

SDMX STANDARDS: SECTION 6

SDMX TECHNICAL NOTES

FOR SDMX VERSION 2.1

Revision 2.0

July 2020

Revision History

Revision	Date	Contents
	April 2011	Initial release
1.0	April 2013	Added section 9 - Transforming between versions of SDMX
2.0	July 2020	Added section 10 – Validation and Transformation Language – before the Annex 1.

Contents

1	Purpose and Structure	1
1.1	Purpose	1
1.2	Structure	1
2	General Notes on This Document	1
3	Guide for SDMX Format Standards	2
3.1	Introduction.....	2
3.2	SDMX Information Model for Format Implementers	2
3.3	SDMX-ML and SDMX-EDI: Comparison of Expressive Capabilities and Function	3
3.4	SDMX-ML and SDMX-EDI Best Practices	6
4	General Notes for Implementers.....	12
4.1	Representations	12
4.2	Time and Time Format	14
4.3	Structural Metadata Querying Best Practices.....	24
4.4	Versioning and External Referencing.....	24
5	Metadata Structure Definition (MSD).....	25
5.1	Scope	25
5.2	Identification of the Object Type to which the Metadata is to be Attached.....	25
5.3	Report Structure	27
5.4	Metadata Set	28
6	Maintenance Agencies	29
7	Concept Roles	31
7.1	Overview	31
7.2	Information Model.....	31
7.3	Technical Mechanism	31

7.4	SDMX-ML Examples in a DSD	32
7.5	SDMX Cross Domain Concept Scheme.....	33
8	Constraints.....	34
8.1	Introduction.....	34
8.2	Types of Constraint.....	34
8.3	Rules for a Content Constraint.....	35
9	Transforming between versions of SDMX	44
9.1	Scope	44
9.2	Groups and Dimension Groups	44
10	Validation and Transformation Language (VTL).....	45
10.1	Introduction	45
10.2	References to SDMX artefacts from VTL statements	46
10.3	Mapping between SDMX and VTL artefacts.....	52
10.4	Mapping between SDMX and VTL Data Types	69
11	Annex I: How to eliminate extra element in the .NET SDMX Web Service	78
11.1	Problem statement.....	78
11.2	Solution.....	79
11.3	Applying the solution	82

1 **1 Purpose and Structure**

2 **1.1 Purpose**

3 The intention of this document is to document certain aspects of SDMX that are
4 important to understand and will aid implementation decisions. The explanations here
5 supplement the information documented in the SDMX XML schema and the
6 Information Model.

7 **1.2 Structure**

8 This document is organized into the following major parts:

9

10 A guide to the SDMX Information Model relating to Data Structure Definitions and
11 Data Sets, statement of differences in functionality supported by the different formats
12 and syntaxes for Data Structure Definitions and Data Sets, and best practices for use
13 of SDMX formats, including the representation for time period

14 A guide to the SDMX Information Model relating to Metadata Structure Definitions,
15 and Metadata Sets

16 Other structural artefacts of interest: agencies, concept role. constraint, partial code
17 list

18 **2 General Notes on This Document**

19 At this version of the standards, the term “Key family” is replaced by Data Structure
20 Definition (also known and referred to as DSD) both in the XML schemas and the
21 Information Model. The term “Key family” is not familiar to many people and its name
22 was taken from the model of SDMX-EDI (previously known as GESMES/TS). The
23 more familiar name “Data Structure Definition” which was used in many documents is
24 now also the technical artefact in the SDMX-ML and Information Model technical
25 specifications. The term “Key family” is still used in the SDMX-EDI specification.

26

27 There has been much work within the SDMX community on the creation of user
28 guides, tutorials, and other aides to implementation and understanding of the
29 standard. This document is not intended to duplicate the function of these
30 documents, but instead represents a short set of technical notes not generally
31 covered elsewhere.

32

33

34

35 **3 Guide for SDMX Format Standards**

36 **3.1 Introduction**

37 This guide exists to provide information to implementers of the SDMX format
38 standards – SDMX-ML and SDMX-EDI – that are concerned with data, i.e. Data
39 Structure Definitions and Data Sets. This section is intended to provide information
40 which will help users of SDMX understand and implement the standards. It is not
41 normative, and it does not provide any rules for the use of the standards, such as
42 those found in *SDMX-ML: Schema and Documentation* and *SDMX-EDI: Syntax and*
43 *Documentation*.
44

45 **3.2 SDMX Information Model for Format Implementers**

46 **3.2.1 Introduction**

47 The purpose of this sub-section is to provide an introduction to the SDMX-IM relating
48 to Data Structure Definitions and Data Sets for those whose primary interest is in the
49 use of the XML or EDI formats. For those wishing to have a deeper understanding of
50 the Information Model, the full SDMX-IM document, and other sections in this guide
51 provide a more in-depth view, along with UML diagrams and supporting explanation.
52 For those who are unfamiliar with DSDs, an appendix to the SDMX-IM provides a
53 tutorial which may serve as a useful introduction.
54

55 The SDMX-IM is used to describe the basic data and metadata structures used in all
56 of the SDMX data formats. The Information Model concerns itself with statistical data
57 and its structural metadata, and that is what is described here. Both structural
58 metadata and data have some additional metadata in common, related to their
59 management and administration. These aspects of the data model are not addressed
60 in this section and covered elsewhere in this guide or in the full SDMX-IM document.
61

62 The Data Structure Definition and Data Set parts of the information model are
63 consistent with the GESMES/TS version 3.0 Data Model (called SDMX-EDI in the
64 SDMX standard), with these exceptions:

65
66 the “sibling group” construct has been generalized to permit any dimension or
67 dimensions to be wildcarded, and not just frequency, as in GESMES/TS. It has been
68 renamed a “group” to distinguish it from the “sibling group” where only frequency is
69 wildcarded. The set of allowable partial “group” keys must be declared in the DSD,
70 and attributes may be attached to any of these group keys;

71 furthermore, whilst the “group” has been retained for compatibility with version 2.0
72 and with SDMX-EDI, it has, at version 2.1, been replaced by the “Attribute
73 Relationship” definition which is explained later

74 the section on data representation is now a convention, to support interoperability
75 with EDIFACT-syntax implementations (see section 3.3.2);

76 DSD-specific data formats are derived from the model, and some supporting features
77 for declaring multiple measures have been added to the structural metadata
78 descriptions

79 Clearly, this is not a coincidence. The GESMES/TS Data Model provides the
80 foundation for the EDIFACT messages in SDMX-EDI, and also is the starting point
81 for the development of SDMX-ML.

82

83 Note that in the descriptions below, text in courier and italicised are the names used
84 in the information model (e.g. *DataSet*).

85 **3.3 SDMX-ML and SDMX-EDI: Comparison of Expressive** 86 **Capabilities and Function**

87 SDMX offers several equivalent formats for describing data and structural metadata,
88 optimized for use in different applications. Although all of these formats are derived
89 directly from the SDM-IM, and are thus equivalent, the syntaxes used to express the
90 model place some restrictions on their use. Also, different optimizations provide
91 different capabilities. This section describes these differences, and provides some
92 rules for applications which may need to support more than one SDMX format or
93 syntax. This section is constrained to the Data Structure Definition and the Data Set.

94 **3.3.1 Format Optimizations and Differences**

95 The following section provides a brief overview of the differences between the
96 various SDMX formats.

97

98 Version 2.0 was characterised by 4 data messages, each with a distinct format:
99 Generic, Compact, Cross-Sectional and Utility. Because of the design, data in some
100 formats could not always be related to another format. In version 2.1, this issue has
101 been addressed by merging some formats and eliminating others. As a result, in
102 SDMX 2.1 there are just two types of data formats: *GenericData* and
103 *StructureSpecificData* (i.e. specific to one Data Structure Definition).

104

105 Both of these formats are now flexible enough to allow for data to be oriented in
106 series with any dimension used to disambiguate the observations (as opposed to
107 only time or a cross sectional measure in version 2.0). The formats have also been
108 expanded to allow for ungrouped observations.

109

110 To allow for applications which only understand time series data, variations of these
111 formats have been introduced in the form of two data messages;
112 *GenericTimeSeriesData* and *StructureSpecificTimeSeriesData*. It is important to note
113 that these variations are built on the same root structure and can be processed in the
114 same manner as the base format so that they do NOT introduce additional
115 processing requirements.

116

117 **Structure Definition**

118 The SDMX-ML Structure Message supports the use of annotations to the structure,
119 which is not supported by the SDMX-EDI syntax.

120 The SDMX-ML Structure Message allows for the structures on which a Data
121 Structure Definition depends – that is, codelists and concepts – to be either included
122 in the message or to be referenced by the message containing the data structure
123 definition. XML syntax is designed to leverage URIs and other Internet-based
124 referencing mechanisms, and these are used in the SDMX-ML message. This option
125 is not available to those using the SDMX-EDI structure message.

126 **Validation**

127 SDMX-EDI – as is typical of EDIFACT syntax messages – leaves validation to
128 dedicated applications (“validation” being the checking of syntax, data typing, and
129 adherence of the data message to the structure as described in the structural
130 definition.)

131 The SDMX-ML Generic Data Message also leaves validation above the XML syntax
132 level to the application.

133 The SDMX-ML DSD-specific messages will allow validation of XML syntax and
134 datatyping to be performed with a generic XML parser, and enforce agreement
135 between the structural definition and the data to a moderate degree with the same
136 tool.

137 **Update and Delete Messages and Documentation Messages**

138 All SDMX data messages allow for both delete messages and messages consisting
139 of only data or only documentation.

140 **Character Encodings**

142 All SDMX-ML messages use the UTF-8 encoding, while SDMX-EDI uses the ISO
143 8879-1 character encoding. There is a greater capacity with UTF-8 to express some
144 character sets (see the “APPENDIX: MAP OF ISO 8859-1 (UNOC) CHARACTER
145 SET (LATIN 1 OR “WESTERN”) in the document “SYNTAX AND
146 DOCUMENTATION VERSION 2.0”.) Many transformation tools are available which
147 allow XML instances with UTF-8 encodings to be expressed as ISO 8879-1-encoded
148 characters, and to transform UTF-8 into ISO 8879-1. Such tools should be used
149 when transforming SDMX-ML messages into SDMX-EDI messages and vice-versa.

150 **Data Typing**

152 The XML syntax and EDIFACT syntax have different data-typing mechanisms. The
153 section below provides a set of conventions to be observed when support for
154 messages in both syntaxes is required. For more information on the SDMX-ML
155 representations of data, see below.

156 **3.3.2 Data Types**

157 The XML syntax has a very different mechanism for data-typing than the EDIFACT
158 syntax, and this difference may create some difficulties for applications which support
159 both EDIFACT-based and XML-based SDMX data formats. This section provides a
160 set of conventions for the expression in data in all formats, to allow for clean
161 interoperability between them.

162
163 It should be noted that this section does not address character encodings – it is
164 assumed that conversion software will include the use of transformations which will
165 map between the ISO 8879-1 encoding of the SDMX-EDI format and the UTF-8
166 encoding of the SDMX-ML formats.

167
168 Note that the following conventions may be followed for ease of interoperation
169 between EDIFACT and XML representations of the data and metadata. For

170 implementations in which no transformation between EDIFACT and XML syntaxes is
171 foreseen, the restrictions below need not apply.

172

173 1. **Identifiers** are:

174 • Maximum 18 characters;

175 • Any of A..Z (upper case alphabetic), 0..9 (numeric), _ (underscore);

176 • The first character is alphabetic.

177 2. **Names** are:

178

179 • Maximum 70 characters.

180 • From ISO 8859-1 character set (including accented characters)

181 3. **Descriptions** are:

182

183 • Maximum 350 characters;

184 • From ISO 8859-1 character set.

185 4. **Code values** are:

186

187 • Maximum 18 characters;

188 • Any of A..Z (upper case alphabetic), 0..9 (numeric), _ (underscore), / (solidus,
189 slash), = (equal sign), - (hyphen);

190 However, code values providing values to a dimension must use only the following
191 characters:

192

193 A..Z (upper case alphabetic), 0..9 (numeric), _ (underscore)

194

195 5. **Observation values** are:

196

197 • Decimal numerics (signed only if they are negative);

198 • The maximum number of significant figures is:

199 • 15 for a positive number

200

201 • 14 for a positive decimal or a negative integer

202

203 • 13 for a negative decimal

204

205 • Scientific notation may be used.

206 6. **Uncoded statistical concept** text values are:

207

208 • Maximum 1050 characters;

- 209
- From ISO 8859-1 character set.

210 **7. Time series keys:**

211

212 In principle, the maximum permissible length of time series keys used in a data
213 exchange does not need to be restricted. However, for working purposes, an effort is
214 made to limit the maximum length to 35 characters; in this length, also (for SDMX-
215 EDI) one (separator) position is included between all successive dimension values;
216 this means that the maximum length allowed for a pure series key (concatenation of
217 dimension values) can be less than 35 characters. The separator character is a
218 colon (":") by conventional usage.

219 **3.4 SDMX-ML and SDMX-EDI Best Practices**

220 **3.4.1 Reporting and Dissemination Guidelines**

221 **3.4.1.1 Central Institutions and Their Role in Statistical Data Exchanges**

222 Central institutions are the organisations to which other partner institutions "report"
223 statistics. These statistics are used by central institutions either to compile
224 aggregates and/or they are put together and made available in a uniform manner
225 (e.g. on-line or on a CD-ROM or through file transfers). Therefore, central institutions
226 receive data from other institutions and, usually, they also "disseminate" data to
227 individual and/or institutions for end-use. Within a country, a NSI or a national central
228 bank (NCB) plays, of course, a central institution role as it collects data from other
229 entities and it disseminates statistical information to end users. In SDMX the role of
230 central institution is very important: every statistical message is based on underlying
231 structural definitions (statistical concepts, code lists, DSDs) which have been devised
232 by a particular agency, usually a central institution. Such an institution plays the role
233 of the reference "structural definitions maintenance agency" for the corresponding
234 messages which are exchanged. Of course, two institutions could exchange data
235 using/referring to structural information devised by a third institution.

236

237 Central institutions can play a double role:

238

- 239
- collecting and further disseminating statistics;

- 240
- devising structural definitions for use in data exchanges.

241 **3.4.1.2 Defining Data Structure Definitions (DSDs)**

242 The following guidelines are suggested for building a DSD. However, it is expected
243 that these guidelines will be considered by central institutions when devising new
244 DSDs.

245

246 Dimensions, Attributes and Code Lists

247

248 ***Avoid dimensions that are not appropriate for all the series in the data***
249 ***structure definition.*** If some dimensions are not applicable (this is evident from the
250 need to have a code in a code list which is marked as "not applicable", "not relevant"
251 or "total") for some series then consider moving these series to a new data structure
252 definition in which these dimensions are dropped from the key structure. This is a

253 judgement call as it is sometimes difficult to achieve this without increasing
254 considerably the number of DSDs.

255 **Devise DSDs with a small number of Dimensions for public viewing of data.** A
256 DSD with the number dimensions in excess 6 or 7 is often difficult for non specialist
257 users to understand. In these cases it is better to have a larger number of DSDs with
258 smaller “cubes” of data, or to eliminate dimensions and aggregate the data at a
259 higher level. Dissemination of data on the web is a growing use case for the SDMX
260 standards: the differentiation of observations by dimensionality which are necessary
261 for statisticians and economists are often obscure to public consumers who may not
262 always understand the semantic of the differentiation.

263 **Avoid composite dimensions.** Each dimension should correspond to a single
264 characteristic of the data, not to a combination of characteristics.

265 **Consider the inclusion of the following attributes.** Once the key structure of a
266 data structure definition has been decided, then the set of (preferably mandatory)
267 attributes of this data structure definition has to be defined. In general, some
268 statistical concepts are deemed necessary across all Data Structure Definitions to
269 qualify the contained information. Examples of these are:

- 270 • A descriptive title for the series (this is most useful for dissemination of data for
271 viewing e.g. on the web)
272
- 273 • Collection (e.g. end of period, averaged or summed over period)
274
- 275 • Unit (e.g. currency of denomination)
276
- 277 • Unit multiplier (e.g. expressed in millions)
278
- 279 • Availability (which institutions can a series become available to)
280
- 281 • Decimals (i.e. number of decimal digits used in numerical observations)
282
- 283 • Observation Status (e.g. estimate, provisional, normal)
284

285 Moreover, additional attributes may be considered as mandatory when a specific
286 data structure definition is defined.

287
288 **Avoid creating a new code list where one already exists.** It is highly
289 recommended that structural definitions and code lists be consistent with
290 internationally agreed standard methodologies, wherever they exist, e.g., System of
291 National Accounts 1993; Balance of Payments Manual, Fifth Edition; Monetary and
292 Financial Statistics Manual; Government Finance Statistics Manual, etc. When
293 setting-up a new data exchange, the following order of priority is suggested when
294 considering the use of code lists:

- 295 • international standard code lists;
- 296 • international code lists supplemented by other international and/or regional
297 institutions;

- 298 • standardised lists used already by international institutions;
- 299 • new code lists agreed between two international or regional institutions;
- 300 • new specific code lists.

301 The same code list can be used for several statistical concepts, within a data
 302 structure definition or across DSDs. Note that SDMX has recognised that these
 303 classifications are often quite large and the usage of codes in any one DSD is only a
 304 small extract of the full code list. In this version of the standard it is possible to
 305 exchange and disseminate a **partial code list** which is extracted from the full code
 306 list and which supports the dimension values valid for a particular DSD.

307
 308 Data Structure Definition Structure

309 The following items have to be specified by a structural definitions maintenance
 310 agency when defining a new data structure definition:

311 Data structure definition (DSD) identification:

- 312 • DSD identifier
- 313 • DSD name

314 A list of metadata concepts assigned as dimensions of the data structure definition.
 315 For each:

- 316 • (statistical) concept identifier
- 317 • ordinal number of the dimension in the key structure (SDMX-EDI only)
- 318 • code list identifier (Id, version, maintenance agency) if the
 319 representation is coded

320 A list of (statistical) concepts assigned as attributes for the data structure definition.
 321 For each:

- 322 • (statistical) concept identifier
- 323 • code list identifier if the concept is coded
- 324 • assignment status: mandatory or conditional
- 325 • attachment level
- 326 • maximum text length for the uncoded concepts
- 327 • maximum code length for the coded concepts

328 A list of the code lists used in the data structure definition. For each:

- 329 • code list identifier

- 330 • code list name
- 331 • code values and descriptions
- 332 Definition of data flow definitions. Two (or more) partners performing data
333 exchanges in a certain context need to agree on:
- 334 • the list of data set identifiers they will be using;
- 335
- 336 • for each data flow:
- 337 • its content and description
- 338 • the relevant DSD that defines the structure of the data reported or
339 disseminated according to the dataflow definition
- 340 **3.4.1.3 Exchanging Attributes**
- 341 **3.4.1.3.1 Attributes on series, sibling and data set level**
- 342 *Static properties.*
- 343 • Upon creation of a series the sender has to provide to the receiver values for all
344 mandatory attributes. In case they are available, values for conditional
345 attributes should also be provided. Whereas initially this information may be
346 provided by means other than SDMX-ML or SDMX-EDI messages (e.g.
347 paper, telephone) it is expected that partner institutions will be in a position to
348 provide this information in SDMX-ML or SDMX-EDI format over time.
- 349
- 350 • A centre may agree with its data exchange partners special procedures for
351 authorising the setting of attributes' initial values.
- 352
- 353 • Attribute values at a data set level are set and maintained exclusively by the
354 centre administrating the exchanged data set.
- 355
- 356 *Communication of changes to the centre.*
- 357 • Following the creation of a series, the attribute values do not have to be
358 reported again by senders, as long as they do not change.
- 359
- 360 • Whenever changes in attribute values for a series (or sibling group) occur, the
361 reporting institutions should report either all attribute values again (this is the
362 recommended option) or only the attribute values which have changed. This
363 applies both to the mandatory and the conditional attributes. For example, if a
364 previously reported value for a conditional attribute is no longer valid, this has
365 to be reported to the centre.
- 366
- 367 • A centre may agree with its data exchange partners special procedures for
368 authorising modifications in the attribute values.
- 369
- 370 Communication of observation level attributes "observation status", "observation
371 confidentiality", "observation pre-break".

- 372 • In SDMX-EDI, the observation level attribute “observation status” is
373 part of the fixed syntax of the ARR segment used for observation reporting.
374 Whenever an observation is exchanged, the corresponding observation
375 status must also be exchanged attached to the observation, regardless of
376 whether it has changed or not since the previous data exchange. This rule
377 also applies to the use of the SDMX-ML formats, although the syntax does
378 not necessarily require this.
379
- 380 • If the “observation status” changes and the observation remains
381 unchanged, both components would have to be reported.
382
- 383 • For Data Structure Definitions having also the observation level
384 attributes “observation confidentiality” and “observation pre-break” defined,
385 this rule applies to these attribute as well: if an institution receives from
386 another institution an observation with an observation status attribute only
387 attached, this means that the associated observation confidentiality and pre-
388 break observation attributes either never existed or from now they do not
389 have a value for this observation.

390 **3.4.2 Best Practices for Batch Data Exchange**

391 **3.4.2.1 Introduction**

392 Batch data exchange is the exchange and maintenance of entire databases between
393 counterparties. It is an activity that often employs SDMX-EDI formats, and might also
394 use the SDMX-ML DSD-specific data set. The following points apply equally to both
395 formats.

396 **3.4.2.2 Positioning of the Dimension "Frequency"**

397 The position of the “frequency” dimension is unambiguously identified in the data
398 structure definition. Moreover, most central institutions devising structural definitions
399 have decided to assign to this dimension the first position in the key structure. This
400 facilitates the easy identification of this dimension, something that it is necessary to
401 frequency's crucial role in several database systems and in attaching attributes at the
402 “sibling” group level.

403 **3.4.2.3 Identification of Data Structure Definitions (DSDs)**

404 In order to facilitate the easy and immediate recognition of the structural definition
405 maintenance agency that defined a data structure definition, most central institutions
406 devising structural definitions use the first characters of the data structure definition
407 identifiers to identify their institution: e.g. BIS_EER, EUROSTAT_BOP_01,
408 ECB_BOP1, etc.

409 **3.4.2.4 Identification of the Data Flows**

410 In order to facilitate the easy and immediate recognition of the institution
411 administrating a data flow definitions, many central institutions prefer to use the first
412 characters of the data flow definition identifiers to identify their institution: e.g.
413 BIS_EER, ECB_BOP1, ECB_BOP1, etc. Note that in GESMES/TS the Data Set
414 plays the role of the data flow definition (see *DataSet* in the SDMX-IM).
415

416 The statistical information in SDMX is broken down into two fundamental parts -
417 structural metadata (comprising the Data Structure Definition, and associated

418 Concepts and Code Lists) - see Framework for Standards -, and observational data
419 (the DataSet). This is an important distinction, with specific terminology associated
420 with each part. Data - which is typically a set of numeric observations at specific
421 points in time - is organized into data sets (*DataSet*) These data sets are structured
422 according to a specific Data Structure Definition (*DataStructureDefinition*) and are
423 described in the data flow definition (*DataflowDefinition*) The Data Structure
424 Definition describes the metadata that allows an understanding of what is expressed
425 in the data set, whilst the data flow definition provides the identifier and other
426 important information (such as the periodicity of reporting) that is common to all of its
427 component data sets.

428
429 Note that the role of the Data Flow (called *DataflowDefinition* in the model) and Data
430 Set is very specific in the model, and the terminology used may not be the same as
431 used in all organisations, and specifically the term Data Set is used differently in
432 SDMX than in GESMES/TS. Essentially the GESMES/TS term "Data Set" is, in
433 SDMX, the "Dataflow Definition" whilst the term "Data Set" in SDMX is used to
434 describe the "container" for an instance of the data.

435 **3.4.2.5 Special Issues**

436 **3.4.2.5.1 "Frequency" related issues**

437 **Special frequencies.** The issue of data collected at special (regular or irregular)
438 intervals at a lower than daily frequency (e.g. 24 or 36 or 48 observations per year,
439 on irregular days during the year) is not extensively discussed here. However, for
440 data exchange purposes:

- 441 • such data can be mapped into a series with daily frequency; this daily series
442 will only hold observations for those days on which the measured event takes
443 place;
- 444
- 445 • if the collection intervals are regular, additional values to the existing frequency
446 code list(s) could be added in the future.
- 447

448 **Tick data.** The issue of data collected at irregular intervals at a higher than daily
449 frequency (e.g. tick-by-tick data) is not discussed here either. However, for data
450 exchange purposes, such series can already be exchanged in the SDMX-EDI format
451 by using the option to send observations with the associated time stamp.

452

453 **4 General Notes for Implementers**

454 This section discusses a number of topics other than the exchange of data sets in
 455 SDMX-ML and SDMX-EDI. Supported only in SDMX-ML, these topics include the
 456 use of the reference metadata mechanism in SDMX, the use of Structure Sets and
 457 Reporting Taxonomies, the use of Processes, a discussion of time and data-typing,
 458 and some of the conventional mechanisms within the SDMX-ML Structure message
 459 regarding versioning and external referencing.

460
 461 This section does not go into great detail on these topics, but provides a useful
 462 overview of these features to assist implementors in further use of the parts of the
 463 specification which are relevant to them.

464 **4.1 Representations**

465 There are several different representations in SDMX-ML, taken from XML Schemas
 466 and common programming languages. The table below describes the various
 467 representations which are found in SDMX-ML, and their equivalents.

468

SDMX-ML Data Type	XML Schema Data Type	.NET Framework Type	Java Data Type
String	xsd:string	System.String	java.lang.String
Big Integer	xsd:integer	System.Decimal	java.math.BigInteger
Integer	xsd:int	System.Int32	int
Long	xsd:long	System.Int64	long
Short	xsd:short	System.Int16	short
Decimal	xsd:decimal	System.Decimal	java.math.BigDecimal
Float	xsd:float	System.Single	float
Double	xsd:double	System.Double	double
Boolean	xsd:boolean	System.Boolean	boolean
URI	xsd:anyURI	System.Uri	Java.net.URI or java.lang.String
DateTime	xsd:dateTime	System.DateTime	javax.xml.datatype.XMLGregorianCalendar
Time	xsd:time	System.DateTime	javax.xml.datatype.XMLGregorianCalendar
GregorianYear	xsd:gYear	System.DateTime	javax.xml.datatype.XMLGregorianCalendar
GregorianMonth	xsd:gYearMonth	System.DateTime	javax.xml.datatype.XMLGregorianCalendar
GregorianDay	xsd:date	System.DateTime	javax.xml.datatype.XMLGregorianCalendar
Day, MonthDay, Month	xsd:g*	System.DateTime	javax.xml.datatype.XMLGregorianCalendar
Duration	xsd:duration	System.TimeSpan	javax.xml.datatype

SDMX-ML Data Type	XML Schema Data Type	.NET Framework Type	Java Data Type
		n	.Duration

469

470 There are also a number of SDMX-ML data types which do not have these direct
 471 correspondences, often because they are composite representations or restrictions
 472 of a broader data type. For most of these, there are simple types which can be
 473 referenced from the SDMX schemas, for others a derived simple type will be
 474 necessary:

475

- 476 • AlphaNumeric (common:AlphaNumericType, string which only allows A-z and
 477 0-9)
- 478 • Alpha (common:AlphaType, string which only allows A-z)
- 479 • Numeric (common:NumericType, string which only allows 0-9, but is not
 480 numeric so that is can having leading zeros)
- 481 • Count (xs:integer, a sequence with an interval of "1")
- 482 • InclusiveValueRange (xs:decimal with the minValue and maxVale facets
 483 supplying the bounds)
- 484 • ExclusiveValueRange (xs:decimal with the minValue and maxVale facets
 485 supplying the bounds)
- 486 • Incremental (xs:decimal with a specified interval; the interval is typically
 487 enforced outside of the XML validation)
- 488 • TimeRange (common:TimeRangeType, start DateTime + Duration,)
- 489 • ObservationalTimePeriod (common: ObservationalTimePeriodType, a union
 490 of StandardTimePeriod and TimeRange).
- 491 • StandardTimePeriod (common: StandardTimePeriodType, a union of
 492 BasicTimePeriod and TimeRange).
- 493 • BasicTimePeriod (common: BasicTimePeriodType, a union of
 494 GregorianTimePeriod and DateTime)
- 495 • GregorianTimePeriod (common:GregorianTimePeriodType, a union of
 496 GregorianYear, GregorianMonth, and GregorianDay)
- 497 • ReportingTimePeriod (common:ReportingTimePeriodType, a union of
 498 ReportingYear, ReportingSemester, ReportingTrimester, ReportingQuarter,
 499 ReportingMonth, ReportingWeek, and ReportingDay).
- 500 • ReportingYear (common:ReportingYearType)
- 501 • ReportingSemester (common:ReportingSemesterType)
- 502 • ReportingTrimester (common:ReportingTrimesterType)
- 503 • ReportingQuarter (common:ReportingQuarterType)
- 504 • ReportingMonth (common:ReportingMonthType)
- 505 • ReportingWeek (common:ReportingWeekType)
- 506 • ReportingDay (common:ReportingDayType)
- 507 • XHTML (common:StructuredText, allows for multi-lingual text content that has
 508 XHTML markup)
- 509 • KeyValues (common:DataKeyType)
- 510 • IdentifiableReference (types for each identifiable object)
- 511 • DataSetReference (common:DataSetReferenceType)
- 512 • AttachmentConstraintReference
 513 (common:AttachmentConstraintReferenceType)

514

515

516 Data types also have a set of facets:

517

- 518 • isSequence = true | false (indicates a sequentially increasing value)
- 519 • minLength = positive integer (# of characters/digits)
- 520 • maxLength = positive integer (# of characters/digits)
- 521 • startValue = decimal (for numeric sequence)
- 522 • endValue = decimal (for numeric sequence)
- 523 • interval = decimal (for numeric sequence)
- 524 • timeInterval = duration
- 525 • startTime = BasicTimePeriod (for time range)
- 526 • endTime = BasicTimePeriod (for time range)
- 527 • minValue = decimal (for numeric range)
- 528 • maxValue = decimal (for numeric range)
- 529 • decimal = Integer (# of digits to right of decimal point)
- 530 • pattern = (a regular expression, as per W3C XML Schema)
- 531 • isMultiLingual = boolean (for specifying text can occur in more than one
- 532 language)

533

534 Note that code lists may also have textual representations assigned to them, in
535 addition to their enumeration of codes.s

536 **4.2 Time and Time Format**

537 **4.2.1 Introduction**

538 First, it is important to recognize that most observation times are a period. SDMX
539 specifies precisely how Time is handled.

540

541 The representation of time is broken into a hierarchical collection of representations.
542 A data structure definition can use of any of the representations in the hierarchy as
543 the representation of time. This allows for the time dimension of a particular data
544 structure definition allow for only a subset of the default representation.

545

546 The hierarchy of time formats is as follows (**bold** indicates a category which is made
547 up of multiple formats, *italic* indicates a distinct format):

548

- 549 • **Observational Time Period**
 - 550 ○ **Standard Time Period**
 - 551 ▪ **Basic Time Period**
 - 552 • **Gregorian Time Period**
 - 553 • *Date Time*
 - 554 ▪ **Reporting Time Period**
 - 555 ○ *Time Range*

556

557 The details of these time period categories and of the distinct formats which make
558 them up are detailed in the sections to follow.

559 **4.2.2 Observational Time Period**

560 This is the superset of all time representations in SDMX. This allows for time to be
561 expressed as any of the allowable formats.

562 **4.2.3 Standard Time Period**

563 This is the superset of any predefined time period or a distinct point in time. A time
564 period consists of a distinct start and end point. If the start and end of a period are
565 expressed as date instead of a complete date time, then it is implied that the start of
566 the period is the beginning of the start day (i.e. 00:00:00) and the end of the period is
567 the end of the end day (i.e. 23:59:59).

568 **4.2.4 Gregorian Time Period**

569 A Gregorian time period is always represented by a Gregorian year, year-month, or
570 day. These are all based on ISO 8601 dates. The representation in SDMX-ML
571 messages and the period covered by each of the Gregorian time periods are as
572 follows:

573

574 **Gregorian Year:**

575 Representation: xs:gYear (YYYY)

576 Period: the start of January 1 to the end of December 31

577 **Gregorian Year Month:**

578 Representation: xs:gYearMonth (YYYY-MM)

579 Period: the start of the first day of the month to end of the last day of the month

580 **Gregorian Day:**

581 Representation: xs:date (YYYY-MM-DD)

582 Period: the start of the day (00:00:00) to the end of the day (23:59:59)

583 **4.2.5 Date Time**

584 This is used to unambiguously state that a date-time represents an observation at a
585 single point in time. Therefore, if one wants to use SDMX for data which is measured
586 at a distinct point in time rather than being reported over a period, the date-time
587 representation can be used.

588 Representation: xs:dateTime (YYYY-MM-DDThh:mm:ss)¹

589 **4.2.6 Standard Reporting Period**

590 Standard reporting periods are periods of time in relation to a reporting year. Each of
591 these standard reporting periods has a duration (based on the ISO 8601 definition)
592 associated with it. The general format of a reporting period is as follows:

593

594 [REPORTING_YEAR]-[PERIOD_INDICATOR][PERIOD_VALUE]

595

596 Where:

597 REPORTING_YEAR represents the reporting year as four digits (YYYY)

598 PERIOD_INDICATOR identifies the type of period which determines the
599 duration of the period

600 PERIOD_VALUE indicates the actual period within the year

601

602 The following section details each of the standard reporting periods defined in SDMX:

603

604 **Reporting Year:**

605 Period Indicator: A

¹ The seconds can be reported fractionally

606 Period Duration: P1Y (one year)
607 Limit per year: 1
608 Representation: common:ReportingYearType (YYYY-A1, e.g. 2000-A1)
609 **Reporting Semester:**
610 Period Indicator: S
611 Period Duration: P6M (six months)
612 Limit per year: 2
613 Representation: common:ReportingSemesterType (YYYY-Ss, e.g. 2000-S2)
614 **Reporting Trimester:**
615 Period Indicator: T
616 Period Duration: P4M (four months)
617 Limit per year: 3
618 Representation: common:ReportingTrimesterType (YYYY-Tt, e.g. 2000-T3)
619 **Reporting Quarter:**
620 Period Indicator: Q
621 Period Duration: P3M (three months)
622 Limit per year: 4
623 Representation: common:ReportingQuarterType (YYYY-Qq, e.g. 2000-Q4)
624 **Reporting Month:**
625 Period Indicator: M
626 Period Duration: P1M (one month)
627 Limit per year: 1
628 Representation: common:ReportingMonthType (YYYY-Mmm, e.g. 2000-M12)
629 Notes: The reporting month is always represented as two digits, therefore 1-9
630 are 0 padded (e.g. 01). This allows the values to be sorted chronologically
631 using textual sorting methods.
632 **Reporting Week:**
633 Period Indicator: W
634 Period Duration: P7D (seven days)
635 Limit per year: 53
636 Representation: common:ReportingWeekType (YYYY-Www, e.g. 2000-W53)
637 Notes: There are either 52 or 53 weeks in a reporting year. This is based on the
638 ISO 8601 definition of a week (Monday - Saturday), where the first week of a
639 reporting year is defined as the week with the first Thursday on or after the
640 reporting year start day.² The reporting week is always represented as two
641 digits, therefore 1-9 are 0 padded (e.g. 01). This allows the values to be sorted
642 chronologically using textual sorting methods.
643 **Reporting Day:**
644 Period Indicator: D
645 Period Duration: P1D (one day)
646 Limit per year: 366
647 Representation: common:ReportingDayType (YYYY-Dddd, e.g. 2000-D366)
648 Notes: There are either 365 or 366 days in a reporting year, depending on
649 whether the reporting year includes leap day (February 29). The reporting day
650 is always represented as three digits, therefore 1-99 are 0 padded (e.g. 001).

² ISO 8601 defines alternative definitions for the first week, all of which produce equivalent results. Any of these definitions could be substituted so long as they are in relation to the reporting year start day.

651 This allows the values to be sorted chronologically using textual sorting
652 methods.

653

654 The meaning of a reporting year is always based on the start day of the year and
655 requires that the reporting year is expressed as the year at the start of the period.
656 This start day is always the same for a reporting year, and is expressed as a day and
657 a month (e.g. July 1). Therefore, the reporting year 2000 with a start day of July 1
658 begins on July 1, 2000.

659

660 A specialized attribute (reporting year start day) exists for the purpose of
661 communicating the reporting year start day. This attribute has a fixed identifier
662 (REPORTING_YEAR_START_DAY) and a fixed representation (xs:gMonthDay) so
663 that it can always be easily identified and processed in a data message. Although
664 this attribute exists in specialized sub-class, it functions the same as any other
665 attribute outside of its identification and representation. It must takes its identity from
666 a concept and state its relationship with other components of the data structure
667 definition. The ability to state this relationship allows this reporting year start day
668 attribute to exist at the appropriate levels of a data message. In the absence of this
669 attribute, the reporting year start date is assumed to be January 1; therefore if the
670 reporting year coincides with the calendar year, this Attribute is not necessary.

671

672 Since the duration and the reporting year start day are known for any reporting
673 period, it is possible to relate any reporting period to a distinct calendar period. The
674 actual Gregorian calendar period covered by the reporting period can be computed
675 as follows (based on the standard format of [REPROTING_YEAR]-
676 [PERIOD_INDICATOR][PERIOD_VALUE] and the reporting year start day as
677 [REPORTING_YEAR_START_DAY]):

678

679 **1. Determine [REPORTING_YEAR_BASE]:**

680 Combine [REPORTING_YEAR] of the reporting period value (YYYY) with
681 [REPORTING_YEAR_START_DAY] (MM-DD) to get a date (YYYY-MM-DD).

682

This is the [REPORTING_YEAR_START_DATE]

683

a) If the [PERIOD_INDICATOR] is W:

684

**1.If [REPORTING_YEAR_START_DATE] is a Friday, Saturday,
685 or Sunday:**

686

Add³ (P3D, P2D, or P1D respectively) to the

687

[REPORTING_YEAR_START_DATE]. The result is the
688 [REPORTING_YEAR_BASE].

689

**2.If [REPORTING_YEAR_START_DATE] is a Monday,
690 Tuesday, Wednesday, or Thursday:**

691

Add³ (P0D, -P1D, -P2D, or -P3D respectively) to the

692

[REPORTING_YEAR_START_DATE]. The result is the
693 [REPORTING_YEAR_BASE].

694

b) Else:

695

The [REPORTING_YEAR_START_DATE] is the
696 [REPORTING_YEAR_BASE].

697

2. Determine [PERIOD_DURATION]:

698

a) If the [PERIOD_INDICATOR] is A, the [PERIOD_DURATION] is P1Y.

699

b) If the [PERIOD_INDICATOR] is S, the [PERIOD_DURATION] is P6M.

700

c) If the [PERIOD_INDICATOR] is T, the [PERIOD_DURATION] is P4M.

701

d) If the [PERIOD_INDICATOR] is Q, the [PERIOD_DURATION] is P3M.

702

e) If the [PERIOD_INDICATOR] is M, the [PERIOD_DURATION] is P1M.

- 703 f) If the [PERIOD_INDICATOR] is W, the [PERIOD_DURATION] is P7D.
704 g) If the [PERIOD_INDICATOR] is D, the [PERIOD_DURATION] is P1D.

705 **3. Determine [PERIOD_START]:**

706 Subtract one from the [PERIOD_VALUE] and multiply this by the
707 [PERIOD_DURATION]. Add³ this to the [REPORTING_YEAR_BASE]. The
708 result is the [PERIOD_START].

709 **4. Determine the [PERIOD_END]:**

710 Multiply the [PERIOD_VALUE] by the [PERIOD_DURATION]. Add³ this to
711 the [REPORTING_YEAR_BASE] add³ -P1D. The result is the
712 [PERIOD_END].
713

714 For all of these ranges, the bounds include the beginning of the [PERIOD_START]
715 (i.e. 00:00:00) and the end of the [PERIOD_END] (i.e. 23:59:59).
716

717 **Examples:**

718
719 **2010-Q2, REPORTING_YEAR_START_DAY = --07-01 (July 1)**

- 720 1. [REPORTING_YEAR_START_DATE] = 2010-07-01
721 b) [REPORTING_YEAR_BASE] = 2010-07-01
722 2. [PERIOD_DURATION] = P3M
723 3. (2-1) * P3M = P3M
724 2010-07-01 + P3M = 2010-10-01
725 [PERIOD_START] = 2010-10-01
726 4. 2 * P3M = P6M
727 2010-07-01 + P6M = 2010-13-01 = 2011-01-01
728 2011-01-01 + -P1D = 2010-12-31
729 [PERIOD_END] = 2010-12-31
730

731 The actual calendar range covered by 2010-Q2 (assuming the reporting year
732 begins July 1) is 2010-10-01T00:00:00/2010-12-31T23:59:59
733

734 **2011-W36, REPORTING_YEAR_START_DAY = --07-01 (July 1)**

- 735 1. [REPORTING_YEAR_START_DATE] = 2010-07-01
736 a) 2011-07-01 = Friday
737 2011-07-01 + P3D = 2011-07-04
738 [REPORTING_YEAR_BASE] = 2011-07-04
739 2. [PERIOD_DURATION] = P7D
740 3. (36-1) * P7D = P245D
741 2011-07-04 + P245D = 2012-03-05
742 [PERIOD_START] = 2012-03-05
743 4. 36 * P7D = P252D
744 2011-07-04 + P252D = 2012-03-12
745 2012-03-12 + -P1D = 2012-03-11
746 [PERIOD_END] = 2012-03-11
747

³ The rules for adding durations to a date time are described in the W3C XML Schema specification. See <http://www.w3.org/TR/xmlschema-2/#adding-durations-to-dateTimes> for further details.

748 The actual calendar range covered by 2011-W36 (assuming the reporting year
 749 begins July 1) is 2012-03-05T00:00:00/2012-03-11T23:59:59
 750

751 **4.2.7 Distinct Range**

752 In the case that the reporting period does not fit into one of the prescribe periods
 753 above, a distinct time range can be used. The value of these ranges is based on the
 754 ISO 8601 time interval format of start/duration. Start can be expressed as either an
 755 ISO 8601 date or a date-time, and duration is expressed as an ISO 8601 duration.
 756 However, the duration can only be positive.
 757

758 **4.2.8 Time Format**

759 In version 2.0 of SDMX there is a recommendation to use the time format attribute to
 760 gives additional information on the way time is represented in the message.
 761 Following an appraisal of its usefulness this is no longer required. However, it is still
 762 possible, if required , to include the time format attribute in SDMX-ML.
 763

Code	Format
OTP	Observational Time Period: Superset of all SDMX time formats (Gregorian Time Period, Reporting Time Period, and Time Range)
STP	Standard Time Period: Superset of Gregorian and Reporting Time Periods
GTP	Superset of all Gregorian Time Periods and date-time
RTP	Superset of all Reporting Time Periods
TR	Time Range: Start time and duration (YYYY-MM-DD(Thh:mm:ss)?/<duration>)
GY	Gregorian Year (YYYY)
GTM	Gregorian Year Month (YYYY-MM)
GD	Gregorian Day (YYYY-MM-DD)
DT	Distinct Point: date-time (YYYY-MM-DDThh:mm:ss)
RY	Reporting Year (YYYY-A1)
RS	Reporting Semester (YYYY-Ss)
RT	Reporting Trimester (YYYY-Tt)
RQ	Reporting Quarter (YYYY-Qq)
RM	Reporting Month (YYYY-Mmm)

Code	Format
RW	Reporting Week (YYYY-Www)
RD	Reporting Day (YYYY-Dddd)

764

Table 1: SDMX-ML Time Format Codes

765

4.2.9 Transformation between SDMX-ML and SDMX-EDI

766

When converting SDMX-ML data structure definitions to SDMX-EDI data structure definitions, only the identifier of the time format attribute will be retained. The representation of the attribute will be converted from the SDMX-ML format to the fixed SDMX-EDI code list. If the SDMX-ML data structure definition does not define a time format attribute, then one will be automatically created with the identifier "TIME_FORMAT".

772

773

When converting SDMX-ML data to SDMX-EDI, the source time format attribute will be irrelevant. Since the SDMX-ML time representation types are not ambiguous, the target time format can be determined from the source time value directly. For example, if the SDMX-ML time is 2000-Q2 the SDMX-EDI format will always be 608/708 (depending on whether the target series contains one observation or a range of observations)

779

780

When converting a data structure definition originating in SDMX-EDI, the time format attribute should be ignored, as it serves no purpose in SDMX-ML.

781

782

When converting data from SDMX-EDI to SDMX-ML, the source time format is only necessary to determine the format of the target time value. For example, a source time format of 604 will result in a target time in the format YYYY-Ss whereas a source format of 608 will result in a target time value in the format YYYY-Qq.

783

784

785

786

4.2.10 Time Zones

787

In alignment with ISO 8601, SDMX allows the specification of a time zone on all time periods and on the reporting year start day. If a time zone is provided on a reporting year start day, then the same time zone (or none) should be reported for each reporting time period. If the reporting year start day and the reporting period time zone differ, the time zone of the reporting period will take precedence. Examples of each format with time zones are as follows (time zone indicated in bold):

788

789

790

791

792

793

794

- Time Range (start date): 2006-06-05-**05:00**/P5D

795

- Time Range (start date-time): 2006-06-05T00:00:00-**05:00**/P5D

796

- Gregorian Year: 2006-**05:00**

797

- Gregorian Month: 2006-06-**05:00**

798

- Gregorian Day: 2006-06-05-**05:00**

799

- Distinct Point: 2006-06-05T00:00:00-**05:00**

800

- Reporting Year: 2006-A1-**05:00**

801

- Reporting Semester: 2006-S2-**05:00**

802

- Reporting Trimester: 2006-T2-**05:00**

803

- Reporting Quarter: 2006-Q3-**05:00**

804

- Reporting Month: 2006-M06-**05:00**

805

- Reporting Week: 2006-W23-**05:00**

- 806 • Reporting Day: 2006-D156-05:00
807 • Reporting Year Start Day: --07-01-05:00

808 According to ISO 8601, a date without a time-zone is considered "local time". SDMX
809 assumes that local time is that of the sender of the message. In this version of
810 SDMX, an optional field is added to the sender definition in the header for specifying
811 a time zone. This field has a default value of 'Z' (UTC). This determination of local
812 time applies for all dates in a message.

813 **4.2.11 Representing Time Spans Elsewhere**

814 It has been possible since SDMX 2.0 for a Component to specify a representation of
815 a time span. Depending on the format of the data message, this resulted in either an
816 element with 2 XML attributes for holding the start time and the duration or two
817 separate XML attributes based on the underlying Component identifier. For example
818 if REF_PERIOD were given a representation of time span, then in the Compact data
819 format, it would be represented by two XML attributes; REF_PERIODStartTime
820 (holding the start) and REF_PERIOD (holding the duration). If a new simple type is
821 introduced in the SDMX schemas that can hold ISO 8601 time intervals, then this will
822 no longer be necessary. What was represented as this:

```
823                   <Series REF_PERIODStartTime="2000-01-01T00:00:00" REF_PERIOD="P2M"/>
```

825
826 can now be represented with this:

```
827                   <Series REF_PERIOD="2000-01-01T00:00:00/P2M"/>
```

829 **4.2.12 Notes on Formats**

830 There is no ambiguity in these formats so that for any given value of time, the
831 category of the period (and thus the intended time period range) is always clear. It
832 should also be noted that by utilizing the ISO 8601 format, and a format loosely
833 based on it for the report periods, the values of time can easily be sorted
834 chronologically without additional parsing.

835 **4.2.13 Effect on Time Ranges**

836 All SDMX-ML data messages are capable of functioning in a manner similar to
837 SDMX-EDI if the Dimension at the observation level is time: the time period for the
838 first observation can be stated and the rest of the observations can omit the time
839 value as it can be derived from the start time and the frequency. Since the frequency
840 can be determined based on the actual format of the time value for everything but
841 distinct points in time and time ranges, this makes it even simpler to process as the
842 interval between time ranges is known directly from the time value.

843

844 **4.2.14 Time in Query Messages**

845 When querying for time values, the value of a time parameter can be provided as any
846 of the Observational Time Period formats and must be paired with an operator. In
847 addition, an explicit value for the reporting year start day can be provided, or this can
848 be set to "Any". This section will detail how systems processing query messages
849 should interpret these parameters.

850

851 Fundamental to processing a time value parameter in a query message is
 852 understanding that all time periods should be handled as a distinct range of time.
 853 Since the time parameter in the query is paired with an operator, this is also
 854 effectively represents a distinct range of time. Therefore, a system processing the
 855 query must simply match the data where the time period for requested parameter is
 856 encompassed by the time period resulting from value of the query parameter. The
 857 following table details how the operators should be interpreted for any time period
 858 provided as a parameter.
 859

Operator	Rule
Greater Than	Any data after the last moment of the period
Less Than	Any data before the first moment of the period
Greater Than or Equal To	Any data on or after the first moment of the period
Less Than or Equal To	Any data on or before the last moment of the period
Equal To	Any data which falls on or after the first moment of the period and before or on the last moment of the period

860
 861 Reporting Time Periods as query parameters are handled based on whether the
 862 value of the reportingYearStartDay XML attribute is an explicit month and day or
 863 "Any":
 864

865 If the time parameter provides an explicit month and day value for the
 866 reportingYearStartDay XML attribute, then the parameter value is converted to
 867 a distinct range and processed as any other time period would be processed.
 868

869 If the reportingYearStartDay XML attribute has a value of "Any", then any data
 870 within the bounds of the reporting period for the year is matched, regardless of
 871 the actual start day of the reporting year. In addition, data reported against a
 872 normal calendar period is matched if it falls within the bounds of the time
 873 parameter based on a reporting year start day of January 1. When determining
 874 whether another reporting period falls within the bounds of a report period
 875 query parameter, one will have to take into account the actual time period to
 876 compare weeks and days to higher order report periods. This will be
 877 demonstrated in the examples to follow.
 878

879 Note that the reportingYearStartDay XML attribute on the time value parameter is
 880 only used to qualify a reporting period value for the given time value parameter. The
 881 usage of this is different than using the attribute value parameter for the actual
 882 reporting year start day attribute. In the case that the attribute value parameters is
 883 used for the reporting year start day data structure attribute, it will be treated as any
 884 other attribute value parameter; data will be filtered to that which matches the values
 885 specified for the given attribute. For example, if the attribute value parameter
 886 references the reporting year start day attribute and specifies a value of "--07-01",
 887 then only data which has this attribute with the value "--07-01" will be returned. In
 888 terms of processing any time value parameters, the value supplied in the attribute
 889 value parameter will be irrelevant.
 890

891 **Examples:**

892

893 **Gregorian Period**

894 Query Parameter: Greater than 2010

895 Literal Interpretation: Any data where the start period occurs after 2010-12-31T23:59:59.

896 Example Matches:

- 897
- 898 • 2011 or later
 - 899 • 2011-01 or later
 - 900 • 2011-01-01 or later
 - 901 • 2011-01-01/P[Any Duration] or any later start date
 - 902 • 2011-[Any reporting period] (any reporting year start day)
 - 903 • 2010-S2 (reporting year start day --07-01 or later)
 - 904 • 2010-T3 (reporting year start day --07-01 or later)
 - 905 • 2010-Q3 or later (reporting year start day --07-01 or later)
 - 906 • 2010-M07 or later (reporting year start day --07-01 or later)
 - 907 • 2010-W28 or later (reporting year start day --07-01 or later)
 - 908 • 2010-D185 or later (reporting year start day --07-01 or later)

909

910 **Reporting Period with explicit start day**

911 Query Parameter: Greater than or equal to 2009-Q3, reporting year start day = "--07-01"

912

913 Literal Interpretation: Any data where the start period occurs on after 2010-01-01T00:00:00 (Note that in this case 2009-Q3 is converted to the explicit date range of 2010-01-01/2010-03-31 because of the reporting year start day value).

914 Example Matches: Same as previous example

915

916

917

918 **Reporting Period with "Any" start day**

919 Query Parameter: Greater than or equal to 2010-Q3, reporting year start day = "Any"

920

921 Literal Interpretation: Any data with a reporting period where the start period is on or after the start period of 2010-Q3 for the same reporting year start day, or and data where the start period is on or after 2010-07-01.

922 Example Matches:

- 923
- 924 • 2011 or later
 - 925 • 2010-07 or later
 - 926 • 2010-07-01 or later
 - 927 • 2010-07-01/P[Any Duration] or any later start date
 - 928 • 2011-[Any reporting period] (any reporting year start day)
 - 929 • 2010-S2 (any reporting year start day)
 - 930 • 2010-T3 (any reporting year start day)
 - 931 • 2010-Q3 or later (any reporting year start day)
 - 932 • 2010-M07 or later (any reporting year start day)
 - 933 • 2010-W27 or later (reporting year start day --01-01)⁴
 - 934 • 2010-D182 or later (reporting year start day --01-01)
 - 935 • 2010-W28 or later (reporting year start day --07-01)⁵
- 936

⁴ 2010-Q3 (with a reporting year start day of --01-01) starts on 2010-07-01. This is day 4 of week 26, therefore the first week matched is week 27.

- 937
- 2010-D185 or later (reporting year start day --07-01)

938 **4.3 Structural Metadata Querying Best Practices**

939 When querying for structural metadata, the ability to state how references should be
940 resolved is quite powerful. However, this mechanism is not always necessary and
941 can create an undue burden on the systems processing the queries if it is not used
942 properly.

943
944 Any structural metadata object which contains a reference to an object can be
945 queried based on that reference. For example, a categorisation references both a
946 category and the object it is categorising. As this is the case, one can query for
947 categorisations which categorise a particular object or which categorise against a
948 particular category or category scheme. This mechanism should be used when the
949 referenced object is known.

950
951 When the referenced object is not known, then the reference resolution mechanism
952 could be used. For example, suppose one wanted to find all category schemes and
953 the related categorisations for a given maintenance agency. In this case, one could
954 query for the category scheme by the maintenance agency and specify that parent
955 and sibling references should be resolved. This would result in the categorisations
956 which reference the categories in the matched schemes to be returned, as well as
957 the object which they categorise.

958 **4.4 Versioning and External Referencing**

959 Within the SDMX-ML Structure Message, there is a pattern for versioning and
960 external referencing which should be pointed out. The identifiers are qualified by their
961 version numbers – that is, an object with an Agency of “A”, and ID of “X” and a
962 version of “1.0” is a different object than one with an Agency of “A”, an ID of “X”, and
963 a version of “1.1”.

964
965 The production versions of identifiable objects/resources are assumed to be static –
966 that is, they have their isFinal attribute set to “true”. Once in production, an object
967 cannot change in any way, or it must be versioned. For cases where an object is not
968 static, the isFinal attribute must have a value of “false”, but non-final objects should
969 not be used outside of a specific system designed to accommodate them. For most
970 purposes, all objects should be declared final before use in production.

971
972 This mechanism is an “early binding” one – everything with a versioned identity is a
973 known quantity, and will not change. It is worth pointing out that in some cases
974 relationships are essentially one-way references: an illustrative case is that of
975 Categories. While a Category may be referenced by many dataflows and metadata
976 flows, the addition of more references from flow objects does not version the
977 Category. This is because the flows are not properties of the Categories – they
978 merely make references to it. If the name of a Category changed, or its sub-
979 Categories changed, then versioning would be necessary.

980

⁵ 2010-Q3 (with a reporting year start day of --07-01) starts on 2011-01-01. This is day 6 of week 27, therefore the first week matched is week 28.

981 Versioning operates at the level of versionable and maintainable objects in the SDMX
982 information model. If any of the children of objects at these levels change, then the
983 objects themselves are versioned.

984

985 One area which is much impacted by this versioning scheme is the ability to
986 reference external objects. With the many dependencies within the various structural
987 objects in SDMX, it is useful to have a scheme for external referencing. This is done
988 at the level of maintainable objects (DSDs, code lists, concept schemes, etc.) In an
989 SDMX-ML Structure Message, whenever an “isExternalReference” attribute is set to
990 true, then the application must resolve the address provided in the associated “uri”
991 attribute and use the SDMX-ML Structure Message stored at that location for the full
992 definition of the object in question. Alternately, if a registry “urn” attribute has been
993 provided, the registry can be used to supply the full details of the object.

994

995 Because the version number is part of the identifier for an object, versions are a
996 necessary part of determining that a given resource is the one which was called for. It
997 should be noted that whenever a version number is not supplied, it is assumed to be
998 “1.0”. (The “x.x” versioning notation is conventional in practice with SDMX, but not
999 required.)

1000 **5 Metadata Structure Definition (MSD)**

1001 **5.1 Scope**

1002 The scope of the MSD is enhanced in this version to better support the types of
1003 construct to which metadata can be attached. In particular it is possible to specify an
1004 attachment to any key or partial key of a data set. This is particularly useful for web
1005 dissemination where metadata may be present for the data, but is not stored with the
1006 data but is related to it. For this use case to be supported it is necessary to be able to
1007 specify in the MSD that metadata is attached to a key or partial key, and the actual
1008 key or partial key to be identified in the Metadata Set.

1009

1010 In addition to the increase in the scope of objects that can be included in an MSD,
1011 the way the identifier mechanism works in this version, and the terminology used, is
1012 much simpler.

1013

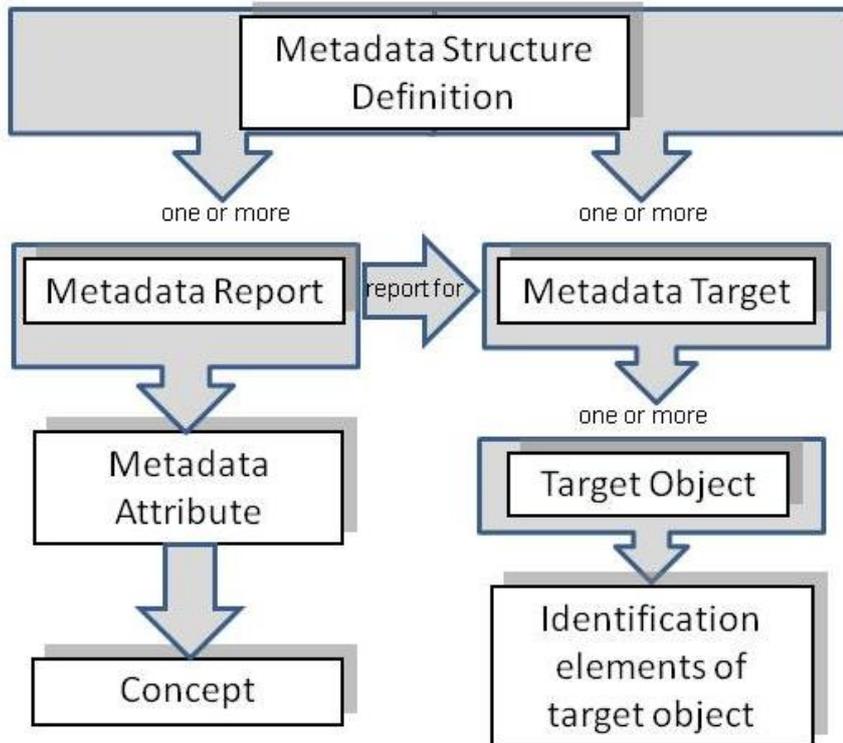
1014 **5.2 Identification of the Object Type to which the Metadata is** 1015 **to be Attached**

1016 The following example shows the structure and naming of the MSD components for
1017 the use case of defining full and partial keys.

1018

1019 The schematic structure of an MSD is shown below.

1020



1021
1022
1023

Figure 1: Schematic of the Metadata Structure Definition

1024 The MSD comprises the specification of the object types to which metadata can be
 1025 reported in a Metadata Set (Metadata Target(s)), and the Report Structure(s)
 1026 comprising the Metadata Attributes that identify the Concept for which metadata may
 1027 be reported in the Metadata Set. Importantly, one Report Structure references the
 1028 Metadata Target for which it is relevant. One Report Structure can reference many
 1029 Metadata Target i.e. the same Report Structure can be used for different target
 1030 objects.

```

1 | <structure:MetadataStructureDefinition id="WEBMETADATA" agencyID="WEBMASTER">
2 |   <!--this enables metadata to be attached to Dimensions, Keys, Partial Keys , and Observations
3 |   relating to data structured according to any DSD -->
4 |   <common.Name xml:lang="en">Web Metadata</common.Name>
5 |   <structure:MetadataTarget id="DATA_KEY_TARGET">
6 |     <structure:DataStructureTarget>
7 |       <structure:ObjectReference/>
8 |     </structure:DataStructureTarget>
9 |     <structure:KeyDescriptorValuesTarget>
10 |      <structure>DataKey/>
11 |    </structure:KeyDescriptorValuesTarget>
12 |  </structure:MetadataTarget>
13 |  <structure:MetadataTarget id="DATASET_TARGET">
14 |    <structure:DataSetTarget>
15 |      <structure:ObjectReference idOnly="true"/>
16 |    </structure:DataSetTarget>
17 |  </structure:MetadataTarget>
18 |  <structure:ReportStructure id="METADATA_REPORT">
19 |    <common.Name xml:lang="en">Metadata Report</common.Name>
20 |    <structure:MetadataTargets>
21 |      <structure:MetadataTargetRef id="DATA_KEY_TARGET"/>
22 |      <structure:MetadataTargetRef id="DATASET_TARGET"/>
23 |    </structure:MetadataTargets>
24 |    <structure:MetadataAttributes>
25 |  </structure:ReportStructure>
26 | </structure:MetadataStructureDefinition>
  
```

Annotations in the code block:

- Target object is Data Structure Definition identified by an Object Reference (points to line 7)
- Target object is Series Key identified by a Data Key (points to line 10)
- Target object is Dataset identified by an Object Reference by its id (points to line 15)
- Metadata Targets for which the Report is valid (points to lines 21-22)

1031
1032

Figure 2: Example MSD showing Metadata Targets

1033 Note that the SDMX-ML schemas have explicit XML elements for each identifiable
 1034 object type because identifying, for instance, a Maintainable Object has different
 1035 properties from an Identifiable Object which must also include the agencyId, version,
 1036 and id of the Maintainable Object in which it resides.

1037 **5.3 Report Structure**

1038 An example is shown below.

```

<structure:MetadataStructureDefinition id="WEBMETADATA" agencyID="WEBMASTER">
  <!--this enables metadata to be attached to Dimensions, Keys, Partial Keys , and Observations
  relating to data structured according to any DSD -->
  <common:Name xml:lang="en">Web Metadata</common:Name>
  <structure:MetadataTarget id="DATA_KEY_TARGET">
  <structure:MetadataTarget id="DATASET_TARGET">
  <structure:ReportStructure id="METADATA_REPORT">
    <common:Name xml:lang="en">Metadata Report</common:Name>
    <structure:MetadataTargets>
      <structure:MetadataTargetRef id="DATA_KEY_TARGET"/>
      <structure:MetadataTargetRef id="DATASET_TARGET"/>
    </structure:MetadataTargets>
    <structure:MetadataAttributes>
      <structure:MetadataAttribute isPresentational="true">
        <structure:ConceptReference>
          <common:ConceptSchemeRef>
            <common:Ref id="IMETADATA_CONCEPTS" agencyID="WEBMASTER" version="1.0"/>
          </common:ConceptSchemeRef>
          <common:ConceptRef id="SOURCE"/>
        </structure:ConceptReference>
        <structure:MetadataAttribute>
          <structure:ConceptReference>
            <common:ConceptSchemeRef>
              <common:Ref id="IMETADATA_CONCEPTS" agencyID="WEBMASTER" version="1.0"/>
            </common:ConceptSchemeRef>
            <common:ConceptRef id="SOURCE_TYPE"/>
          </structure:ConceptReference>
        </structure:MetadataAttribute>
        <structure:MetadataAttribute>
          <structure:ConceptReference>
            <common:ConceptSchemeRef>
              <common:Ref id="METADATA_CONCEPTS" agencyID="WEBMASTER" version="1.0"/>
            </common:ConceptSchemeRef>
            <common:ConceptRef id="COLLECTION_SOURCE_NAME"/>
          </structure:ConceptReference>
        </structure:MetadataAttribute>
        <structure:MetadataAttribute>
          <!-- and so on for the remaining metadata attribute -->
        </structure:MetadataAttribute>
      </structure:MetadataAttributes>
    </structure:ReportStructure>
  </structure:MetadataStructureDefinition>
  
```

1039

1040 **Figure 3: Example MSD showing specification of three Metadata Attributes**

1041 This example shows the following hierarchy of Metadata Attributes:

1042 Source – this is presentational and no metadata is expected to be reported at this
 1043 level

- 1044 ○ Source Type
- 1045 ○ Collection Source Name

1046 **5.4 Metadata Set**

1047 An example of reporting metadata according to the MSD described above, is shown
1048 below.

```
1049 ] <g:MetadataSet>
  <c:MetadataStructureDefinitionReference>
    <c:Ref id="WEBMETADATA" agencyID="WEBMASTER" version="1.0"/>
  </c:MetadataStructureDefinitionReference>
  <g:AttributeValueSet>
    <g:ReportRef>METADATA_REPORT</g:ReportRef>
    <!-- This is a partial key report (combination of codes from different dimensions) -->
    <g:TargetValues>
      <g:MetadataTargetValue id="DATA_KEY_TARGET">
        <g:ReferenceValue>
          <c:DataStructureReference>
            <c:Ref id="FINANCE_DSD" agencyID="WEBMASTER" version="1.0"/>
          </c:DataStructureReference>
          </g:ReferenceValue>
          <g:ReferenceValue>
            <c>DataKey>
              <c>DataKeyValue dimensionID="ECONOMICCONCEPT">
                <c:DimensionValue>ENDA</c:DimensionValue>
              </c>DataKeyValue>
              <c>DataKeyValue dimensionID="DATASOURCE">
                <c:DimensionValue>IFS</c:DimensionValue>
              </c>DataKeyValue>
            </c>DataKey>
          </g:ReferenceValue>
        </g:MetadataTargetValue>
      </g:TargetValues>
      <g:ReportedAttribute id="SOURCE">
        <g:ReportedAttribute id="SOURCE_TYPE">
          <g:Value>Market Values</g:Value>
        </g:ReportedAttribute>
        <g:ReportedAttribute id="COLLECTION_SOURCE_NAME">
          <g:Value>These series are typically the monthly average of market rates or official rates of the reporting countr
are not available, they are the monthly average rates in New York. Or if the latter are not available, they are estimates basec
averages of the end-of-month market rates quoted in the reporting country.
        </g:Value>
        </g:ReportedAttribute>
      </g:ReportedAttribute>
    </g:AttributeValueSet>
  </g:MetadataSet>
```

1051 **Figure 4: Example Metadata Set**

1052 This example shows:

- 1053 1. The reference to the MSD, Metadata Report, and Metadata Target
1054 (MetadataTargetValue)
- 1055 2. The reported metadata attributes (AttributeValueSet)

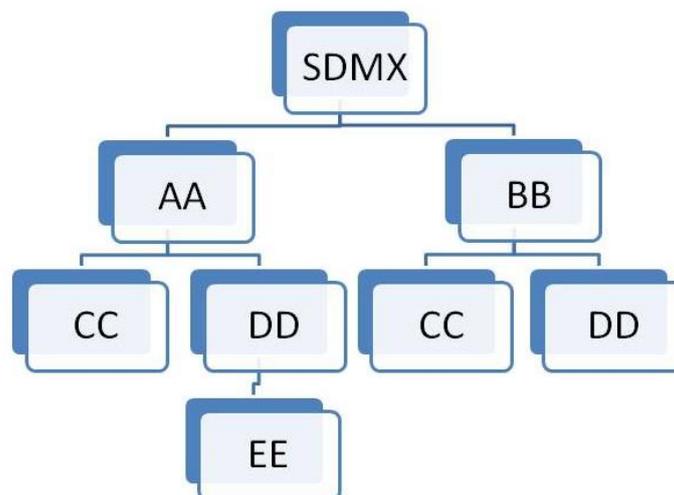
1056 6 Maintenance Agencies

1057 All structural metadata in SDMX is owned and maintained by a maintenance agency
 1058 (Agency identified by `agencyID` in the schemas). It is vital to the integrity of the
 1059 structural metadata that there are no conflicts in `agencyID`. In order to achieve this
 1060 SDMX adopts the following rules:

- 1061
- 1062 1. Agencies are maintained in an Agency Scheme (which is a sub class of
 1063 Organisation Scheme)
- 1064 2. The maintenance agency of the Agency Scheme must also be declared in a
 1065 (different) Agency Scheme.
- 1066 3. The “top-level” agency is SDMX and this agency scheme is maintained by
 1067 SDMX.
- 1068 4. Agencies registered in the top-level scheme can themselves maintain a single
 1069 Agency Scheme. SDMX is an agency in the SDMX agency scheme. Agencies
 1070 in this scheme can themselves maintain a single Agency Scheme and so on.
- 1071 5. The `AgencyScheme` cannot be versioned and so take a default version
 1072 number of 1.0 and cannot be made “final”.
- 1073 6. There can be only one `AgencyScheme` maintained by any one Agency. It has
 1074 a fixed Id of `AgencyScheme`.
- 1075 7. The format of the agency identifier is `agencyId.agencyID` etc. The top-level
 1076 agency in this identification mechanism is the agency registered in the SDMX
 1077 agency scheme. In other words, SDMX is not a part of the hierarchical ID
 1078 structure for agencies. SDMX is, itself, a maintenance agency.

1079
 1080 This supports a hierarchical structure of `agencyID`.

1081
 1082 An example is shown below.
 1083



1084
 1085 **Figure 5: Example of Hierarchic Structure of Agencies**

1086 Each agency is identified by its full hierarchy excluding SDMX.

1087
 1088 The XML representing this structure is shown below.
 1089

```

<structure:Organisations>
  <structure:AgencyScheme agencyID="SDMX" id="AGENCY_SCHEME">
    <common:Name>name</common:Name>
    <structure:Agency id="AA">
      <common:Name>AA Name</common:Name>
    </structure:Agency>
    <structure:Agency id="BB">
      <common:Name>BB Name</common:Name>
    </structure:Agency>
  </structure:AgencyScheme>
  <structure:AgencyScheme agencyID="AA" id="AGENCY_SCHEME">
    <common:Name>name</common:Name>
    <structure:Agency id="CC">
      <common:Name>CC Name</common:Name>
    </structure:Agency>
    <structure:Agency id="DD">
      <common:Name>DD Name</common:Name>
    </structure:Agency>
  </structure:AgencyScheme>
  <structure:AgencyScheme agencyID="BB" id="AGENCY_SCHEME">
    <common:Name>name</common:Name>
    <structure:Agency id="CC">
      <common:Name>CC Name</common:Name>
    </structure:Agency>
    <structure:Agency id="DD">
      <common:Name>DD Name</common:Name>
    </structure:Agency >
  </structure:AgencyScheme>
  <structure:AgencyScheme agencyID="AA.CC" id="AGENCY_SCHEME">
    <common:Name>name</common:Name>
    <structure:Agency id="EE">
      <common:Name>EE Name</common:Name>
    </structure:Agency >
  </structure:AgencyScheme>
</structure:Organisations>

```

1090
1091

Figure 6: Example Agency Schemes Showing a Hierarchy

1092 Example of Structure Definitions:
1093

```

<structure:Codelists>
  <structure:Codelist id="CL_BOP" agencyID="SDMX" version="1.0"
    urn="urn:sdmx:org.sdmx.infomodel.codelist.Codelist=SDMX:CL_BOP[1.0]">
    <common:Name>name</common:Name>
  </structure:Codelist>
  <structure:Codelist id="CL_BOP" agencyID="AA" version="1.0"
    urn="urn:sdmx:org.sdmx.infomodel.codelist.Codelist=AA:CL_BOP[1.0]" >
    <common:Name>name</common:Name>
  </structure:Codelist>
  <structure:Codelist id="CL_BOP" agencyID="AA.CC" version="1.0"
    urn="urn:sdmx:org.sdmx.infomodel.codelist.Codelist=AA.CC:CL_BOP[1.0]" >
    <common:Name>name</common:Name>
  </structure:Codelist>
  <structure:Codelist id="CL_BOP" agencyID="BB.CC" version="1.0"
    urn="urn:sdmx:org.sdmx.infomodel.codelist.Codelist=BB.CC:CL_BOP[1.0]">
    <common:Name>name</common:Name>
  </structure:Codelist>
</structure:Codelists>

```

1094
1095

Figure 7: Example Showing Use of Agency Identifiers

1096
1097 Each of these maintenance agencies has an identical Codelist with the Id CL_BOP.
1098 However, each is uniquely identified by means of the hierarchic agency structure.

1099 7 Concept Roles

1100 7.1 Overview

1101 The DSD Components of Dimension and Attribute can play a specific role in the DSD
1102 and it is important to some applications that this role is specified. For instance, the
1103 following roles are some examples:

1104
1105 **Frequency** – in a data set the content of this Component contains information on the
1106 frequency of the observation values

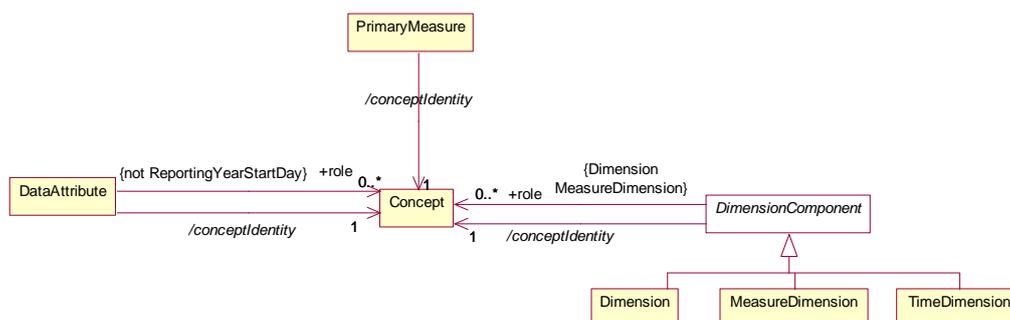
1107 **Geography** - in a data set the content of this Component contains information on the
1108 geographic location of the observation values

1109 **Unit of Measure** - in a data set the content of this Component contains information
1110 on the unit of measure of the observation values

1111
1112 In order for these roles to be extensible and also to enable user communities to
1113 maintain community-specific roles, the roles are maintained in a controlled
1114 vocabulary which is implemented in SDMX as Concepts in a Concept Scheme. The
1115 Component optionally references this Concept if it is required to declare the role
1116 explicitly. Note that a Component can play more than one role and therefore multiple
1117 “role” concepts can be referenced.

1118 7.2 Information Model

1119 The Information Model for this is shown below:
1120



1121
1122

Figure 8: Information Model Extract for Concept Role

1123 It is possible to specify zero or more concept roles for a Dimension, Measure
1124 Dimension and Data Attribute (but not the ReportingYearStartDay). The Time
1125 Dimension, Primary Measure, and the Attribute ReportingYearStartDay have
1126 explicitly defined roles and cannot be further specified with additional concept roles.

1127 7.3 Technical Mechanism

1128 The mechanism for maintain and using concept roles is as follows:
1129

- 1130 1. Any recognized Agency can have a concept scheme that contains concepts
 1131 that identify concept roles. Indeed, from a technical perspective any agency
 1132 can have more than one of these schemes, though this is not recommended.
 1133
- 1134 2. The concept scheme that contains the “role” concepts can contain concepts
 1135 that do not play a role.
 1136
- 1137 3. There is no explicit indication on the Concept whether it is a “role” concept.
 1138
- 1139 4. Therefore, any concept in any concept scheme is capable of being a “role”
 1140 concept.
 1141
- 1142 5. It is the responsibility of Agencies to ensure their community knows which
 1143 concepts in which concept schemes play a “role” and the significance and
 1144 interpretation of this role. In other words, such concepts must be known by
 1145 applications, there is no technical mechanism that can inform an application
 1146 on how to process such a “role”.
 1147
- 1148 6. If the concept referenced in the Concept Identity in a DSD component
 1149 (Dimension, Measure Dimension, Attribute) is contained in the concept
 1150 scheme containing concept roles then the DSD component could play the role
 1151 implied by the concept, if this is understood by the processing application.
 1152
- 1153 7. If the concept referenced in the Concept Identity in a DSD component
 1154 (Dimension, Measure Dimension, Attribute) is not contained in the concept
 1155 scheme containing concept roles, and the DSD component is playing a role,
 1156 then the concept role is identified by the Concept Role in the schema.
 1157

1158 **7.4 SDMX-ML Examples in a DSD**

1159
 1160 The Cross-Domain Concept Scheme maintained by SDMX contains concept role
 1161 concepts (FREQ chosen as an example).

```

1162 <structure:Dimension id="FREQ">
  <structure:ConceptIdentity>
    | <URN>
    urn:sdmx:org.sdmx.infomodel.conceptscheme.Concept=SDMX:CROSS_DOMAIN_CONCEPTS[1.0].FREQ
    </URN>
  </structure:ConceptIdentity>
</structure:Dimension>

```

1163

1164 Whether this is a role or not depends upon the application understanding that FREQ
 1165 in the Cross-Domain Concept Scheme is a role of Frequency.

1166 Using a Concept Scheme that is not the Cross-Domain Concept Scheme where it is
 1167 required to assign a role using the Cross-Domain Concept Scheme. Again FREQ is
 1168 chosen as the example.

```

1169 ]<structure:Dimension id="FREQ">
1170   > <structure:ConceptIdentity>
1171     > <URN>
1172       urn:sdmx:org.sdmx.infomodel.conceptscheme.Concept=JBG:MY_CONCEPTS[1.0].FREQ
1173     </URN>
1174   </structure:ConceptIdentity>
1175   > <structure:ConceptRole>
1176     > <URN>
1177       urn:sdmx:org.sdmx.infomodel.conceptscheme.Concept=SDMX:CROSS_DOMAIN_CONCEPTS[1.0].FREQ
1178     </URN>
1179   </structure: ConceptRole >
1180 </structure:Dimension>

```

1171 This explicitly states that this Dimension is playing a role identified by the FREQ
 1172 concept in the Cross-Domain Concept Scheme. Again the application needs to
 1173 understand what FREQ in the Cross-Domain Concept Scheme implies in terms of a
 1174 role.

1175 This is all that is required for interoperability within a community. The important point
 1176 is that a community must recognise a specific Agency as having the authority to
 1177 define concept roles and to maintain these "role" concepts in a concept scheme
 1178 together with documentation on the meaning of the role and any relevant processing
 1179 implications. This will then ensure there is interoperability between systems that
 1180 understand the use of these concepts.

1181 Note that each of the Components (Data Attribute, Primary Measure, Dimension,
 1182 Measure Dimension, Time Dimension) has a mandatory identity association
 1183 (Concept Identity) and if this Concept also identifies the role then it is possible to
 1184 state this by
 1185
 1186

1187 **7.5 SDMX Cross Domain Concept Scheme**

1188 All concepts in the SDMX Cross Domain Concept Scheme are capable of playing a
 1189 role and this scheme will contain all of the roles that were allowed at version 2.0 and
 1190 will be maintained with new roles that are agreed at the level of the community using
 1191 the Cross Domain Concept Scheme.

1192 The table below lists the Concepts that need to be in this scheme either for
 1193 compatibility with version 2.0 or because of requests for additional roles at version
 1194 2.1 which have been accepted.

1195 Note that each of the Components (Data Attribute, Primary Measure, Dimension,
 1196 Measure Dimension, Time Dimension) has a mandatory identity association
 1197 (Concept Identity) and if this Concept also identifies the role then it is possible to
 1198 state this by means of the `isRole` attribute (`isRole=true`) Additional roles can still
 1199 be specified by means of the `+role` association to additional Concepts that identify
 1200 the role.
 1201
 1202

1203 **8 Constraints**

1204 **8.1 Introduction**

1205 In this version of SDMX the Constraints is a Maintainable Artefact can be associated
1206 to one or more of:

- 1207
- 1208 • Data Structure Definition
- 1209 • Metadata Structure Definition
- 1210 • Dataflow
- 1211 • Metadataflow
- 1212 • Provision Agreement
- 1213 • Data Provider (this is restricted to a Release Calendar Constraint)
- 1214 • Simple or Queryable Datasources
- 1215

1216 Note that regardless of the artifact to which the Constraint is associated, it is
1217 constraining the contents of code lists in the DSD to which the constrained object is
1218 related. This does not apply, of course, to a Data Provider as the Data Provider can
1219 be associated, via the Provision Agreement, to many DSDs. Hence the reason for
1220 the restriction on the type of Constraint that can be attached to a Data Provider.

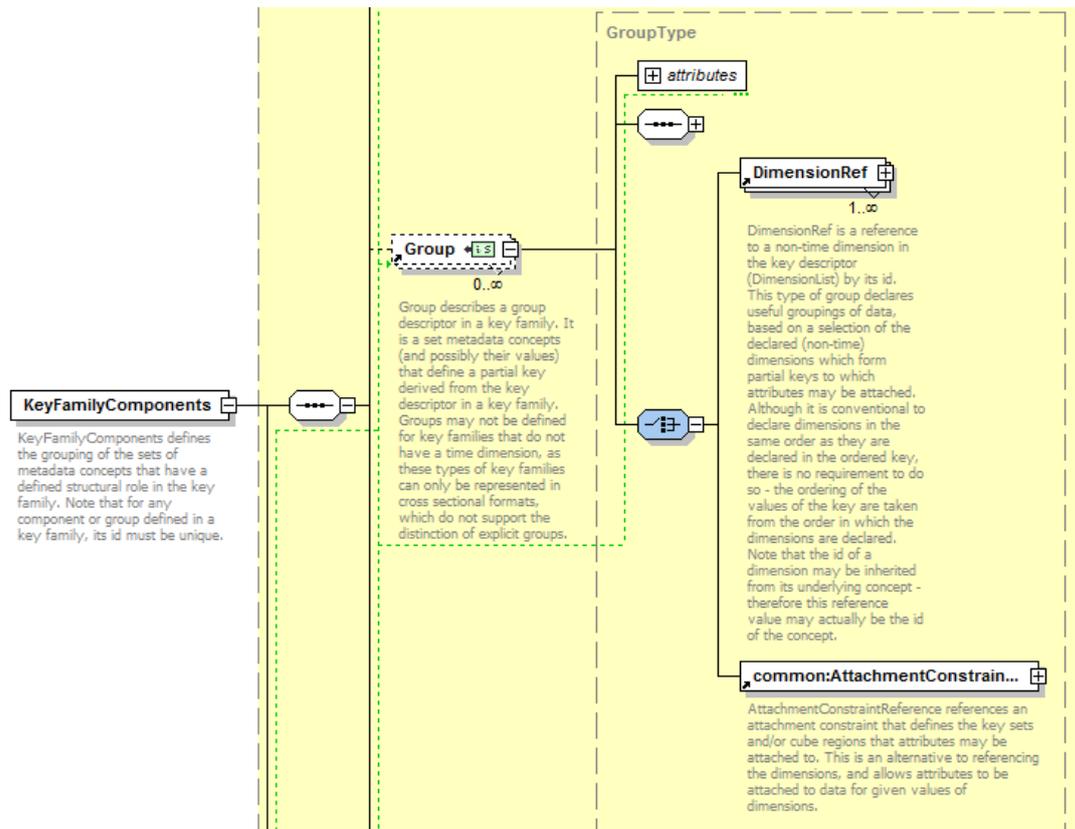
1221 **8.2 Types of Constraint**

1222 The Constraint can be of one of two types:

- 1223
- 1224 • Content constraint
- 1225 • Attachable constraint
- 1226

1227 The attachable constraint is used to define “cube slices” which identify sub sets of
1228 data in terms of series keys or dimension values. The purpose of this is to enable
1229 metadata to be attached to the constraint, and thereby to the cube slices defined in
1230 the Constraint. The metadata can be attached via the “reference metadata”
1231 mechanism – MSD and Metadata Set – or via a Group in the DSD. Below is snippet
1232 of the schema for a DSD that shows the constructs that enable the Constraint to
1233 referenced from a Group in a DSD.

1234



1235
1236

Figure 9: Extract from the SDMX-ML Schema showing reference to Attachment Constraint

1237 For the Content Constraint specific “inheritance” rules apply and these are detailed
1238 below.

1241 **8.3 Rules for a Content Constraint**

1242 **8.3.1 Scope of a Content Constraint**

1243 A Content Constraint is used specify the content of a data or metadata source in
1244 terms of the component values or the keys.

1245

1246 In terms of data the components are:

1247

- 1248 • Dimension
- 1249 • Measure Dimension
- 1250 • Time Dimension
- 1251 • Data Attribute
- 1252 • Primary Measure

1253

1254 And the keys are the content of the KeyDescriptor – i.e. the series keys composed,
1255 for each key, by a value for each Dimension and Measure Dimension

1256

1257 In terms of reference metadata the components are:

1258

- 1259 • Target Object which is one of:

- 1260 ○ Key Descriptor Values
- 1261 ○ Data Set
- 1262 ○ Report Period
- 1263 ○ IdentifiableObject

1264

- 1265 ● Metadata Attribute

1266

1267 The “key” is therefore the combination of the Target Objects that are defined for the
1268 Metadata Target.

1269

1270 For a Constraint based on a DSD the Content Constraint can reference one or more
1271 of:

1272

- 1273 ● Data Structure Definition
- 1274 ● Dataflow
- 1275 ● Provision Agreement

1276

1277 For a Constraint based on an MSD the Content Constraint can reference one or
1278 more of:

1279

- 1280 ● Metadata Structure Definition
- 1281 ● Metadataflow
- 1282 ● Provision Agreement

1283

1284 Furthermore, there can be more than one Content Constraint specified for a specific
1285 object e.g. more than one Constraint for a specific DSD.

1286

1287 In view of the flexibility of constraints attachment, clear rules on their usage are
1288 required. These are elaborated below.

1289 **8.3.2 Multiple Content Constraints**

1290 There can be many Content Constraints for any Constraining Artefact (e.g. DSD),
1291 subject to the following restrictions:

1292 **8.3.2.1 Cube Region**

- 1293 1. The constraint can contain multiple Member Selections (e.g. Dimension) but:
- 1294 2. A specific Member Selection (e.g. Dimension FREQ) can only be contained in
1295 one Content Constraint for any one attached object (e.g. a specific DSD or
1296 specific Dataflow)

1297 **8.3.2.2 Key Set**

1298 Key Sets will be processed in the order they appear in the Constraint and wildcards
1299 can be used (e.g. any key position not reference explicitly is deemed to be “all
1300 values”). As the Key Sets can be “included” or “excluded” it is recommended that Key
1301 Sets with wildcards are declared before KeySets with specific series keys. This will
1302 minimize the risk that keys are inadvertently included or excluded.

1303 **8.3.3 Inheritance of a Content Constraint**

1304 **8.3.3.1 Attachment levels of a Content Constraint**

1305 There are three levels of constraint attachment for which these inheritance rules
1306 apply:

- 1307 • DSD/MSD – top level
1308 ○ Dataflow/Metadataflow – second level
1309 ▪ Provision Agreement – third level
1310

1311 Note that these rules do not apply to the Simple Datasource or Queryable
1312 Datasource: the Content Constraint(s) attached to these artefacts are resolved for
1313 this artefact only and do not take into account Constraints attached to other artefacts
1314 (e.g. Provision Agreement, Dataflow, DSD).

1315 It is not necessary for a Content Constraint to be attached to higher level artifact. e.g.
1316 it is valid to have a Content Constraint for a Provision Agreement where there are no
1317 constraints attached the relevant dataflow or DSD.

1318 **8.3.3.2 Cascade rules for processing Constraints**

1319 The processing of the constraints on either Dataflow/Metadataflow or Provision
1320 Agreement must take into account the constraints declared at higher levels. The
1321 rules for the lower level constraints (attached to Dataflow/ Metadataflow and
1322 Provision Agreement) are detailed below.

1323 Note that there can be a situation where a constraint is specified at a lower level
1324 before a constraint is specified at a higher level. Therefore, it is possible that a higher
1325 level constraint makes a lower level constraint invalid. SDMX makes no rules on how
1326 such a conflict should be handled when processing the constraint for attachment.
1327 However, the cascade rules on evaluating constraints for usage are clear - the higher
1328 level constraint takes precedence in any conflicts that result in a less restrictive
1329 specification at the lower level.

1330 **8.3.3.3 Cube Region**

- 1331 1. It is not necessary to have a constraint on the higher level artifact (e.g. DSD
1332 referenced by the Dataflow) but if there is such a constraint at the higher
1333 level(s) then:
1334 a. The lower level constraint cannot be less restrictive than the constraint
1335 specified for the same Member Selection (e.g. Dimension) at the next
1336 higher level which constraints that Member Selection (e.g. if the
1337 Dimension FREQ is constrained to A, Q in a DSD then the constraint
1338 at the Dataflow or Provision Agreement cannot be A, Q, M or even just
1339 M – it can only further constrain A,Q).
1340 b. The constraint at the lower level for any one Member Selection further
1341 constrains the content for the same Member Selection at the higher
1342 level(s).
1343 2. Any Member Selection which is not referenced in a Content Constraint is
1344 deemed to be constrained according to the Content Constraint specified at
1345 the next higher level which constraints that Member Selection.

1346 3. If there is a conflict when resolving the constraint in terms of a lower-level
 1347 constraint being less restrictive than a higher-level constraint then the
 1348 constraint at the higher-level is used.
 1349

1350 Note that it is possible for a Content Constraint at a higher level to constrain, say,
 1351 four Dimensions in a single constraint, and a Content Constraint at a lower level to
 1352 constrain the same four in two, three, or four Content Constraints.

1353 8.3.3.4 Key Set

1354 1. It is not necessary to have a constraint on the higher level artefact (e.g. DSD
 1355 referenced by the Dataflow) but if there is such a constraint at the higher
 1356 level(s) then:
 1357
 1358 a. The lower level constraint cannot be less restrictive than the constraint
 1359 specified at the higher level.
 1360 b. The constraint at the lower level for any one Member Selection further
 1361 constrains the keys specified at the higher level(s).
 1362 2. Any Member Selection which is not referenced in a Content Constraint is
 1363 deemed to be constrained according to the Content Constraint specified at
 1364 the next higher level which constraints that Member Selection.
 1365 3. If there is a conflict when resolving the keys in the constraint at two levels, in
 1366 terms of a lower-level constraint being less restrictive than a higher-level
 1367 constraint, then the offending keys specified at the lower level are not
 1368 deemed part of the constraint.
 1369

1370 Note that a Key in a Key Set can have wildcarded Components. For instance the
 1371 constraint may simply constrain the Dimension FREQ to "A", and all keys where the
 1372 FREQ=A are therefore valid.
 1373

1374 The following logic explains how the inheritance mechanism works. Note that this is
 1375 conceptual logic and actual systems may differ in the way this is implemented.
 1376

1377 1. Determine all possible keys that are valid at the higher level.
 1378 2. These keys are deemed to be inherited by the lower level constrained object,
 1379 subject to the constraints specified at the lower level.
 1380 3. Determine all possible keys that are possible using the constraints specified at
 1381 the lower level.
 1382 4. At the lower level inherit all keys that match with the higher level constraint.
 1383 5. If there are keys in the lower level constraint that are not inherited then the key
 1384 is invalid (i.e. it is less restrictive).

1385 8.3.4 Constraints Examples

1386 The following scenario is used.

1387 DSD

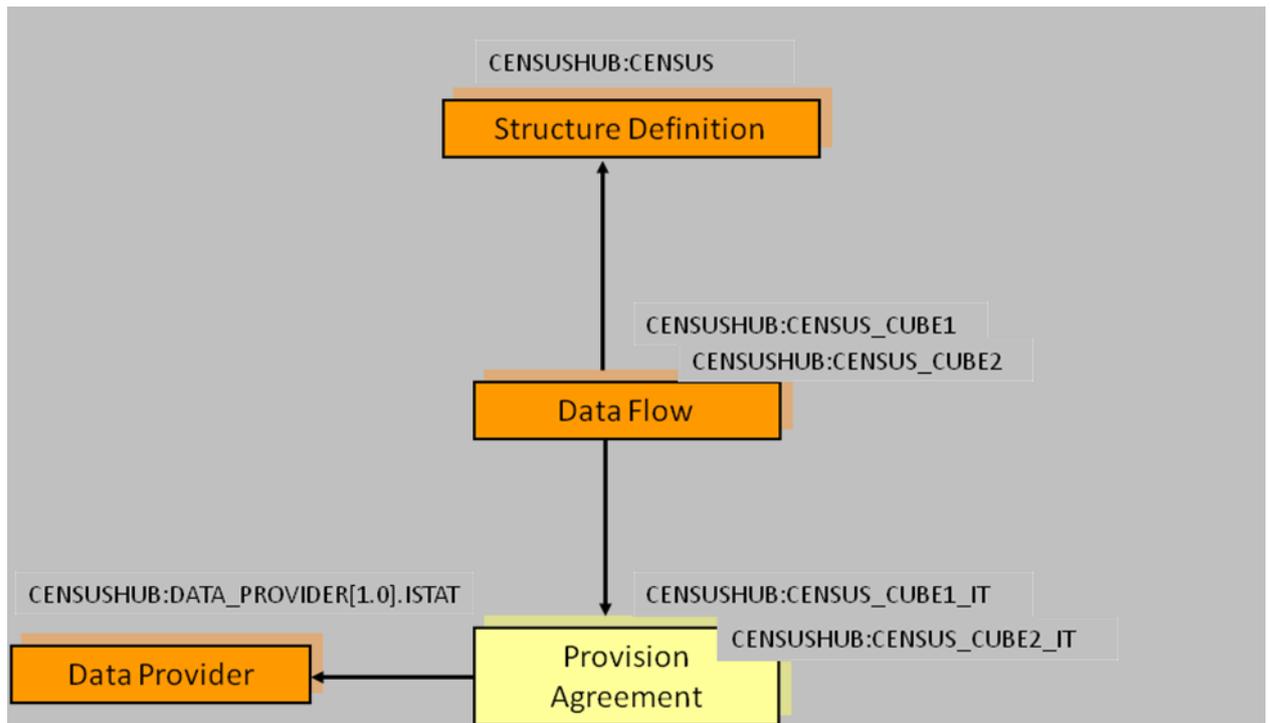
1388 This contains the following Dimensions:

- 1389 • GEO – Geography
- 1390 • SEX – Sex

1391 • AGE – Age

1392 • CAS – Current Activity Status

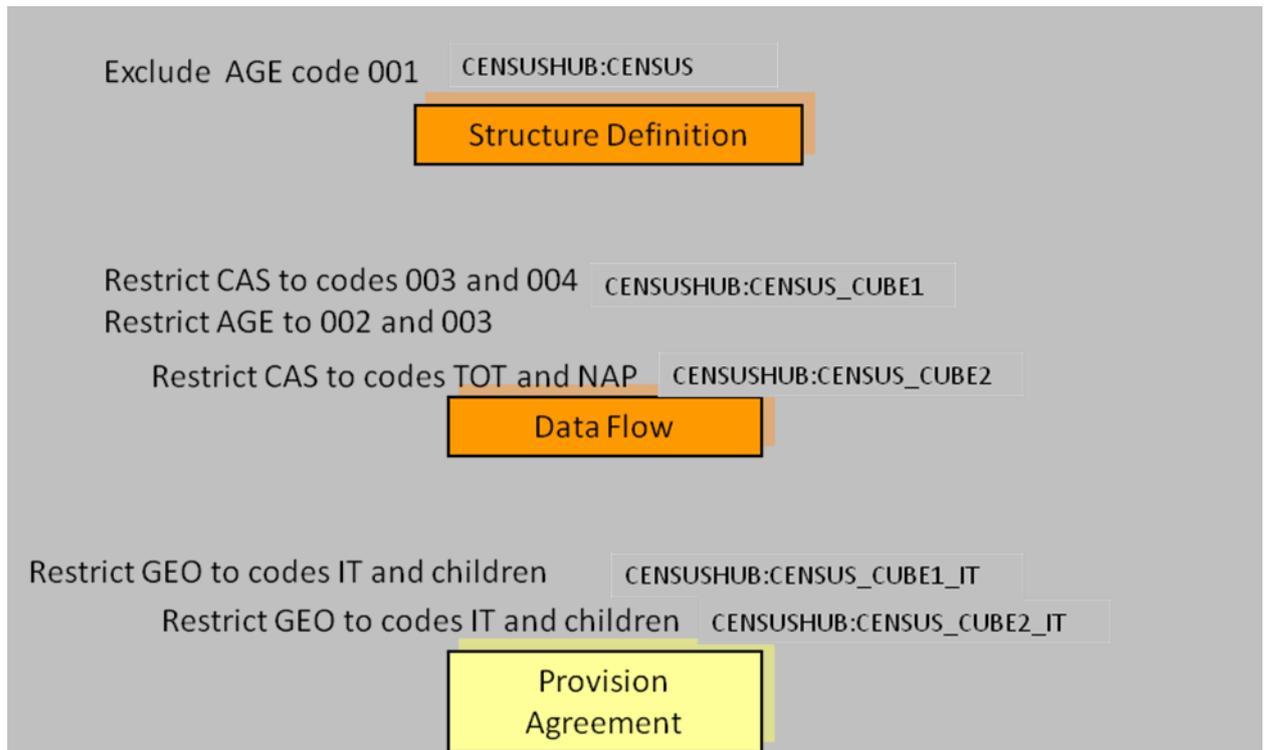
1393 In the DSD common code lists are used and the requirement is to restrict these at
 1394 various levels to specify the actual code that are valid for the object to which the
 1395 Content Constraint is attached.



1396

1397 **Figure 10: Example Scenario for Constraints**

1398 Constraints are declared as follows:



1399

1400

Figure 11: Example Content Constraints

1401

Notes:

1402

1. AGE is constrained for the DSD and is further restricted for the Dataflow

1403

CENSUS_CUBE1.

1404

2. The same Constraint applies to both Provision Agreements.

1405

1406

The cascade rules elaborated above result as follows:

1407

DSD

1408

1. Constrained by eliminating code 001 from the code list for the AGE Dimension.

1409

1410

Dataflow CENSUS_CUBE1

1411

1. Constrained by restricting the code list for the AGE Dimension to codes 002 and

1412

003(note that this is a more restrictive constraint than that declared for the DSD which specifies all codes except code 001).

1413

2. Restricts the CAS codes to 003 and 004.

1414

1415

1416

Dataflow CENSUS_CUBE2

1417

1. Restricts the code list for the CAS Dimension to codes TOT and NAP.

1418

2. Inherits the AGE constraint applied at the level of the DSD.

1419

1420 Provision Agreements CENSUS_CUBE1_IT

- 1421 1. Restricts the codes for the GEO Dimension to IT and its children.
 1422 2. Inherits the constraints from Dataflow CENSUS_CUBE1 for the AGE and CAS
 1423 Dimensions.
 1424

1425 Provision Agreements CENSUS_CUBE2_IT

- 1426 1. Restricts the codes for the GEO Dimension to IT and its children.
 1427 2. Inherits the constraints from Dataflow CENSUS_CUBE2 for the CAS Dimension.
 1428 3. Inherits the AGE constraint applied at the level of the DSD.
 1429

1430 The constraints are defined as follows:

1431 DSD Constraint

```
<structure:ContentConstraint id="CONSTRAINT1" agencyID="CENSUSHUB" type="Allowed" >
  <common:Name>name</common:Name>
  <structure:ConstraintAttachment>
    <structure:DataStructure>
      <Ref agencyID="CENSUSHUB" id="CENSUS"></Ref>
    </structure:DataStructure>
  </structure:ConstraintAttachment>
  <structure:CubeRegion include="true">
    <!-- note uses the ability of exclude values - i.e all values valid except this one -->
    <common:KeyValue id="AGE" include="false">
      <common:Value>001</common:Value>
    </common:KeyValue>
  </structure:CubeRegion>
</structure:ContentConstraint>
```

1432

1433

1434 Dataflow Constraints

```

} <structure:ContentConstraint id="CONSTRAINT2" agencyID="CENSUSHUB" type="Allowed" >
  <common:Name>name</common:Name>
} <structure:ConstraintAttachment>
} <structure:Dataflow>
  ..... <Ref agencyID="CENSUSHUB" id="CENSUS_CUBE1"></Ref>
} </structure:Dataflow>
} </structure:ConstraintAttachment>
} <structure:CubeRegion include="true">
} <common:KeyValue id="AGE" include="true">
  ..... <common:Value>002</common:Value>
  ..... <common:Value>003</common:Value>
} </common:KeyValue>
} <common:KeyValue id="CAS">
  ..... <common:Value>003</common:Value>
  ..... <common:Value>004</common:Value>
} </common:KeyValue>
} </structure:CubeRegion>
} </structure:ContentConstraint>

```

1435

```

} <structure:ContentConstraint id="CONSTRAINT3" agencyID="CENSUSHUB" type="Allowed" >
  <common:Name>name</common:Name>
} <structure:ConstraintAttachment>
} <structure:Dataflow>
  ..... <Ref agencyID="CENSUSHUB" id="CENSUS_CUBE2"></Ref>
} </structure:Dataflow>
} </structure:ConstraintAttachment>
} <structure:CubeRegion include="true">
} <common:KeyValue id="CAS" include="true">
  ..... <common:Value>TOT</common:Value>
  ..... <common:Value>NAP</common:Value>
} </common:KeyValue>
} </structure:CubeRegion>
} </structure:ContentConstraint>

```

1436

1437 Provision Agreement Constraint

```
<structure:ContentConstraint id="CONSTRAINT4" agencyID="CENSUSHUB" type="Allowed" >
  <common:Name>name</common:Name>
  <structure:ConstraintAttachment>
    <structure:ProvisionAgreement>
      ..... <Ref agencyID="CENSUSHUB" id="CENSUS_CUBE1_IT"></Ref>
    </structure:ProvisionAgreement>
    <structure:ProvisionAgreement>
      ..... <Ref agencyID="CENSUSHUB" id="CENSUS_CUBE2_IT"></Ref>
    </structure:ProvisionAgreement>
  </structure:ConstraintAttachment>
  <structure:CubeRegion include="true">
    <common:KeyValue id="GEO" include="true">
      ..... <common:Value cascadeValues="true">IT</common:Value>
    </common:KeyValue>
  </structure:CubeRegion>
</structure:ContentConstraint>
```

1438

1439

1440 **9 Transforming between versions of SDMX**

1441 **9.1 Scope**

1442 The scope of this section is to define both best practices and mandatory behaviour
1443 for specific aspects of transformation between different formats of SDMX.

1444 **9.2 Groups and Dimension Groups**

1445 **9.2.1 Issue**

1446 Version 2.1 introduces a more granular mechanism for specifying the relationship
1447 between a Data Attribute and the Dimensions to which the attribute applies. The
1448 technical construct for this is the Dimension Group. This Dimension Group has no
1449 direct equivalent in versions 2.0 and 1.0 and so the application transforming data
1450 from a version 2.1 data set to a version 2.0 or version 1.0 data set must decide to
1451 which construct the attribute value, whose Attribute is declared in a Dimension
1452 Group, should be attached. The closest construct is the “Series” attachment level and
1453 in many cases this is the correct construct to use.

1454 However, there is one case where the attribute **MUST** be attached to a Group in the
1455 version 2.0 and 1.0 message. The conditions of this case are:

- 1456 1. A Group is defined in the DSD with exactly the same Dimensions as a Dimension
1457 Group in the same DSD.
- 1458 2. The Attribute is defined in the DSD with an Attribute Relationship to the Dimension
1459 Group. This attribute is **NOT** defined as having an Attribute Relationship to the
1460 Group.

1461 **9.2.2 Structural Metadata**

1462 If the conditions defined in 9.2.1 are true then on conversion to a version 2.0 or 1.0
1463 DSD (Key Family) the Component/Attribute.attachmentLevel must be set to “Group”
1464 and the Component/Attribute/AttachmentGroup” is used to identify the Group. Note
1465 that under rule(1) in 1.2.1 this group will have been defined in the V 2.1 DSD and so
1466 will be present in the V 2.0 transformation.

1467 **9.2.3 Data**

1468 If the conditions defined in 9.2.1 are true then, on conversion from a 2.1 data set to a
1469 2.0 or 1.0 dataset the attribute value will be placed in the relevant <Group>. If these
1470 conditions are not true then the attribute value will be placed in the <Series>.

1471 **9.2.4 Compact Schema**

1472 If the conditions defined in 9.2.1 are true then the Compact Schema must be
1473 generated with the Group present and the Attribute(s) present in that group definition.

1474 **10 Validation and Transformation Language (VTL)**

1475 **10.1 Introduction**

1476 The Validation and Transformation Language (VTL) supports the definition of
1477 Transformations, which are algorithms to calculate new data starting from already
1478 existing ones⁶. The purpose of the VTL in the SDMX context is to enable the:

- 1479
- 1480 • definition of validation and transformation algorithms, in order to specify how to
1481 calculate new data from existing ones;
- 1482 • exchange of the definition of VTL algorithms, also together the definition of the
1483 data structures of the involved data (for example, exchange the data
1484 structures of a reporting framework together with the validation rules to be
1485 applied, exchange the input and output data structures of a calculation task
1486 together with the VTL Transformations describing the calculation algorithms);
- 1487 • compilation and execution of VTL algorithms, either interpreting the VTL
1488 transformations or translating them in whatever other computer language is
1489 deemed as appropriate.

1490

1491 It is important to note that the VTL has its own information model (IM), derived from
1492 the Generic Statistical Information Model (GSIM) and described in the VTL User
1493 Guide. The VTL IM is designed to be compatible with more standards, like SDMX,
1494 DDI (Data Documentation Initiative) and GSIM, and includes the model artefacts that
1495 can be manipulated (inputs and/or outputs of transformations, e.g. “Data Set”, “Data
1496 Structure”) and the model artefacts that allow the definition of the transformation
1497 algorithms (e.g. “Transformation”, “Transformation Scheme”).

1498

1499 The VTL language can be applied to SDMX artefacts by mapping the SDMX IM
1500 model artefacts to the model artefacts that VTL can manipulate. Thus, the SDMX
1501 artefacts can be used in VTL as inputs and/or outputs of transformations. It is
1502 important to be aware that the artefacts do not always have the same names in the
1503 SDMX and VTL IMs, nor do they always have the same meaning. The more evident
1504 example is given by the SDMX `Dataset` and the VTL “Data Set”, which do not
1505 correspond one another: as a matter of fact, the VTL “Data Set” maps to the SDMX
1506 “Dataflow”, while the SDMX “Dataset” has no explicit mapping to VTL (such an
1507 abstraction is not needed in the definition of VTL transformations). A SDMX
1508 “Dataset”, however, is an instance of a SDMX “Dataflow” and can be the artefact
1509 on which the VTL transformations are executed (i.e., the transformations are defined
1510 on `Dataflows` and are applied to `Dataflow` instances that can be `Datasets`).

1511

1512 The VTL programs (Transformation Schemes) are represented in SDMX through the
1513 `TransformationScheme` maintainable class which is composed of
1514 `Transformation` (nameable artefact). Each `Transformation` assigns the
1515 outcome of the evaluation of a VTL expression to a result.

1516

⁶ The Validation and Transformation Language is a standard language designed and published under the SDMX initiative. VTL is described in the VTL User and Reference Guides available on the SDMX website <https://sdmx.org>.

1517 This section does not explain the VTL language or any of the content published in the
1518 VTL guides. Rather, this is a description of how the VTL can be used in the SDMX
1519 context and applied to SDMX artefacts.

1520 **10.2 References to SDMX artefacts from VTL statements**

1521 **10.2.1 Introduction**

1522 The VTL can manipulate SDMX artefacts (or objects) by referencing them through
1523 pre-defined conventional names (aliases).

1524 The alias of a SDMX artefact can be its URN (Universal Resource Name), an
1525 abbreviation of its URN or another user-defined name.

1526 In any case, the aliases used in the VTL transformations have to be mapped to the
1527 SDMX artefacts through the `VtlMappingScheme` and `VtlMapping` classes (see
1528 the section of the SDMX IM relevant to the VTL). A `VtlMapping` allows specifying
1529 the aliases to be used in the VTL transformations, rulesets⁷ or user defined
1530 operators⁸ to reference SDMX artefacts. A `VtlMappingScheme` is a container for
1531 zero or more `VtlMapping`.

1532 The correspondence between an alias and a SDMX artefact must be one-to-one,
1533 meaning that a generic alias identifies one and just one SDMX artefact while a
1534 SDMX artefact is identified by one and just one alias. In other words, within a
1535 `VtlMappingScheme` an artefact can have just one alias and different artefacts
1536 cannot have the same alias.

1537 The references through the URN and the abbreviated URN are described in the
1538 following paragraphs.

1539 **10.2.2 References through the URN**

1540 This approach has the advantage that in the VTL code the URN of the referenced
1541 artefacts is directly intelligible by a human reader but has the drawback that the
1542 references are verbose.

1543
1544 The SDMX URN⁹ is the concatenation of the following parts, separated by special
1545 symbols like dot, equal, asterisk, comma, and parenthesis:

- 1546 • SDMXprefix
- 1547 • SDMX-IM-package-name
- 1548 • class-name
- 1549 • agency-id

⁷ See also the section “VTL-DL Rulesets” in the VTL Reference Manual.

⁸ The `VtlMapping` are used also for User Defined Operators (UDO). Although UDO operators are envisaged to be defined on generic operands, so that the specific artefacts to be manipulated are passed as parameters at their invocation, it is also possible that an UDO operator invokes directly some specific SDMX artefacts. These SDMX artefacts have to be mapped to the corresponding aliases used in the definition of the UDO through the `VtlMappingScheme` and `VtlMapping` classes as well.

⁹ For a complete description of the structure of the URN see the SDMX 2.1 Standards - Section 5 - Registry Specifications, paragraph 6.2.2 (“Universal Resource Name (URN)”).

- 1550 • maintainedobject-id
- 1551 • maintainedobject-version
- 1552 • container-object-id ¹⁰
- 1553 • object-id

1554 The generic structure of the URN is the following:

1555
1556 SDMXprefix.SDMX-IM-package-name.class-name=agency-id:maintainedobject-id
1557 (maintainedobject-version).*container-object-id.object-id

1558 The **SDMX prefix** is “urn:sdmx:org”, always the same for all SDMX artefacts.

1559 The **SDMX-IM-package-name** is the concatenation of the string “sdmx.infomodel.”
1560 with the package-name which the artefact belongs to. For example, for referencing a
1561 dataflow the SDMX-IM-package-name is “sdmx.infomodel.datastructure”, because
1562 the class `Dataflow` belongs to the package “datastructure”.

1563 The **class-name** is the name of the SDMX object class which the SDMX object
1564 belongs to (e.g., for referencing a dataflow the class-name is “Dataflow”). The VTL
1565 can reference SDMX artefacts that belong to the classes `Dataflow`, `Dimension`,
1566 `MeasureDimension`, `TimeDimension`, `PrimaryMeasure`, `DataAttribute`,
1567 `Concept`, `ConceptScheme`, `Codelist`.

1568 The **agency-id** is the acronym of the agency that owns the definition of the artefact,
1569 for example for the Eurostat artefacts the agency-id is “ESTAT”). The agency-id can
1570 be composite (for example `AgencyA.Dept1.Unit2`).

1571 The **maintainedobject-id** is the name of the maintained object which the artefact
1572 belongs to, and in case the artefact itself is maintainable¹¹, coincides with the name
1573 of the artefact. Therefore the maintainedobject-id depends on the class of the
1574 artefact:

- 1575 • if the artefact is a `Dataflow`, which is a maintainable class, the
1576 maintainedobject-id is the `Dataflow` name (`dataflow-id`);
- 1577 • if the artefact is a `Dimension`, `MeasureDimension`, `TimeDimension`,
1578 `PrimaryMeasure` or `DataAttribute`, which are not maintainable and belong
1579 to the `DataStructure` maintainable class, the maintainedobject-id is the
1580 name of the `DataStructure` (`dataStructure-id`) which the artefact belongs
1581 to;
- 1582 • if the artefact is a `Concept`, which is not maintainable and belongs to the
1583 `ConceptScheme` maintainable class, the maintainedobject-id is the name of
1584 the `ConceptScheme` (`conceptScheme-id`) which the artefact belongs to;
- 1585 • if the artefact is a `ConceptScheme`, which is a maintainable class, the
1586 maintainedobject-id is the name of the `ConceptScheme` (`conceptScheme-id`);
- 1587 • if the artefact is a `Codelist`, which is a maintainable class, the
1588 maintainedobject-id is the `Codelist` name (`codelist-id`).

¹⁰ The container-object-id can repeat and may not be present.

¹¹ i.e., the artefact belongs to a maintainable class

1589 The **maintainedobject-version** is the version of the maintained object which the
1590 artefact belongs to (for example, possible versions are 1.0, 2.1, 3.1.2).

1591 The **container-object-id** does not apply to the classes that can be referenced in
1592 VTL transformations, therefore is not present in their URN

1593 The **object-id** is the name of the non-maintainable artefact (when the artefact is
1594 maintainable its name is already specified as the maintainedobject-id, see above),
1595 in particular it has to be specified:

- 1596 • if the artefact is a `Dimension`, `MeasureDimension`, `TimeDimension`,
1597 `PrimaryMeasure` or `DataAttribute` (the object-id is the name of one of
1598 the artefacts above, which are data structure components)
- 1599 • if the artefact is a `Concept` (the object-id is the name of the `Concept`)

1600 For example, by using the URN, the VTL transformation that sums two SDMX
1601 dataflows DF1 and DF2 and assigns the result to a third persistent dataflow DFR,
1602 assuming that DF1, DF2 and DFR are the maintainedobject-id of the three
1603 dataflows, that their version is 1.0 and their Agency is AG, would be written as¹²:

```
1604
1605 'urn:sdmx:org.sdmx.infomodel.datastructure.Dataflow=AG:DFR(1.0)' <-
1606 'urn:sdmx:org.sdmx.infomodel.datastructure.Dataflow=AG:DF1(1.0)' +
1607 'urn:sdmx:org.sdmx.infomodel.datastructure.Dataflow=AG:DF2(1.0)'
```

1608 10.2.3 Abbreviation of the URN

1609 The complete formulation of the URN described above is exhaustive but verbose,
1610 even for very simple statements. In order to reduce the verbosity through a simplified
1611 identifier and make the work of transformation definers easier, proper abbreviations
1612 of the URN are possible. Using this approach, the referenced artefacts remain
1613 intelligible in the VTL code by a human reader.

1614 The URN can be abbreviated by omitting the parts that are not essential for the
1615 identification of the artefact or that can be deduced from other available information,
1616 including the context in which the invocation is made. The possible abbreviations are
1617 described below.

- 1618 • The **SDMXPrefix** can be omitted for all the SDMX objects, because it is a
1619 prefixed string (urn:sdmx:org), always the same for SDMX objects.
- 1620 • The **SDMX-IM-package-name** can be omitted as well because it can be
1621 deduced from the class-name that follows it (the table of the SDMX-IM
1622 packages and classes that allows this deduction is in the SDMX 2.1
1623 Standards - Section 5 - Registry Specifications, paragraph 6.2.3). In
1624 particular, considering the object classes of the artefacts that VTL can
1625 reference, the package is:
 - 1626 ○ "datastructure" for the classes `Dataflow`, `Dimension`,
1627 `MeasureDimension`, `TimeDimension`, `PrimaryMeasure`,
1628 `DataAttribute`,

¹² Since these references to SDMX objects include non-permitted characters as per the VTL ID notation, they need to be included between single quotes, according to the VTL rules for irregular names.

- 1629 ○ “conceptscheme” for the classes `Concept` and `ConceptScheme`
1630 ○ “codelist” for the class `Codelist`.
- 1631 • The **class-name** can be omitted as it can be deduced from the VTL invocation.
1632 In particular, starting from the VTL class of the invoked artefact (e.g. dataset,
1633 component, identifier, measure, attribute, variable, valuedomain), which is
1634 known given the syntax of the invoking VTL operator¹³, the SDMX class can
1635 be deduced from the mapping rules between VTL and SDMX (see the section
1636 “Mapping between VTL and SDMX” hereinafter)¹⁴.
- 1637 • If the **agency-id** is not specified, it is assumed by default equal to the agency-id
1638 of the `TransformationScheme`, `UserDefinedOperatorScheme` or
1639 `RulesetScheme` from which the artefact is invoked. For example, the
1640 agency-id can be omitted if it is the same as the invoking
1641 `TransformationScheme` and cannot be omitted if the artefact comes from
1642 another agency.¹⁵ Take also into account that, according to the VTL
1643 consistency rules, the agency of the result of a `Transformation` must be the
1644 same as its `TransformationScheme`, therefore the agency-id can be omitted
1645 for all the results (left part of `Transformation` statements).
- 1646 • As for the **maintainedobject-id**, this is essential in some cases while in other
1647 cases it can be omitted:
1648 ○ if the referenced artefact is a `Dataflow`, which is a maintainable class,
1649 the `maintainedobject-id` is the `dataflow-id` and obviously cannot be
1650 omitted;
1651 ○ if the referenced artefact is a `Dimension`, `MeasureDimension`,
1652 `TimeDimension`, `PrimaryMeasure`, `DataAttribute`, which are not
1653 maintainable and belong to the `DataStructure` maintainable class,
1654 the `maintainedobject-id` is the `dataStructure-id` and can be omitted,
1655 given that these components are always invoked within the invocation
1656 of a `Dataflow`, whose `dataStructure-id` can be deduced from the
1657 SDMX structural definitions;
1658 ○ if the referenced artefact is a `Concept`, which is not maintainable and
1659 belong to the `ConceptScheme` maintainable class, the `maintained`
1660 object is the `conceptScheme-id` and cannot be omitted;
1661 ○ if the referenced artefact is a `ConceptScheme`, which is a
1662 maintainable class, the `maintained object` is the `conceptScheme-id`
1663 and obviously cannot be omitted;

¹³ For the syntax of the VTL operators see the VTL Reference Manual

¹⁴ In case the invoked artefact is a VTL component, which can be invoked only within the invocation of a VTL data set (SDMX dataflow), the specific SDMX class-name (e.g. `Dimension`, `MeasureDimension`, `TimeDimension`, `PrimaryMeasure` or `DataAttribute`) can be deduced from the data structure of the SDMX `Dataflow` which the component belongs to.

¹⁵ If the Agency is composite (for example `AgencyA.Dept1.Unit2`), the agency is considered different even if only part of the composite name is different (for example `AgencyA.Dept1.Unit3` is a different Agency than the previous one). Moreover the agency-id cannot be omitted in part (i.e., if a `TransformationScheme` owned by `AgencyA.Dept1.Unit2` references an artefact coming from `AgencyA.Dept1.Unit3`, the specification of the agency-id becomes mandatory and must be complete, without omitting the possibly equal parts like `AgencyA.Dept1`)

1664 ○ if the referenced artefact is a `Codelist`, which is a maintainable
1665 class, the `maintainedobject-id` is the `codelist-id` and obviously
1666 cannot be omitted.

1667 • When the `maintainedobject-id` is omitted, the **maintainedobject-version** is
1668 omitted too. When the `maintainedobject-id` is not omitted and the
1669 `maintainedobject-version` is omitted, the version 1.0 is assumed by default.

1670 • As said, the **container-object-id** does not apply to the classes that can be
1671 referenced in VTL transformations, therefore is not present in their URN

1672 • The **object-id** does not exist for the artefacts belonging to the `Dataflow`,
1673 `ConceptScheme` and `Codelist` classes, while it exists and cannot be omitted
1674 for the artefacts belonging to the classes `Dimension`, `MeasureDimension`,
1675 `TimeDimension`, `PrimaryMeasure`, `DataAttribute` and `Concept`, as for
1676 them the `object-id` is the main identifier of the artefact

1677 The simplified object identifier is obtained by omitting all the first part of the URN,
1678 including the special characters, till the first part not omitted.

1679

1680 For example, the full formulation that uses the complete URN shown at the end of the
1681 previous paragraph:

1682

1683 'urn:sdmx:org.sdmx.infomodel.datastructure.Dataflow=AG:DFR(1.0)' :=

1684 'urn:sdmx:org.sdmx.infomodel.datastructure.Dataflow=AG:DF1(1.0)' +

1685 'urn:sdmx:org.sdmx.infomodel.datastructure.Dataflow=AG:DF2(1.0)'

1686

1687 by omitting all the non-essential parts would become simply:

1688

DFR := DF1 + DF2

1689 The references to the `Codelists` can be simplified similarly. For example, given the
1690 non-abbreviated reference to the `Codelist` `AG:CL_FREQ(1.0)`, which is¹⁶:

1691

'urn:sdmx:org.sdmx.infomodel.codelist.Codelist=AG:CL_FREQ(1.0)'

1692

if the `Codelist` is referenced from a ruleset scheme belonging to the agency AG,
1693 omitting all the optional parts, the abbreviated reference would become simply¹⁷:

1694

CL_FREQ

1695

As for the references to the components, it can be enough to specify the `component-
1696 id`, given that the `dataStructure-Id` can be omitted. An example of non-abbreviated
1697 reference, if the data structure is `DST1` and the component is `SECTOR`, is the
1698 following:

1699

'urn:sdmx:org.sdmx.infomodel.datastructure.DataStructure=AG:DST1(1.0).SECTOR'

¹⁶ Single quotes are needed because this reference is not a VTL regular name.

¹⁷ Single quotes are not needed in this case because `CL_FREQ` is a VTL regular name.

1700 The corresponding fully abbreviated reference, if made from a transformation
1701 scheme belonging to AG, would become simply:

1702 SECTOR

1703 For example, the transformation for renaming the component SECTOR of the
1704 dataflow DF1 into SEC can be written as¹⁸:

1705 'DFR(1.0)' := 'DF1(1.0)' [rename SECTOR to SEC]

1706 In the references to the Concepts, which can exist for example in the definition of the
1707 VTL Rulesets, at least the `conceptScheme-id` and the `concept-id` must be
1708 specified.

1709 An example of non-abbreviated reference, if the `conceptScheme-id` is CS1 and the
1710 `concept-id` is SECTOR, is the following:

1711 'urn:sdmx:org.sdmx.infomodel.conceptscheme.Concept=AG:CS1(1.0).SECTOR'

1712 The corresponding fully abbreviated reference, if made from a `RulesetScheme`
1713 belonging to AG, would become simply:

1714 CS1(1.0).SECTOR

1715 The Codes and in general all the Values can be written without any other
1716 specification, for example, the transformation to check if the values of the measures
1717 of the dataflow DF1 are between 0 and 25000 can be written like follows:

1718 'DFR(1.0)' := between ('DF1(1.0)', 0, 25000)

1719 The artefact (component, concept, codelist ...) which the Values are referred to can
1720 be deduced from the context in which the reference is made, taking also into account
1721 the VTL syntax. In the transformation above, for example, the values 0 and 2500 are
1722 compared to the values of the measures of DF1(1.0).

1723 **10.2.4 User-defined alias**

1724 The third possibility for referencing SDMX artefacts from VTL statements is to use
1725 user-defined aliases not related to the SDMX URN of the artefact.

1726 This approach gives preference to the use of symbolic names for the SDMX
1727 artefacts. As a consequence, in the VTL code the referenced artefacts would become
1728 not directly intelligible by a human reader. In any case, the VTL aliases are
1729 associated to the SDMX URN through the `VtlMappingScheme` and `VtlMapping`
1730 classes. These classes provide for structured references to SDMX artefacts whatever
1731 kind of reference is used in VTL statements (URN, abbreviated URN or user-defined
1732 aliases).

¹⁸ The result DFR(1.0) is be equal to DF1(1.0) save that the component SECTOR is called SEC

1733 **10.2.5 References to SDMX artefacts from VTL Rulesets**

1734 The VTL Rulesets allow defining sets of reusable rules that can be applied by some
1735 VTL operators, like the ones for validation and hierarchical roll-up. A “rule” consists in
1736 a relationship between Values belonging to some Value Domains or taken by some
1737 Variables, for example: (i) when the Country is USA then the Currency is USD; (ii)
1738 the Benelux is composed by Belgium, Luxembourg, Netherlands.

1739 The VTL Rulesets have a signature, in which the Value Domains or the Variables on
1740 which the Ruleset is defined are declared, and a body, which contains the rules.

1741 In the signature, given the mapping between VTL and SDMX better described in the
1742 following paragraphs, a reference to a VTL Value Domain becomes a reference to a
1743 SDMX `Codelist` or to a SDMX `ConceptScheme` (for SDMX measure dimensions),
1744 while a reference to a VTL Represented Variable becomes a reference to a SDMX
1745 `Concept`, assuming for it a definite representation¹⁹.

1746 In general, for referencing SDMX `Codelists` and `Concepts`, the conventions
1747 described in the previous paragraphs apply. In the Ruleset syntax, the elements that
1748 reference SDMX artefacts are called “valueDomain” and “variable” for the Datapoint
1749 Rulesets and “ruleValueDomain”, “ruleVariable”, “condValueDomain” “condVariable”
1750 for the Hierarchical Rulesets). The syntax of the Ruleset signature allows also to
1751 define aliases of the elements above, these aliases are valid only within the specific
1752 ruleset definition statement and cannot be mapped to SDMX.²⁰

1753 In the body of the Rulesets, the Codes and in general all the Values can be written
1754 without any other specification, because the artefact which the Values are referred
1755 (`Codelist`, `ConceptScheme`, `Concept`) to can be deduced from the Ruleset
1756 signature.

1757 **10.3 Mapping between SDMX and VTL artefacts**

1758 **10.3.1 When the mapping occurs**

1759 The mapping methods between the VTL and SDMX object classes allow
1760 transforming a SDMX definition in a VTL one and vice-versa for the artefacts to be
1761 manipulated.

1762 It should be remembered that VTL programs (i.e. Transformation Schemes) are
1763 represented in SDMX through the `TransformationScheme` maintainable class
1764 which is composed of `Transformations` (nameable artefacts). Each
1765 `Transformation` assigns the outcome of the evaluation of a VTL expression to a
1766 result: the input operands of the expression and the result can be SDMX artefacts.

1767 Every time a SDMX object is referenced in a VTL Transformation as an input
1768 operand, there is the need to generate a VTL definition of the object, so that the VTL
1769 operations can take place. This can be made starting from the SDMX definition and
1770 applying a SDMX-VTL mapping method in the direction from SDMX to VTL. The

¹⁹ Rulesets of this kind cannot be reused when the referenced Concept has a different representation.

²⁰ See also the section “VTL-DL Rulesets” in the VTL Reference Manual.

1771 possible mapping methods from SDMX to VTL are described in the following
1772 paragraphs and are conceived to allow the automatic deduction of the VTL definition
1773 of the object from the knowledge of the SDMX definition.

1774 In the opposite direction, every time an object calculated by means of VTL must be
1775 treated as a SDMX object (for example for exchanging it through SDMX), there is the
1776 need of a SDMX definition of the object, so that the SDMX operations can take place.
1777 The SDMX definition is needed for the VTL objects for which a SDMX use is
1778 envisaged²¹.

1779
1780 The mapping methods from VTL to SDMX are described in the following paragraphs
1781 as well, however they do not allow the complete SDMX definition to be automatically
1782 deduced from the VTL definition, more than all because the former typically contains
1783 additional information in respect to the latter. For example, the definition of a SDMX
1784 DSD includes also some mandatory information not available in VTL (like the concept
1785 scheme to which the SDMX components refer, the assignmentStatus and
1786 attributeRelationship for the DataAttributes and so on). Therefore the mapping
1787 methods from VTL to SDMX provide only a general guidance for generating SDMX
1788 definitions properly starting from the information available in VTL, independently of
1789 how the SDMX definition it is actually generated (manually, automatically or part and
1790 part).

1791 **10.3.2 General mapping of VTL and SDMX data structures**

1792 This section makes reference to the VTL “Model for data and their structure”²² and
1793 the correspondent SDMX “Data Structure Definition”²³.

1794 The main type of artefact that the VTL can manipulate is the VTL Data Set, which in
1795 general is mapped to the SDMX *Dataflow*. This means that a VTL Transformation,
1796 in the SDMX context, expresses the algorithm for calculating a derived *Dataflow*
1797 starting from some already existing *Dataflows* (either collected or derived).²⁴

1798 While the VTL Transformations are defined in term of *Dataflow* definitions, they are
1799 assumed to be executed on instances of such *Dataflows*, provided at runtime to
1800 the VTL engine (the mechanism for identifying the instances to be processed are not
1801 part of the VTL specifications and depend on the implementation of the VTL-based
1802 systems). As already said, the SDMX *Datasets* are instances of SDMX
1803 *Dataflows*, therefore a VTL Transformation defined on some SDMX *Dataflows*
1804 can be applied on some corresponding SDMX *Datasets*.

²¹ If a calculated artefact is persistent, it needs a persistent definition, i.e. a SDMX definition in a SDMX environment. Also possible calculated artefact that are not persistent may require a SDMX definition, for example when the result of a non-persistent calculation is disseminated through SDMX tools (like an inquiry tool).

²² See the VTL 2.0 User Manual

²³ See the SDMX 2.1 Section 2 – Information Model

²⁴ Besides the mapping between one SDMX *Dataflow* and one VTL Data Set, it is also possible to map distinct parts of a SDMX *Dataflow* to different VTL Data Set, as explained in a following paragraph.

1805 A VTL Data Set is structured by one and just one Data Structure and a VTL Data
1806 Structure can structure any number of Data Sets. Correspondingly, in the SDMX
1807 context a SDMX Dataflow is structured by one and just one
1808 DataStructureDefinition and one DataStructureDefinition can
1809 structure any number of Dataflows.

1810 A VTL Data Set has a Data Structure made of Components, which in turn can be
1811 Identifiers, Measures and Attributes. Similarly, a SDMX DataflowDefinition has
1812 a DataStructureDefinition made of components that can be
1813 DimensionComponents, PrimaryMeasure and DataAttributes. In turn, a
1814 SDMX DimensionComponent can be a Dimension, a TimeDimension or a
1815 MeasureDimension. Correspondingly, in the SDMX implementation of the VTL, the
1816 VTL Identifiers can be (optionally) distinguished in three sub-classes (Simple
1817 Identifier, Time Identifier, Measure Identifier) even if such a distinction is not
1818 evidenced in the VTL IM.

1819 However, a VTL Data Structure can have any number of Identifiers, Measures and
1820 Attributes, while a SDMX 2.1 DataStructureDefinition can have any number
1821 of Dimensions and DataAttributes but just one PrimaryMeasure²⁵. This is
1822 due to a difference between SDMX 2.1 and VTL in the possible representation
1823 methods of the data that contain more measures.

1824 As for SDMX, because the data structure cannot contain more than one measure
1825 component (i.e., the primaryMeasure), the representation of data having more
1826 measures is possible only by means of a particular dimension, called
1827 MeasureDimension, which is aimed at containing the name of the measure
1828 concepts, so that for each observation the value contained in the PrimaryMeasure
1829 component is the value of the measure concept reported in the MeasureDimension
1830 component.

1831 Instead VTL allows either the method above (an identifier containing the name of the
1832 measure together with just one measure component) or a more generic method that
1833 consists in defining more measure components in the data structure, one for each
1834 measure.

1835 Therefore for multi-measure data more mapping options are possible, as described in
1836 more detail in the following sections.

1837 **10.3.3 Mapping from SDMX to VTL data structures**

1838 **10.3.3.1 Basic Mapping**

1839 The main mapping method from SDMX to VTL is called **Basic** mapping. This is
1840 considered as the default mapping method and is applied unless a different method
1841 is specified through the VtlMappingScheme and VtlDataflowMapping classes.

²⁵ The SDMX community is evaluating the opportunity of allowing more than one measure component in a DataStructureDefinition in the next SDMX major version.

1842 When transforming **from SDMX to VTL**, this method consists in leaving the
 1843 components unchanged and maintaining their names and roles, according to the
 1844 following table:

SDMX	VTL
Dimension	(Simple) Identifier
Time Dimension	(Time) Identifier
Measure Dimension	(Measure) Identifier
Primary Measure	Measure
Data Attribute	Attribute

1845
 1846 According to this method, the resulting VTL structures are always mono-measure
 1847 (i.e., they have just one measure component) and their Measure is the SDMX
 1848 PrimaryMeasure. Nevertheless, if the SDMX data structure has a
 1849 MeasureDimension, which can convey the name of one or more measure
 1850 concepts, such unique measure component can contain the value of more
 1851 (conceptual) measures (one for each observation).

1852 As for the SDMX DataAttributes, in VTL they are all considered “at data point /
 1853 observation level” (i.e. dependent on all the VTL Identifiers), because VTL does not
 1854 have the SDMX AttributeRelationships, which defines the construct to which
 1855 the DataAttribute is related (e.g. observation, dimension or set or group of
 1856 dimensions, whole data set).

1857 With the Basic mapping, one SDMX observation generates one VTL data point.

1858 10.3.3.2 Pivot Mapping

1859 An alternative mapping method from SDMX to VTL is the **Pivot** mapping, which is
 1860 different from the Basic method only for the SDMX data structures that contain a
 1861 MeasureDimension, which are mapped to multi-measure VTL data structures.

1862 The SDMX structures that do not contain a MeasureDimension are mapped like in
 1863 the Basic mapping (see the previous paragraph).

1864 The SDMX structures that contain a MeasureDimension are mapped as follows
 1865 (this mapping is equivalent to a pivoting operation):

- 1866 • A SDMX simple dimension becomes a VTL (simple) identifier and a SDMX
 1867 TimeDimension becomes a VTL (time) identifier;
- 1868 • Each possible Concept C_j of the SDMX MeasureDimension is mapped to a
 1869 VTL Measure, having the same name as the SDMX Concept (i.e. C_j); the VTL
 1870 Measure C_j is a new VTL component even if the SDMX data structure has not
 1871 such a Component;
- 1872 • The SDMX MeasureDimension is not mapped to VTL (it disappears in the
 1873 VTL Data Structure);
- 1874 • The SDMX PrimaryMeasure is not mapped to VTL as well (it disappears in
 1875 the VTL Data Structure);
- 1876 • A SDMX DataAttribute is mapped in different ways according to its
 1877 AttributeRelationship:

- 1878 ○ If, according to the SDMX `AttributeRelationship`, the values of
 1879 the `DataAttribute` do not depend on the values of the
 1880 `MeasureDimension`, the SDMX `DataAttribute` becomes a VTL
 1881 Attribute having the same name. This happens if the
 1882 `AttributeRelationship` is not specified (i.e. the `DataAttribute`
 1883 does not depend on any `DimensionComponent` and therefore is at
 1884 data set level), or if it refers to a set (or a group) of dimensions which
 1885 does not include the `MeasureDimension`;
 1886 ○ Otherwise if, according to the SDMX `AttributeRelationship`, the
 1887 values of the `DataAttribute` depend on the `MeasureDimension`,
 1888 the SDMX `DataAttribute` is mapped to one VTL Attribute for each
 1889 possible `Concept` of the SDMX `MeasureDimension`; by default, the
 1890 names of the VTL Attributes are obtained by concatenating the name
 1891 of the SDMX `DataAttribute` and the names of the correspondent
 1892 `Concept` of the `MeasureDimension` separated by underscore; for
 1893 example, if the SDMX `DataAttribute` is named DA and the
 1894 possible concepts of the SDMX `MeasureDimension` are named C1,
 1895 C2, ..., Cn, then the corresponding VTL Attributes will be named
 1896 DA_C1, DA_C2, ..., DA_Cn (if different names are desired, they can
 1897 be achieved afterwards by renaming the Attributes through VTL
 1898 operators).
 1899 ○ Like in the Basic mapping, the resulting VTL Attributes are considered
 1900 as dependent on all the VTL identifiers (i.e. “at data point / observation
 1901 level”), because VTL does not have the SDMX notion of `Attribute`
 1902 `Relationship`.

1903 The summary mapping table of the “pivot” mapping from SDMX to VTL for the SDMX
 1904 data structures that contain a `MeasureDimension` is the following:

SDMX	VTL
<code>Dimension</code>	(Simple) Identifier
<code>TimeDimension</code>	(Time) Identifier
<code>MeasureDimension</code> & <code>PrimaryMeasure</code>	One Measure for each Concept of the SDMX Measure Dimension
<code>DataAttribute</code> not depending on the <code>MeasureDimension</code>	Attribute
<code>DataAttribute</code> depending on the <code>MeasureDimension</code>	One Attribute for each Concept of the SDMX Measure Dimension

1905
 1906 Using this mapping method, the components of the data structure can change in the
 1907 conversion from SDMX to VTL and it must be taken into account that the VTL
 1908 statements can reference only the components of the resulting VTL data structure.
 1909

1910 At observation / data point level, calling C_j (j=1, ... n) the jth Concept of the
 1911 `MeasureDimension`:

- 1912 • The set of SDMX observations having the same values for all the Dimensions
 1913 except than the `MeasureDimension` become one multi-measure VTL Data
 1914 Point, having one Measure for each Concept C_j of the SDMX
 1915 `MeasureDimension`;

- 1916 • The values of the SDMX simple Dimensions, TimeDimension and
1917 DataAttributes not depending on the MeasureDimension (these
1918 components by definition have always the same values for all the
1919 observations of the set above) become the values of the corresponding VTL
1920 (simple) Identifiers, (time) Identifier and Attributes.
1921 • The value of the PrimaryMeasure of the SDMX observation belonging to the
1922 set above and having MeasureDimension=Cj becomes the value of the VTL
1923 Measure Cj
1924 • For the SDMX DataAttributes depending on the MeasureDimension, the
1925 value of the DataAttribute DA of the SDMX observation belonging to the
1926 set above and having MeasureDimension=Cj becomes the value of the VTL
1927 Attribute DA_Cj

1928 **10.3.3.3 From SDMX DataAttributes to VTL Measures**

- 1929 • In some cases it may happen that the DataAttributes of the SDMX
1930 DataStructure need to be managed as Measures in VTL. Therefore, a
1931 variant of both the methods above consists in transforming all the SDMX
1932 DataAttributes in VTL Measures. When DataAttributes are converted
1933 to Measures, the two methods above are called Basic_A2M and Pivot_A2M
1934 (the suffix “A2M” stands for Attributes to Measures). Obviously, the resulting
1935 VTL data structure is, in general, multi-measure and does not contain
1936 Attributes.

1937 The Basic_A2M and Pivot_A2M behaves respectively like the Basic and Pivot
1938 methods, except that the final VTL components, which according to the Basic and
1939 Pivot methods would have had the role of Attribute, assume instead the role of
1940 Measure.

1941 Proper VTL features allow changing the role of specific attributes even after the
1942 SDMX to VTL mapping: they can be useful when only some of the
1943 DataAttributes need to be managed as VTL Measures.

1944 **10.3.4 Mapping from VTL to SDMX data structures**

1945 **10.3.4.1 Basic Mapping**

1946 The main mapping method **from VTL to SDMX** is called **Basic** mapping as well.

1947 This is considered as the default mapping method and is applied unless a different
1948 method is specified through the VtlMappingScheme and VtlDataflowMapping
1949 classes.

1950 The method consists in leaving the components unchanged and maintaining their
1951 names and roles in SDMX, according to the following mapping table, which is the
1952 same as the basic mapping from SDMX to VTL, only seen in the opposite direction.
1953

1954 This mapping method cannot be applied for SDMX 2.1 if the VTL data structure has
1955 more than one measure component, given that the SDMX 2.1
1956 DataStructureDefinition allows just one measure component (the
1957 PrimaryMeasure). In this case it becomes mandatory to specify a different

1958 mapping method through the `VtlMappingScheme` and `VtlDataflowMapping`
 1959 classes.²⁶

1960 Please note that the VTL measures can have any name while in SDMX 2.1 the
 1961 `MeasureComponent` has the mandatory name “obs_value”, therefore the name of
 1962 the VTL measure name must become “obs_value” in SDMX 2.1.

1963

1964 Mapping table:

1965

VTL	SDMX
(Simple) Identifier	Dimension
(Time) Identifier	TimeDimension
(Measure) Identifier	MeasureDimension
Measure	PrimaryMeasure
Attribute	DataAttribute

1966

1967

1968 If the distinction between simple identifier, time identifier and measure identifier is not
 1969 maintained in the VTL environment, the classification between `Dimension`,
 1970 `TimeDimension` and `MeasureDimension` exists only in SDMX, as declared in the
 1971 relevant `DataStructureDefinition`.

1972

1973 Regarding the Attributes, because VTL considers all of them “at observation level”,
 1974 the corresponding SDMX `DataAttributes` should be put “at observation level” as
 1975 well (`AttributeRelationships` referred to the `PrimaryMeasure`), unless some
 1976 other information about their `AttributeRelationship` is available.

1977

1978 Note that the basic mappings in the two directions (from SDMX 2.1 to VTL 2.0 and
 1979 vice-versa) are (almost completely) reversible. In fact, if a SDMX 2.1 structure is
 1980 mapped to a VTL structure and then the latter is mapped back to SDMX 2.1, the
 1981 resulting data structure is like the original one (apart for the
 1982 `AttributeRelationship`, that can be different if the original SDMX 2.1 structure
 1983 contains `DataAttributes` that are not at observation level). In reverse order, if a
 1984 VTL 2.0 mono-measure structure is mapped to SDMX 2.1 and then the latter is
 1985 mapped back to VTL 2.0, the original data structure is obtained (apart from the name
 1986 of the VTL measure, that in SDMX 2.1 must become “obs_value”).

1987

1988 As said, the resulting SDMX definitions must be compliant with the SDMX
 1989 consistency rules. For example, the SDMX DSD must have the
 1990 `assignmentStatus`, which does not exist in VTL, the `AttributeRelationship`
 1991 for the `DataAttributes` and so on.

1992 10.3.4.2 Unpivot Mapping

1993 An alternative mapping method from VTL to SDMX is the **Unpivot** mapping.

1994

1995 Although this mapping method can be used in any case, it makes major sense in
 1996 case the VTL data structure has more than one measure component (multi-measures

²⁶ If future SDMX major versions will allow multi-measures data structures, this method is expected to become applicable even if the VTL data structure has more than one measure

1997 VTL structure). For such VTL structures, in fact, the basic method cannot be applied,
 1998 given that by maintaining the data structure unchanged the resulting SDMX data
 1999 structure would have more than one measure component, which is not allowed by
 2000 SDMX 2.1 (it allows just one measure component, the `PrimaryMeasure`, called
 2001 “obs_value”).

2002
 2003 The multi-measures VTL structures have not a Measure Identifier (because the
 2004 Measures are separate components) and need to be converted to SDMX
 2005 dataflows having an added `MeasureDimension` which disambiguates the
 2006 multiple measures, and an added `PrimaryMeasure`, in which the measures’ values
 2007 are maintained.

2008
 2009 The **unpivot** mapping behaves like follows:

- 2010 • like in the basic mapping, a VTL (simple) identifier becomes a SDMX
 2011 `Dimension` and a VTL (time) identifier becomes a SDMX `TimeDimension`
 2012 (as said, a measure identifier cannot exist in multi-measure VTL structures);
- 2013 • a `MeasureDimension` component called “measure_name” is added to the
 2014 SDMX `DataStructure`;
- 2015 • a `PrimaryMeasure` component called “obs_value” is added to the SDMX
 2016 `DataStructure`;
- 2017 • each VTL Measure is mapped to a Concept of the SDMX `MeasureDimension`
 2018 having the same name as the VTL Measure (therefore all the VTL Measure
 2019 Components do not originate Components in the SDMX `DataStructure`);
- 2020 • a VTL Attribute becomes a SDMX `DataAttribute` having
 2021 `AttributeRelationship` referred to all the SDMX
 2022 `DimensionComponents` including the `TimeDimension` and except the
 2023 `MeasureDimension`.

2024
 2025 The summary mapping table of the **unpivot** mapping method is the following:
 2026

VTL	SDMX
(Simple) Identifier	<code>Dimension</code>
(Time) Identifier	<code>TimeDimension</code>
All Measure Components	<code>MeasureDimension</code> (having one Measure Concept for each VTL measure component) & <code>PrimaryMeasure</code>
Attribute	<code>DataAttribute</code> depending on all SDMX Dimensions including the <code>TimeDimension</code> and except the <code>MeasureDimension</code>

2027
 2028
 2029

At observation / data point level:

- 2030 • a multi-measure VTL Data Point becomes a set of SDMX observations, one for
 2031 each VTL measure
- 2032 • the values of the VTL identifiers become the values of the corresponding
 2033 SDMX `Dimensions`, for all the observations of the set above

- 2034 • the name of the j^{th} VTL measure (e.g. “C j ”) becomes the value of the SDMX
2035 MeasureDimension of the j^{th} observation of the set (i.e. the Concept C j)
2036 • the value of the j^{th} VTL measure becomes the value of the SDMX
2037 PrimaryMeasure of the j^{th} observation of the set
2038 • the values of the VTL Attributes become the values of the corresponding SDMX
2039 DataAttributes (in principle for all the observations of the set above)

2040 If desired, this method can be applied also to mono-measure VTL structures,
2041 provided that none of the VTL components has already the role of measure identifier.
2042 Like in the general case, a MeasureDimension component called “measure_name”
2043 would be added to the SDMX DataStructure and would have just one possible
2044 measure concept, corresponding to the unique VTL measure. The original VTL
2045 measure component would not become a Component in the SDMX data structure.
2046 The value of the VTL measure would be assigned to the SDMX PrimaryMeasure
2047 called “obs_value”.

2048 In any case, the resulting SDMX definitions must be compliant with the SDMX
2049 consistency rules. For example, the possible Concepts of the SDMX
2050 MeasureDimension need to be listed in a SDMX ConceptScheme, with proper id,
2051 agency and version; moreover, the SDMX DSD must have the assignmentStatus,
2052 which does not exist in VTL, the attributeRelationship for the
2053 DataAttributes and so on.

2054 **10.3.4.3 From VTL Measures to SDMX Data Attributes**

2055 For the multi-measure VTL structures (having more than one Measure Component),
2056 it may happen that the Measures of the VTL Data Structure need to be managed as
2057 DataAttributes in SDMX. Therefore a third mapping method consists in
2058 transforming one VTL measure in the SDMX primaryMeasure and all the other
2059 VTL Measures in SDMX DataAttributes. This method is called M2A (“M2A”
2060 stands for “Measures to DataAttributes”).

2061 When applied to mono-measure VTL structures (having one Measure component),
2062 the M2A method behaves like the Basic mapping (the VTL Measure component
2063 becomes the SDMX primary measure “obs_value”, there is no additional VTL
2064 measure to be converted to SDMX DataAttribute). Therefore the mapping table is the
2065 same as for the Basic method:

VTL	SDMX
(Simple) Identifier	Dimension
(Time) Identifier	TimeDimension
(Measure) Identifier (if any)	MeasureDimension
Measure	PrimaryMeasure
Attribute	DataAttribute

2066

2067 For multi-measure VTL structures (having more than one Measure component), one
2068 VTL Measure becomes the SDMX PrimaryMeasure while the other VTL Measures
2069 maintain their names and values but assume the role of DataAttribute in SDMX.
2070 The choice of the VTL Measure that correspond to the SDMX PrimaryMeasure is
2071 left to the definer of the SDMX data structure definition.

2072 Taking into account that the multi-measure VTL structures do not have a measure
 2073 identifier, the mapping table is the following:

VTL	SDMX
(Simple) Identifier	Dimension
(Time) Identifier	TimeDimension
One of the Measures	PrimaryMeasure
Other Measures	DataAttribute
Attribute	DataAttribute

2074

2075 Even in this case, the resulting SDMX definitions must be compliant with the SDMX
 2076 consistency rules. For example, the SDMX DSD must have the
 2077 assignmentStatus, which does not exist in VTL, the attributeRelationship
 2078 for the DataAttributes and so on. In particular, the primaryMeasure of the
 2079 SDMX 2.1 DSD must be called “obs_value” and must be one of the VTL Measures,
 2080 chosen by the DSD definer.

2081 **10.3.5 Declaration of the mapping methods between data structures**

2082 In order to define and understand properly VTL transformations, the applied mapping
 2083 methods must be specified in the SDMX structural metadata. If the default mapping
 2084 method (Basic) is applied, no specification is needed.
 2085

2086 A customized mapping can be defined through the `VtlMappingScheme` and
 2087 `VtlDataflowMapping` classes (see the section of the SDMX IM relevant to the
 2088 VTL). A `VtlDataflowMapping` allows specifying the mapping methods to be used
 2089 for a specific `dataflow`, both in the direction from SDMX to VTL
 2090 (`toVtlMappingMethod`) and from VTL to SDMX (`fromVtlMappingMethod`); in
 2091 fact a `VtlDataflowMapping` associates the structured URN that identifies a SDMX
 2092 `dataflow` to its VTL alias and its mapping methods.

2093 It is possible to specify the `toVtlMappingMethod` and `fromVtlMappingMethod`
 2094 also for the conventional `dataflow` called “generic_dataflow”: in this case the
 2095 specified mapping methods are intended to become the default ones, overriding the
 2096 “Basic” methods. In turn, the `toVtlMappingMethod` and
 2097 `fromVtlMappingMethod` declared for a specific `Dataflow` are intended to
 2098 override the default ones for such a `Dataflow`.

2099 The `VtlMappingScheme` is a container for zero or more `VtlDataflowMapping`
 2100 (besides possible mappings to artefacts other than `dataflows`).

2101 **10.3.6 Mapping dataflow subsets to distinct VTL data sets²⁷**

2102 Until now it has been assumed to map one SDMX `Dataflow` to one VTL dataset and
 2103 vice-versa. This mapping one-to-one is not mandatory according to VTL because a

²⁷ The kind of mapping explained here works in combination with a SDMX specific naming convention that requires pre-processing before parsing the VTL expressions. As highlighted below, the identifiers of the VTL datasets are a shortcut of some specific VTL operators applied to the SDMX `Dataflows`. This is not safe to use outside an SDMX context, as the naming convention may have no meaning there.

2104 VTL data set is meant to be a set of observations (data points) on a logical plane,
2105 having the same logical data structure and the same general meaning, independently
2106 of the possible physical representation or storage (see VTL 2.0 User Manual page
2107 24), therefore a `SDMX Dataflow` can be seen either as a unique set of data
2108 observations (corresponding to one VTL data set) or as the union of many sets of
2109 data observations (each one corresponding to a distinct VTL data set).

2110 As a matter of fact, in some cases it can be useful to define VTL operations involving
2111 definite parts of a `SDMX Dataflow` instead than the whole.²⁸

2112 Therefore, in order to make the coding of VTL operations simpler when applied on
2113 parts of `SDMX Dataflows`, it is allowed to map distinct parts of a `SDMX Dataflow`
2114 to distinct VTL data sets according to the following rules and conventions. This kind
2115 of mapping is possible both from `SDMX` to VTL and from VTL to `SDMX`, as better
2116 explained below.²⁹

2117 Given a `SDMX Dataflow` and some predefined `Dimensions` of its
2118 `DataStructure`, it is allowed to map the subsets of observations that have the
2119 same combination of values for such `Dimensions` to correspondent VTL datasets.

2120 For example, assuming that the `SDMX dataflow DF1(1.0)` has the `Dimensions`
2121 `INDICATOR`, `TIME_PERIOD` and `COUNTRY`, and that the user declares the
2122 `Dimensions` `INDICATOR` and `COUNTRY` as basis for the mapping (i.e. the
2123 mapping dimensions): the observations that have the same values for `INDICATOR`
2124 and `COUNTRY` would be mapped to the same VTL dataset (and vice-versa).

2125 In practice, this kind mapping is obtained like follows:

2126 • For a given `SDMX dataflow`, the user (VTL definer) declares the dimension
2127 components on which the mapping will be based, in a given order.³⁰ Following
2128 the example above, imagine that the user declares the `dimensions`
2129 `INDICATOR` and `COUNTRY`.

2130 • The VTL dataset is given a name using a special notation also called “ordered
2131 concatenation” and composed of the following parts:

²⁸ A typical example of this kind is the validation, and more in general the manipulation, of individual time series belonging to the same `Dataflow`, identifiable through the dimension components of the `Dataflow` except the time `Dimension`. The coding of these kind of operations might be simplified by mapping distinct time series (i.e. different parts of a `SDMX Dataflow`) to distinct VTL data sets.

²⁹ Please note that this kind of mapping is only an option at disposal of the definer of VTL Transformations; in fact it remains always possible to manipulate the needed parts of `SDMX Dataflows` by means of VTL operators (e.g. “sub”, “filter”, “calc”, “union” ...), maintaining a mapping one-to-one between `SDMX Dataflows` and VTL datasets.

³⁰ This definition is made through the `ToVtlSubspace` and `ToVtlSpaceKey` classes and/or the `FromVtlSuperspace` and `FromVtlSpaceKey` classes, depending on the direction of the mapping (“key” means “dimension”). The mapping of `Dataflow` subsets can be applied independently in the two directions, also according to different `Dimensions`. When no `Dimension` is declared for a given direction, it is assumed that the option of mapping different parts of a `SDMX Dataflow` to different VTL datasets is not used.

- 2132 ○ The reference to the SDMX *dataflow* (expressed according to the
2133 rules described in the previous paragraphs, i.e. URN, abbreviated
2134 URN or another alias); for example DF(1.0);
2135 ○ a slash (“/”) as a separator;³¹
2136 ○ The reference to a specific part of the SDMX *dataflow* above,
2137 expressed as the concatenation of the values that the SDMX
2138 dimensions declared above must have, separated by dots (“.”) and
2139 written in the order in which these dimensions are defined³². For
2140 example POPULATION.USA would mean that such a VTL dataset is
2141 mapped to the SDMX observations for which the dimension
2142 *INDICATOR* is equal to POPULATION and the dimension *COUNTRY*
2143 is equal to USA.

2144 In the VTL transformations, this kind of dataset name must be referenced between
2145 single quotes because the slash (“/”) is not a regular character according to the VTL
2146 rules.

2147 Therefore, the generic name of this kind of VTL datasets would be:

2148 ‘DF(1.0)/*INDICATORvalue*.*COUNTRYvalue*’

2149 Where DF(1.0) is the *Dataflow* and *INDICATORvalue* and *COUNTRYvalue* are
2150 placeholders for one value of the *INDICATOR* and *COUNTRY* dimensions.

2151 Instead the specific name of one of these VTL datasets would be:

2152 ‘DF(1.0)/POPULATION.USA’

2153 In particular, this is the VTL dataset that contains all the observations of the
2154 *dataflow* DF(1.0) for which *INDICATOR* = POPULATION and *COUNTRY* = USA.

2155 Let us now analyse the different meaning of this kind of mapping in the two mapping
2156 directions, i.e. from SDMX to VTL and from VTL to SDMX.

2157 As already said, the mapping from SDMX to VTL happens when the VTL datasets
2158 are operand of VTL transformations, instead the mapping from VTL to SDMX
2159 happens when the VTL datasets are result of VTL transformations³³ and need to be
2160 treated as SDMX objects. This kind of mapping can be applied independently in the
2161 two directions and the *Dimensions* on which the mapping is based can be different
2162 in the two directions: these *Dimensions* are defined in the *ToVtlSpaceKey* and in
2163 the *FromVtlSpaceKey* classes respectively.

³¹ As a consequence of this formalism, a slash in the name of the VTL dataset assumes the specific meaning of separator between the name of the *Dataflow* and the values of some of its *Dimensions*.

³² This is the order in which the dimensions are defined in the *ToVtlSpaceKey* class or in the *FromVtlSpaceKey* class, depending on the direction of the mapping.

³³ It should be remembered that, according to the VTL consistency rules, a given VTL dataset cannot be the result of more than one VTL transformation.

2164 First, let us see what happens in the mapping direction from SDMX to VTL, i.e. when
 2165 parts of a SDMX *dataflow* (e.g. DF1(1.0)) need to be mapped to distinct VTL
 2166 datasets that are operand of some VTL transformations.

2167 As already said, each VTL dataset is assumed to contain all the observations of the
 2168 SDMX *dataflow* having *INDICATOR=INDICATORvalue* and *COUNTRY=COUNTRYvalue*. For example, the VTL dataset 'DF1(1.0)/POPULATION.USA' would
 2169 contain all the observations of DF1(1.0) having *INDICATOR = POPULATION* and
 2170 *COUNTRY = USA*.
 2171

2172 In order to obtain the data structure of these VTL datasets from the SDMX one, it is
 2173 assumed that the SDMX dimensions on which the mapping is based are dropped, i.e.
 2174 not maintained in the VTL data structure; this is possible because their values are
 2175 fixed for each one of the invoked VTL datasets³⁴. After that, the mapping method
 2176 from SDMX to VTL specified for the *dataflow* DF1(1.0) is applied (i.e. basic, pivot ...).

2177 In the example above, for all the datasets of the kind
 2178 'DF1(1.0)/*INDICATORvalue.COUNTRYvalue*', the dimensions *INDICATOR* and
 2179 *COUNTRY* would be dropped so that the data structure of all the resulting VTL data
 2180 sets would have the identifier *TIME_PERIOD* only.

2181 It should be noted that the desired VTL datasets (i.e. of the kind 'DF1(1.0)/
 2182 *INDICATORvalue.COUNTRYvalue*') can be obtained also by applying the VTL
 2183 operator "**sub**" (subspace) to the *dataflow* DF1(1.0), like in the following VTL
 2184 expression:

```
2185     'DF1(1.0)/POPULATION.USA' :=
2186     DF1(1.0) [ sub INDICATOR="POPULATION", COUNTRY="USA" ];
2187
2188     'DF1(1.0)/POPULATION.CANADA' :=
2189     DF1(1.0) [ sub INDICATOR="POPULATION", COUNTRY="CANADA" ];
2190
2191     ... ..
```

2192 In fact the VTL operator "**sub**" has exactly the same behaviour. Therefore, mapping
 2193 different parts of a SDMX *dataflow* to different VTL datasets in the direction from
 2194 SDMX to VTL through the ordered concatenation notation is equivalent to a proper
 2195 use of the operator "**sub**" on such a *dataflow*.³⁵

2196 In the direction from SDMX to VTL it is allowed to omit the value of one or more
 2197 Dimensions on which the mapping is based, but maintaining all the separating dots

³⁴ If these dimensions would not be dropped, taking into account that the typical binary VTL operations at dataset level (+, -, *, / and so on) are executed on the observations having matching identifiers, the VTL datasets resulting from this kind of mapping would have non-matching values for the mapping dimensions (e.g. *POPULATION* and *COUNTRY*), therefore it would not be possible to compose the resulting VTL datasets one another (e.g. it would not be possible to calculate the population ratio between USA and CANADA).

³⁵ In case the ordered concatenation notation is used, the VTL Transformation described above, e.g. 'DF1(1.0)/POPULATION.USA' := DF1(1.0) [sub INDICATOR="POPULATION", COUNTRY="USA"], is implicitly executed and, in order to test the overall compliance of the VTL program to the VTL consistency rules, it has to be considered as part of the VTL program even if it is not explicitly coded.

2198 (therefore it may happen to find two or more consecutive dots and dots in the
2199 beginning or in the end). The absence of value means that for the corresponding
2200 Dimension all the values are kept and the Dimension is not dropped.

2201 For example, 'DF(1.0)/POPULATION.' (note the dot in the end of the name) is the
2202 VTL dataset that contains all the observations of the `dataflow` DF(1.0) for which
2203 `INDICATOR = POPULATION` and `COUNTRY = any value`.

2204 This is equivalent to the application of the VTL "sub" operator only to the identifier
2205 `INDICATOR`:

```
2206 'DF1(1.0)/POPULATION.' :=  
2207 DF1(1.0) [ sub INDICATOR="POPULATION" ];
```

2208
2209 Therefore the VTL dataset 'DF1(1.0)/POPULATION.' would have the identifiers
2210 `COUNTRY` and `TIME_PERIOD`.

2211 Heterogeneous invocations of the same `Dataflow` are allowed, i.e. omitting different
2212 `Dimensions` in different invocations.

2213 Let us now analyse the mapping direction from VTL to SDMX.

2214 In this situation, distinct parts of a SDMX `dataflow` are calculated as distinct VTL
2215 datasets, under the constraint that they must have the same VTL data structure.

2216 For example, let us assume that the VTL programmer wants to calculate the SDMX
2217 `dataflow` DF2(1.0) having the `Dimensions` `TIME_PERIOD`, `INDICATOR`, and
2218 `COUNTRY` and that such a programmer finds it convenient to calculate separately
2219 the parts of DF2(1.0) that have different combinations of values for `INDICATOR` and
2220 `COUNTRY`:

- 2221 • each part is calculated as a VTL derived dataset, result of a dedicated VTL
2222 transformation;³⁶
- 2223 • the data structure of all these VTL datasets has the `TIME_PERIOD` identifier
2224 and does not have the `INDICATOR` and `COUNTRY` identifiers.³⁷

2225 Under these hypothesis, such derived VTL datasets can be mapped to DF2(1.0) by
2226 declaring the `Dimensions` `INDICATOR` and `COUNTRY` as mapping dimensions³⁸.

2227 The corresponding VTL transformations, assuming that the result needs to be
2228 persistent, would be of this kind:³⁹

³⁶ If the whole DF2(1.0) is calculated by means of just one VTL transformation, then the mapping between the SDMX `dataflow` and the corresponding VTL dataset is one-to-one and this kind of mapping (one SDMX `Dataflow` to many VTL datasets) does not apply..

³⁷ This is possible as each VTL dataset corresponds to one particular combination of values of `INDICATOR` and `COUNTRY`

³⁸ The mapping dimensions are defined as `FromVtlSpaceKeys` of the `FromVtlSuperSpace` of the `VtlDataflowMapping` relevant to DF2(1.0)

³⁹ the symbol of the VTL persistent assignment is used (<-)

2229 'DF2(1.0)/INDICATORvalue.COUNTRYvalue' <- expression

2230

2231 Some examples follow, for some specific values of INDICATOR and COUNTRY:

2232

2233 'DF2(1.0)/GDPPERCAPITA.USA' <- expression11;

2234 'DF2(1.0)/GDPPERCAPITA.CANADA' <- expression12;

2235

2236 'DF2(1.0)/POPGROWTH.USA' <- expression21;

2237 'DF2(1.0)/POPGROWTH.CANADA' <- expression22;

2238

2239

2240 As said, it is assumed that these VTL derived datasets have the TIME_PERIOD as
2241 the only identifier. In the mapping from VTL to SMDX, the *Dimensions* INDICATOR
2242 and COUNTRY are added to the VTL data structure on order to obtain the SDMX
2243 one, with the following values respectively:

2244

<i>VTL dataset</i>	<i>INDICATOR value</i>	<i>COUNTRY value</i>
'DF2(1.0)/GDPPERCAPITA.USA'	GDPPERCAPITA	USA
'DF2(1.0)/GDPPERCAPITA.CANADA'	GDPPERCAPITA	CANADA
... ..		
'DF2(1.0)/POPGROWTH.USA'	POPGROWTH	USA
'DF2(1.0)/POPGROWTH.CANADA'	POPGROWTH	CANADA
... ..		

2253

2254 It should be noted that the application of this many-to-one mapping from VTL to
2255 SDMX is equivalent to an appropriate sequence of VTL Transformations. These use
2256 the VTL operator "calc" to add the proper VTL identifiers (in the example,
2257 INDICATOR and COUNTRY) and to assign to them the proper values and the
2258 operator "union" in order to obtain the final VTL dataset (in the example DF2(1.0)),
2259 that can be mapped one-to-one to the homonymous SDMX *Dataflow*. Following
2260 the same example, these VTL transformations would be:

2261

2262 DF2bis_GDPPERCAPITA_USA := 'DF2(1.0)/GDPPERCAPITA.USA'
2263 [calc identifier INDICATOR := "GDPPERCAPITA",
2264 identifier COUNTRY := "USA"];

2265

2266 DF2bis_GDPPERCAPITA_CANADA := 'DF2(1.0)/GDPPERCAPITA.CANADA'
2267 [calc identifier INDICATOR:= "GDPPERCAPITA",
2268 identifier COUNTRY:= "CANADA"];

2269

2270 DF2bis_POPGROWTH_USA := 'DF2(1.0)/POPGROWTH.USA'
2271 [calc identifier INDICATOR := "POPGROWTH",
2272 identifier COUNTRY := "USA"];

2273 DF2bis_POPGROWTH_CANADA := 'DF2(1.0)/POPGROWTH.CANADA'
2274 [calc identifier INDICATOR := "POPGROWTH",
2275 identifier COUNTRY := "CANADA"];

2276

2277 DF2(1.0) <- UNION (DF2bis_GDPPERCAPITA_USA',
2278 DF2bis_GDPPERCAPITA_CANADA',

2279 ... ,
 2280 DF2bis_POPGROWTH_USA',
 2281 DF2bis_POPGROWTH_CANADA'
 2282 ...);
 2283

2284 In other words, starting from the datasets explicitly calculated through VTL (in the
 2285 example 'DF2(1.0)/GDPPERCAPITA.USA' and so on), the first step consists in
 2286 calculating other (non-persistent) VTL datasets (in the example
 2287 DF2bis_GDPPERCAPITA_USA and so on) by adding the identifiers INDICATOR and
 2288 COUNTRY with the desired values (*INDICATORvalue* and *COUNTRYvalue*). Finally,
 2289 all these non-persistent data sets are united and give the final result DF2(1.0)⁴⁰,
 2290 which can be mapped one-to-one to the homonymous SDMX dataflow having the
 2291 dimension components TIME_PERIOD, INDICATOR and COUNTRY.

2292 Therefore, mapping different VTL datasets having the same data structure to
 2293 different parts of a SDMX *dataflow*, i.e. in the direction from VTL to SDMX, through
 2294 the ordered concatenation notation is equivalent to a proper use of the operators
 2295 "calc" and "union" on such datasets.⁴¹

2296 It is worth noting that in the direction from VTL to SDMX it is mandatory to specify the
 2297 value for every Dimension on which the mapping is based (in other word, in the name
 2298 of the calculated VTL dataset is not possible to omit the value of some of the
 2299 Dimensions).

2300 **10.3.7 Mapping variables and value domains between VTL and SDMX**

2301 With reference to the VTL "model for Variables and Value domains", the following
 2302 additional mappings have to be considered:

VTL	SDMX
Data Set Component	Although this abstraction exists in SDMX, it does not have an explicit definition and correspond to a Component (either a Dimension or a PrimaryMeasure or a DataAttribute) belonging to one specific Dataflow ⁴²
Represented Variable	Concept with a definite Representation

⁴⁰ The result is persistent in this example but it can be also non persistent if needed.

⁴¹ In case the ordered concatenation notation from VTL to SDMX is used, the set of transformations described above is implicitly performed; therefore, in order to test the overall compliance of the VTL program to the VTL consistency rules, these implicit transformations have to be considered as part of the VTL program even if they are not explicitly coded.

⁴² Through SDMX *Constraints*, it is possible to specify the values that a Component of a Dataflow can assume.

Value Domain	Representation (see the Structure Pattern in the Base Package)
Enumerated Value Domain / Code List	CodeList (for enumerated Dimension, PrimaryMeasure, DataAttribute) or ConceptScheme (for MeasureDimension)
Code	Code (for enumerated Dimension, PrimaryMeasure, DataAttribute) or Concept (for MeasureDimension)
Described Value Domain	non-enumerated Representation (having Facets / ExtendedFacets, see the Structure Pattern in the Base Package)
Value	Although this abstraction exists in SDMX, it does not have an explicit definition and correspond to a Code of a CodeList (for enumerated Representations) or to a valid value (for non-enumerated Representations) or to a Concept (for MeasureDimension)
Value Domain Subset / Set	This abstraction does not exist in SDMX
Enumerated Value Domain Subset / Enumerated Set	This abstraction does not exist in SDMX
Described Value Domain Subset / Described Set	This abstraction does not exist in SDMX
Set list	This abstraction does not exist in SDMX

2303

2304 The main difference between VTL and SDMX relies on the fact that the VTL artefacts
 2305 for defining subsets of Value Domains do not exist in SDMX, therefore the VTL
 2306 features for referring to predefined subsets are not available in SDMX. These
 2307 artefacts are the Value Domain Subset (or Set), either enumerated or described, the
 2308 Set List (list of values belonging to enumerated subsets) and the Data Set
 2309 Component (aimed at defining the set of values that the Component of a Data Set
 2310 can take, possibly a subset of the codes of Value Domain).

2311 Another difference consists in the fact that all Value Domains are considered as
 2312 identifiable objects in VTL either if enumerated or not, while in SDMX the CodeList
 2313 (corresponding to a VTL enumerated Value Domain) is identifiable, while the SDMX
 2314 non-enumerated Representation (corresponding to a VTL non-enumerated Value
 2315 Domain) is not identifiable. As a consequence, the definition of the VTL rulesets,

2316 which in VTL can refer either to enumerated or non-enumerated value domains, in
2317 SDMX can refer only to enumerated Value Domains (i.e. to SDMX Codelists).

2318 As for the mapping between VTL variables and SDMX Concepts, it should be noted
2319 that these artefacts do not coincide perfectly. In fact, the VTL variables are
2320 represented variables, defined always on the same Value Domain (“Representation”
2321 in SDMX) independently of the data set / data structure in which they appear⁴³, while
2322 the SDMX Concepts can have different Representations in different
2323 DataStructures.⁴⁴ This means that one SDMX Concept can correspond to many
2324 VTL Variables, one for each representation the Concept has.

2325 Therefore, it is important to be aware that some VTL operations (for example the
2326 binary operations at data set level) are consistent only if the components having the
2327 same names in the operated VTL data sets have also the same representation (i.e.
2328 the same Value Domain as for VTL). For example, it is possible to obtain correct
2329 results from the VTL expression

2330 $DS_c := DS_a + DS_b$ (where DS_a, DS_b, DS_c are VTL Data Sets)

2331 if the matching components in DS_a and DS_b (e.g. ref_date, geo_area, sector ...)
2332 refer to the same general representation. In simpler words, DS_a and DS_b must
2333 use the same values/codes (for ref_date, geo_area, sector ...), otherwise the
2334 relevant values would not match and the result of the operation would be wrong.

2335 As mentioned, the property above is not enforced by construction in SDMX, and
2336 different representations of the same Concept can be not compatible one another
2337 (for example, it may happen that geo_area is represented by ISO-alpha-3 codes in
2338 DS_a and by ISO alpha-2 codes in DS_b). Therefore, it will be up to the definer of
2339 VTL transformations to ensure that the VTL expressions are consistent with the
2340 actual representations of the correspondent SDMX Concepts.

2341 It remains up to the SDMX-VTL definer also the assurance of the consistency
2342 between a VTL Ruleset defined on Variables and the SDMX Components on which
2343 the Ruleset is applied. In fact, a VTL Ruleset is expressed by means of the values of
2344 the Variables (i.e. SDMX Concepts), i.e. assuming definite representations for them
2345 (e.g. ISO-alpha-3 for country). If the Ruleset is applied to SDMX Components that
2346 have the same name of the Concept they refer to but different representations (e.g.
2347 ISO-alpha-2 for country), the Ruleset cannot work properly.

2348 **10.4 Mapping between SDMX and VTL Data Types**

2349 **10.4.1 VTL Data types**

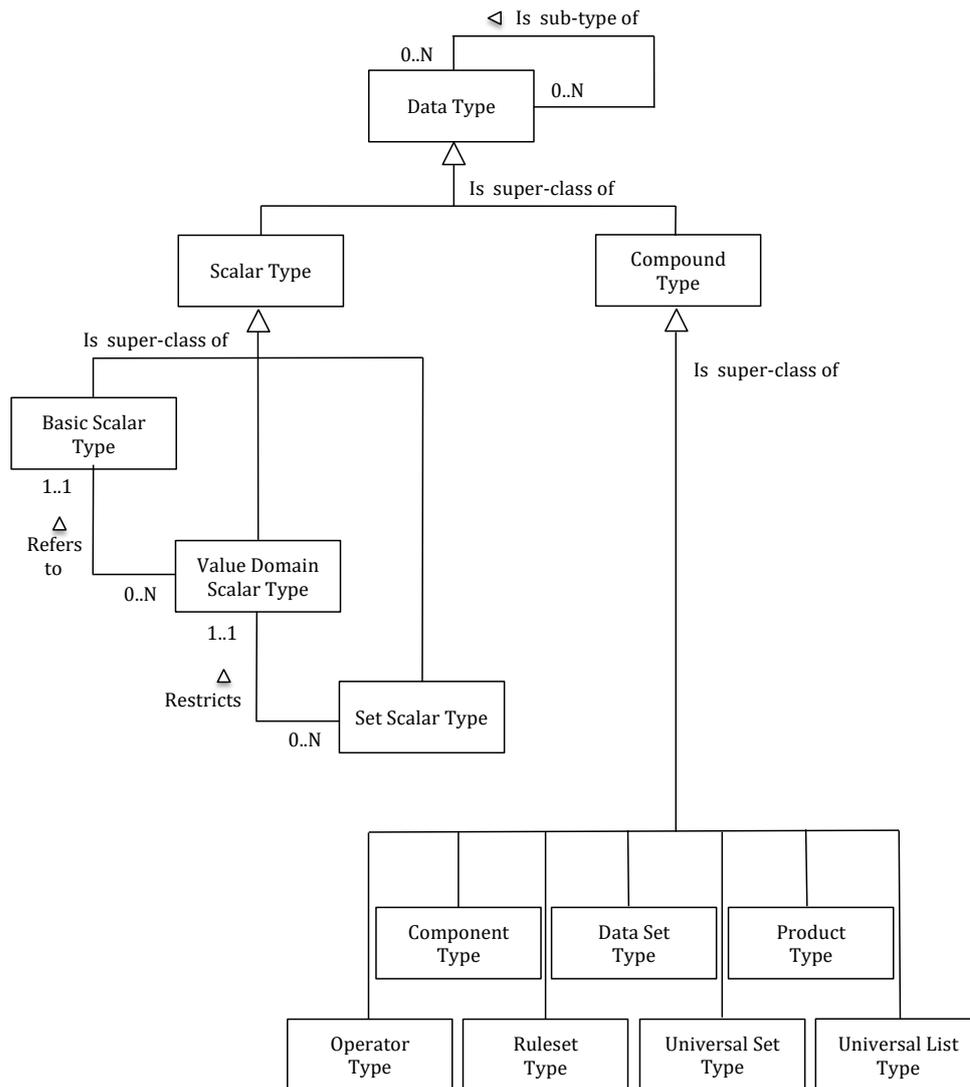
2350 According to the VTL User Guide the possible operations in VTL depend on the data
2351 types of the artefacts. For example, numbers can be multiplied but text strings

⁴³ By using represented variables, VTL can assume that data structures having the same variables as identifiers can be composed one another because the correspondent values can match.

⁴⁴ A Concept becomes a Component in a DataStructureDefinition, and Components can have different LocalRepresentations in different DataStructureDefinitions, also overriding the (possible) base representation of the Concept.

2352 cannot. In the VTL Transformations, the compliance between the operators and the
 2353 data types of their operands is statically checked, i.e., violations result in compile-
 2354 time errors.

2355 The VTL data types are sub-divided in scalar types (like integers, strings, etc.), which
 2356 are the types of the scalar values, and compound types (like data sets, components,
 2357 rulesets, etc.), which are the types of the compound structures. See below the
 2358 diagram of the VTL data types, taken from the VTL User Manual:

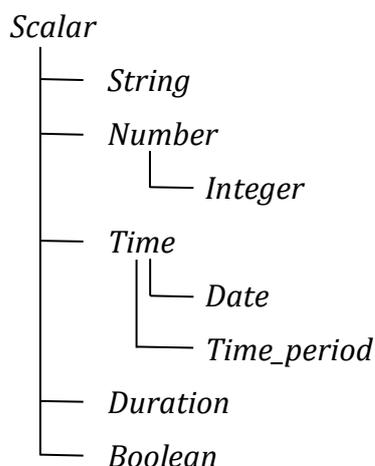


2359
 2360
 2361

Figure 12 – VTL Data Types

2362 The VTL scalar types are in turn subdivided in basic scalar types, which are
 2363 elementary (not defined in term of other data types) and Value Domain and Set
 2364 scalar types, which are defined in terms of the basic scalar types.

2365 The VTL basic scalar types are listed below and follow a hierarchical structure in
 2366 terms of supersets/subsets (e.g. “scalar” is the superset of all the basic scalar types):



2367

2368

Figure 13 – VTL Basic Scalar Types

2369

10.4.2 VTL basic scalar types and SDMX data types

2370

The VTL assumes that a basic scalar type has a unique internal representation and can have more external representations.

2371

2372

The internal representation is the format used within a VTL system to represent (and process) all the scalar values of a certain type. In principle, this format is hidden and not necessarily known by users. The external representations are instead the external formats of the values of a certain basic scalar type, i.e. the formats known by the users. For example, the internal representation of the dates can be an integer counting the days since a predefined date (e.g. from 01/01/4713 BC up to 31/12/5874897 AD like in Postgres) while two possible external representations are the formats YYYY-MM-GG and MM-GG-YYYY (e.g. respectively 2010-12-31 and 12-31-2010).

2373

2374

2375

2376

2377

2378

2379

2380

2381

The internal representation is the reference format that allows VTL to operate on more values of the same type (for example on more dates) even if such values have different external formats: these values are all converted to the unique internal representation so that they can be composed together (e.g. to find the more recent date, to find the time span between these dates and so on).

2382

2383

2384

2385

2386

The VTL assumes that a unique internal representation exists for each basic scalar type but does not prescribe any particular format for it, leaving the VTL systems free to using they preferred or already existing internal format. By consequence, in VTL the basic scalar types are abstractions not associated to a specific format.

2387

2388

2389

2390

SDMX data types are conceived instead to support the data exchange, therefore they do have a format, which is known by the users and correspond, in VTL terms, to external representations. Therefore, for each VTL basic scalar type there can be more SDMX data types (the latter are explained in the section “General Notes for Implementers” of this document and are actually much more numerous than the former).

2391

2392

2393

2394

2395

2396

2397

The following paragraphs describe the mapping between the SDMX data types and the VTL basic scalar types. This mapping shall be presented in the two directions of possible conversion, i.e. from SDMX to VTL and vice-versa.

2398

2399

2400

2401 The conversion from SDMX to VTL happens when an SDMX artefact acts as inputs
2402 of a VTL transformation. As already said, in fact, at compile time the VTL needs to
2403 know the VTL type of the operands in order to check their compliance with the VTL
2404 operators and at runtime it must convert the values from their external (SDMX)
2405 representations to the corresponding internal (VTL) ones.

2406

2407 The opposite conversion, i.e. from VTL to SDMX, happens when a VTL result, i.e. a
2408 VTL data set output of a transformation, must become a SDMX artefact (or part of it).
2409 The values of the VTL result must be converted into the desired (SDMX) external
2410 representations (data types) of the SDMX artefact.

2411 **10.4.3 Mapping SDMX data types to VTL basic scalar types**

2412 The following table describes the default mapping for converting from the SDMX data
2413 types to the VTL basic scalar types.

SDMX data type (BasicComponentDataType)	Default VTL basic scalar type
String (string allowing any character)	string
Alpha (string which only allows A-z)	string
AlphaNumeric (string which only allows A-z and 0-9)	string
Numeric (string which only allows 0-9, but is not numeric so that is can having leading zeros)	string
BigInteger (corresponds to XML Schema xs:integer datatype; infinite set of integer values)	integer
Integer (corresponds to XML Schema xs:int datatype; between -2147483648 and +2147483647 (inclusive))	integer
Long (corresponds to XML Schema xs:long datatype; between -9223372036854775808 and +9223372036854775807 (inclusive))	integer
Short (corresponds to XML Schema xs:short datatype; between -32768 and -32767 (inclusive))	integer
Decimal (corresponds to XML Schema xs:decimal datatype; subset of real numbers that can be represented as decimals)	number
Float (corresponds to XML Schema xs:float datatype; patterned after the IEEE single-precision 32-bit floating point type)	number
Double (corresponds to XML Schema xs:double datatype; patterned after the IEEE double-precision 64-bit floating point type)	number
Boolean (corresponds to the XML Schema xs:boolean datatype; support the mathematical concept of binary-valued logic:	boolean

{true, false}}	
URI (corresponds to the XML Schema xs:anyURI; absolute or relative Uniform Resource Identifier Reference)	string
Count (an integer following a sequential pattern, increasing by 1 for each occurrence)	integer
InclusiveValueRange (decimal number within a closed interval, whose bounds are specified in the SDMX representation by the facets minValue and maxValue)	number
ExclusiveValueRange (decimal number within an open interval, whose bounds are specified in the SDMX representation by the facets minValue and maxValue)	number
Incremental (decimal number the increased by a specific interval (defined by the interval facet), which is typically enforced outside of the XML validation)	number
ObservationalTimePeriod (superset of StandardTimePeriod and TimeRange)	time
StandardTimePeriod (superset of BasicTimePeriod and ReportingTimePeriod)	time
BasicTimePeriod (superset of GregorianTimePeriod and DateTime)	date
GregorianTimePeriod (superset of GregorianYear, GregorianYearMonth, and GregorianDay)	date
GregorianYear (YYYY)	date
GregorianYearMonth / GregorianMonth (YYYY-MM)	date
GregorianDay (YYYY-MM-DD)	date
ReportingTimePeriod (superset of RepostingYear, ReportingSemester, ReportingTrimester, ReportingQuarter, ReportingMonth, ReportingWeek, ReportingDay)	time_period
ReportingYear (YYYY-A1 – 1 year period)	time_period
ReportingSemester (YYYY-Ss – 6 month period)	time_period
ReportingTrimester (YYYY-Tt – 4 month period)	time_period
ReportingQuarter (YYYY-Qq – 3 month period)	time_period
ReportingMonth (YYYY-Mmm – 1 month period)	time_period
ReportingWeek (YYYY-Www – 7 day period; following ISO 8601 definition of a week in a year)	time_period
ReportingDay (YYYY-Dddd – 1 day period)	time_period
DateTime (YYYY-MM-DDThh:mm:ss)	date

TimeRange (YYYY-MM-DD(Th:mm:ss)?/<duration>)	time
Month (--MM; specifies a month independent of a year; e.g. February is black history month in the United States)	string
MonthDay (--MM-DD; specifies a day within a month independent of a year; e.g. Christmas is December 25 th ; used to specify reporting year start day)	string
Day (---DD; specifies a day independent of a month or year; e.g. the 15 th is payday)	string
Time (hh:mm:ss; time independent of a date; e.g. coffee break is at 10:00 AM)	string
Duration (corresponds to XML Schema xs:duration datatype)	duration
XHTML KeyValues IdentifiableReference DataSetReference AttachmentConstraintReference	Metadata type – not applicable Metadata type – not applicable Metadata type – not applicable Metadata type – not applicable Metadata type – not applicable

2414 **Figure 14 – Mappings from SDMX data types to VTL Basic Scalar Types**

2415 When VTL takes in input SDMX artefacts, it is assumed that a type conversion
 2416 according to the table above always happens. In case a different VTL basic scalar
 2417 type is desired, it can be achieved in the VTL program taking in input the default VTL
 2418 basic scalar type above and applying to it the VTL type conversion features (see the
 2419 implicit and explicit type conversion and the “cast” operator in the VTL Reference
 2420 Manual).

2421 **10.4.4 Mapping VTL basic scalar types to SDMX data types**

2422 The following table describes the default conversion from the VTL basic scalar types
 2423 to the SDMX data types .

VTL basic scalar type	Default SDMX data type (BasicComponentDataType)	Default output format
String	String	Like XML (xs:string)
Number	Float	Like XML (xs:float)
Integer	Integer	Like XML (xs:int)
Date	DateTime	YYYY-MM-DDT00:00:00Z
Time	StandardTimePeriod	<date>/<date> (as defined above)
time_period	ReportingTimePeriod (StandardReportingPeriod)	YYYY-Pppp (according to SDMX)
Duration	Duration	Like XML (xs:duration) PnYnMnDTnHnMnS
Boolean	Boolean	Like XML (xs:boolean) with the values “true” or “false”

2424 **Figure 14 – Mappings from SDMX data types to VTL Basic Scalar Types**

2425 In case a different default conversion is desired, it can be achieved through the
2426 `CustomTypeScheme` and `CustomType` artefacts (see also the section
2427 Transformations and Expressions of the SDMX information model).

2428 The custom output formats can be specified by means of the VTL formatting mask
2429 described in the section “Type Conversion and Formatting Mask” of the VTL
2430 Reference Manual. Such a section describes the masks for the VTL basic scalar
2431 types “number”, “integer”, “date”, “time”, “time_period” and “duration” and gives
2432 examples. As for the types “string” and “boolean” the VTL conventions are extended
2433 with some other special characters as described in the following table.

VTL special characters for the formatting masks	
Number	
D	one numeric digit (if the scientific notation is adopted, D is only for the mantissa)
E	one numeric digit (for the exponent of the scientific notation)
.	(dot) possible separator between the integer and the decimal parts.
,	(comma) possible separator between the integer and the decimal parts.
Time and duration	
C	century
Y	year
S	semester
Q	quarter
M	month
W	week
D	day
h	hour digit (by default on 24 hours)
M	minute
S	second
D	decimal of second
P	period indicator (representation in one digit for the duration)
P	number of the periods specified in the period indicator
AM/PM	indicator of AM / PM (e.g. am/pm for “am” or “pm”)
MONTH	uppercase textual representation of the month (e.g., JANUARY for January)
DAY	uppercase textual representation of the day (e.g., MONDAY for Monday)
Month	lowercase textual representation of the month (e.g., january)
Day	lowercase textual representation of the month (e.g., monday)
Month	First character uppercase, then lowercase textual representation of the month (e.g., January)
Day	First character uppercase, then lowercase textual representation of the day using (e.g. Monday)
String	
X	any string character
Z	any string character from “A” to “z”

9	any string character from “0” to “9”
Boolean	
B	Boolean using “true” for True and “false” for False
1	Boolean using “1” for True and “0” for False
0	Boolean using “0” for True and “1” for False
Other qualifiers	
*	an arbitrary number of digits (of the preceding type)
+	at least one digit (of the preceding type)
()	optional digits (specified within the brackets)
\	prefix for the special characters that must appear in the mask
N	fixed number of digits used in the preceding textual representation of the month or the day

2434

2435 The default conversion, either standard or customized, can be used to deduce
2436 automatically the representation of the components of the result of a VTL
2437 transformation. In alternative, the representation of the resulting SDMX Dataflow
2438 can be given explicitly by providing its DataStructureDefinition. In other
2439 words, the representation specified in the DSD, if available, overrides any default
2440 conversion⁴⁵.

2441 **10.4.5 Null Values**

2442 In the conversions from SDMX to VTL it is assumed by default that a missing value in
2443 SDMX becomes a NULL in VTL. After the conversion, the NULLs can be
2444 manipulated through the proper VTL operators.

2445 On the other side, the VTL programs can produce in output NULL values for
2446 Measures and Attributes (Null values are not allowed in the Identifiers). In the
2447 conversion from VTL to SDMX, it is assumed that a NULL in VTL becomes a missing
2448 value in SDMX.

2449 In the conversion from VTL to SDMX, the default assumption can be overridden,
2450 separately for each VTL basic scalar type, by specifying which the value that
2451 represents the NULL in SDMX is. This can be specified in the attribute “nullValue”
2452 of the CustomType artefact (see also the section Transformations and Expressions
2453 of the SDMX information model). A CustomType belongs to a CustomTypeScheme,
2454 which can be referenced by one or more TransformationScheme (i.e. VTL
2455 programs). The overriding assumption is applied for all the SDMX Dataflows
2456 calculated in the TransformationScheme.

2457 **10.4.6 Format of the literals used in VTL transformations**

2458 The VTL programs can contain literals, i.e. specific values of certain data types
2459 written directly in the VTL definitions or expressions. The VTL does not prescribe a
2460 specific format for the literals and leave the specific VTL systems and the definers of
2461 VTL transformations free of using their preferred formats.

⁴⁵ The representation given in the DSD should obviously be compatible with the VTL data type.

2462 Given this discretion, it is essential to know which are the external representations
2463 adopted for the literals in a VTL program, in order to interpret them correctly. For
2464 example, if the external format for the dates is YYYY-MM-DD the date literal 2010-
2465 01-02 has the meaning of 2nd January 2010, instead if the external format for the
2466 dates is YYYY-DD-MM the same literal has the meaning of 1st February 2010.

2467 Hereinafter, i.e. in the SDMX implementation of the VTL, it is assumed that the
2468 literals are expressed according to the “default output format” of the table of the
2469 previous paragraph (“Mapping VTL basic scalar types to SDMX data types”) unless
2470 otherwise specified.

2471 A different format can be specified in the attribute “vtlLiteralFormat” of the
2472 CustomType artefact (see also the section Transformations and Expressions of the
2473 SDMX information model).

2474 Like in the case of the conversion of NULLs described in the previous paragraph, the
2475 overriding assumption is applied, for a certain VTL basic scalar type, if a value is
2476 found for the vtlLiteralFormat attribute of the CustomType of such VTL basic
2477 scalar type. The overriding assumption is applied for all the literals of a related VTL
2478 TransformationScheme.

2479 In case a literal is operand of a VTL Cast operation, the format specified in the Cast
2480 overrides all the possible otherwise specified formats.

2481

2482 **11 Annex I: How to eliminate extra element in the .NET**
 2483 **SDMX Web Service**

2484 **11.1 Problem statement**

2485 For implementing an SDMX compliant Web Service the standardised WSDL file
 2486 should be used that describes the expected request/response structure. The request
 2487 message of the operation contains a wrapper element (e.g. “GetGenericData”) that
 2488 wraps a tag called “GenericDataQuery”, which is the actual SDMX query XML
 2489 message that contains the query to be processed by the Web Service. In the same
 2490 way the response is formulated in a wrapper element “GetGenericDataResponse”.

2491 As defined in the SOAP specification, the root element of a SOAP message is the
 2492 Envelope, which contains an optional Header and a mandatory Body. These are
 2493 illustrated below along with the Body contents according to the WSDL:

```

XML
<SOAP-ENV:Envelope
  <SOAP-ENV:Body>
    <GetGenericData>
      <sdmx:GenericDataQuery>
        ...
      </sdmx:GenericDataQuery>
    </GetGenericData>
  </SOAP-ENV:Body>
</SOAP-ENV:Envelope>
  
```

2494

2495 The problem that initiated the present analysis refers to the difference in the way
 2496 SOAP requests are when trying to implement the aforementioned Web Service in
 2497 .NET framework.

2498 Building such a Web Service using the .NET framework is done by exposing a
 2499 method (i.e. the getGenericData in the example) with an XML document argument
 2500 (lets name it “Query”). **The difference that appears in Microsoft .Net**
 2501 **implementations is that there is a need for an extra XML container around the**
 2502 **SDMX GenericDataQuery.** This is the expected behavior since the framework is let
 2503 to publish automatically the Web Service as a remote procedure call, thus wraps
 2504 each parameter into an extra element. The .NET request is illustrated below:

```

XML
<SOAP-ENV:Envelope
  
```

```

<SOAP-ENV:Body>

  <GetGenericData>

    <Query>          <!-- MS .Net implementation -->

      <GenericDataQuery>

        ...

      </GenericDataQuery>

    </Query>        <!-- MS .Net implementation -->

  </GetGenericData>

</SOAP-ENV:Body>

</SOAP-ENV:Envelope>

```

2505

2506 Furthermore this extra element is also inserted in the automatically generated WSDL
 2507 from the framework. Therefore this particularity requires custom clients for the .NET
 2508 Web Services that is not an interoperable solution.

2509

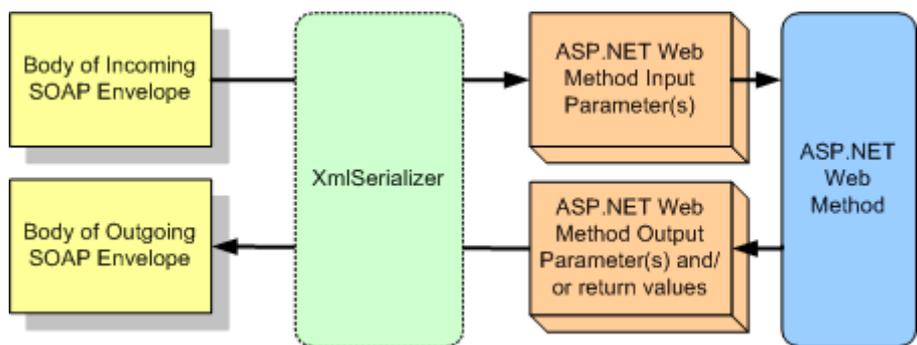
2510 **11.2 Solution**

2511

2512 The solution proposed for conforming the .NET implementation to the envisioned
 2513 SOAP requests has to do with the manual intervention to the serialisation and
 2514 deserialisation of the XML payloads. Since it is a Web Service of already prepared
 2515 XML messages requests/responses this is the indicate way so as to have full control
 2516 on the XML messages. This is the way the Java implementation (using Apache Axis)
 2517 of the SDMX Web Service has adopted.

2518 As regards the .NET platform this is related with the usage of **XmlAnyElement**
 2519 parameter for the .NET web methods.

2520 Web methods use XmlSerializer in the .NET Framework to invoke methods and build
 2521 the response.



2522

2523

2524 The XML is passed to the XmlSerializer to de-serialize it into the instances of classes
2525 in managed code that map to the input parameters for the Web method. Likewise,
2526 the output parameters and return values of the Web method are serialized into XML
2527 in order to create the body of the SOAP response message.

2528 In case the developer wants more control over the serialization and de-serialization
2529 process a solution is represented by the usage of **XmlElement** parameters. This
2530 offers the opportunity of validating the XML against a schema before de-serializing it,
2531 avoiding de-serialization in the first place, analyzing the XML to determine how you
2532 want to de-serialize it, or using the many powerful XML APIs that are available to
2533 deal with the XML directly. This also gives the developer the control to handle errors
2534 in a particular way instead of using the faults that the XmlSerializer might generate
2535 under the covers.

2536 In order to control the de-serialization process of the XmlSerializer for a Web method,
2537 **XmlAnyElement** is a simple solution to use.

2538 To understand how the **XmlAnyElement** attribute works we present the following two
2539 web methods:

C#

```
// Simple Web method using XmlElement parameter  
  
[WebMethod]  
  
public void SubmitXml(XmlElement input)  
  
{ return; }
```

2540

2541 In this method the **input** parameter is decorated with the **XmlAnyElement**
2542 parameter. This is a hint that this parameter will be de-serialized from an **xsd:any**
2543 element. Since the attribute is not passed any parameters, it means that the entire
2544 XML element for this parameter in the SOAP message will be in the Infoset that is
2545 represented by this **XmlElement** parameter.

2546

C#

```
// Simple Web method...using the XmlAnyElement attribute  
  
[WebMethod]  
  
public void SubmitXmlAny([XmlAnyElement] XmlElement input)  
  
{ return; }
```

2547

2548 The difference between the two is that for the first method, **SubmitXml**, the
2549 XmlSerializer will expect an element named **input** to be an immediate child of the

2550 **SubmitXml** element in the SOAP body. The second method, **SubmitXmlAny**, will
2551 not care what the name of the child of the **SubmitXmlAny** element is. It will plug
2552 whatever XML is included into the input parameter. The message style from
2553 ASP.NET Help for the two methods is shown below. First we look at the message for
2554 the method without the **XmlAnyElement** attribute.

2555

XML

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

<soap:Envelope

  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"

  xmlns:xsd="http://www.w3.org/2001/XMLSchema"

  xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">

  <soap:Body>

    <SubmitXml xmlns="http://msdn.microsoft.com/AYS/XEService">

      <input>xml</input>

    </SubmitXml>

  </soap:Body>

</soap:Envelope>
```

2556 Now we look at the message for the method that uses the **XmlAnyElement** attribute.

XML

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

<!-- SOAP message for method using XmlAnyElement -->

<soap:Envelope

  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"

  xmlns:xsd="http://www.w3.org/2001/XMLSchema"

  xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">

  <soap:Body>

    <SubmitXmlAny xmlns="http://msdn.microsoft.com/AYS/XEService">

      Xml

    </SubmitXmlAny>

  </soap:Body>
```

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

2557 The method decorated with the **XmlAnyElement** attribute has one fewer wrapping
2558 elements. Only an element with the name of the method wraps what is passed to the
2559 **input** parameter.

2560 For more information please consult:

2561 <http://msdn.microsoft.com/en-us/library/aa480498.aspx>

2562 Furthermore at this point the problem with the different requests has been solved.
2563 However there is still the difference in the produced WSDL that has to be taken care.
2564 The automatic generated WSDL now doesn't insert the extra element, but defines the
2565 content of the operation wrapper element as "xsd:any" type.

XML

```
<xs:element name="GetGenericData">
  <xs:complexType>
    <xs:sequence>
      <xs:any minOccurs="0" maxOccurs="1" />
    </xs:sequence>
  </xs:complexType>
</xs:element>
```

2566 Without a common WSDL still the solution doesn't enforce interoperability. In order to
2567 "fix" the WSDL, there two approaches. The first is to intervene in the generation
2568 process. This is a complicated approach, compared to the second approach, which
2569 overrides the generation process and returns the envisioned WSDL for the SDMX
2570 Web Service.

2571 This is done by redirecting the request to the "/Service?WSDL" to the envisioned
2572 WSDL stored locally into the application. To do this, from the project add a "Global
2573 Application Class" item (.asax file) and override the request in the
2574 "Application_BeginRequest" method. This is demonstrated in detail in the next
2575 section.

2576 This approach has the disadvantage that for each deployment the WSDL end point
2577 has to be changed to reflect the current URL. However this inconvenience can be
2578 easily eliminated if a developer implements a simple rewriting module for changing
2579 the end point to the one of the current deployment.

2580 **11.3 Applying the solution**

2581 In the context of the SDMX Web Service, applying the above solution translates into
2582 the following:

C#

```
[return: XmlAnyElement]

public XmlDocument GetGenericData([XmlAnyElement]XmlDocument Query)

{ return; }
```

2583 The SOAP request/response will then be as follows:

2584 **GenericData Request**

2585

XML

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

<soap:Envelope xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">

  <soap:Body>

    <GetGenericData xmlns="http://www.sdmx.org/resources/webservices">

      Xml

    </GetGenericData>

  </soap:Body>

</soap:Envelope>
```

2586

2587 **GenericData Response**

2588

XML

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

<soap:Envelope xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">

  <soap:Body>

    <GetGenericDataResponse
xmlns="http://www.sdmx.org/resources/webservices">

      Xml

    </GetGenericDataResponse>

  </soap:Body>

</soap:Envelope>
```

2589 For overriding the automatically produced WSDL, in the solution explorer right click
2590 the project and select "Add" -> "New item...". Then select the "Global Application
2591 Class". This will create ".asax" class file in which the following code should replace
2592 the existing empty method:

C#

```
protected void Application_BeginRequest(object sender, EventArgs e)
{
    System.Web.HttpApplication app = (System.Web.HttpApplication)sender;
    if (Request.RawUrl.EndsWith("/Service1.asmx?WSDL"))
    {
        app.Context.RewritePath("/SDMX_WSDL.wsdl", false);
    }
}
```

2593

2594 The SDMX_WSDL.wsdl should reside in the in the root directory of the application.
2595 After applying this solution the returned WSDL is the envisioned. Thus in the request
2596 message definition contains:

XML

```
<xs:element name="GetGenericData">
  <xs:complexType>
    <xs:sequence>
      <xs:element ref="sdmx:GenericQueryData"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>
```

2597