePeriod constraints would also be specified to indicate the first and last time
859 periods for which data is available.

860

861

862 **8.2.3 Registration Request: Functions and Behaviour**

863

Function	Comprises object types	Data attributes	Status	Notes
Submit Registration Request	Registration(s)	lastUpdated	C	The date of when the dataset was last updated
		validFrom	C	The date from which the registry can broadcast the dataset
		validTo	C	The date at which the registry must retire the dataset
		Action	M	Append, Replace or Delete
	Datasource		C	Specifies the data-source for the dataset. If not specified, a data source must have been set up during provisioning
		dataUrl	M	Data access URL for data source
		WSDLUrl	C	WSDL URL for data source
	ConstrainableRef	URN	M	Reference to a constrainable object: one of Data Provider, Provision Agreement or Dataflow Definition
	Constraint		C	Identical to section 6.3.2

864

Table 8: Table of functions and data for Registration Request

865

866 **8.2.4 Registration Registry Service**

867 As mentioned earlier, a registry service must be provided which can take a URL of an
 868 SDMX-ML file (in generic format) and create a Registration Request from the
 869 contents of the file, then perform registration with the request. This service takes
 870 information to construct the Registration Request from the SDMX-ML message
 871 header and from the dataset itself. The mapping is shown in the table below:
 872

SDMX-ML Element/@attribute	Registration Request class and attribute
Sender/@id	ProvisionAgreementRef.dataProviderId
DataSetAgency	ProvisionAgreementRef.dataflowAgencyId
DataSetID	ProvisionAgreementRef.dataflowId
Extracted	Registration.lastUpdated
ReportingBegin	TimePeriod.startDate
ReportingEnd	TimePeriod.endDate
Value/@concept	Key.ConceptRef.conceptId
Value/@value	Key.ConceptRef.CodeRef.codeValue

873
 874 The last two rows of the table represent the values for time-series keys: each
 875 complete key is built and assembled into a key set constraint.


876 **8.3 Registration Response**

877 **8.3.1 Registration Response Model**

878 After a registration request has been submitted to the registry, a response is returned
 879 to the submitter indicating success or failure. Given that a registration request can
 880 hold many Registrations, then the response must give a result for each. The
 881 Registration Result class has a status field which is either set to "Success" or
 882 "Failure".

883
 884 If the registration has succeeded, a ProvisionAgreementRef (subclass of
 885 ConstrainableRef) will be returned - this holds the URN of the newly registered
 886 dataset plus a Datasource holding the URL to access the dataset or metadataset.

887
 888 In the event of registration failure, a set of Registry Exceptions are returned, giving
 889 the error messages that occurred during registration. It is entirely possible when
 890 registering a batch of datasets, that the response will contain some successful and
 891 some failed statuses. The UML for the response is shown below:
 892



893

894

895

Figure 9: Class diagram showing the registration response



896 **8.3.2 Registration Response: Functions and Behaviour**

Function	Comprises object types	Data attributes	Status	Notes
Consume Response	Registration	Registration Response	M	
	Registration Result(s)		M	One per dataset registered
		Status	M	urn:sdmx:registry:status:success or urn:sdmx:registry:status:failure.
	Registry Exception(s)		C	On failure
		Exception	M	The error message
	Datasource		C	On success, this gives the datasource that is registered. Note this could have been set up at provisioning time.
		dataUrl	M	Data access URL for data source
		WSDLUrl	C	WSDL URL for data source
	ConstrainableRef		C	On success
		URN	M	Reference to a constrainable object: usually Provision Agreement Ref. This handle can be used to replace or delete the registration later.

Table 9: Table of functions and data for Registration Response

897

898

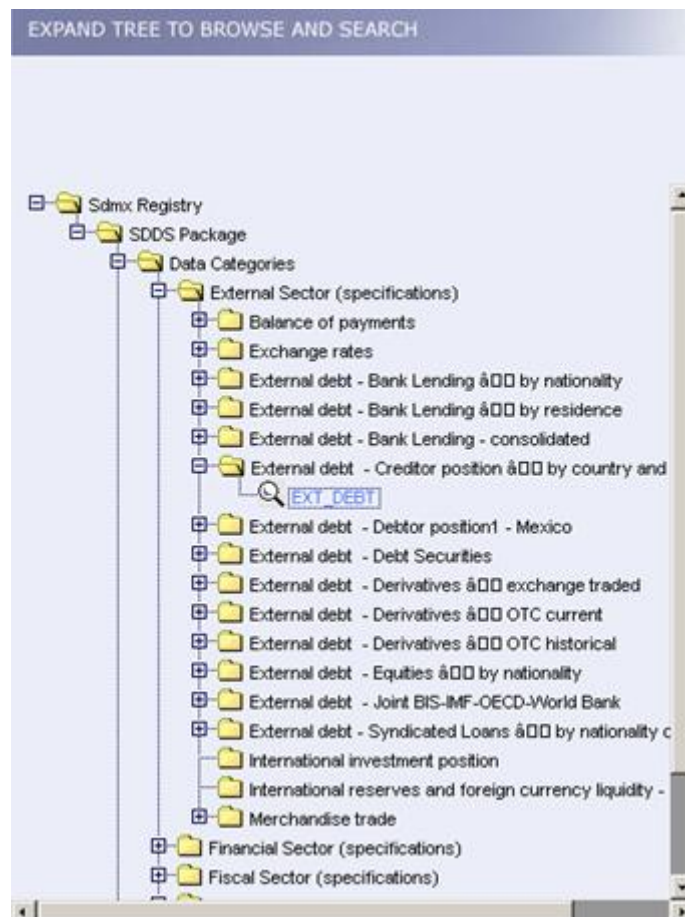
9 DATASET AND METADATA SET DISCOVERY

So far, we have described the ability to set up the SDMX-RR with structural and provisioning metadata, and how to register datasets and metadata sets against this framework. Now it's time to discuss the process of searching for datasets in the registry, referred to as discovery.

One of the main functions of the SDMX-RR is to aid in discovering data. This process may be instigated by either human users via a web interface, or via other application computer processes. The steps involved are outlined below.

9.1 Finding the Data and Metadata We Want

The user (or client program) may not know in advance the dataflow definition or key family that has been used to structure the dataset required. Browsing through subject matter domain categories will lead to dataflows and metadataflows which have been linked to the categories. The image below, taken from the SDMX case-study demonstration, shows how a dataflow called EXT_DEBT has been located by browsing through subject matter domain categories.

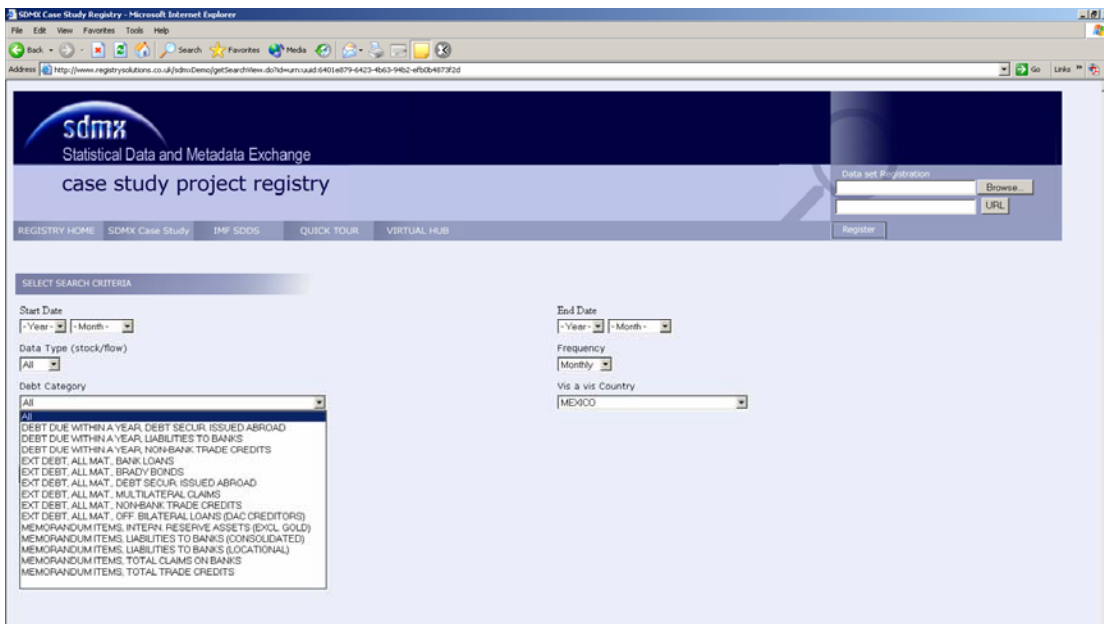


Alternatively, the end-user may be aware of some concepts that they are looking for. A simple application could suggest key families which use these concepts, and thus suggest dataflows which use these key families.

922 **9.2 Building a Query**

923 The registry supports queries at the key family/metadatastructure definition,
 924 dataflow/metadataflow, and provision agreement level. Key family queries are the
 925 most general and will search amongst all dataflows that use the specified key family.
 926 (The parallel is metadata structure definition queries which search amongst all
 927 metadataflows). A dataflow or metadataflow query will only target a specific dataflow
 928 or metadataflow, and will include datasets or metadatasets from any data provider
 929 (this is also true for the key family/metadata structure definition query). A provision
 930 agreement query is limited to a specific dataflow/metadataflow from a specific data
 931 provider.

932
 933 An end-user application could construct a search screen based on the underlying key
 934 family, with drop-down list for each dimension to allow dimension values to be
 935 specified for the query. The following image (again from the SDMX case-study
 936 demonstration) shows a screen which has been constructed around the EXT_DEBT
 937 dataflow and allows users to select values to build a query.
 938

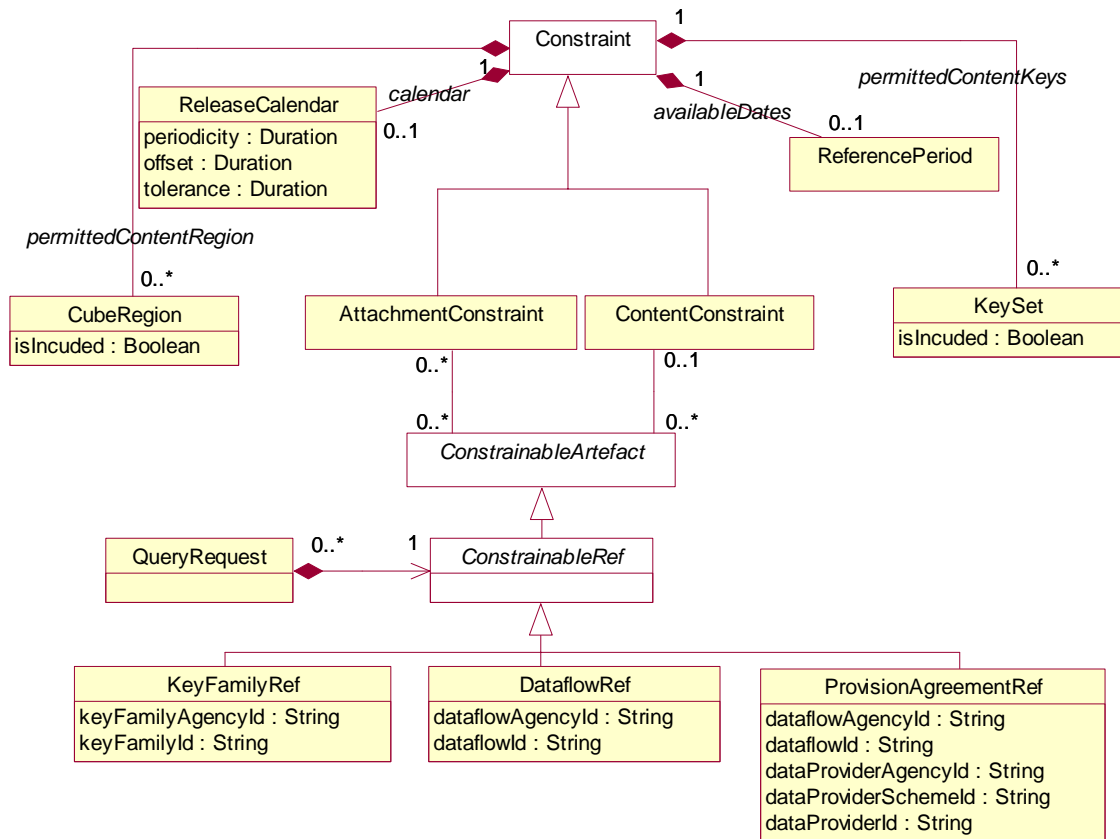


939
 940
 941 In the preceding image, one can see that the drop-down lists have been populated
 942 with all the possible values that a dimension can take. Because dataflows are
 943 constrained views of key families, the drop-down lists can actually be populated with
 944 code values that are legal for the selected dataflow. This sub-setting of code-lists via
 945 constraints on the dataflow reduces the possibility of searching for data which does
 946 not exist (i.e. sparse areas).

947
 948 The user input from such a selection screen is then used to construct either cube
 949 region or key set content constraints (as described earlier in the data provisioning
 950 section). Key set constraints are used when searching for an actual time-series key
 951 and hence specify values for all dimensions. A cube region constraint is specified for
 952 a more general search. In the screen above only the vis-à-vis (partner) country has
 953 been specified so this would indicate a cube region constraint.
 954

955 The query may also wish to restrict the search in temporal terms. Start and end dates
 956 from the query screen would be transformed into a reference period constraint.

957 **9.2.1 The Data/Metadata Query Model**



958

959

Figure 10: Class diagram showing the query request

960 The preceding UML class diagram shows how the query request is constructed from
 961 classes already used in data provisioning and registration.

962



963 **9.2.2 Data/Metadata Query Functions and Behaviour**

Function	Comprises object types	Data attributes	Status	Notes
Submit Query Request	QueryRequest		M	The date of when the dataset/metadata set was last updated
	ConstrainableRef	URN	M	Reference to a constrainable object: one of Data Provider, Provision Agreement or Dataflow/Metadataflow Definition
	Constraint		C	Identical to section 6.3.2

964 **Table 10: Table of functions and data for Query Functions**

964

965

966 **9.3 Processing the Query**

967 There are three types of query that can be presented to the SDMX-RR v1.0,
968 according to the way that constraints have been used within the QueryRequest.

969 **9.3.1 No Content (Key Set or Cube Region) Constraints**

970 This is the simplest case. The registry just has to find all datasets and datasources
971 that have been registered against the key family, dataflow or provision agreement
972 specified in the query. There may be some trimming of results by the registry to
973 comply with any specified ReferencePeriod constraint. This scenario is useful when
974 the registry is being used to feed a database or cache mechanism. Content based
975 queries will then be handled by the database or cache.

976

977 It should be remembered that datasources can be registered against data providers
978 and provision agreements during the provisioning process, and these should be
979 returned in the query result.

980 **9.3.2 Cube Region Content Constraint**

981 The process is as above, except all datasets and datasources must satisfy the cube
982 region constraint. That means any concept values in the cube region must be found
983 in the key set or cube region associated with the datasource or dataset in the
984 registry.

985 **9.3.3 Key Set Content Constraint**

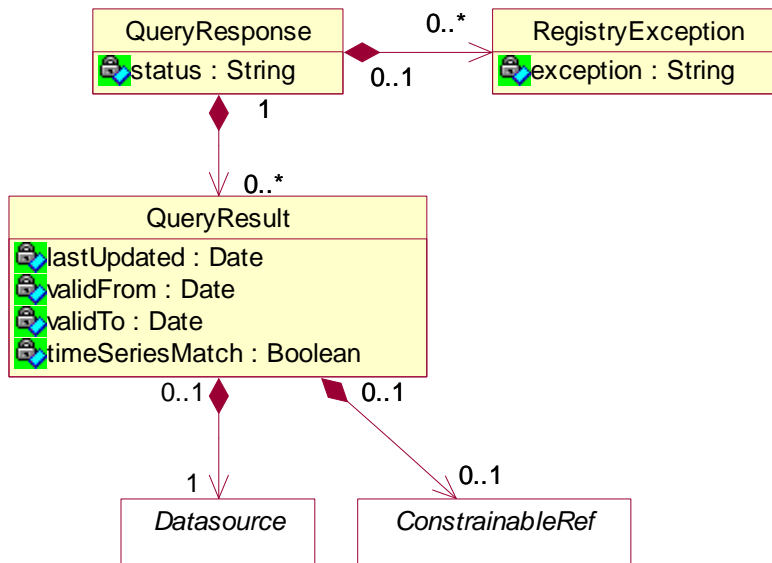
986 This process is as above, except that all concept values found in the individual keys
987 comprising the key set must be present in the cube region or key set constraints
988 registered with the datasource or dataset. In the case of a dataset registered with a
989 key set constraint, when a match occurs then it will be an exact match on the key of a
990 time-series, so the “timeSeriesMatch” attribute can be set to true in the query result.
991 If the match is with datasources registered with only a cube region constraint, it is not
992 certain that the particular combination of dimension values occurs in the same time-
993 series, but only that it is possible. Thus the “timeSeriesMatch” must be set to false.

994 **9.4 Handling the Results**

995 After the query request has been processed by the registry, the result is a set of
996 registry objects which correspond to either QueryDatasources which were registered
997 at provisioning time, or SimpleDatasources that were registered along with datasets.

998 **9.4.1 The Results Model**

999 SDMX-RR v1.0 constructs a response to the query which conforms to the UML
1000 model shown below. The registry must make sure that no duplicate datasources are
1001 returned in the result set. For each datasource in the resultset, a QueryResult object
1002 is created, which contains both datasource information containing the URL of the
1003 data file or web-service, and whether the datasource was registered against a key
1004 family, dataflow or provision agreement (the ConstrainingRef).
1005



1006
1007

Figure 11: Class diagram showing the query response

1008 Any results which matched a key set query with a key set constrained datasource are
 1009 marked as an exact time-series match. Information regarding when the last update to
 1010 the dataset was made, plus the validity period of the datasource are also returned.
 1011 This information allows consumers to make sure that they are using the most up-to-
 1012 date information.
 1013

1014 9.4.2 Query Response: Functions and Behaviour

1015

Function	Comprises object types	Data attributes	Status	Notes
Consume Query Response	QueryResponse		M	
	QueryResult(s)		C	On success
		Status	M	urn:sdmx:registry:status:success or urn:sdmx:registry:status:failure.
	RegistryException(s)		C	On failure
		Exception	M	The error message
	Datasource		C	On success, this gives the datasource that is registered. Note this could have been set up at provisioning time.
		dataUrl	M	Data access URL for data source
		WSDLUrl	C	WSDL URL for data source
	ConstrainableRef		C	On success
		URN	M	Reference to a constrainable object: usually Provision Agreement Ref. This handle can be used to replace or delete the registration later.

Table 11: Table of functions and data for Query Response

1016

1017

1018 **10 SUBSCRIPTION AND NOTIFICATION SERVICE**

1019 The contents of the SDMX Registry/Repository will change regularly: new code lists
 1020 and key families will be published, new datasets and metadata-sets will be
 1021 registered. To obviate the need for users to repeatedly query the registry to see when
 1022 new information is available, a mechanism is provided to allow users to be notified
 1023 when these events happen.

1024
 1025 A user can store a subscription in the registry that defines which events are of
 1026 interest, and either an email and/or an HTTP address to which qualifying events will
 1027 be delivered. The subscription will be identified in the registry by a URN which is
 1028 returned to the user when the subscription is created. If the user wants to delete the
 1029 subscription at a later point, the URN is used as identification. Subscriptions have a
 1030 validity period expressed as a date range (startDate, endDate) and the registry is free
 1031 to delete any expired subscriptions, and will notify the subscriber on expiry.

1032
 1033 When a registry/repository artefact is modified, any subscriptions which are
 1034 observing the object are activated, and either an email or HTTP POST is instigated to
 1035 report details of the changes to the user specified in the subscription. This is called a
 1036 “notification”.

1037

1038 **10.1 Subscription**

1039 **10.1.1 Common Information**

1040 Regardless of the type of Registry/Repository events being observed, a subscription
 1041 always contains:

- 1042 • A set of URIs describing the end-points to which notifications must be sent if
 1043 the subscription is activated. The URIs can be either mailto: or http: protocol.
 1044 In the former case an email notification is sent; in the latter an HTTP POST
 1045 notification is sent.
- 1046 • A user-defined identifier which is returned in the response to the subscription
 1047 request. This helps with asynchronous processing.
- 1048 • A validity period which defines both when the subscription becomes active
 1049 and expires. The subscriber will be sent a notification on expiration of the
 1050 subscription.
- 1051 • A selector which specifies which type of events are of interest. The set of
 1052 event types is:

Event Type	Comment
STRUCTURAL_REPOSITORY_EVENTS	Life-cycle changes to Maintainable Artefacts in the structural metadata repository. These include: KeyFamily, CategoryScheme, ConceptScheme, CodeList, MetadataStructureDefinition, DataflowDefinition, OrganisationScheme, MetadataflowDefinition
PROVISIONING_REPOSITORY_EVENTS	Life-cycle changes to ProvisionAgreement. Addition or deletion of Constraints and Datasources to ProvisionAgreement and to

Event Type	Comment
	DataProvider and DataflowDefinition by reference.
ALL_REPOSITORY_EVENTS	Both of the above
DATA_REGISTRATION_EVENTS	Whenever a published dataset is registered. This can be either a SDMX-ML data file or an SDMX conformant database.
METADATA_REGISTRATION_EVENTS	Whenever a published metadataset is registered. This can be either a SDMX-ML reference metadata file or an SDMX conformant database.
ALL_REGISTRATION_EVENTS	Both of the above
ALL_EVENTS	All of the above

1053 **10.1.2 Structural Metadata Events**

1054 Whenever a maintainable artefact (key family, concept scheme, code list, metadata
 1055 structure definition, category scheme, etc.) is added to, deleted from, or modified in
 1056 the structural metadata repository, a structural metadata event is triggered.
 1057 Subscriptions may observe all such events, or focus on specific artefacts such as a
 1058 key family and may use a wildcard text pattern match. This is specified in the same
 1059 way as for SQL, so that JEDH% would match any artefact whose ID begins with
 1060 JEDH.

1061 **10.1.3 Provisioning Metadata Events**

1062 Whenever a Provision Agreement is created, updated or deleted and whenever a
 1063 datasource or constraint is added to or removed from a data provider or dataflow (via
 1064 a reference) a provisioning metadata event occurs. As for structural metadata events,
 1065 wildcard expressions can be used to limit the scope of the subscription.
 1066

1067 **10.1.4 Registration Events**

1068 Whenever a dataset or metadata-set is registered a registration event is created. A
 1069 subscription may be observing all data or metadata registrations, or it may focus on
 1070 specific registrations as shown in the table below:
 1071

Selector	Comment
DataProvider	Any datasets or metadata-sets registered by the specified data provider will activate the subscription
ProvisionAgreement	Any datasets or metadata-sets registered by the specified data provider and dataflow (or metadataflow) will activate the subscription
Dataflow (&Metadataflow)	Any datasets or metadata-sets registered for the specified dataflow (or metadataflow) will activate the subscription
KeyFamily (&MetadataStructureDefinition)	Any datasets or metadata-sets registered for the those dataflows (or metadataflows) that are based on the



Selector	Comment
	specified key family will activate the subscription
Category	Any datasets or metadata-sets registered for the those dataflows (or metadataflows) that fall within the specified category (statistical domain) will activate the subscription

1072
1073
1074
1075
1076

The event will also capture the semantic of the registration; is this a deletion or replacement of an existing registration or a new registration. As for other subscription selection criteria, the use of the SQL wildcard syntax (at least %) should be supported.

1077

10.2 Notification

1078
1079
1080

A notification is an XML document that is sent to a user via email or http POST whenever a subscription is activated. It is an asynchronous one-way message. It has three forms according to the type of event that generated it:

1081
1082
1083

- structural metadata notification
- provisioning metadata notification
- registration notification

1084

1085

All three share some common information:

1086
1087
1088
1089
1090

- Date and time that the event occurred
- The URN of the artefact that caused the event
- The URN of the subscription that produced the notification
- Notification type: Structural, Provisioning or Registration.
- Event type: Add, Replace, or Delete.

1091

1092

Additionally, supplementary information will be contained in the notification as detailed below.

1093

1094

10.2.1 Structural Metadata Notification

1095
1096

The notification will contain the MaintainableArtefact that triggered the event in a form similar to the SDMX-ML structural message (using elements from that namespace).

1097

10.2.2 Provisioning Metadata Notification

1098
1099
1100
1101
1102

The notification will contain the ConstrainingArtefact that triggered the event in the form of an SDMX-ML provisioning message. The ConstrainingArtefacts in question are DataProviderRef, DataflowRef (&MetadataflowRef), and ProvisionAgreement. Any specified datasources and constraints may also be returned (except in the case of a deletion).

1103

10.2.3 Registration Notification

1104
1105
1106
1107
1108
1109

The notification will contain the Registration which has attributes of lastUpdated, validFrom and validTo. It will also contain any Datasource associated with the registration in URL form. Constraints may optionally be returned (in some cases they could be very large).

1110 10.2.4 Subscription and Notification: Functions and Behaviour

Function	Comprises object types	Data attributes	Status	Notes
Submit Subscription Request	SubscriptionRequest		M	
		validFrom	M	Date when subscription comes into effect.
		validTo	M	Date when subscription expires.
		eventType	M	What events are being observed: urn:sdmx:registry:subscription:event:repository:structure, urn:sdmx:registry:subscription:event:repository:provisioning, urn:sdmx:registry:subscription:event:repository:all, urn:sdmx:registry:subscription:event:registration:data, urn:sdmx:registry:subscription:event:registration:metadata, urn:sdmx:registry:subscription:event:registration:all, urn:sdmx:registry:subscription:event:all.
		userToken	C	A user-defined token that is returned in the Submit Subscription Response
	EndPoint(s)		M	Where to send the notification, may be repeated
		endPointType	M	urn:sdmx:registry:subscription:event:endpoint:mailto or urn:sdmx:registry:subscription:event:endpoint:http:post
	MaintainableArtefact		C	For subscribing to structural metadata events
		URN	C	Either URN or objectType, agencyId and id must be specified. Wildcards in SQL form may be used (i.e. %)
		objectType	C	KeyFamily, CodeList, ConceptScheme etc.
		agencyId	C	To identify a MaintainableArtefact
		id	C	To identify a MaintainableArtefact

Function	Comprises object types	Data attributes	Status	Notes
	ConstrainableArtefact		C	For subscribing to structural metadata events and registration events
		URN		Either URN or fully qualified identifying attributes must be specified. Wildcards in SQL form may be used (i.e. %)
		objectType	C	Constrainable type: one of DataProvider, ProvisionAgreement, DataflowDefinition, Metadataflow Definition
		dataProviderAgencyId	C	To identify a DataProvider or ProvisionAgreement
		dataProviderSchemeId	C	To identify a DataProvider or ProvisionAgreement
		dataProviderId	C	To identify a DataProvider or ProvisionAgreement
		dataflowAgencyId	C	To identify a (meta)Dataflow or ProvisionAgreement
		dataflowId	C	To identify a (meta)Dataflow or ProvisionAgreement
		keyFamilyAgencyId		To identify a (meta)Dataflow or ProvisionAgreement
		keyFamilyId		To identify a (meta)Dataflow or ProvisionAgreement
		CategorySchemeAgencyId		To identify a (meta)Dataflow or ProvisionAgreement
		CategorySchemeId		To identify a (meta)Dataflow or ProvisionAgreement
		CategoryId		To identify a (meta)Dataflow or ProvisionAgreement

Table 12: Table of functions and data for Subscription and Notification

1111

1112



1113 **11 SDMX-RR LOGICAL MODEL AND SDMX-ML**

1114 There is a natural correspondence between the logical interfaces described in the
1115 preceding sections and the SDMX-ML implementation of them. This section
1116 describes the basic document flows from a functional perspective.

1117
1118 Note that all registry interfaces are contained in a single XML element
1119 “RegistryInterface”, and thus always appear inside a standard SDMX-ML Message
1120 as the contents of this element. The message types used for interactions with the
1121 SDMX-RR are characterized by which of the child elements of RegistryInterface they
1122 contain – each message is only allowed to contain a single child element/registry
1123 interface.

1124
1125 Thus, there are a limited number of SDMX-ML Registry/Repository interfaces, or
1126 “message types”:

- 1127
- 1128 • SubmitSubscriptionRequest
- 1129 • SubmitSubscriptionResponse
- 1130 • NotifyRegistryEvent
- 1131 • SubmitRegistrationRequest
- 1132 • SubmitRegistrationResponse
- 1133 • QueryRegistrationRequest
- 1134 • QueryRegistrationResponse
- 1135 • SubmitStructureRequest
- 1136 • SubmitStructureResponse
- 1137 • QueryStructureRequest
- 1138 • QueryStructureResponse
- 1139 • SubmitProvisioningRequest
- 1140 • SubmitProvisioningResponse
- 1141 • QueryProvisioningRequest
- 1142 • QueryProvisioningResponse

1143
1144 Several patterns are evident from the naming conventions of the message
1145 types/interfaces. First, all interactions with the registry – with the exception of
1146 NotifyRegistryEvent – are designed in pairs. The first document – the one which
1147 invokes the SDMX-RR interface, is a “Request” document. The message returned by
1148 the interface is a “Response” document.

1149
1150 It should be noted that all interactions are assumed to be synchronous, with the
1151 exception of NotifyRegistryEvent. This document is sent by the SDMX-RR to all
1152 subscribers whenever an even occurs to which any users have subscribed. Thus, it
1153 does not conform to the request-response pattern, because it is inherently
1154 asynchronous.

1155
1156 The second pattern which is evident is that the SDMX-ML registry interface
1157 messages correspond to the layers of the logical SDMX-RR architecture. For each
1158 layer - the structural repository, the provisioning repository, and the data and
1159 metadata registry - there are two sets of SDMX-ML registry interface messages:

1160
1161
1162



1163 **Structural Repository:**

1164
1165 SubmitStructureRequest & SubmitStructureResponse

1166
1167 QueryStructureRequest & QueryStructureResponse

1168
1169 **Provisioning Repository:**

1170
1171 SubmitProvisioningRequest & SubmitProvisioningResponse

1172
1173 QueryProvisioningRequest & QueryProvisioningResponse

1174
1175 **Data and Metadata Registry:**

1176
1177 SubmitRegistrationRequest & SubmitRegistrationResponse

1178
1179 QueryRegistrationRequest & QueryRegistrationResponse

1180
1181
1182 Additionally, there is a set of messages which apply to all three layers of the
1183 architecture, related to subscription/notification:

1184
1185 **Subscription/Notification:**

1186
1187 SubmitSubscriptionRequest & SubmitSubscriptionResponse

1188
1189 NotifyRegistryEvent

1190
1191 The third pattern is the division between “submit” and “query” messages. All “submit”
1192 messages are used to register or submit objects to the appropriate level of the
1193 architecture. Thus, structural objects are submitted to the structural repository, while
1194 data sets are registered in the data and metadata registry layer. Submit messages
1195 can be understood as those which are used for maintenance activities – adding,
1196 modifying, or deleting maintained objects.

1197 Query messages are those which are used to discover the objects in each of the
1198 architectural layers. Again, the subscription and notification is a special case,
1199 because – although subscriptions are submitted in a normal fashion – the
1200 notifications are outside the basic model.