



Revision History

Revision	Date	Contents
DRAFT 1.0	May 2021	Draft release updated for SDMX 3.0 for public consultation



Contents

1	Pur	pos	e and Structure	6
	1.1	Pur	pose	6
	1.2	Stru	ucture	6
2	Gei	nera	I Notes on This Document	7
3	Gui	ide f	or SDMX Format Standards	8
	3.1	Intro	oduction	8
	3.2	SDI	MX Information Model for Format Implementers	8
	3.2.	.1	Introduction	8
	3.3	SDI	MX Formats: Expressive Capabilities and Function	8
	3.3.	.1	Format Optimizations and Differences	8
	3.4	SDI	MX Best Practices	10
	3.4.	.1	Reporting and Dissemination Guidelines	10
	3.4.	.2	Best Practices for Batch Data Exchange	13
4	Gei	nera	Notes for Implementers	15
	4.1	Rep	presentations	15
	4.1.	.1	Data Types	17
	4.2	Tim	e and Time Format	18
	4.2.	.1	Introduction	18
	4.2.	.2	Observational Time Period	18
	4.2.	.3	Standard Time Period	18
	4.2.	.4	Gregorian Time Period	19
	4.2.	.5	Date Time	19
	4.2.	.6	Standard Reporting Period	19
	4.2.	.7	Distinct Range	22
	4.2.	.8	Time Format	22
	4.2.	.9	Time Zones	23
	4.2.	.10	Representing Time Spans Elsewhere	24
	4.2.	.11	Notes on Formats	24
	4.2.	.12	Effect on Time Ranges	24
	4.2.	.13	Time in Query Messages	24
	4.3	Stru	ctural Metadata Querying Best Practices	26
	4.4	Ver	sioning	27
5	Ref	eren	ce Metadata	29
	5.1	Sco	pe of the Metadata Structure Definition (MSD)	29
	5.2	lder	ntification of the Object(s) to which the Metadata is to be attached	29
	5.3	Met	adata Structure Definition	30
	5.4	Met	adata Set	31
	5.5	Ref	erence Metadata in Data Structure Definition and Dataset	32
6	Coo	delis	t	33
	6.1	Geo	ospatial Codelist	33
	6.1.	.1	Indirect Reference to Geospatial Information.	33
	6.1.	.2	Geographic Coordinates	34
	6.1.	.3	A Geographic Code List	37



	6.2	Codelist extension and discriminated unions	39
	6.2.	1 Prefixing Code Ids	40
	6.2.2	2 Including / Excluding Specific Codes	40
	6.2.3	3 Parent Ids	41
	6.2.4	4 Discriminated Unions	43
	6.3	Linking Hierarchies	44
7	Mai	ntenance Agencies and Metadata Providers	46
8	Con	cept Roles	49
	8.1	Overview	49
	8.2	Information Model	49
	8.3	Technical Mechanism	49
	8.4	SDMX-ML Examples in a DSD	50
	8.5	SDMX standard roles Concept Scheme	51
9	Con	straints	53
	9.1	Introduction	53
	9.2	Types of Constraint	53
	9.3	Rules for a Constraint	54
	9.3.	1 Scope of a Constraint	54
	9.3.2	2 Multiple Content Constraints	54
	9.3.3	3 Inheritance of a Content Constraint	55
	9.3.4	4 Constraints Examples	57
1() Tra	ansforming between versions of SDMX	65
	10.1	Scope	65
	10.2	Compatibility and new DSD features	65
1	l Va	lidation and Transformation Language (VTL)	66
	11.1	Introduction	66
	11.2	References to SDMX artefacts from VTL statements	67
	11.2	2.1 Introduction	67
	11.2	2.2 References through the URN	67
	11.2	2.3 Abbreviation of the URN	69
	11.2	2.4 User-defined alias	72
	11.2	2.5 References to SDMX artefacts from VTL Rulesets	72
	11.3	Mapping between SDMX and VTL artefacts	73
	11.3	8.1 When the mapping occurs	73
	11.3	3.2 General mapping of VIL and SDMX data structures	74
	11.3	3.3 Mapping from SDMX to VIL data structures	74
	11.3	8.4 Mapping from VIL to SDMX data structures	11
	11.3	B.5 Declaration of the mapping methods between data structures	80
	11.3	8.6 Mapping dataflow subsets to distinct VIL Data Sets	81
	11.3	Mapping variables and value domains between VIL and SDMX	86
	11.4	Mapping between SDMX and VIL Data Types	88
	11.4	A.1 VIL Data types	88
	11.4	H.2 VIL basic scalar types and SDIMX data types	90
	11.4	H.3 IVIAPPING SUIVIX data types to VIL basic scalar types	91
	11.4	1.4 Inapping VIL basic scalar types to SDMX data types	93



11		05
11.	4.5 Null Values	
11.	.4.6 Format of the literals used in VIL Transformations	
12 S	structure Mapping	
12.1	Introduction	97
12.2	1-1 structure maps	97
12.3	N-n structure maps	98
12.4	Ambiguous mapping rules	99
12.5	Representation maps	99
12.6	Regular expression and substring rules	100
12.	.6.1 Regular expressions	101
12.	.6.2 Substrings	101
12.7	Mapping non-SDMX time formats to SDMX formats	102
12.	.7.1 Pattern based dates	103
12.	.7.2 Numerical based datetime	106
12.	.7.3 Mapping more complex time inputs	107
12.8	Using TIME PERIOD in mapping rules	107
12.9	Time span mapping rules using validity periods	107
12.10	0 Mapping examples	108
12	10.1 Many to one mapping (N-1)	
12	10.2 Mapping other data types to Code Id	108
12	10.3 Observation Attributes for Time Period	109
12	10.4 Time manning	109
13 Δ	ANNEX Semantic Versioning	111
13 1	Introduction to Semantic Versioning	111
13.1	Semantic Versioning Specification for SDMX 3.0(.0)	111
13.2	Backus, Naur Form Grammar for Valid SDMX 3.0(.0) Somantic V	/orsions113
12.0	Dependency Management in SDMX 2.0(.0):	11/
10.4	Liperade and environment in SDIVIA S.U(.U).	114
IJ.J Vorei	opyrade and conversions of arteracts defined with previous SDM	
12 6	EAO for Somentic Versioning	
13.0		



1 **1 Purpose and Structure**

2 **1.1 Purpose**

The intention of this document is to document certain aspects of SDMX that are important to understand and will aid implementation decisions. The explanations here supplement the information documented in the SDMX XML/JSON schemas and the Information Model.

7 1.2 Structure

- 8 This document is organized into the following major parts:
- 9
- A guide to the SDMX Information Model relating to Data Structure Definitions and Data Sets, statement of differences in functionality supported by the different formats and syntaxes for Data Structure Definitions and Data Sets, and best practices for use of SDMX formats, including the representation for time period.
- A guide to the SDMX Information Model relating to Metadata Structure
 Definitions, and Metadata Sets.
- Other structural artefacts of interest: Agencies, Concept Role, Constraint, Codelist.



18 2 General Notes on This Document

As of version SDMX 2.1, the term "Key family" has been replaced by Data Structure 19 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 specifications. The SDMX-EDI specification, that was using the term "Key family", is 25 26 deprecated in this version of the specification.

27

There has been much work within the SDMX community on the creation of user guides, tutorials, and other aides to implementation and understanding of the standard. This document is not intended to duplicate the function of these documents, but instead

31 represents a short set of technical notes not generally covered elsewhere.

32



33 Guide for SDMX Format Standards

34 **3.1** Introduction

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

41

42 **3.2 SDMX Information Model for Format Implementers**

43 **3.2.1 Introduction**

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

51

52 The SDMX-IM is used to describe the basic data and metadata structures used in all 53 of the SDMX data formats. The Information Model concerns itself with statistical data 54 and its structural metadata, and that is what is described here. Both structural 55 metadata and data have some additional metadata in common, related to their 56 management and administration. These aspects of the data model are not addressed 57 in this section and covered elsewhere in this guide or in the full SDMX-IM document.

58

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

61 **3.3** SDMX Formats: Expressive Capabilities and Function

62 SDMX offers several equivalent formats for describing data and structural metadata, 63 optimized for use in different applications. Although all of these formats are derived 64 directly from the SDMX-IM, and are thus equivalent, the syntaxes used to express the 65 model place some restrictions on their use. Also, different optimizations provide 66 different capabilities. This section describes these differences and provides some rules 67 for applications which may need to support more than one SDMX format or syntax. 68 This section is constrained to the Data Structure Definition and the Date Set.

69 **3.3.1 Format Optimizations and Differences**

The following section provides a brief overview of the differences between the variousSDMX formats.

72

Version 2.0 was characterised by 4 data messages, each with a distinct format: Generic, Compact, Cross-Sectional and Utility. Because of the design, data in some formats could not always be related to another format. In version 2.1, this issue has been addressed by merging some formats and eliminating others. As a result, in SDMX

2.1 there were just two types of data formats: *GenericData* and *StructureSpecificData*



(i.e., specific to one Data Structure Definition). As of SDMX 3.0, based also on the reallife usage of 2.1 XML formats but also the new formats introduced (JSON and CSV),
only one XML format remains, i.e., *StructureSpecificData*. Furthermore, the time
specific sub-formats have also been deprecated due to the lack of usage.

82

SDMX-JSON and SDMX-CSV feature also only one flavour, each. It should be noted,
 though, that both XML and JSON messages allow for series oriented as well as flat
 representations.

86

87 Structure Definition

- The SDMX-ML Structure Message is currently the main way of modelling a DSD.
 The SDMX-JSON version follows the same principles, while the SDMX-CSV does not support structures, yet.
- The SDMX-ML Structure Message allows for the structures on which a Data Structure Definition depends – that is, codelists and concepts – to be either included in the message or to be referenced by the message containing the data structure definition. XML syntax is designed to leverage URIs and other Internetbased referencing mechanisms, and these are used in the SDMX-ML message. This option is also available in SDMX-JSON. The latter, though, further supports conveying data with some structural metadata within a single message.

98 Validation

- The SDMX-ML structure specific messages will allow validation of XML syntax
 and data typing to be performed with a generic XML parser and enforce
 agreement between the structural definition and the data to a moderate degree
 with the same tool.
- Similarly the SDMX-JSON message can be validated using JSON Schema and hence may also be generically parsed and validated.
- The SDMX-CSV format cannot be validated by generic tools.

106 Update and Delete Messages and Documentation Messages

 All messages allow for both append/replace/delete messages and messages consisting of only data or only documentation.

109 Character Encodings

All formats use the UTF-8 encoding. The SDMX-CSV may use a different encoding if this is reported properly in the mime type of a web service response.

112

113 Data Typing

114 The XML syntax and JSON syntax have similar data-typing mechanisms. Hence, there 115 is no need for conventions in order to allow transition from one format to another, like those required for EDIFACT in SDMX 2.1. On the other hand, JSON schema has a 116 117 simpler set of data types (as explained in section 2, paragraph "3.6.3.3 Representation 118 Constructs") but complements its data types with a fixed set of formats or regular 119 expressions. In addition, the JSON schema has also types that are not natively 120 supported in XML schema and need to be implemented as complex types in the latter. 121 The section below provides examples of those cases that are not natively supported 122 by either the XML or JSON data types. More details on the data mapping between 123 XML and JSON schemas are also explained in section "4.1.1 Data Types". 124



125 3.4 SDMX Best Practices

126 3.4.1 Reporting and Dissemination Guidelines

127 3.4.1.1 Central Institutions and Their Role in Statistical Data Exchanges

128 Central institutions are the organisations to which other partner institutions "report" statistics. These statistics are used by central institutions either to compile aggregates 129 130 and/or they are put together and made available in a uniform manner (e.g., on-line or 131 on a CD-ROM or through file transfers). Therefore, central institutions receive data from other institutions and, usually, they also "disseminate" data to individual and/or 132 institutions for end-use. Within a country, a NSI or a national central bank (NCB) plays, 133 134 of course, a central institution role as it collects data from other entities and it 135 disseminates statistical information to end users. In SDMX the role of central institution 136 is very important: every statistical message is based on underlying structural definitions (statistical concepts, code lists, DSDs) which have been devised by a particular 137 138 agency, usually a central institution. Such an institution plays the role of the reference 139 "structural definitions maintenance agency" for the corresponding messages which are 140 exchanged. Of course, two institutions could exchange data using/referring to 141 structural information devised by a third institution.

- 142
- 143 Central institutions can play a double role:
- collecting and further disseminating statistics;
- devising structural definitions for use in data exchanges.

146 3.4.1.2 Defining Data Structure Definitions (DSDs)

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

150

151 Dimensions, Attributes and Codelists

152

167

- Avoid dimensions that are not appropriate for all the series in the data structure definition. If some dimensions are not applicable (this is evident from the need to have a code in a code list which is marked as "not applicable", "not relevant" or "total") for some series then consider moving these series to a new data structure definition in which these dimensions are dropped from the key structure. This is a judgement call as it is sometimes difficult to achieve this without increasing considerably the number of DSDs.
- Devise DSDs with a small number of Dimensions for public viewing of data.
 A DSD with the number dimensions in excess 6 or 7 is often difficult for non-specialist users to understand. In these cases, it is better to have a larger number of DSDs with smaller "cubes" of data, or to eliminate dimensions and aggregate the data at a higher level. Dissemination of data on the web is a growing use case for the SDMX standards: the differentiation of observations by dimensionality, which are necessary for statisticians and economists, are often obscure to public
- Avoid composite dimensions. Each dimension should correspond to a single characteristic of the data, not to a combination of characteristics.

consumers who may not always understand the semantic of the differentiation.



170 171 172 173 174	• C s a s t	Consider the inclusion of the following attributes. Once the key structure of a data tructure definition has been decided, then the set of (preferably mandatory) ttributes of this data structure definition has to be defined. In general, some tatistical concepts are deemed necessary across all Data Structure Definitions of qualify the contained information. Examples of these are:
175 176	0	A descriptive title for the series (this is most useful for dissemination of data for viewing e.g., on the web).
177	0	Collection (e.g., end of period, averaged or summed over period).
178	0	Unit (e.g., currency of denomination).
179	0	Unit multiplier (e.g., expressed in millions).
180	0	Availability (which institutions can a series become available to).
181	0	Decimals (i.e., number of decimal digits used in numerical observations).
182	0	Observation Status (e.g., estimate, provisional, normal).
183 184 185 186	Moreo structu	ver, additional attributes may be considered as mandatory when a specific data are definition is defined.
187 188 189 190 191 192 193	• A re ir o a V V	Avoid creating a new code list where one already exists . It is highly ecommended that structural definitions and code lists be consistent with internationally agreed standard methodologies, wherever they exist, e.g., System of National Accounts 1993; Balance of Payments Manual, Fifth Edition; Monetary and Financial Statistics Manual; Government Finance Statistics Manual, etc. When setting-up a new data exchange, the following order of priority is suggested when considering the use of code lists:
194	0	international standard code lists;
195 196	0	international code lists supplemented by other international and/or regional institutions;
197	0	standardised lists used already by international institutions;
198	0	new code lists agreed between two international or regional institutions;
199	0	new code lists which extend existing code lists, by adding only missing codes;
200	0	new specific code lists.
201 202 203 204 205 206 207 208 209	The sa definit are off the fu dissen suppo	ame code list can be used for several statistical concepts, within a data structure ion or across DSDs. Note that SDMX has recognised that these classifications en quite large and the usage of codes in any one DSD is only a small extract of Il code list. In this version of the standard, it is possible to exchange and ninate a partial code list which is extracted from the full code list and which rts the dimension values valid for a particular DSD.
209	Dala	

- The following items have to be specified by a structural definitions maintenance agency when defining a new data structure definition:
- Data structure definition (DSD) identification:



213	DSD identifier
214	DSD name
215 216	 A list of metadata concepts assigned as dimensions of the data structure definition. For each:
217	(statistical) concept identifier
218 219	 code list identifier (id, version, maintenance agency) if the representation is coded
220 221	 A list of (statistical) concepts assigned as attributes for the data structure definition. For each:
222	(statistical) concept identifier
223	 code list identifier if the concept is coded
224	 assignment status: mandatory, conditional
225	 relationship to dimensions and measures
226	 maximum text length for the uncoded concepts
227	 maximum code length for the coded concepts
228	A list of the code lists used in the data structure definition. For each:
229	code list identifier
230	code list name
231	code values and descriptions
232 233	 Definition of Dataflow. Two (or more) partners performing data exchanges in a certain context need to agree on:
234	 the list of dataset identifiers they will be using;
235	for each Dataflow:
236	 its content (e.g., by Constraints) and description
237 238	 the relevant DSD that defines the structure of the data reported or disseminated according the Dataflow
239	3.4.1.3 Exchanging Attributes

240 3.4.1.3.1 Attributes on series and group levels

- Static properties.
- Upon creation of a series the sender has to provide to the receiver values for all mandatory attributes. In case they are available, values for conditional attributes should also be provided. Whereas initially this information may be provided by means other than SDMX-ML/JSON/CSV messages (e.g., paper, telephone) it is expected that partner institutions will be in a position to provide this information in the available formats over time.
- A centre may agree with its data exchange partners special procedures for authorising the setting of attributes' initial values.



- Communication of changes to the centre.
- Following the creation of a series, the attribute values do not have to be reported again by senders, as long as they do not change.
- Whenever changes in attribute values for a series (or group) occur, the reporting institutions should report either all attribute values again (this is the recommended option) or only the attribute values which have changed. This applies both to the mandatory and the conditional attributes. For example, if a previously reported value for a conditional attribute is no longer valid, this has to be reported to the centre.
- A centre may agree with its data exchange partners special procedures for authorising modifications in the attribute values.
- Communication of observation level attributes "observation status", "observation confidentiality", "observation pre-break" is recommended.
- Whenever an observation is exchanged, the corresponding observation status is
 recommended to also be exchanged attached to the observation, regardless of
 whether it has changed or not since the previous data exchange.
- If the "observation status" changes and the observation remains unchanged, both components would have to be reported (unless the observation is deleted).
- For Data Structure Definitions having also the observation level attributes "observation confidentiality" and "observation pre-break" defined, this rule applies to these attributes as well: if an institution receives from another institution an observation with an observation status attribute only attached, this means that the associated observation confidentiality and pre-break observation attributes either never existed or from now they do not have a value for this observation.

275 **3.4.2 Best Practices for Batch Data Exchange**

276 **3.4.2.1** Introduction

Batch data exchange is the exchange and maintenance of entire databases between
counterparties. It is an activity that often employs SDMX-CSV format, and might also
use the SDMX-ML dataset. The following points apply equally to both formats.

280 3.4.2.2 Positioning of the Dimension "Frequency"

The position of the "frequency" dimension is unambiguously identified in the data structure definition. Moreover, most central institutions devising structural definitions have decided to assign to this dimension the first position in the key structure. Nevertheless, a standard role (i.e., that of 'Frequency') would facilitate the easy identification of this dimension, something that it is necessary to frequency's crucial role in several database systems and in attaching attributes at the "sibling" group level.

287 **3.4.2.3** Identification of Data Structure Definitions (DSDs)

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



293 **3.4.2.4** Identification of the Dataflows

In order to facilitate the easy and immediate recognition of the institution administrating
a Dataflow, many central institutions prefer to use the first characters of the Dataflow
identifiers to identify their institution: e.g. BIS_EER, ECB_BOP1, ECB_BOP1, etc.

297

The statistical information in SDMX is broken down into two fundamental parts -298 299 structural metadata (comprising the DataStructureDefinition, and associated Concepts and Codelists) - see Framework for Standards - and observational data 300 301 (the DataSet). This is an important distinction, with specific terminology associated with each part. Data, which is typically a set of numeric observations at specific points 302 303 in time, is organised into data sets (DataSet). These data sets are structured 304 according to a specific DataStructureDefinition and are described in the 305 Dataflow (via Constraints). The DataStructureDefinition describes the 306 metadata that allows an understanding of what is expressed in the DataSet, whilst 307 the Dataflow provides the identifier and other important information (such as the 308 periodicity of reporting) that is common to all of its Components.

309

Note that the role of the Dataflow and DataSet is very specific in the model, and the terminology used may not be the same as used in all organisations, and specifically the term DataSet is used differently in SDMX than in GESMES/TS. Essentially the GESMES/TS term "Data Set" is, in SDMX, the "Dataflow" whilst the term "Data Set" in SDMX is used to describe the "container" for an instance of the data.

- 315 **3.4.2.5 Special Issues**
- 316 3.4.2.5.1 "Frequency" related issues
- Special frequencies. The issue of data collected at special (regular or irregular)
 intervals at a lower than daily frequency (e.g., 24 or 36 or 48 observations per year, on irregular days during the year) is not extensively discussed here.
 However, for data exchange purposes:
- such data can be mapped into a series with daily frequency; this daily series
 will only hold observations for those days on which the measured event takes
 place;
- if the collection intervals are regular, additional values to the existing
 frequency code list(s) could be added in the future.
- *Tick data.* The issue of data collected at irregular intervals at a higher than daily frequency (e.g., tick-by-tick data) is not discussed here either.



328 4 General Notes for Implementers

This section discusses a number of topics other than the exchange of data sets in SDMX formats. Supported only in SDMX-ML (and some in SDMX-JSON), these topics include the use of the reference metadata mechanism in SDMX, the use of Structure Sets and Reporting Taxonomies, the use of Processes, a discussion of time and datatyping, and some of the conventional mechanisms within the SDMX-ML Structure message regarding versioning and external referencing.

335

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

339 4.1 Representations

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

343

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	<pre>xsd:decimal</pre>	System.Decimal	java.math.BigDecimal
Float	<pre>xsd:float</pre>	System.Single	float
Double	xsd:double	System.Double	double
Boolean	<pre>xsd:boolean</pre>	System.Boolean	boolean
URI	xsd:anyURI	System.Uri	Java.net.URI or java.lang.String
DateTime	xsd:dateTime	System.DateTime	javax.xml.datatype.XMLG regorianCalendar
Time	xsd:time	System.DateTime	javax.xml.datatype.XMLG regorianCalendar
GregorianYear	xsd:gYear	System.DateTime	javax.xml.datatype.XMLG regorianCalendar
GregorianMonth	xsd:gYearMonth	System.DateTime	javax.xml.datatype.XMLG regorianCalendar
GregorianDay	xsd:date	System.DateTime	javax.xml.datatype.XMLG regorianCalendar
Day, MonthDay, Month	xsd:g*	System.DateTime	javax.xml.datatype.XMLG regorianCalendar
Duration	xsd:duration	System.TimeSpan	javax.xml.datatype.Dura tion

344

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



349		
350 351	•	AlphaNumeric (common:AlphaNumericType, string which only allows A-z and 0-9)
352	•	Alpha (common: AlphaType, string which only allows A-z)
353	•	Numeric (common: Numeric Type, String which only allows 0-9, but is not numeric
354		so that is can having leading zeros)
355	•	Count (xs:integer, a sequence with an interval of "1")
356	•	InclusiveValueRange (xs:decimal With the minValue and maxValue facets
357		supplying the bounds)
358	•	ExclusiveValueRange (xs:decimal With the minValue and maxValue facets
359		supplying the bounds)
360	•	Incremental (xs:decimal with a specified interval: the interval is typically
361		enforced outside of the XML validation)
362	•	TimeRange (common:TimeRangeType, startDateTime + Duration)
363	•	ObservationalTimePeriod (common:ObservationalTimePeriodType. a Union Of
364		StandardTimePeriod and TimeRange).
365	•	StandardTimePeriod (common:StandardTimePeriodType. a UNION Of
366		BasicTimePeriod and ReportingTimePeriod).
367	•	BasicTimePeriod (common:BasicTimePeriodType, a Union of
368		GregorianTimePeriod and DateTime)
369	•	GregorianTimePeriod (common:GregorianTimePeriodType, a Union of
370		GregorianYear, GregorianMonth, And GregorianDay)
371	•	ReportingTimePeriod (common:ReportingTimePeriodType, a UNION Of
372		ReportingYear, ReportingSemester, ReportingTrimester, ReportingQuarter,
373		ReportingMonth, ReportingWeek, and ReportingDay).
374	•	ReportingYear (common:ReportingYearType)
375	•	ReportingSemester (common:ReportingSemesterType)
376	•	ReportingTrimester (common:ReportingTrimesterType)
377	•	ReportingQuarter (common:ReportingQuarterType)
378	•	ReportingMonth (common:ReportingMonthType)
379	•	ReportingWeek (common:ReportingWeekType)
380	•	ReportingDay (common:ReportingDayType)
381	•	XHTML (common:StructuredText, allows for multi-lingual text content that has
382		XHTML markup)
383	•	KeyValues (common:DataKeyType)
384	•	<pre>IdentifiableReference (types for each IdentifiableObject)</pre>
385	•	GeographicalInformation (a geo feature set, according to the pattern in
386		section 6.1.2)
387		
388	Data ty	ypes also have a set of facets:
389		
390	•	isSequence = true false (indicates a sequentially increasing value)
391	•	<pre>minLength = positive integer (# Of characters/digits)</pre>
392	•	<pre>maxLength = positive integer (# of characters/digits)</pre>
393	•	startValue = decimal (for numeric sequence)
394	•	endValue = decimal (for numeric sequence)
395	•	interval = decimal (for numeric sequence)
396	•	timeInterval = duration
397	•	<pre>startTime = BasicTimePeriod (for time range)</pre>



Statistical Data and Metadata eXchange

398	• endTime = BasicTimePeriod (for time range)
399	 minValue = decimal (for numeric range)
400	• maxValue = decimal (for numeric range)
401	 decimal = Integer (# of digits to right of decimal point)
402	 pattern = (a regular expression as per W3C XML Schema)
403	• isMultiLingual = boolean (for specifying text can occur in more than one
404	language)
405	languago)
406	Note that code lists may also have textual representations assigned to them, in addition
407	to their enumeration of codes.
408	4.1.1 Data Types
409	XML and JSON schemas support a variety of data types that, although rich, are not
410	mapped one-to-one in all cases. This section provides an explanation of the mapping
411	performed in SDMX 3.0, between such cases.
412	
413	For identifiers, text fields and Codes there are no restriction from either side, since a
414	generic type (e.g., that of string) accompanied by the proper regular expression works
415	equally well for both XML and JSON.
416	
417	For example, for the id type, this is the XML schema definition:
418	<xs:simpletype name="IDType"></xs:simpletype>
419	<pre><xs:restriction base="NestedIDType"></xs:restriction></pre>
420	<pre><xs:pattern value="[A-Za-z0-9_@\$\-]+"></xs:pattern></pre>
421	
422	
423	Where the NestedIDType is also a restriction of string.
424	
425	The above looks like this, in JSON schema:
420	
427 199	vype : Suring",
420 120	pattern : [A-2a-20-9_09-]+9"
429	j
43U 121	There are also cases, though that data types cannot be menned like above. One such
431	THETE ALE AISO CASES, LIDUUT, LIAL UALA LYPES CALINOL DE MAPPEU IRE ADOVE. ONE SUCH

There are also cases, though, that data types cannot be mapped like above. One such case is the array data type, which was introduced in SDMX 3.0 as a new representation. In JSON schema an array is already natively foreseen, while in the XML schema, this has to be defined as a complex type, with an SDMX specific definition (i.e., specific element/attribute names for SDMX). Beyond that, the minimum and/or maximum number of items within an array is possible in both cases.

437

Further to the above, the mapping between the non-native data types is presented in the table below:

SDMX Facet	XML Schema	JSON schema "pattern" ¹ for "string" type
GregorianYear	xsd:gYear	"^-?([1-9][0-9]{3,} 0[0-9]{3})(Z (\+ -)((0[0- 9] 1[0-3]):[0-5][0-9] 14:00))?\$"
GregorianMonth	xsd:gYearMonth	"^-?([1-9][0-9]{3,}]0[0-9]{3})-(0[1-9] 1[0- 2])(Z (\+ -)((0[0-9] 1[0-3]):[0-5][0- 9] 14:00))?\$"

¹ Regular expressions, as specified in <u>W3C XML Schema Definition Language (XSD)</u> <u>1.1 Part 2: Datatypes</u>.



GregorianDay	xsd:date	"^-?([1-9][0-9]{3,} 0[0-9]{3})-(0[1-9] 1[0-2])- (0[1-9] [12][0-9] 3[01])(Z (\+ -)((0[0-9] 1[0- 3]):[0-5][0-9] 14:00))?\$"
Day	xsd:gDay	"^(0[1-9] [12][0-9] 3[01])(Z (\+ -)((0[0-9] 1[0-3]):[0-5][0-9] 14:00))?\$"
MonthDay	xsd:gMonthDay	"^(0[1-9] 1[0-2])-(0[1-9] [12][0- 9] 3[01])(Z (\+ -)((0[0-9] 1[0-3]):[0-5][0- 9] 14:00))?\$"
Month	xsd:Month	"^(0[1-9] 1[0-2])(Z (\+ -)((0[0-9] 1[0- 3]):[0-5][0-9] 14:00))?\$"
Duration	xsd:duration	"^-?P[0-9]+Y?([0-9]+M)?([0-9]+D)?(T([0- 9]+H)?([0-9]+M)?([0-9]+(\.[0-9]+)?S)?)?\$"

440

441

442 4.2 Time and Time Format

443 4.2.1 Introduction

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

446

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

451

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

454 455

456

- **Observational Time Period**
- Standard Time Period
 Basic Time Period
- 457
- 458 459

462

- Gregorian Time Period
 Data Time
- Date Time Reporting Time Period
- 460 461 ○
- Time Range
- The details of these time period categories and of the distinct formats which make them up are detailed in the sections to follow.

465 4.2.2 Observational Time Period

466 This is the superset of all time representations in SDMX. This allows for time to be 467 expressed as any of the allowable formats.

468 4.2.3 Standard Time Period

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



474 **4.2.4 Gregorian Time Period**

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

- 479 **Gregorian Year:**
- 480 Representation: xs:gYear (YYYY)
- 481 Period: the start of January 1 to the end of December 31

482 **Gregorian Year Month**:

- 483 Representation: xs:gYearMonth (YYYY-MM)
- 484 Period: the start of the first day of the month to end of the last day of the month

485 **Gregorian Day**:

- 486 Representation: xs:date (YYYY-MM-DD)
- 487 Period: the start of the day (00:00:00) to the end of the day (23:59:59)

488 **4.2.5 Date Time**

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

493 Representation: xs:dateTime (YYYY-MM-DDThh:mm:ss)²

494 **4.2.6 Standard Reporting Period**

495 Standard reporting periods are periods of time in relation to a reporting year. Each of
496 these standard reporting periods has a duration (based on the ISO 8601 definition)
497 associated with it. The general format of a reporting period is as follows:
498

499 [REPORTING_YEAR]-[PERIOD_INDICATOR][PERIOD_VALUE]

500 501 Where:

- 502 REPORTING_YEAR represents the reporting year as four digits (YYYY) 503 PERIOD_INDICATOR identifies the type of period which determines the duration
- 504 of the period

505

506

- PERIOD VALUE indicates the actual period within the year
- 507 The following section details each of the standard reporting periods defined in SDMX:

508 509 **Reporting Year**:

- 510 Period Indicator: A
- 511 Period Duration: P1Y (one year)
- 512 Limit per year: 1
- 513 Representation: common:ReportingYearType (YYYY-A1, e.g. 2000-A1)
- 514 **Reporting Semester:**
- 515 Period Indicator: S
- 516 Period Duration: P6M (six months)
- 517 Limit per year: 2
- 518 Representation: common:ReportingSemesterType (YYYY-Ss, e.g. 2000-S2)
- 519 **Reporting Trimester:**
- 520 Period Indicator: T

² The seconds can be reported fractionally



521	Period Duration: P4M (four months)
522	Limit per year: 3
523	Representation: common:ReportingTrimesterType (YYYY-Tt, e.g. 2000-T3)
524	Reporting Quarter:
525	Period Indicator: Q
526	Period Duration: P3M (three months)
527	Limit per year: 4
528	Representation: common:ReportingQuarterType (YYYY-Qq, e.g. 2000-Q4)
529	Reporting Month:
530	Period Indicator: M
531	Period Duration: P1M (one month)
532	Limit per year: 1
533	Representation: common:ReportingMonthType (YYYY-Mmm, e.g. 2000-M12)
534	Notes: The reporting month is always represented as two digits, therefore 1-9
535	are 0 padded (e.g. 01). This allows the values to be sorted chronologically using
536	textual sorting methods.
537	Reporting Week:
538	Period Indicator: W
539	Period Duration: P7D (seven days)
540	Limit per year: 53
541	Representation: common:ReportingWeekType (YYYY-Www, e.g. 2000-W53)
542	Notes: There are either 52 or 53 weeks in a reporting year. This is based on the
543	ISO 8601 definition of a week (Monday - Saturday), where the first week of a
544	reporting year is defined as the week with the first Thursday on or after the
545	reporting year start day. ³ The reporting week is always represented as two digits,
546	therefore 1-9 are 0 padded (e.g. 01). This allows the values to be sorted
547	chronologically using textual sorting methods.
548	Reporting Day:
549	Period Indicator: D
550	Period Duration: P1D (one day)
551	Limit per year: 366
552	Representation: common:ReportingDayType (YYYY-Dddd, e.g. 2000-D366)
553	Notes: There are either 365 or 366 days in a reporting year, depending on
554	whether the reporting year includes leap day (February 29). The reporting day is
555	always represented as three digits, therefore 1-99 are 0 padded (e.g. 001). This
556	allows the values to be sorted chronologically using textual sorting methods.
557	
558	The meaning of a reporting year is always based on the start day of the year and
559	requires that the reporting year is expressed as the year at the start of the period. This
560	start day is always the same for a reporting year. and is expressed as a day and a
561	month (e.g. July 1). Therefore, the reporting year 2000 with a start day of July 1 begins
562	on July 1, 2000.

562 563

A specialized attribute (reporting year start day) exists for the purpose of communicating the reporting year start day. This attribute has a fixed identifier (REPORTING_YEAR_START_DAY) and a fixed representation (xs:gMonthDay) so that it can always be easily identified and processed in a data message. Although this attribute exists in specialized sub-class, it functions the same as any other attribute

³ 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.



outside of its identification and representation. It must takes its identity from a concept and state its relationship with other components of the data structure definition. The ability to state this relationship allows this reporting year start day attribute to exist at the appropriate levels of a data message. In the absence of this attribute, the reporting year start date is assumed to be January 1; therefore if the reporting year coincides with the calendar year, this Attribute is not necessary.

576 Since the duration and the reporting year start day are known for any reporting period, it is possible to relate any reporting period to a distinct calendar period. The actual 577 578 Gregorian calendar period covered by the reporting period can be computed as follows [REPROTING YEAR]-579 (based on the standard format of 580 [PERIOD_INDICATOR][PERIOD_VALUE] and the reporting year start day as [REPORTING_YEAR_START_DAY]): 581

582		
583	1.	Determine [REPORTING_YEAR_BASE]:
584	(Combine [REPORTING_YEAR] of the reporting period value (YYYY) with
585	[REPORTING_YEAR_START_DAY] (MM-DD) to get a date (YYYY-MM-DD).
586	Ī	This is the [REPORTING_YEAR_START_DATE]
587		a) If the [PERIOD_INDICATOR] is W:
588		1. If [REPORTING_YEAR_START_DATE] is a Friday,
589		Saturday, or Sunday:
590		Add ⁴ (P3D, P2D, or P1D respectively) to the
591		[REPORTING_YEAR_START_DATE]. The result is the
592		[REPORTING_YEAR_BASE].
593		2. If [REPORTING YEAR START DATE] is a Monday,
594		Tuesday, Wednesday, or Thursday:
595		Add ⁴ (P0D, -P1D, -P2D, or -P3D respectively) to the
596		[REPORTING_YEAR_START_DATE]. The result is the
597		[REPORTING_YEAR_BASE].
598		b) Else:
599		The [REPORTING_YEAR_START_DATE] is the
600		[REPORTING_YEAR_BASE].
601	2.	Determine [PERIOD_DURATION]:
602		a) If the [PERIOD_INDICATOR] is A, the [PERIOD_DURATION] is P1Y.
603		b) If the [PERIOD_INDICATOR] is S, the [PERIOD_DURATION] is P6M.
604		c) If the [PERIOD_INDICATOR] is T, the [PERIOD_DURATION] is P4M.
605		d) If the [PERIOD_INDICATOR] is Q, the [PERIOD_DURATION] is P3M.
606		e) If the [PERIOD_INDICATOR] is M, the [PERIOD_DURATION] is P1M.
607		f) If the [PERIOD_INDICATOR] is W, the [PERIOD_DURATION] is P7D.
608		g) If the [PERIOD_INDICATOR] is D, the [PERIOD_DURATION] is P1D.
609	3.	Determine [PERIOD_START]:
610		Subtract one from the [PERIOD_VALUE] and multiply this by the
611		[PERIOD_DURATION]. Add ⁴ this to the [REPORTING_YEAR_BASE]. The
612		result is the [PERIOD_START].
613	4.	Determine the [PERIOD_END]:
614		Multiply the [PERIOD_VALUE] by the [PERIOD_DURATION]. Add ⁴ this to
615		the [REPORTING_YEAR_BASE] add ⁴ -P1D. The result is the
616		[PERIOD_END].

⁴ The rules for adding durations to a date time are described in the W3C XML Schema specification. See <u>http://www.w3.org/TR/xmlschema-2/#adding-durations-to-dateTimes</u> for further details.



617	
618	For all of these ranges, the bounds include the beginning of the [PERIOD_START]
619	(i.e. 00:00:00) and the end of the [PERIOD_END] (i.e. 23:59:59).
620	
621	Examples:
622	
623	2010-Q2, REPORTING YEAR START DAY =07-01 (July 1)
624	1. IREPORTING YEAR START DATE1 = 2010-07-01
625	b) [REPORTING YEAR BASE] = 2010-07-01
626	2. [PERIOD DURATION] = $P3M$
627	3. $(2-1) * P3M = P3M$
628	2010-07-01 + P3M = 2010-10-01
629	[PERIOD_START] = 2010-10-01
630	4. $2 * P3M = P6M$
631	2010-07-01 + P6M = 2010-13-01 = 2011-01-01
632	2011-01-01 + -P1D = 2010-12-31
633	[PERIOD END] = 2010-12-31
634	
635	The actual calendar range covered by 2010-Q2 (assuming the reporting year
636	begins July 1) is 2010-10-01T00:00:00/2010-12-31T23:59:59
637	5,,
638	2011-W36, REPORTING YEAR START DAY =07-01 (July 1)
639	1. [RÉPORTING YEAR START DATE] = 2010-07-01
640	a) 2011-07-01 = Friday
641	2011-07-01 + P3D = 2011-07-04
642	[REPORTING YEAR BASE] = 2011-07-04
643	2. $[PERIOD_DURATION] = P7D$
644	3. (36-1) * P7D = P245D
645	2011-07-04 + P245D = 2012-03-05
646	[PERIOD_START] = 2012-03-05
647	4. 36 * P7D = P252D
648	2011-07-04 + P252D =2012-03-12
649	2012-03-12 + -P1D = 2012-03-11
650	[PERIOD_END] = 2012-03-11
651	
652	The actual calendar range covered by 2011-W36 (assuming the reporting year
653	begins July 1) is 2012-03-05T00:00/2012-03-11T23:59:59
654	

655 4.2.7 Distinct Range

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

661

662 **4.2.8 Time Format**

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



667

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>)</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)
RW	Reporting Week (YYYY-Www)
RD	Reporting Day (YYYY-Dddd)

668

Table 1: SDMX-ML Time Format Codes

669 4.2.9 Time Zones

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):

- 676
- Time Range (start date): 2006-06-05-05:00/P5D
- Time Range (start date-time): 2006-06-05T00:00:00-05:00/P5D
- Gregorian Year: 2006-05:00
- 680 Gregorian Month: 2006-06-05:00
- 681 Gregorian Day: 2006-06-05-05:00
- Distinct Point: 2006-06-05T00:00:00-05:00
- 683 Reporting Year: 2006-A1-05:00
- Reporting Semester: 2006-S2-05:00
- Reporting Trimester: 2006-T2-05:00
- Reporting Quarter: 2006-Q3-05:00
- 687 Reporting Month: 2006-M06-05:00
- Reporting Week: 2006-W23-05:00



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

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

696 4.2.10 Representing Time Spans Elsewhere

697 It has been possible since SDMX 2.0 for a Component to specify a representation of a 698 time span. Depending on the format of the data message, this resulted in either an 699 element with 2 XML attributes for holding the start time and the duration or two 700 separate XML attributes based on the underlying Component identifier. For example, 701 if REF PERIOD were given a representation of time span, then in the Compact data 702 format, it would be represented by two XML attributes; REF_PERIODStartTime 703 (holding the start) and REF PERIOD (holding the duration). If a new simple type is 704 introduced in the SDMX schemas that can hold ISO 8601 time intervals, then this will 705 no longer be necessary. What was represented as this:

706

690

707 708

711

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

can now be represented with this:710

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

712 4.2.11 Notes on Formats

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

718 4.2.12 Effect on Time Ranges

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

726 4.2.13 Time in Query Messages

When querying for time values, the value of a time parameter can be provided as any of the Observational Time Period formats and must be paired with an operator. This section will detail how systems processing query messages should interpret these parameters.

731

Fundamental to processing a time value parameter in a query message is understanding that all time periods should be handled as a distinct range of time. Since the time parameter in the query is paired with an operator, this also effectively



represents a distinct range of time. Therefore, a system processing the query must
simply match the data where the time period for requested parameter is encompassed
by the time period resulting from value of the query parameter. The following table
details how the operators should be interpreted for any time period provided as a
parameter.

740

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

741

742 Reporting Time Periods as query parameters are handled like this: any data within the 743 bounds of the reporting period for the year is matched, regardless of the actual start 744 day of the reporting year. In addition, data reported against a normal calendar period 745 is matched if it falls within the bounds of the time parameter based on a reporting year 746 start day of January 1. When determining whether another reporting period falls within 747 the bounds of a report period query parameter, one will have to take into account the 748 actual time period to compare weeks and days to higher order report periods. This will 749 be demonstrated in the examples to follow.

750

769

770 771

773

751 Note that the reportingYearStartDay attribute on the time value parameter is only 752 used to qualify a reporting period value for the given time value parameter. The usage 753 of this is different than using the attribute value parameter for the actual reporting year 754 start day attribute. In the case that the attribute value parameters is used for the 755 reporting year start day data structure attribute, it will be treated as any other attribute 756 value parameter; data will be filtered to that which matches the values specified for the 757 given attribute. For example, if the attribute value parameter references the reporting 758 year start day attribute and specifies a value of "--07-01", then only data which has this 759 attribute with the value "--07-01" will be returned. In terms of processing any time value 760 parameters, the value supplied in the attribute value parameter will be irrelevant. 761

762 Examples:

763 764 **Gregorian Period**

- 765 Query Parameter: Greater than 2010
- Literal Interpretation: Any data where the start period occurs after 2010-12-31T23:59:59.
- 768 Example Matches:
 - 2011 or later
 - 2011-01 or later
 - 2011-01-01 or later
- 2011-01-01/P[Any Duration] or any later start date
 - 2011-[Any reporting period] (any reporting year start day)
- 2010-S2 (reporting year start day --07-01 or later)
- 2010-T3 (reporting year start day --07-01 or later)
- 2010-Q3 or later (reporting year start day --07-01 or later)



777	•	2010-M07 or later (reporting year start day07-01 or later)
778	•	2010-W28 or later (reporting year start day07-01 or later)

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

781 **Reporting Period**

779

780

788

789

790

791

792

796

797

798

799

Query Parameter: Greater than or equal to 2010-Q3 782

- 783 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 784 data where the start period is on or after 2010-07-01. 785
- 786 Example Matches: 787
 - 2011 or later •
 - 2010-07 or later •
 - 2010-07-01 or later •
 - 2010-07-01/P[Any Duration] or any later start date
 - 2011-[Any reporting period] (any reporting year start day) •
 - 2010-S2 (any reporting year start day) •
- 793 2010-T3 (any reporting year start day) •
- 2010-Q3 or later (any reporting year start day) 794 •
- 2010-M07 or later (any reporting year start day) 795 •
 - 2010-W27 or later (reporting year start day --01-01)⁵ •
 - 2010-D182 or later (reporting year start day --01-01) •
 - 2010-W28 or later (reporting year start day --07-01)⁶ •
 - 2010-D185 or later (reporting year start day --07-01) •

4.3 Structural Metadata Querying Best Practices 800

801 When guerying for structural metadata, the ability to state how references should be resolved is quite powerful. However, this mechanism is not always necessary and can 802 803 create an undue burden on the systems processing the queries if it is not used properly.

804

Any structural metadata object which contains a reference to an object can be queried 805 806 based on that reference. For example, a categorisation references both a category and 807 the object is it categorising. As this is the case, one can guery for categorisations which categorise a particular object or which categorise against a particular category or 808 809 category scheme. This mechanism should be used when the referenced object is 810 known.

811

When the referenced object is not known, then the reference resolution mechanism 812 813 could be used. For example, suppose one wanted to find all category schemes and the related categorisations for a given maintenance agency. In this case, one could 814 815 query for the category scheme by the maintenance agency and specify that parent and sibling references should be resolved. This would result in the categorisations which 816 817 reference the categories in the matched schemes to be returned, as well as the object

818 which they categorise.

⁵ 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.

⁶ 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.



819 4.4 Versioning

Within the SDMX Structure Message and Medataset, there is a pattern for versioning and external referencing, which should be pointed out. The identifiers are qualified by their version numbers – that is, an object with an Agency of "A", and ID of "X" and a version of "1.0.0" is a different object than one with an Agency of "A", an ID of "X", and a version of "1.1.0".

825

826 As of SDMX 3.0, versioning following the rules of the well-known practice called 827 "Semantic Versioning". Despite that, some use cases do not need or are incompatible with versioning for some or all of their structural artefacts. Therefore, SDMX 3.0 allows 828 829 the 'version' property being nulled for non-versioned artefacts. This means that 830 whenever a structural artefact has no version number than it is not versioned and vice-831 versa. Non-versioned structural artefacts are allowed changing in any way. SDMX 3.0 832 ceases to define a default for the version number, which then needs to be provided by 833 the modeller, even if it is nulled. The OrganisationScheme family of artefacts 834 become truly non-versioned. In addition, the "isFinal" property is removed from 835 MaintainableArtefact.

836

Since the purpose of SDMX versioning is to allow communicating the structural artefact
changes to data exchange partners and connected systems, SDMX 3.0 offers
Semantic Versioning (aka SemVer) with a clear and unambiguous syntax to all
semantically versioned SDMX 3.0 structural artefacts.

841

The semantic version number consists of four parts: MAJOR, MINOR, PATCH and
EXTENSION, the first three parts being separated by a dot (.), the last two parts being
separated by a hyphen (-): MAJOR.MINOR.PATCH-EXTENSION. All versions are
ordered.

847 Given a version number MAJOR.MINOR.PATCH (without EXTENSION), when making
848 changes to that semantically versioned SDMX artefact, then one must increment the:

- 8491. MAJOR version when backwards incompatible artefact changes are850made,
- 851
 852
 2. MINOR version when artefact elements are added in a backwards compatible manner, or
- 8533. PATCH version when backwards compatible artefact property changes854are made.
- 855

When incrementing a version part, the right-hand side parts are 0-ed. Extensions can
be added, changed or dropped.

Given a version number MAJOR.MINOR.PATCH-EXTENSION, when making changes to that semantically versioned SDMX artefact, then one is not required to increment the version if those changes are within the allowed scope of the version increment from the previous version (if that existed); otherwise, the above version increment rules apply. EXTENSIONS can be used e.g., for drafting or a pre-release.

864

New flexible dependency specifications through wildcarding allow for easier data
 model maintenance and enhancements for semantically versioned SDMX artefacts.

507



Semantically versioned SDMX artefacts will be safe to use. Specific version patterns 868 allow them to become either immutable, i.e., the maintainer commits to never change 869 870 their content, or changeable only within a well-defined scope. If any further change is 871 required, a new version must be created first. Furthermore, the impact of the further 872 change is communicated using a clear version increment. This allows implementing a 873 smart referencing mechanism, whereby an artefact may reference, for example, any 874 backward compatible version of another artefact. The built-in version extension facility 875 allows for eased drafting of new SDMX artefact versions.

876

Organisations wishing to keep a maximum of backwards compatibility can continue
using the previous 2-digit convention for version numbers (MAJOR.MINOR), such as
'2.3'. The new SDMX 3.0 standard, though, does not add any guarantees about
changes in those artefacts.

881

The production versions of identifiable objects/resources are assumed to be those that are stable, i.e., they do not have an EXTENSION. This is because once in production, an object cannot change in any way, or it must be versioned. For cases where an object is not static, the version must indicate this by including an EXTENSION. Details on the rules may be found in the section 13. Draft objects should not be used outside of a specific system designed to accommodate them. For most purposes, all objects should become stable before use in production.

889

890 This mechanism is an "early binding" one - everything with a versioned identity is a 891 known quantity and will not change. Nevertheless, Semantic Versioning allows also a "late binding" mechanism where wildcards may be used in references (more in §13). It 892 893 is worth pointing out that in some cases relationships are essentially one-way 894 references: an illustrative case is that of Constraints and flows. While a Constraint may 895 reference many Dataflows or Metadataflows, the addition of more references to flow 896 objects does not version the Constraint. This is because the Constraints are not 897 properties of the flows - they merely make references to them.

898

Versioning operates at the level of versionable and maintainable objects in the SDMX
information model. The Semantic Versioning rules apply here, in order to identify the
version change of an Artefact, as explained in section "13 ANNEX Semantic
Versioning".

904 One area which is much impacted by this versioning scheme is the ability to reference 905 external objects. With the many dependencies within the various structural objects in 906 SDMX, it is useful to have a scheme for external referencing. This is done at the level 907 of maintainable objects (DSDs, Codelists, Concept Schemes, etc.) In an SDMX 908 Structure Message, whenever an "isExternalReference" attribute is set to true, 909 then the application must resolve the address provided in the associated "uri" 910 attribute and use the SDMX Structure Message stored at that location for the full definition of the object in question. Alternately, if a registry "urn" attribute has been 911 provided, the registry can be used to supply the full details of the object. 912

913

Because the version number is part of the identifier for an object, versions are a necessary part of determining that a given resource is the one which was called for. It should be noted that whenever a version number is not supplied, this indicates a nonversioned Artefact.

918



919 5 Reference Metadata

920 **5.1 Scope of the Metadata Structure Definition (MSD)**

The scope of the MSD is reduced in SDMX 3.0, by means of simplifying its structure, but also in the way referenced Artefacts are targeted. In fact, the MSD is restricted to play the role of a single container, without targeting any specific Artefact. The possible targets of Metadata Set are specified in the Metadataflows or Metadata Provision Agreements relating to that MSD. To achieve that, the structure of the Metadataflow has changed in this version of the standard. Moreover, the Metadata Provision Agreement Artefact is introduced to include this feature.

928

929 Two more changes, introduced in this version, are the following:

- The Metadata Set becomes a Maintainable Artefact but maintained by a Metadata
 Provider (another new Artefact in this version).
- Metadata Attributes may also be used in Data Structure Definitions, as long as the latter reference the Metadata Structure Definition that specify those Metadata Attributes.

935

936 5.2 Identification of the Object(s) to which the Metadata is to 937 be attached

The following example shows the structure and naming of the MSD and related
components for creating reference metadata.

941 The schematic structure of an MSD is shown below.

942



943 944

Figure 1: Schematic of the Metadata Structure Definition

945 The MSD contains one Metadata Attribute Descriptor comprising the Metadata 946 Attributes that identify the Concepts for which metadata may be reported in the 947 Metadata Set. The Metadataflow and Metadata Provision Agreement comprise the



948 specification of the objects to which metadata can be reported in a Metadata Set 949 (Metadata Target(s)).

950

951 The high-level view of the MSD, as well as the way the Metadataflow and Metadata 952 Provision Agreement specify the Targets:

953 954

```
<str:MetadataStructure agencyID="SDMX" id="MSD" version="1.0.0-draft">
955
956
         <com:Name>MSD 3.0 sample</com:Name>
         <str:MetadataAttributeDescriptor id="MetadataAttributeDescriptor">
957
958
         </str: MetadataAttributeDescriptor>
959
       </str:MetadataStructure>
```

960 961

```
962
       <str:Metadataflow agencyID="OECD" id="GENERAL METADATA" version="1.0.0-</pre>
963
      draft">
964
        <com:Name xml:lang="en">Metadataflow 3.0 sample</com:Name>
965
        <str:Structure>
966
           <Ref agencyID="OECD" id="MSD" version="1.0.0-draft" />
967
        </str:Structure>
968
        <str:Targets> <!-- Attach to any Dataflows maintained by the OECD -->
969
           <Ref package="dataflow" class="Dataflow" agencyID="OECD" />
970
         </str:Targets>
```

971 </str:Metadataflow>

972

Figure 3: Wildcarded Target Objects as specified in a Metadataflow

Figure 2: The high-level view of the MSD containing one Metadata Attribute Descriptor

973	
974	<pre><str:metadataprovisionagreement <="" agencyid="OECD" id="ABS_INDICATORS" pre=""></str:metadataprovisionagreement></pre>
975	version="1.0.0-draft">
976	<pre><com:name xml:lang="en">Metadata Provision Agreement 3.0 sample</com:name></pre>
977	<str:structureusage></str:structureusage>
978	<pre><ref <="" agencyid="OECD" class="Metadataflow" package="metadatastructure" pre=""></ref></pre>
979	id="GENERAL_METADATA" version="1.0.0-draft" />
980	
981	<pre><str:metadataprovider></str:metadataprovider></pre>
982	<pre><ref <="" agencyid="OECD" id="ABS" maintainableparentid="METADATA_PROVIDERS" pre=""></ref></pre>
983	/>
984	
985	<pre><str:targets> <!-- Attach to specific Dataflows maintained by the OECD--></str:targets></pre>
986	<pre><ref agencyid="OECD" class="Dataflow" id="GDP" package="dataflow"></ref></pre>
987	<pre><ref agencyid="OECD" class="Dataflow" id="EXR" package="dataflow"></ref></pre>
988	<pre><ref agencyid="OECD" class="Dataflow" id="ABC" package="dataflow"></ref></pre>
989	
990	

991

Figure 4: Specific Target Objects as specified in a Metadata Provision Agreement

992 Note that the SDMX-ML schemas have specific XML elements for each identifiable object type because identifying, for instance, a Maintainable Object has different 993 properties from an Identifiable Object which must also include the agencyld, version, 994 995 and id of the Maintainable Object in which it resides.

5.3 Metadata Structure Definition 996

An example is shown below.

1002

<str:MetadataStructure agencyID="SDMX" id="MSD" version="1.0.0-draft"> <com:Name>MSD 3.0 sample</com:Name> <str:MetadataAttributeDescriptor id="MetadataAttributeDescriptor"> <str:MetadataAttribute id="CONTACT" isPresentational="true">



Statistical Data and Metadata eXchange

1003	<pre><str:conceptidentity></str:conceptidentity></pre>
1004	<pre><ref <="" agencyid="SDMX" id="CONTACT" maintainableparentid="CONCEPTS" pre=""></ref></pre>
1005	<pre>maintainableParentVersion="1.0.0"/></pre>
1006	
1007	<pre><str:metadataattribute id="CONTACT_NAME" maxoccurs="1" minoccurs="1"></str:metadataattribute></pre>
1008	<pre><str:conceptidentity></str:conceptidentity></pre>
1009	<pre><ref <="" agencyid="SDMX" maintainableparentid="CONCEPTS" pre=""></ref></pre>
1010	id="CONTACT_NAME" maintainableParentVersion="1.0.0"/>
1011	
1012	<pre><str:localrepresentation> </str:localrepresentation></pre>
1013	<pre><str:textformat texttype="string"></str:textformat> </pre>
1014	<pre> </pre>
1015	<pre></pre> <pre></pre> <pre></pre> <pre></pre> <pre>//str:MetadataAttribute</pre> <pre>//str:MetadataAttribute</pre> <pre>//str:MetadataAttribute</pre>
1017	isPresentational="true">
1018	<pre><str:concentidentity></str:concentidentity></pre>
1019	<pre><ref <="" agencyid="SDMX" id="ADDRESS" maintainableparentid="CONCEPTS" pre=""></ref></pre>
1020	maintainableParentVersion="1.0.0"/>
1021	
1022	<pre><str:metadataattribute <="" id="HOUSE NUMBER" minoccurs="1" pre=""></str:metadataattribute></pre>
1023	maxOccurs="1">
1024	<str:conceptidentity></str:conceptidentity>
1025	<pre><ref <="" agencyid="SDMX" maintainableparentid="CONCEPTS" pre=""></ref></pre>
1026	<pre>id="HOUSE_NUMBER" maintainableParentVersion="1.0.0"/></pre>
1027	
1028	<pre><str:localrepresentation></str:localrepresentation></pre>
1029	<pre><str:textformat texttype="Integer"></str:textformat></pre>
1030	
1031	
1032	
1034	<pre></pre>
1035	<pre></pre>
1036	Figure 5: Example MSD showing specification of some Metadata Attributes
1037	This example shows the following hierarchy of Metadata Attributes:
1038 1039	 Contact – this is presentational; no metadata is expected to be reported at this level
1040	 Contact Name
1041	 Address – this is also presentational; up to 3 addresses are allowed
1042	 House Number
1043	5.4 Metadata Set
4044	An example of examples and a late examplice to the MOD decay's the late is the
1044	An example of reporting metadata according to the MSD described above, is shown

1045 1046

10-0	
1047	<pre><msg:metadataset id="ALB" metadataproviderid="OECD" version="1.0.0"></msg:metadataset></pre>
1048	<str:metadataprovision></str:metadataprovision>
1049	<pre><ref agencyid="OECD" id="ABS_INDICATORS" version="1.0.0-draft"></ref></pre>
1050	
1051	<str:target></str:target>
1052	<pre><ref <="" agencyid="OECD" class="Dataflow" id="GDP" package="dataflow" pre=""></ref></pre>
1053	version="1.0.0" />
1054	
1055	<md:attributeset></md:attributeset>
1056	<md:reportedattribute id="CONTACT"></md:reportedattribute>
1057	<md:attributeset></md:attributeset>
1058	<pre><md:reportedattribute id="CONTACT_NAME">John Doe</md:reportedattribute></pre>
	· _

below.



Statistical Data and Metadata eXchange

1059	
1060	<md:reportedattribute id="ADDRESS"></md:reportedattribute>
1061	<md:attributeset></md:attributeset>
1062	<pre><md:reportedattribute id="STREET NAME"></md:reportedattribute></pre>
1063	<com:text xml:lang="en">5th Avenue</com:text>
1064	
1065	<pre><md:reportedattribute id="HOUSE NUMBER">12</md:reportedattribute></pre>
1066	
1067	
1068	
1069	<md:reportedattribute id="HTML ATTR"></md:reportedattribute>
1070	<pre><com:structuredtext xml:lang="en"></com:structuredtext></pre>
1071	<pre><div xmlns="http://www.w3.org/1999/xhtml"></div></pre>
1072	Lorem Ipsum
1073	
1074	
1075	<pre></pre>
1076	
1077	
1078	
1079	
1080	Figure 6: Example Metadata Set

1081 This example shows:

1082

1083

1089

1092

1093

1096

1097

- 1. The reference to the Metadata Provision Agreement and Metadata Target
 - 2. The reported metadata attributes (AttributeSet)

10845.5ReferenceMetadatainDataStructureDefinitionand1085Dataset

An important change of SDMX 3.0 is the ability to reference an MSD within a DSD, in
order to report any Metadata Attributes defined in the former to Datasets of the latter.
This is achieved by the following:

- In a DSD, the user may add a reference to one MSD.
- In the Attribute Descriptor of the DSD, the user may include any Metadata
 Attributes defined in the linked MSD.
 - For each link to a Metadata Attribute, an Attribute Relationship may be specified (similarly to that for Data Attributes).
- In any Dataset complying with this DSD, Metadata Attributes may be reported according to the specified Attribute Relationship.
 - Any hierarchy of the Metadata Attributes defined in the MSD is ignored; Metadata Attributes are reported as Data Attributes.
- In Data Constraints, the user is allowed to restrict values for Metadata Attributes, in the same way as Data Attributes (more on this in section "9 Constraints").



1101 6 Codelist

As of SDMX 3.0, Codelists have gained new features like geospatial properties, inheritance and extension. Moreover, hierarchies (used to build complex hierarchies of one or more Codelists) are now linked to other Artefacts in order to facilitate the formers' usage in dissemination or other scenarios.

1106

1120

1121 1122

1123

1107 6.1 Geospatial Codelist

SDMX recognizes that statistics refers to units or facts sited in places or areas that may be referenced to geodesic coordinates. This section presents the technical specifications to "geo-reference" those objects and facts in SDMX, by establishing ways to make relations to geographic features over the Earth using a defined coordinates system.

- 1114 SDMX can support three different ways for referencing geospatial data:
- Indirect Reference to Geospatial Information. Including a link to an external file containing the geospatial information. This is the only backwards compatible approach. Since this representation of geospatial information is not included inside the data message, the main use case would be connecting dissemination systems for making use of external tools, like GIS software.
 - Geographic Coordinates. Including the coordinates of a specific geospatial feature as a set of coordinates. This is suitable for any statistical information that needs to be georeferenced especially for the exchange of microdata.
- 1124 5. A Geographic Codelist. Includes a type of Codelist, listing predefined geographies that are represented by geospatial information. These 1125 geographies could be administrative (including administrative 1126 1127 boundaries or enumeration areas), lines, points, or gridded geographies. Regardless, the geospatial information used to represent 1128 1129 the geography would contain both dimensions and/or attributes; therefore, representing an advantage for the data modellers as it 1130 1131 provides a clear way to identify those dimensions developing a 1132 "Geospatial" role.

1133 **6.1.1 Indirect Reference to Geospatial Information.**

1134 This option provides a way to include external references to geospatial information 1135 through a file containing it. The external content may be geographical or thematic maps 1136 with different levels of precision. All the processing of geospatial data is made through 1137 external applications that can interpret the information in different formats.

1138

1139 The reference to the external files containing geospatial information is made using 1140 some recommended SDMX Attributes, with the following content:

- **GEO_INFO_TEXT**. A description of the kind of information being referenced.
- GEO_INFO_URL. A URL which points to the resource containing the referred geospatial information. The resource might be a file with static geodesic information or a web service providing dynamic construction of geometries.
- **GEO_INFO_TYPE**. Coded information describing a standard format of the file that contains the geospatial information. The format types are taken from the list



1147 of Format descriptions for Geospatial Data managed by the US Library of the 1148 Congress

1149

1150

(https://www.loc.gov/preservation/digital/formats/fdd/gis_fdd.shtml). Allowed types in SDMX are listed in the Geographical Formats code list (CL GEO FORMATS). Examples of the codes contained in the document are:

Code	Description
GML	Geography Markup Language
GeoTIFF	GeoTIFF
KML_2_2	KML Version 2.2
GEOJSON_1_1	GeoJSON Version 1.1

1151

1152 Depending on the intended use, these attributes may be attached at the dataflow level, 1153 the series level or the observation level.

1154

1155 The indirect reference to geospatial information in SDMX is recommended to be used 1156 for dissemination purposes, where the statistical information is complemented by geographical representations of places or regions. 1157

1158

6.1.2 Geographic Coordinates 1159

1160 This option to represent geospatial information in SDMX provides an efficient way for including geographic information with different levels of granularity, due to its flexibility. 1161 Geospatial information is represented using the Geospatial Information type, as 1162 defined in the data types of the SDMX Information Model. A "GEO FEATURE SET" role 1163 1164 should be assigned to any Component of this type.

1165

1166 The GeospatialInformation data type can be assigned to a Dimension, with 1167 DataAttribute, MetadataAttribute or а Measure the 1168 "GEO FEATURE SET" role assigned; it can be included in a data message or a metadata report. 1169 1170

1171 Any Component used for representing a Geographical Feature Set, i.e., used to 1172 describe geographical characteristics, must have a "GEO FEATURE SET" role. Its Representation would be of textType="GeospatialInformation". The 1173 1174 GeospatialInformation type is not intended to replace standard geospatial 1175 information formats, but instead provide a simple way to describe precise geographical characteristics in SDMX data sets agnostic of any particular geospatial software 1176 1177 product or use case.

1178

1179 The GeospatialInformation type should be used to describe geographical features like points (e.g., locations of places, landmarks, buildings, etc.), lines (e.g., 1180 rivers, roads, streets, etc.), or areas (e.g., geographical regions, countries, islands, 1181 land lots, etc.). A mix of different features is possible too, e.g., combining polygons and 1182 geographical points to describe a country and the location of its capital. 1183 1184

1185 The components that conform to the structure of the GeospatialInformation type 1186 are:



1187 1188 1189 1190 1191 1192 1193 1194	 X_COORDINATE: The horizontal (longitude) value of a pair of coordinates expressed in the Coordinate Reference System (CRS), mandatory. Y_COORDINATE: The vertical value (latitude) of a pair of coordinates expressed in the CRS units, mandatory. ALT: The height (altitude) from the reference surface is expressed in meters, optional. CRS: The code of the Coordinate Reference System is used to reference the coordinates in the flow, optional.
1195 1196 1197 1198 1199 1200	 The code of the CRS will be as it appears in the EPSG Geodetic Parameter Registry (<u>http://www.epsg-registry.org/</u>) maintained by the International Association of Oil and Gas Producers. If this element is omitted, by default, the CRS will be the World Geodetic System 1984 (WGS 84, EPSG:4326). PRECISION: Precision of the coordinates, expressing the possible deviation in meters from the exact geodesic point, optional.
1201 1202 1203 1204 1205	 This component is only allowed if the CRS is specified too. If omitted, it will be interpreted as limited it to the expected measurement accuracy (e. g. a standard GPS has an accuracy of ~ 10m). GEO_DESCRIPTION: Text for additional information about the place, geographical feature, or set of geographical features, optional.
1206 1207 Ge 1208 int 1209 ob 1210 etc 1211 1212 (G 1213	eographical features (GEO_FEATURES) are collections of geographical features tended to be used to represent geographical areas like countries, regions, etc., or jects, like water bodies (e. g. rivers, lakes, oceans, etc.), roads (streets, highways, c.), hospitals, schools, and the like. They are represented in the following way: GEO_FEATURE, GEO_FEATURE): GEO_DESCRIPTION
1214 1215 1216 1217 1218	 GEO_FEATURE represents a set of points defining a feature following the ISO/IEC 13249-3:2016 standard to conform Well-known Text (WKT) for the representation of geometries in a format defined in the following way: GEOMETY_TYPE (GEOMETRY_REP)
1219 1220 1221 1222 1223 1224 1225 1226 1227 1228 1229	 GEOMETRY_TYPE: A string with a closed vocabulary defining the type of the geometry that represents a geographical component of the GEO_FEATURES collection, mandatory. Three types are allowed: Point, a specific geodesic point, like the centroid of a city or a hospital. It is represented with the string "POINT" Line, a feature defining a line like a road, a river, or similar. It is represented with the string "LINESTRING" Area, a polygon defining a closed area. It is represented with the string
1230 1231	"POLYGON"



	If the GEOMETRY_REP is going to be including the height (ALT) component, a
	"Z" must be added after the string qualifying the GEOMETRY_TYPE. In this way,
	we will have: "POINT Z", "LINESTRING Z" and "POLYGON Z"
	Other feature types (a.g. Triangular irregular petworks, "TIN") are not supported
	vet directly except grids that are detailed in 6.1.3
•	GEOMETRY REP: Representation of each of the types The way to represent
	each GEO FEATURE TYPE will be:
	• A point (POINT): "COORDINATES"
	• A line (LINESTRING): "COORDINATES, COORDINATES,"
	 An area (POLYGON): "(COORDINATES, COORDINATES,),
	(COORDINATES, COORDINATES,)"
Where	
٠	COORDINATES: Represents an individual set of coordinates composed by the
	X_COORDINATE (X), Y_COORDINATE (Y), and ALT (Z) in the following
	way "X Y Z" or "X Y" defining a single point of the polygon. Altitude is to be
	reported in meters.
ام مم د	wooddd wow and an any be represented in the following wover
nane	spanded way, GEO_FEATORE may be represented in the following ways.
LINES Y_COC LINES Y_COC	Z (X_COORDINATE Y_COORDINATE ALT): GEO_DESCRIPTION TRING (X_COORDINATE Y_COORDINATE, X_COORDINATE PRDINATE,): GEO_DESCRIPTION TRING Z (X_COORDINATE Y_COORDINATE ALT, X_COORDINATE RDINATE ALT,): GEO_DESCRIPTION ON ((X_COORDINATE X_COORDINATE X_COORDINATE
LINES Y_COC POLYG Y_COC Y_COC Y_COC Y_COC Y_COC Y_COC X_COC	Z (X_COORDINATE Y_COORDINATE ALT): GEO_DESCRIPTION STRING (X_COORDINATE Y_COORDINATE, X_COORDINATE ORDINATE,): GEO_DESCRIPTION STRING Z (X_COORDINATE Y_COORDINATE ALT, X_COORDINATE ORDINATE ALT,): GEO_DESCRIPTION SON ((X_COORDINATE Y_COORDINATE, X_COORDINATE ORDINATE,), (X_COORDINATE Y_COORDINATE, X_COORDINATE ORDINATE,); GEO_DESCRIPTION SON Z ((X_COORDINATE Y_COORDINATE ALT, X_COORDINATE ORDINATE,), (X_COORDINATE ALT, X_COORDINATE ORDINATE ALT,), (X_COORDINATE ALT, X_COORDINATE ORDINATE ALT,), (X_COORDINATE Y_COORDINATE ALT, ORDINATE ALT,), (X_COORDINATE Y_COORDINATE ALT, ORDINATE ALT,), (X_COORDINATE Y_COORDINATE ALT, ORDINATE Y_COORDINATE ALT,);): GEO_DESCRIPTION
LINES Y_COC LINES Y_COC POLYG Y_COC Y_COC Y_COC Y_COC X_COC	Z (X_COORDINATE Y_COORDINATE ALT): GEO_DESCRIPTION STRING (X_COORDINATE Y_COORDINATE, X_COORDINATE PRDINATE,): GEO_DESCRIPTION STRING Z (X_COORDINATE Y_COORDINATE ALT, X_COORDINATE PRDINATE ALT,): GEO_DESCRIPTION SON ((X_COORDINATE Y_COORDINATE, X_COORDINATE PRDINATE,), (X_COORDINATE Y_COORDINATE, X_COORDINATE PRDINATE,); GEO_DESCRIPTION SON Z ((X_COORDINATE Y_COORDINATE ALT, X_COORDINATE PRDINATE ALT,), (X_COORDINATE ALT, X_COORDINATE PRDINATE ALT,), (X_COORDINATE Y_COORDINATE ALT, PRDINATE ALT,), (X_COORDINATE Y_COORDINATE ALT, PRDINATE ALT,), (X_COORDINATE Y_COORDINATE ALT, PRDINATE Y_COORDINATE ALT,); GEO_DESCRIPTION Ample of how GEO_FEATURES may be represented in an expanded way would
LINES Y_COC POLYG Y_COC Y_COC Y_COC Y_COC Y_COC X_COC X_COC X_COC (Y_COC X_COC (X_CC Y_COC X_COC (X_CC Y_COC Y_COC (X_CC (X_CC)	<pre>2 Z (X_COORDINATE Y_COORDINATE ALT): GEO_DESCRIPTION STRING (X_COORDINATE Y_COORDINATE, X_COORDINATE PRDINATE,): GEO_DESCRIPTION STRING Z (X_COORDINATE Y_COORDINATE ALT, X_COORDINATE PRDINATE ALT,): GEO_DESCRIPTION SON ((X_COORDINATE Y_COORDINATE, X_COORDINATE PRDINATE,), (X_COORDINATE Y_COORDINATE, X_COORDINATE PRDINATE,),): GEO_DESCRIPTION SON Z ((X_COORDINATE Y_COORDINATE ALT, X_COORDINATE PRDINATE ALT,), (X_COORDINATE ALT, X_COORDINATE PRDINATE ALT,), (X_COORDINATE Y_COORDINATE ALT, PRDINATE ALT,), (X_COORDINATE Y_COORDINATE ALT, PRDINATE Y_COORDINATE ALT,); GEO_DESCRIPTION SON Z ((X_COORDINATE ALT,),): GEO_DESCRIPTION STRING GEO_FEATURES may be represented in an expanded way would GON Z ((X_COORDINATE ALT,),), POLYGON Z COORDINATE Y_COORDINATE ALT, X_COORDINATE ALT, PRDINATE Y_COORDINATE ALT, X_COORDINATE ALT, PRDINATE Y_COORDINATE ALT, X_COORDINATE Y_COORDINATE PRDINATE Y_COORDINATE ALT, X_COORDINATE Y_COORDINATE PRDINATE Y_COORDINATE ALT, X_COORDINATE Y_COORDINATE PRDINATE ALT,),), POLYGON Z COORDINATE Y_COORDINATE ALT, X_COORDINATE PRDINATE ALT, X_COORDINATE ALT, X_COORDINATE ALT,), ORDINATE Y_COORDINATE ALT, X_COORDINATE ALT,), ORDINATE Y_COORDINATE ALT, X_COORDINATE ALT,), ORDINATE Y_COORDINATE ALT, X_COORDINATE ALT,); ORDINATE Y_</pre>


Statistical Data and Metadata eXchange

1282 "CRS, PRECISION: { (POLYGON Z (X COORDINATE Y COORDINATE ALT, X_COORDINATE Y_COORDINATE ALT, ...), (X_COORDINATE Y_COORDINATE ALT, X_COORDINATE Y_COORDINATE ALT, ...), ...), POLYGON Z 1283 1284 1285 ((X COORDINATE Y COORDINATE ALT, X COORDINATE Y COORDINATE 1286 ALT, ...), (X COORDINATE Y COORDINATE ALT, X COORDINATE Y_COORDINATE ALT, ...), ...), POLYGON Z ((X COORDINATE 1287 Y COORDINATE ALT, X COORDINATE Y COORDINATE ALT, ...), 1288 (X COORDINATE Y COORDINATE ALT, X COORDINATE Y COORDINATE ALT, 1289 1290 ...), ...); GEO DESCRIPTION}, {(POLYGON Z ((X COORDINATE Y COORDINATE ALT, X COORDINATE Y COORDINATE ALT, ...), 1291 (X COORDINATE Y COORDINATE ALT, X COORDINATE Y COORDINATE ALT, 1292 1293 ...), POLYGON Z ((X COORDINATE Y COORDINATE ALT, 1294 X COORDINATE Y COORDINATE ALT, ...), (X COORDINATE Y COORDINATE 1295 ALT, X COORDINATE Y COORDINATE ALT, ...), POLYGON Z 1296 ((X COORDINATE Y COORDINATE ALT, X COORDINATE Y COORDINATE 1297 ALT, ...), (X COORDINATE Y COORDINATE ALT, X COORDINATE 1298 Y COORDINATE ALT, ...), ...): GEO DESCRIPTION}, ...: 1299 GEO DESCRIPTION" 1300

Validation rules must be added to the XML Schema to ensure the integrity of thespecification according to the proposed syntax.

1303

1304 6.1.3 A Geographic Code List

1305 Geography is represented by geospatial information. Within SDMX, geospatial 1306 information is conceptually represented by the "GEO FEATURE SET" 1307 role/specification. This approach uses a specialized form of SDMX Codelist, named 1308 "GeoCodelist", which is a Codelist containing the Geography used to demarcate the geographic extent. This is implemented in two ways: 1309

- 1310 1. Geographic. It is a regular codelist that has been extended to add a 1311 geographical feature set to each of its items, typically, this would include all 1312 types of administrative geographies;
- 13132. Grid. As a codelist that has defined a geographical grid composed of cells1314 representing regular squared portions of the Earth.

A GeoCodelist is a Codelist as defined in the SDMX Information Model that has the
GeoType property added. GeoType can take one of two values "Geographic" or
"GeoGrid".

1318

1323

"Geographic" corresponds to the first way to implement a GeoCodelist. When the
GeoCodelist includes a GeoType="Geographic" property, a GeoFeatureSet
property is added to each of the items in the code list to implement a Geographic
GeoCodelist.

1324 On the other hand, when GeoType="GeoGrid" it is defining a gridded 1325 GeoCodelist, and it is necessary to add a grid definition to the Codelist identifier 1326 using the gridDefinition property. The components needed to define a 1327 geographical grid are the following:

CRS: The code of the Coordinate Reference System is used to reference the coordinates in the flow, optional. The code of the CRS will be as it appears in the EPSG Geodetic Parameter Registry (<u>http://www.epsg-registry.org/</u>) maintained



- by the International Association of Oil and Gas Producers. If this component is
 omitted, by default the CRS will be the World Geodetic System 1984 (WGS 84,
 EPSG:4326).
- REFERENCE_CORNER: A code composed of two characters to define the position of the coordinates that will be used as a starting reference to locate the cells. The possible values of this code can be UL (Upper Left), UR (Upper Right), LL (Lower 1337 Left), or LR (Lower Right). If this component is omitted the value LL (Lower Left) will be taken by default. This element is optional.
- REFERENCE_COORDINATES: Represents the starting point to reference the cells
 of the grid, accordingly to the CRS and the REFERENCE_CORNER. It is represented
 by an individual set of coordinates composed by the X_COORDINATE (X) and
 Y_COORDINATE (Y) in the following way "X, Y". This element is mandatory if
 GEO_STD is omitted.
- CELL_WIDTH: The size in meters of a horizontal side of the cells in the grid. This
 element is mandatory if GEO_STD is omitted.
- CELL_HEIGHT: The size in meters of a vertical side of the cells in the grid. This
 element is mandatory if GEO_STD is omitted.
- GEO_STD: A restricted text value expressing that the cells in the grid will provide information about matching codes existing in another reference system that establishes a mechanism to define the grid. This element is optional.
- 1351Accepted values for this component are included in the Geographical Grids1352Codelist (CL_GEO_GRIDS). Examples contained in the mentioned document1353are:1354

Value	Description
GEOHASH	GeoHash
GEOREF	World Geographic Reference System
MGRS	Military Grid Reference System
OLC	Open Location Code / Plus Code
QTH	Maidenhead Locator System /QTH Locator / IAURU Locator
W3W	What3words™
WOEID	Where On Earth Identifier

1356The GRID_DEFINITION element will contain a regular expression string1357corresponding to the following format:

- 1358 "CRS: REFERENCE_CORNER; REFERENCE_COORDINATES; CELL_WIDTH,
- 1359 Cell_height: geo_std"
- 1360

1361 In an expanded way we would have:

1362 "CRS:REFERENCE_CORNER; X_COORDINATE, Y_COORDINATE; CELL_WIDTH,

1363 CELL HEIGHT: GEO_STD"

1364

1365 If the grid will be fully adhering to a standard declared in the GEO_STD, the definition 1366 of each code in the code list will be optional. In other case, each item in the code list



must be assigned to one cell in the grid, which is made by adding the GEO_CELL
element to each item of the code list that will contain a regular expression string
composed of the following components:

- GEO COL: The number of the column in the grid starting by zero.
- GEO_ROW: The number of the row in the grid starting by zero.
- **1372** GEO_TAG: An optional text to include additional information to the cell.
- GEO_CELL will have values with the following format: "GEO_COL, GEO_ROW:
 GEO_TAG"

When using a gridded GeoCodelist we may use the GEO_TAG to integrate the cells in the grid to the codes used by other standard defined grids. As an example, GEO_TAG can take the values of the Open Location Codes, GeoHash, etc. If this is done, the GEO_STD component must have been added to the definition of the grid. If the GEO_STD is omitted, the GEO_TAG contents will be taken just as free text.

1381 6.2 Codelist extension and discriminated unions

A Codelist can extend one or more Codelists. Codelist extensions are defined as a list of references to parent Codelists. The order of the references is important when it comes to conflict resolution on Code Ids. When two Codelists have the same Code Id, the Codelist referenced later takes priority. In the example below, the code 'A', exists in both CL_INDICATOR and CL_SERIES. The Codelist CL_INDICATOR_EX will contain the code 'A' from CL_SERIES as this was the second Codelist to be referenced in the sequence of references.



1389 1390

Figure 7: Codelist extension

As the extended Codelist, CL_INDICATOR_EX in this example, may also define its own Codes, these take the ultimate priority over any referenced Codelists. If CL_INDICATOR_EX defines Code 'A', then this will be used instead of Code 'A' from CL_INDICATOR and CL_SERIES, as shown below:





Figure 8: Codelist extension with new Codes

1397

1398 6.2.1 Prefixing Code Ids

A reference to a Codelist may contain a prefix. If a prefix is provided, this prefix will
be applied to all the codes in the Codelist before they are imported into the extended
Codelist. In the above example if the CL_INDICATOR reference includes a prefix of
'IND_' then the resulting Codelist would contain 6 codes, IND_A, IND_X, IND_Y, A,
B, C.



1404 1405

Figure 9: Extended Codelist with prefix

1406 6.2.2 Including / Excluding Specific Codes

The default behaviour of extending another Codelist is to import all of the Codes.
However, an explicit list of Code Ids may be provided for explicit inclusion or exclusion.
This list of Ids may contain wildcards using the same notation as Constraints (%).
Cascading values is also supported using the same syntax as the Constraints. The
list of Ids is either a list of excluded items, or included items, exclusion and inclusion
is not supported against a single Codelist.





1414 Fig

Figure 10: Extended Codelist with include/exclude terms

1415 6.2.3 Parent Ids

Parent Ids are preserved in the extended Codelist if they can be. If a Code is inherited but its parent Code is excluded, then the Code's parent Id will be removed. This rule is also true if the parent Code is excluded because it is overwritten by another Code with the same Id from another Codelist. This ensures the parent Ids do not inadvertently link to Codes originating from different Codelists, and also prevents circular references from occurring.



Figure 11: Parent Code included







Figure 12: Parent Code from different extended Codelist





Figure 13: Parent Code overridden by local Code





Figure 14: Parent Code not included



1430 6.2.4 Discriminated Unions

A common use case solved in SDMX 3.0 is that of discriminated unions, i.e., dealing
with classification or breakdown "variants" which are all valid but mutually exclusive.
For example, there are many versions of the international classification for economic
activities "ISIC". In SDMX, classifications are enumerated concepts, normally modelled
as dimensions corresponding to breakdowns. Each enumerated concept is associated
to one and only one code list.

- 1437
- 1438 To support this use case, the following have to be considered:
- Independent Codelists per variant: Having each variant in a separate Codelist facilitates the maintenance and allows keeping the original codes, even if different versions of the classification have the same code for different concepts. For example, in ISIC Rev. 4 the code "A" represents "Agriculture, forestry and fishing", while in ISIC 3.1 "A" means "Agriculture, hunting and forestry".
- Prefixing Code Ids: When extending Codelists, the reference to an extension Codelist may contain a prefix. If a prefix is provided, this prefix will be applied to all the codes in the Codelist before they are imported into the extended Codelist. In this case, the reference to CL_ISIC4 includes a prefix of "ISIC4_" and the reference to ISIC3 includes "ISIC3_", so the resulting Codelist will have no conflict for the "A" items which will become "ISIC3_A" and "ISIC4 A".
- Including / Excluding Specific Codes: As explained above, there will be independent DFs/PAs with specific Constraint attached, in order to keep the proper items according to the variant in use by each data provider.
- 1455 For example, assuming:
- 1456 DSD DSD_EXDU contains a Dimension: ACTIVITY enumerated by 1457 CL_ACTIVITY.
- 1458 CL_ACTIVITY has no items and is extended by:
- CL_ISIC4, prefix="ISIC4_"
- CL_ISIC3, prefix="ISIC3_"
- 1461 CL NACE2, prefix="NACE2 "
- 1462 CL_AGGR, prefix="AGGR_"
- Dataflow DF1, with a DataConstraint CC_NACE2, CubeRegion for ACTIVITY
 and Value="NACE2_%"
- Dataflow DF2, with a DataConstraint CC_ISIC3, CubeRegion for ACTIVITY
 and Value="ISIC3 %"
- Dataflow DF3, with a DataConstraint CC_ISIC4, CubeRegion for ACTIVITY
 and Value="ISIC4 %", Value="AGGR TOTAL", Value="AGGR Z"

1469

1470 The discriminated unions are achieved, by requesting any of the above Dataflows 1471 references="all" with and detail="referencepartial": returns 1472 with corresponding extensions resolved and the CL ACTIVITY the



1473 DataConstraint, referencing the Dataflow, applied. Thus, the CL_ACTIVITY will 1474 only include Codes prefixed according to the Dataflow, i.e.:

- **Prefix** "NACE2 %" for DF1;
- Prefix "ISIC3 %" for DF2;
- Prefix "ISIC4_%" for DF3; note that Codes "AGGR_TOTAL" and "AGGR_Z" are also included in this case.

1479

1484

1480 6.3 Linking Hierarchies

1481 To facilitate the usage of Hierarchy within other SDMX Artefacts, the
1482 HierarchyAssociation is defined to link any Hierarchy with any
1483 IdentifiableArtefact within a specific context.

1485The HierarchyAssociation is a simple Artefact operating like a1486Categorisation. The former specifies three references:

- The link to a Hierarchy;
- The link to the IdentifiableArtefact that the Hierarchy is linked (e.g., a Dimension);
- The link to the context that the linking is taking place (e.g., a DSD).

1491 As an example, let's assume:

- A DSD with a COUNTRY Dimension that uses Codelist CL_AREA as representation.
- A Hierarchy (e.g., EU_COUNTRIES) that builds a hierarchy for the CL_AREA
 Codelist.

1496 In order to use this Hierarchy for data of a Dataflow (e.g., EU_INDICATORS), we 1497 need to build the following HierarchyAssociation:

- 1498 Links to the Hierarchy EU COUNTRIES (what is associated?)
- Links to the Dimension COUNTRY (where is it associated?)
- 1500 Links to the context: Dataflow EU_INDICATORS (when is it 1501 associated?)
- 1502 The above are also shown in the schematic below:







1506 7 Maintenance Agencies and Metadata Providers

All structural metadata in SDMX is owned and maintained by a maintenance agency (Agency identified by agencyID in the schemas). Similarly, all reference metadata (i.e., MetadataSets) is owned and maintained by a metadata provider organisation (MetadataProvider identified by metadataProviderID in the schemas). It is vital to the integrity of the structural metadata that there are no conflicts in agencyID and metadataProviderID. In order to achieve this, SDMX adopts the following rules:

1513 1514

1515

1516

1517

1518

1525

1526

1527

1528

1533

1535

1537

- 1. Agencies are maintained in an AgencyScheme (which is a sub class of *OrganisationScheme*); Data Providers are maintained in a MetadataProviderScheme.
 - 2. The maintenance agency of the Agency/Metadata Provider Scheme must also be declared in a (different) AgencyScheme.
- 15193. The "top-level" agency is SDMX and this agency scheme is maintained by1520SDMX.
- 4. Agencies registered in the top-level scheme can themselves maintain a single AgencyScheme and a single MetadataProviderScheme. SDMX is an agency in the SDMX AgencyScheme. Agencies in any AgencyScheme can themselves maintain a single AgencyScheme and so on.
 - 5. The AgencyScheme cannot be versioned.
 - 6. There can be only one AgencyScheme maintained by any one Agency. It has a fixed Id of 'AGENCIES'. Similarly, only one MetadataProvideScheme is maintained by one Agency and has a fixed id of 'METADATA PROVIDERS'.
- 7. The format of the agency identifier is agencyId.agencyID etc. The top-level agency in this identification mechanism is the agency registered in the SDMX agency scheme. In other words, SDMX is not a part of the hierarchical ID structure for agencies. SDMX is, itself, a maintenance agency.
- 1534 This supports a hierarchical structure of agencyID.
- 1536 An example is shown below.
 - AA BB CCC DD CCC DD EE



1539	Figure 16: Example of Hierarchic Structure of Agencies
1540 1541	Each agency is identified by its full hierarchy excluding SDMX.
1542 1543	The XML representing this structure is shown below.
1544	<str:organisationschemes></str:organisationschemes>
1545	<pre><str:agencyscheme agencyid="SDMX" id="AGENCIES"></str:agencyscheme></pre>
1546	<pre><com:name xml:lang="en">Top-level Agency Scheme</com:name></pre>
1547	<str:agency id="AA"></str:agency>
1548	<pre><com:name xml:lang="en">AA Name</com:name></pre>
1549	
1550	<str:agency id="BB"></str:agency>
1551	<pre><com:name xml:lang="en">BB Name</com:name></pre>
1552	
1553	
1554	
1555	<pre><str:agencyscheme agencyid="AA" id="AGENCIES"></str:agencyscheme></pre>
1556	<pre><com:name xml:lang="en">AA Agencies</com:name></pre>
1557	<str:agency id="CC"></str:agency>
1558	<pre><com:name xml:lang="en">CC Name</com:name></pre>
1559	
1560	<str:agency id="DD"></str:agency>
1561	<pre><com:name xml:lang="en">DD Name</com:name></pre>
1562	
1563	
1564	
1565	<pre><str:agencyscheme agencyid="BB" id="AGENCIES"></str:agencyscheme></pre>
1566	<pre><com:name xml:lang="en">BB Agencies</com:name></pre>
1567	<str:agency id="CC"></str:agency>
1568	<pre><com:name xml:lang="en">CC Name</com:name></pre>
1569	
1570	<str:agency id="DD"></str:agency>
1571	<pre><com:name xml:lang="en">DD Name</com:name></pre>
1572	
1573	
1574	
1575	<pre><str:agencyscheme agencyid="AA.DD" id="AGENCIES"></str:agencyscheme></pre>
1576	<pre><com:name xml:lang="en">AA.DD Agencies</com:name></pre>
1577	<pre><str:agency id="EE"></str:agency></pre>
1578	<com:name xml:lang="en">EE Name</com:name>
1579	
1580	
1581	
1582	

Figure 17: Example Agency Schemes Showing a Hierarchy

1584 Examples of Structure definitions that show how Agencies are used, are presented 1585 below:

1586	<pre><str:codelist <="" agencyid="SDMX" id="CL_FREQ" pre="" version="1.0.0"></str:codelist></pre>
1587	<pre>urn="urn:sdmx:org.sdmx.infomodel.codelist.Codelist=SDMX:CL_FREQ(1.0.0)"></pre>
1588	<pre><com:name xml:lang="en">Frequency</com:name></pre>
1589	
1590	<pre><str:codelist <="" agencyid="AA" id="CL FREQ" pre="" version="1.0.0"></str:codelist></pre>
1591	urn="urn:sdmx:org.sdmx.infomodel.codelist.Codelist=AA:CL FREQ(1.0.0)">
1592	<pre><com:name xml:lang="en">Frequency</com:name></pre>
1593	
1594	<pre><str:codelist <="" agencyid="AA.CC" id="CL FREQ" pre="" version="1.0.0"></str:codelist></pre>
1595	urn="urn:sdmx:org.sdmx.infomodel.codelist.Codelist=AA.CC:CL FREQ(1.0.0)">
1596	<pre><com:name xml:lang="en">Frequency</com:name></pre>
1597	
1598	<pre><str:codelist <="" agencyid="BB.CC" id="CL FREQ" pre="" version="1.0.0"></str:codelist></pre>
1599	urn="urn:sdmx:org.sdmx.infomodel.codelist.Codelist=BB.CC:CL FREQ(1.0.0)">



1600 1601	<com:name xml:lang="en">Frequency</com:name>
1602	Figure 18: Example Showing Use of Agency Identifiers
1603 1604 1605	Each of these maintenance agencies has a Codelist with an identical id 'CL_FREQ'. However, each is uniquely identified by means of the hierarchic agency structure.



1606 8 Concept Roles

1607 **8.1 Overview**

1608 The DSD Components of Dimension and Attribute can play a specific role in the DSD 1609 and it is important to some applications that this role is specified. For instance, the 1610 following roles are some examples:

- Frequency in a data set the content of this Component contains information on the frequency of the observation values.
- Geography in a data set the content of this Component contains information on the geographic location of the observation values.

1615 8.2 Information Model

1616 The Information Model for this is shown below:



1618 1619

1617

Figure 19: Information Model Extract for Concept Role

1620 It is possible to specify zero or more concept roles for a Dimension, Measure and Data
1621 Attribute (but not the ReportingYearStartDay). The Time Dimension and the Attribute
1622 ReportingYearStartDay have explicitly defined roles and cannot be further specified
1623 with additional concept roles.

1624 8.3 Technical Mechanism

- 1625 The mechanism for maintain and using concept roles is as follows:
- 1626 1627

- 6. A standard Concept Scheme maintained in the Global Registry, with the following identification: SDMX:CONCEPT_ROLES(1.0.0), shall include the default roles, specified by the SDMX SWG (as detailed in 8.5).
- 1629
 1630
 1630
 1631
 1632
 7. Any recognized Agency can have a concept scheme that contains concepts that identify concept roles. Indeed, from a technical perspective any agency can have more than one of these schemes, though this is not recommended.



1633	8.	The	concept	scheme	that	contains	the	"role"	concepts	can	contain
1634		cond	cepts that	do not p	lay a	role.					

- 16359. There is no explicit indication on the Concept whether it is a 'role'1636concept.
- 163710. Therefore, any concept in any concept scheme is capable of being a
'role' concept.
- 163911. It is the responsibility of Agencies to ensure their community knows1640which concepts in which concept schemes play a 'role' and the1641significance and interpretation of this role. In other words, such1642concepts must be known by applications, there is no technical1643mechanism that can inform an application on how to process such a1644'role'.
- 164512. If the concept referenced in the Concept Identity in a DSD component1646(Dimension, Measure Dimension, Attribute) is contained in the concept1647scheme containing concept roles then the DSD component could play1648the role implied by the concept, if this is understood by the processing1649application.
- 165013. If the concept referenced in the Concept Identity in a DSD component
(Dimension, Measure Dimension, Attribute) is not contained in the
concept scheme containing concept roles, and the DSD component is
playing a role, then the concept role is identified by the Concept Role in
the schema.

1655 8.4 SDMX-ML Examples in a DSD

1656 The standard roles Concept Scheme, is still a normal Concept Scheme, thus it may be 1657 used also for the concept identity of a Component, e.g., the 'FREQ':

```
1658
1659
1659
1660
1661
1662
1663
1664

<pr
```

1665 Given this is the standard roles Concept Scheme, any application should interpret the 1666 above Dimension to have the role of Frequency.

1668 Using a Concept Scheme that is not the standard roles Concept Scheme where it is 1669 required to assign a role using the standard roles Concept Scheme. Again, FREQ is 1670 chosen as the example.

1671	<pre><str:dimension id="FREQ"></str:dimension></pre>
1672	<str:conceptidentity></str:conceptidentity>
1673	<pre><ref <="" agencyid="SDMX" id="FREQ" maintainableparentid="CONCEPTS" pre=""></ref></pre>
1674	maintainableParentVersion="1.0.0"/>
1675	
1676	<str:conceptrole></str:conceptrole>
1677	<pre><ref <="" agencyid="SDMX" id="FREO" maintainableparentid="CONCEPT ROLE" pre=""></ref></pre>
1678	maintainableParentVersion="1.0.0"/>
1679	
1680	
1000	() 5 CL . D LING HOLDON

1681



1682 This explicitly states that this Dimension is playing a role identified by the FREQ concept in the standard roles Concept Scheme. Again, the application must interpret 1683 1684 this as a Frequency role.

1685

1686 In other cases where a role from a non-standard roles Concept Scheme is used, then 1687 the application has to know how to interpret the provided roles, e.g., like in the case 1688 below:

1689 1690

<str:Dimension id="FREQ"> <str:ConceptIdentity> <Ref id="FREQ" agencyID="SDMX" maintainableParentID="CONCEPTS" 1691 1692 maintainableParentVersion="1.0.0"/> 1693 </str:ConceptIdentity> 1694 <str:ConceptRole> 1695 <Ref id="FREQ" agencyID="SDMX" maintainableParentID="MY CONCEPT ROLES"</pre> 1696 maintainableParentVersion="1.0.0"/> 1697 </str:ConceptRole> </str:Dimension>

1698 1699

This is all that is required for interoperability within a community. Having a standard 1700 roles Concept Scheme, maintained by the SDMX SWG, allows the SDMX community 1701 to have a common understanding of the roles, while also being able to extend the roles 1702 in bilateral (or multilateral) agreements, by maintaining their own roles Concept 1703 1704 Scheme. This will then ensure there is interoperability between systems that 1705 understand the use of these concepts. 1706

1707 Note that each of the Components (Data Attribute, Measure, Dimension, Time Dimension) has a mandatory identity association (Concept Identity) and if this Concept 1708 1709 also identifies the role then it must be interpreted accordingly.

1710

1711 In order for these roles to be extensible and also to enable user communities to 1712 maintain community-specific roles, the roles are maintained in a controlled vocabulary which is implemented in SDMX as Concepts in a Concept Scheme. The Component 1713 1714 optionally references this Concept if it is required to declare the role explicitly. Note 1715 that a Component can play more than one role and therefore multiple "role" concepts can be referenced. 1716

8.5 SDMX standard roles Concept Scheme 1717

1718 As of this version of the specifications, there is a predefined Concept Scheme, with a 1719 set of Concepts that are considered the standard roles for SDMX. Beyond that, a user is free to add other roles, using custom Concept Schemes. This predefined Concept 1720 1721 Scheme is the result of the SWG guidelines for Concept Roles, plus that for Measure, 1722 and includes the following Concepts:

COMMENT	Comment	Descriptive text which can be attached to data or metadata.
ENTITY	Entity	Describes the subject of the data set (e.g., a country).
FLAG	Flag	Coded attribute in a data set that represents qualitative information for the cell or partial key (e.g. series) value.
FREQ	Frequency	Time interval at which the source data are collected.



GEO	Geographical	Geographic area to which the measured statistical phenomenon relates.
OPERATION	Statistical operation	Signifies statistical operations have been done on the observations.
VARIABLE	Variable	Characteristic of a unit being observed that may assume more than one of a set of values to which a numerical measure or a category from a classification can be assigned.
MEASURE	Measure	Used for emulating the functionality of the deprecated MeasureDimension.
GEO_FEATU RE_SET	Geographical Feature Set	Georeferencing information to describe the location or the shape of a statistical unit, recognizable object or geographical area.
PRIMARY	Primary Measure	Used for backwards compatibility with SDMX 2.1 and back, or when the "Primary Measure" concept is needed.



1725 9 Constraints

- 1726 **9.1** Introduction
- 1727 A Constraint is a Maintainable Artefact that can be associated to one or more of:
- 1728 Data Structure Definition
- 1729 Metadata Structure Definition
- 1730 Dataflow
- Metadataflow
- Provision Agreement
- 1733 Metadata Provision Agreement
- Data Provider or Metadata Provider (this is restricted to a Release Calendar Constraint)
- Simple or Queryable Data Sources
- Note that regardless of the Artefact to which the Constraint is associated, it is constraining the contents of code lists in the DSD to which the constrained object is related. This does not apply, of course, to a Metadata/Data Provider as the latter can be associated, via the (Metadata) Provision Agreement, to many MSDs/DSDs. Hence the reason for the restriction on the type of Constraint that can be attached to a Metadata/Data Provider.

1743 **9.2 Types of Constraint**

- 1744 The Constraint can be of one of two types:
- Data constraint
- Metadata constraint
- 1747
- 1748 The Data Constraint may serve two different perspectives, depending on the way the 1749 latter is retrieved. These are:
- Allowed constraint
- Actual constraint
- The former (allowed also valid for Metadata Constraint) is specified by a data or
 metadata provider or consumer for sharing the allowed data and metadata in the
 context of their DSD or MSD exchanges, e.g., only Monthly data for a specific Dataflow.
 The latter (actual) is a dynamic Constraint in response to an availability request (only
 possible for data).
- 1757

1759

1760

1761 1762

1758 For Actual Data Constraints, there a few characteristics that are worth noting:

- They can only be retrieved by the availability requests (as specified in the REST API).
- They depend on the data available in an SDMX Web Service and thus they can only be dynamically generated according to that data.
- Although they are Maintainable Artefacts, they cannot change independently of data; thus, they cannot be versioned (they are non-versioned, as explained in section 13).



Their identifier may also be dynamically generated and thus there is no REST resource based on their identification.

1768 **9.3** *Rules for a Constraint*

- 1769 **9.3.1 Scope of a Constraint**
- A Constraint is used specify the content of a data or metadata source in terms of thecomponent values or the keys.
- 1773 In terms of data the components are:
- 1774 Dimension
- 1775 Time Dimension
- Data Attribute
- 1777 Measure
- Metadata Attribute
- DataKeySets: the keys are the content of the KeyDescriptor i.e., the series keys composed, for each key, by a value for each Dimension.
- 1781
- 1782 In terms of reference metadata the components are:
- 1783 Metadata Attribute
- 1784

1785 For a Constraint based on a DSD the Constraint can reference one or more of:

- Data Structure Definition
- 1787 Dataflow
- 1788 Provision Agreement
- 1789
- 1790 For a Constraint based on an MSD the Constraint can reference one or more of:
- 1791 Metadata Structure Definition
- Metadataflow
- Metadata Provision Agreement
- 1794
- Furthermore, there can be more than one Constraint specified for a specific object e.g.,more than one Constraint for a specific DSD.
- 1797
- 1798 In view of the flexibility of constraints attachment, clear rules on their usage are1799 required. These are elaborated below.
- 1800 9.3.2 Multiple Content Constraints
- 1801 There can be many Content Constraints for any Constrainable Artefact (e.g., DSD),1802 subject to the following restrictions:



1803 9.3.2.1 Cube Region

- 1804 A Constraint can contain multiple Member Selections (e.g., Dimensions).
- A specific Member Selection (e.g., Dimension FREQ) can only be contained in one Cube Region for any one attached object (e.g., a specific DSD or specific Dataflow).
- If a validity period is specified for Cube Regions, then a specific Member Selection may be repeated across Cube Regions, provided their validity periods are not overlapping.
- For partial reference resolution purposes (as per the SDMX REST API), the latest non-draft Constraint must be considered.
- A Member Selection may include wildcarding of values (using character '%' to represent zero or more occurrences of any character), as well as cascading through hierarchic structures (e.g., parents in Codelist), or localised values (e.g., text for English only). Lack of locale, means any language may match. Cascading values are mutual exclusive to localised values, as the former refer to coded values, while the latter refer to uncoded values.
- Any values included in a Member Selection for Components with an array data type (i.e., Measures, Attributes or Metadata Attributes), will be applied as single values and will not be assessed combined with other values to match all possible array values. For example, including the Code 'A' for an Attribute will allow any instance of the Attribute that includes 'A', like ['A', 'B'] or ['A', 'C', 'D']. Similarly, if Code 'A' was excluded, all those arrays of values would also be excluded.

1825

1826 9.3.2.2 Key Set

1827 Key Sets will be processed in the order they appear in the Constraint and wildcards
1828 can be used (e.g., any key position not reference explicitly is deemed to be "all
1829 values"). As the Key Sets can be "included" or "excluded" it is recommended that Key
1830 Sets with wildcards are declared before KeySets with specific series keys. This will
1831 minimize the risk that keys are inadvertently included or excluded.

1832

1833 In addition, Attribute, Measure and Metadata Attribute constraints may accompany
1834 KeySets, in order to specify the allowed values per Key. Those are expressed following
1835 the rules for Cube Regions, as explained above.

1836 9.3.3 Inheritance of a Content Constraint

1837 9.3.3.1 Attachment levels of a Constraint

1838 There are three levels of constraint attachment for which these inheritance rules apply:

- 1839 DSD/MSD top level
- 1840

- 1841
- Dataflow/Metadataflow second level
 - Provision Agreement third level
- 1843 Note that these rules do not apply to the Simple Datasource or Queryable Datasource;
 1844 the Constraint(s) attached to these artefacts are resolved for this artefact only and do
 1845 not take into account Constraints attached to other artefacts (e.g., Provision
 1846 Agreement, Dataflow, DSD).



1847 It is not necessary for a Constraint to be attached to a higher level artefact. e.g., it is 1848 valid to have a Constraint for a Provision Agreement where there are no constraints 1849 attached the relevant dataflow or DSD.

1850 9.3.3.2 Cascade rules for processing Constraints

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

1855

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

1863 **9.3.3.3 Cube Region**

1864 It is not necessary to have a Constraint on the higher-level artefact (e.g., DSD 1865 referenced by the Dataflow), but if there is such a Constraint at the higher level(s) then:

- The lower-level Constraint cannot be less restrictive than the Constraint specified for the same Member Selection (e.g. Dimension) at the next higher level, which constrains that Member Selection. For example, if the Dimension FREQ is constrained to A, Q in a DSD, then the Constraint at the Dataflow or Provision Agreement cannot be A, Q, M or even just M it can only further constrain A, Q.
- The Constraint at the lower level for any one Member Selection further constrains
 the content for the same Member Selection at the higher level(s).
- Any Member Selection, which is not referenced in a Constraint, is deemed to be constrained according to the Constraint specified at the next higher level which constraints that Member Selection.
- If there is a conflict when resolving the Constraint in terms of a lower-level
 Constraint being less restrictive than a higher-level Constraint, then the
 Constraint at the higher-level is used.
- 1879

1880 Note that it is possible for a Constraint at a higher level to constrain, say, four
1881 Dimensions in a single Constraint, and a Constraint at a lower level to constrain the
1882 same four in two, three, or four Constraints.

1884 9.3.3.4 Key Set

1885 It is not necessary to have a Constraint on the higher-level artefact (e.g., DSD referenced by the Dataflow), but if there is such a Constraint at the higher level(s) then:

- The lower-level Constraint cannot be less restrictive than the Constraint specified at the higher level.
- The Constraint at the lower level for any one Member Selection further constrains the keys specified at the higher level(s).



- Any Member Selection, which is not referenced in a Constraint, is deemed to be constrained according to the Constraint specified at the next higher level which constraints that Member Selection.
- If there is a conflict when resolving the keys in the Constraint at two levels, in terms of a lower-level constraint being less restrictive than a higher-level Constraint, then the offending keys specified at the lower level are not deemed part of the Constraint.

1899 Note that a Key in a Key Set can have wildcarded Components. For instance, the
1900 Constraint may simply constrain the Dimension FREQ to "A", and all keys where the
1901 FREQ="A" are therefore valid.

- 1902
 1903 The following logic explains how the inheritance mechanism works. Note that this is
 1904 conceptual logic and actual systems may differ in the way this is implemented.
- 19051. Determine all possible keys that are valid at the higher level.
- These keys are deemed to be inherited by the lower-level constrained object, subject to the Constraints specified at the lower level.
 - 3. Determine all possible keys that are possible using the Constraints specified at the lower level.
 - 4. At the lower level inherit all keys that match with the higher-level Constraint.
 - 5. If there are keys in the lower-level Constraint that are not inherited then the key is invalid (i.e., it is less restrictive).
- 1914 9.3.4 Constraints Examples

1915 9.3.4.1 Data Constraint and Cascading

- 1916 The following scenario is used.
- 1917

1920

1921

1922

1909 1910

1911

1912

1913

- 1918 A DSD contains the following Dimensions:
- GEO Geography
 - SEX Sex
 - AGE Age
 - CAS Current Activity Status

1923 In the DSD, common code lists are used and the requirement is to restrict these at 1924 various levels to specify the actual code that are valid for the object to which the 1925 Constraint is attached.







1935	
1936	The cascade rules elaborated above result as follows:
1937	
1938	DSD
1939 1940	 Constrained by eliminating code 001 from the code list for the AGE Dimension.
1941	Dataflow CENSUS_CUBE1
1942	Constrained by restricting the code list for the AGE Dimension to codes 002
1943	and 003 (note that this is a more restrictive constraint than that declared for the
1944	DSD which specifies all codes except code 001)
1945	\circ Restricts the CAS codes to 003 and 004
1946	
1947	Dataflow CENSUS_CUBE2
1948	Restricts the code list for the CAS Dimension to codes TOT and NAP.
1949	 Inherits the AGE constraint applied at the level of the DSD.
1950	
1951	Provision Agreement CENSUS_CUBE1_IT
1952	Restricts the codes for the GEO Dimension to IT and its children
1953	 Inherits the constraints from Dataflow CENSUS CUBE1 for the AGE
1954	and CAS Dimensions.
1955	
1956	Provision Agreement CENSUS CUBE2 IT
1957	Restricts the codes for the GEO Dimension to IT and its children.
1958	 Inherits the constraints from Dataflow CENSUS CUBE2 for the CAS
1959	Dimension.
1960	 Inherits the AGE constraint applied at the level of the DSD
1300	
1961	
1961 1962	The Constraints are defined as follows:
1960 1961 1962 1963	The Constraints are defined as follows: DSD Constraint
1960 1961 1962 1963 1964	The Constraints are defined as follows: DSD Constraint <str:dataconstraint agencyid="SDMX" id="DATA_CONSTRAINT" sdmx"="" type="Allowed" version="1.0.0-
draft"></str:dataconstraint>
1960 1961 1962 1963 1964 1965 1966 1967	The Constraints are defined as follows: DSD Constraint <str:dataconstraint agencyid="SDMX" id="DATA_CONSTRAINT" type="Allowed" version="1.0.0-
draft"> <com:name xml:lang="en">SDMX 3.0 Data Constraint sample</com:name> <str:constraintattachment></str:constraintattachment></str:dataconstraint>
1961 1962 1963 1964 1965 1966 1967 1968	The Constraints are defined as follows: DSD Constraint <str:dataconstraint agencyid="SDMX" id="DATA_CONSTRAINT" type="Allowed" version="1.0.0-
draft"> <com:name xml:lang="en">SDMX 3.0 Data Constraint sample</com:name> <str:constraintattachment> <str:datastructure></str:datastructure></str:constraintattachment></str:dataconstraint>
1960 1961 1962 1963 1964 1965 1966 1967 1968 1969	The Constraints are defined as follows: DSD Constraint <str:dataconstraint agencyid="SDMX" id="DATA_CONSTRAINT" type="Allowed" version="1.0.0-
draft"> <com:name xml:lang="en">SDMX 3.0 Data Constraint sample</com:name> <str:constraintattachment> <str:datastructure> <ref agencyid="CENSUSHUB" id="CENSUS" version="1.0.0"></ref></str:datastructure></str:constraintattachment></str:dataconstraint>
1960 1961 1962 1963 1964 1965 1966 1967 1968 1969 1970	The Constraints are defined as follows: DSD Constraint <str:dataconstraint agencyid="SDMX" id="DATA_CONSTRAINT" type="Allowed" version="1.0.0-
draft"> <com:name xml:lang="en">SDMX 3.0 Data Constraint sample</com:name> <str:constraintattachment> <str:datastructure> <ref agencyid="CENSUSHUB" id="CENSUS" version="1.0.0"></ref> </str:datastructure></str:constraintattachment></str:dataconstraint>
1960 1961 1962 1963 1964 1965 1966 1967 1968 1969 1970 1971 1972	The Constraints are defined as follows: DSD Constraint <str:dataconstraint agencyid="SDMX" id="DATA_CONSTRAINT" type="Allowed" version="1.0.0-
draft"> <com:name xml:lang="en">SDMX 3.0 Data Constraint sample</com:name> <str:constraintattachment> <str:datastructure> <fstr:datastructure> </fstr:datastructure></str:datastructure></str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:constraint< td=""></str:constraint<></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:dataconstraint>
1960 1961 1962 1963 1964 1965 1966 1967 1968 1969 1970 1971 1972 1973	The Constraints are defined as follows: DSD Constraint <pre> Str:DataConstraint agencyID="SDMX" id="DATA_CONSTRAINT" version="1.0.0- draft" type="Allowed"></pre>
1960 1961 1962 1963 1964 1965 1966 1967 1968 1969 1970 1971 1972 1973 1974	The Constraints are defined as follows: DSD Constraint <pre> Str:DataConstraint agencyID="SDMX" id="DATA_CONSTRAINT" version="1.0.0- draft" type="Allowed"> <com:name xml:lang="en">SDMX 3.0 Data Constraint sample</com:name> <str:constraintattachment></str:constraintattachment></pre>
1960 1961 1962 1963 1964 1965 1966 1967 1968 1969 1970 1971 1972 1973 1974 1975	The Constraints are defined as follows: DSD Constraint <pre> Str:DataConstraint agencyID="SDMX" id="DATA_CONSTRAINT" version="1.0.0- draft" type="Allowed"></pre>
1960 1961 1962 1963 1964 1965 1966 1967 1968 1969 1970 1971 1972 1973 1974 1975 1976 1977	The Constraints are defined as follows: DSD Constraint <pre></pre>
1960 1961 1962 1963 1964 1965 1966 1967 1968 1969 1970 1971 1972 1973 1974 1975 1976 1977 1978	The Constraints are defined as follows: DSD Constraint <pre></pre>
1960 1961 1962 1963 1964 1965 1966 1967 1968 1969 1970 1971 1972 1973 1974 1975 1976 1977 1978 1979	The Constraints are defined as follows: DSD Constraint <str:dataconstraint agencyid="SDMX" id="DATA_CONSTRAINT" type="Allowed" version="1.0.0-
draft"> <com:name xml:lang="en">SDMX 3.0 Data Constraint sample</com:name> <str:constraintattachment> <str:datastructure> <fstr:datastructure> </fstr:datastructure></str:datastructure></str:constraintattachment> <str:cuberegion include="true"> <!-- the ability to exclude values is illustrated - i.e., all values<br-->valid except this one> <com:keyvalue id="AGE" include="false"> </com:keyvalue> </str:cuberegion> </str:dataconstraint>
1960 1961 1962 1963 1964 1965 1966 1967 1968 1969 1970 1971 1972 1973 1974 1975 1976 1977 1978 1979 1980	<pre>Soft Whitehas the AGE constraint applied at the level of the DSD. The Constraints are defined as follows: DSD Constraint <str:dataconstraint agencyid="SDMX" id="DATA_CONSTRAINT" type="Allowed" version="1.0.0- draft"></str:dataconstraint></pre>
1961 1962 1963 1964 1965 1966 1967 1968 1969 1970 1971 1972 1973 1974 1975 1976 1977 1978 1979 1980 1981	The Constraints are defined as follows: DSD Constraint <str:dataconstraint agencyid="SDMX" id="DATA_CONSTRAINT" type="Allowed" version="1.0.0-
draft"> <com:name xml:lang="en">SDMX 3.0 Data Constraint sample</com:name> <str:constraintattachment> <str:datastructure> <fstr:datastructure> <fstr:constraintattachment> <str:cuberegion include="true"> <!-- the ability to exclude values is illustrated - i.e., all values<br-->valid except this one> <com:keyvalue id="AGE" include="false"> <com:keyvalue> </com:keyvalue></com:keyvalue></str:cuberegion> </fstr:constraintattachment></fstr:datastructure></str:datastructure></str:constraintattachment></str:dataconstraint>
1961 1962 1963 1964 1965 1966 1967 1968 1969 1970 1971 1972 1973 1974 1975 1976 1977 1978 1979 1980 1982	The Constraints are defined as follows: DSD Constraint <str:dataconstraint agencyid="SDMX" id="DATA_CONSTRAINT" type="Allowed" version="1.0.0-
draft"> <com:name xml:lang="en">SDMX 3.0 Data Constraint sample</com:name> <str:constraintattachment> <str:datastructure> </str:datastructure> </str:constraintattachment> <str:cuberegion include="true"> <!-- the ability to exclude values is illustrated - i.e., all values<br-->valid except this one> <com:keyvalue id="AGE" include="false"> <com:keyvalue id="AGE" include="false"> <com:keyvalue> </com:keyvalue></com:keyvalue></com:keyvalue></str:cuberegion> </str:dataconstraint> Dataflow Constraints
1961 1962 1963 1964 1965 1966 1967 1968 1969 1970 1971 1972 1973 1974 1975 1976 1977 1978 1977 1978 1979 1980 1981 1982 1983 1984	The Constraints are defined as follows: DSD Constraint <str:dataconstraint agencyid="SDMX" id="DATA_CONSTRAINT" type="Allowed" version="1.0.0-
draft"> <com:name xml:lang="en">SDMX 3.0 Data Constraint sample</com:name> <str:constraintattachment> <str:constraintattachment> <str:constraintattachment> <str:cuberegion include="true"> <fstr:cuberegion include="true"> <fstr:cuberegion include="true"> <fstr:cuberegion include="true"> <fstr:cuberegion include="true"> <fstr:cuberegion include="true"> <fstr:cuberegion include="true"> <fstr:cuberegion include="true"> <fstr:cuberegion include="true"> <for:keyvalue id="AGE" include="false"> <for:keyvalue id="AGE" include="false"> <for:keyvalue> <fstr:cuberegion> <fstr:dataconstraint> Dataflow Constraints </fstr:dataconstraint></fstr:cuberegion></for:keyvalue></for:keyvalue></for:keyvalue></fstr:cuberegion></fstr:cuberegion></fstr:cuberegion></fstr:cuberegion></fstr:cuberegion></fstr:cuberegion></fstr:cuberegion></fstr:cuberegion></str:cuberegion></str:constraintattachment></str:constraintattachment></str:constraintattachment></str:dataconstraint>
1961 1962 1963 1964 1965 1966 1967 1968 1969 1970 1971 1972 1973 1974 1975 1976 1977 1978 1979 1980 1981 1982 1984 1985	The Constraints are defined as follows: DSD Constraint <tr:dataconstraint agencyid="SDMX" id="DATA_CONSTRAINT" type="Allowed" version="1.0.0-
draft"> <com:name xml:lang="en">SDMX 3.0 Data Constraint sample</com:name> <str:constraintattachment> <str:datastructure> </str:datastructure> <itr:cuberegion include="true"> <itr:cuberegion include="true"> <com:keyvalue id="AGE" include="false"> <com:keyvalue id="AGE" include="false"> </com:keyvalue> </com:keyvalue></itr:cuberegion></itr:cuberegion></str:constraintattachment></tr:dataconstraint>
1961 1962 1963 1964 1965 1966 1967 1968 1969 1970 1971 1972 1973 1974 1975 1976 1977 1978 1980 1981 1982 1983 1984 1985 1986	The Constraints are defined as follows: DSD Constraint <tr:dataconstraint agencyid="SDMX" id="DATA_CONSTRAINT" type="Allowed" version="1.0.0-
draft"> <com:name xml:lang="en">SDMX 3.0 Data Constraint sample</com:name> <str:constraintattachment> <str:datastructure> <fstr:constraintattachment> <str:cuberegion include="true"> <fstr:cuberegion include="true"> <form:keyvalue id="NGE" include="false"> <form:keyvalue id="NGE" include="false"> <form:keyvalue id="NGE" include="false"> <form:keyvalue id="NGE" include="false"> <form:keyvalue id="NGE" include="false"> <form:keyvalue id="NGE" include="false"> <form:keyvalue> <fstr:dataconstraint> Dataflow Constraints </fstr:dataconstraint></form:keyvalue></form:keyvalue></form:keyvalue></form:keyvalue></form:keyvalue></form:keyvalue></form:keyvalue></fstr:cuberegion></fstr:cuberegion></fstr:cuberegion></fstr:cuberegion></fstr:cuberegion></fstr:cuberegion></fstr:cuberegion></fstr:cuberegion></fstr:cuberegion></fstr:cuberegion></fstr:cuberegion></str:cuberegion></fstr:constraintattachment></str:datastructure></str:constraintattachment></tr:dataconstraint>
1961 1962 1963 1964 1965 1966 1967 1968 1969 1970 1971 1972 1973 1974 1975 1976 1977 1978 1979 1980 1981 1983 1984 1985 1986 1987 1988 1988 1987 1988 <t< td=""><td>The Constraints are defined as follows: DSD Constraint <pre> Str:DataConstraint agencyID="SDMX" id="DATA_CONSTRAINT" version="1.0.0- draft" type="Allowed"></pre></td></t<>	The Constraints are defined as follows: DSD Constraint <pre> Str:DataConstraint agencyID="SDMX" id="DATA_CONSTRAINT" version="1.0.0- draft" type="Allowed"></pre>
1961 1962 1963 1964 1965 1966 1967 1968 1967 1968 1969 1970 1971 1972 1973 1974 1975 1976 1977 1978 1977 1978 1979 1980 1981 1983 1984 1985 1986 1987 1988 1987 1988 1989	The Constraints are defined as follows: DSD Constraint <pre> Str:DataConstraint agencyID="SDMX" id="DATA_CONSTRAINT" version="1.0.0- draft" type="Allowed"></pre>



Statistical Data and Metadata eXchange

1991 1992 1993 1994 1995 1996 1997 1998	<pre><com:keyvalue id="AGE" include="true"> <com:value>002</com:value> <com:value>003</com:value> </com:keyvalue> <com:keyvalue id="CAS"> <com:value>003</com:value> <com:value>003</com:value> <com:value>004</com:value> </com:keyvalue></pre>
2000 2001	
2002 2003	<pre><str:dataconstraint agencyid="SDMX" id="DATA_CONSTRAINT_3" type="Allowed" version="1.0.0-
draft"></str:dataconstraint></pre>
2004 2005	<pre><com:name xml:lang="en">SDMX 3.0 Data Constraint sample</com:name> <str:constraintattachment></str:constraintattachment></pre>
2006 2007	<pre><str:dataflow></str:dataflow></pre>
2008 2009	
2010 2011	<pre><str:cuberegion include="true"> <com:keyvalue id="CAS" include="true"></com:keyvalue></str:cuberegion></pre>
2012	<pre><com:value>TOT</com:value> <com:value>NAP</com:value></pre>
2014 2015	
2010	
2018	Provision Agreement Constraint
2019	<pre>draft" type="Allowed"></pre>
2021 2022	<com:name xml:lang="en">SDMX 3.0 Data Constraint sample</com:name> <str:constraintattachment></str:constraintattachment>
2023 2024	<pre><str:provisionagreement></str:provisionagreement></pre>
2025 2026	 <str:provisionagreement></str:provisionagreement>

2024	<pre><ref agencyid="CENSUSHUB" id="CENSUS CUBE1 IT" version="1.0.0"></ref></pre>
2025	
2026	<str:provisionagreement></str:provisionagreement>
2027	<pre><ref agencyid="CENSUSHUB" id="CENSUS CUBE2 IT" version="1.0.0"></ref></pre>
2028	
2029	
2030	<pre><str:cuberegion include="true"></str:cuberegion></pre>
2031	<pre><com:keyvalue id="GEO" include="true"></com:keyvalue></pre>
2032	<pre><com:value cascadevalues="true">IT</com:value></pre>
2033	
2034	
2035	

2036

2037 9.3.4.2 Combination of Constraints

2038 The possible combination of constraining terms are explained in this section, following 2039 a few examples.

2040

2041 Let's assume a DSD with the following Components:

Dimension	FREQ
Dimension	JD_TYPE
Dimension	JD_CATEGORY
Dimension	VIS_CTY
TimeDimension	TIME_PERIOD
Attribute	OBS_STATUS



Statistical Data and Metadata eXchange

Attribute	UNIT
Attribute	COMMENT
MetadataAttribute	CONTACT
Measure	MULTISELECT
Measure	CHOICE

2042

2043 On the above, let's assume the following use cases with their constraining 2044 requirements:

2045 9.3.4.2.1 Use Case 1: A Constraint on allowed values for some Dimensions

- 2046 R1: Allow monthly and quarterly data
- 2047 R2: Allow Mexico for vis-à-vis country

2048

2049 This is expressed with the following CubeRegion:

FREQ	M, Q
VIS_CTY	MX

20509.3.4.2.2Use Case 2: A Constraint on allowed combinations for some2051Dimensions

- 2052 R1: Allow monthly data for Germany
- 2053 R2: Allow quarterly data for Mexico
- 2054

2055 This is expressed with the following DataKeySet:

Key1	FREQ	Μ
	VIS_CTY	DE
Key2	FREQ	Q
	VIS_CTY	MX

20569.3.4.2.3Use Case 3: A Constraint on allowed values for some Dimensions2057combined with allowed values for some Attributes

- 2058 R1: Allow monthly and quarterly data
- 2059 R2: Allow Mexico for vis-à-vis country
- 2060 R3: Allow present for status
- 2061

2062 This may be expressed with the following CubeRegion:

FREQ	M, Q
VIS_CTY	MX
OBS_STATUS	A

20639.3.4.2.4Use Case 4: A Constraint on allowed combinations for some2064Dimensions combined with specific Attribute values

- 2065 R1: Allow monthly data, for Germany, with unit euro
- 2066 R2: Allow quarterly data, for Mexico, with unit usd
- 2067

2068 This is may be expressed with the following DataKeySet:

|--|



	VIS_CTY	DE
	UNIT	EUR
Key2	FREQ	Q
	VIS_CTY	MX
	UNIT	USD

20699.3.4.2.5Use Case 5: A Constraint on allowed values for some Dimensions2070together with some combination of Dimension values

2071 R1: For annually and quarterly data, for Mexico and Germany, only A status is 2072 allowed

2073 R2: For monthly data, for Mexico and Germany, only F status is allowed 2074

2074

2075 Considering the above examples, the following CubeRegions would be created:

CubeRegion1	FREQ	Q, A
	VIS_CTY	MX, DE
	OBS_STATUS	А
CubeRegion2	FREQ	Μ
	VIS_CTY	MX, DE
	OBS_STATUS	F

2076

The problem with this approach is that according to the business rule for Constraints, only one should be specified per Component. Thus, if a software would perform some conflict resolution would end up with empty sets for FREQ and OBS STATUS (as they do not share any values).

2082 Nevertheless, there is a much easier approach to that; this is the cascading 2083 mechanism of Constraints (as shown in 9.3.4.1). Hence, these rules would be 2084 expressed into two levels of Constraints, e.g., DSD and Dataflows:

2085

2081

2086 DSD CubeRegion:

FREQ	M, Q, A
VIS_CTY	MX, DE
OBS_STATUS	A, F

2087 2088

Dataflow1 CubeReg	jion:
FREQ	Q, A
VIS_CTY	MX, DE
OBS_STATUS	F

2089 2090

Dataflow2	CubeReg	jion:
FREQ		Μ

FREQ	Μ
VIS_CTY	MX, DE
OBS_STATUS	А

9.3.4.2.6 Use case 6: A Constraint on allowed values for some Dimensions combined with allowed values for Measures

2093 R1: Allow monthly data, for Germany, with unit euro, and measure choice is 'A'



2094 R2: Allow quarterly data, for Mexico, with unit usd, and measure choice is 'B'

2095

2096 This is may be expressed with the following DataKeySet:

Key1	FREQ	Μ
	VIS_CTY	DE
	UNIT	EUR
	CHOICE	А
Key2	FREQ	Q
	VIS_CTY	MX
	UNIT	USD
	CHOICE	В

2097

20989.3.4.2.7Use Case 7: A Constraint with wildcards for Codes and removePrefix2099property

- For this example, we assume that the VIS_CTY representation has been prefixed with prefix 'AREA'. In this Constraint, we need to remove the prefix.
- 2102 R1: Allow monthly and quarterly data
- 2103 R2: Allow vis-à-vis countries that start with M
- 2104 R3: Remove the prefix 'AREA '
- 2105

2106 This may be expressed with the following CubeRegion:

FREQ	M, Q
VIS_CTY (removePrefix='AREA_')	M%

2107

2108 9.3.4.2.8 Use Case 8: A Constraint with multilingual support on Attributes

- 2109 R1: Allow monthly and quarterly data
- 2110 R2: Allow Mexico for vis-à-vis country
- 2111 R3: Allow a comment, in English, which includes the term <code>adjusted</code> for status
- 2112
- 2113 This may be expressed with the following CubeRegion:

FREQ	M, Q
VIS_CTY	MX
COMMENT (lang='en')	%adjusted%

2114

9.3.4.2.9 Use Case 9: A Constraint on allowed values for Dimensions combined with allowed values for Metadata Attributes

- 2117 R1: Allow monthly and quarterly data
- 2118 R2: Allow Mexico for vis-à-vis country
- 2119 R3: Allow John Doe for contact
- 2120

2121 This may be expressed with the following CubeRegion:

FREQ	M, Q
VIS_CTY	MX



CONTACT John Doe

2122

2123 9.3.4.3 Other constraining terms

Beyond the cube regions and keysets, there are two more constraining terms, the ReleaseCalendar and the ReferencePeriod. The latter only applies to Data Constraints (and hence datasets), while the former applies both to Data and Metadata Constraints.

The ReferencePeriod is specifically used for the TimeDimension and is oriented
for time-series datasets. Nevertheless, a Dimension with a Time Representation
can also be constrained within a Cube or MetadataTarget Region.

The ReleaseCalendar is the only term that does not apply on Components; it specifies the schedule of publication or reporting of the dataset or metadataset.

- For example, the Release Calendar for Provider BIS, is specified in the three following terms:
- 2136

2133

2137

- Periodicity: how often data should be reported, e.g., monthly
- Offset: the number of days between the 1st of January and the first release of data, e.g., 10 days
- Tolerance: the maximum allowed of days that data may be considered, without
 being considered as late, e.g., 5 days
- With the above terms, BIS would need to report data between the 10th and 15th of every
 month.



10 Transforming between versions of SDMX

2145 **10.1 Scope**

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

2148 **10.2 Compatibility and new DSD features**

The following table provides an overview of the backwards compatibility between SDMX 3.0 and 2.1.

2151

SDMX 3.0 feature	SDMX 2.1 compatibility	Comments
Multiple Measures	Create a Measure Dimension Or Model Measures as Attributes	For a Measure Dimensions, all Concepts must reside in the same Concept Scheme
Arrays for values	Cannot be supported	Arrays are always reported in a verbose format, even if one value is reported
Measure	Can be ignored, as it does not	
Relationship	affect dataset format	
Metadata Attributes	Can be created as Data Attributes	Not for extended facets
Multilingual Components	Cannot be supported	
No Measure	Can only be emulated by ignoring the Primary Measure value	
Use extended Codelist	A new Codelist with all Codes must be created	
Sentinel values	Cannot be supported in the DSD	Rules may be supported outside the DSD, in bilateral agreements

2152 2153

The following table illustrates forward compatibility issues from SDMX 2.1 to 3.0.

2154

SDMX 2.1 feature	SDMX 3.0 compatibility	Comments
Measure	Create a Dimension with role	If the dataset has
Dimension	'MEASURE'	to resemble that of
	Or	SDMX 2.1
	Create multiple Measures from the	Structure Specific,
	Measure Dimension Concept Scheme	then the first option
		must be used
Primary Measure	Create one Measure with role	
-	'PRIMARY'; use id="OBS_VALUE"	



2156 **11 Validation and Transformation Language (VTL)**

2157 **11.1 Introduction**

The Validation and Transformation Language (VTL) supports the definition of Transformations, which are algorithms to calculate new data starting from already existing ones⁷. The purpose of the VTL in the SDMX context is to enable the:

2161 2162

2163

2164 2165

2166

2167 2168

2169

2170

2171

2172

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

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

2181 The VTL language can be applied to SDMX artefacts by mapping the SDMX IM model artefacts to the model artefacts that VTL can manipulate⁸. Thus, the SDMX artefacts 2182 2183 can be used in VTL as inputs and/or outputs of Transformations. It is important to be 2184 aware that the artefacts do not always have the same names in the SDMX and VTL IMs, nor do they always have the same meaning. The more evident example is given 2185 2186 by the SDMX Dataset and the VTL "Data Set", which do not correspond one another: 2187 as a matter of fact, the VTL "Data Set" maps to the SDMX "Dataflow", while the SDMX "Dataset" has no explicit mapping to VTL (such an abstraction is not needed 2188 in the definition of VTL Transformations). A SDMX "Dataset", however, is an instance 2189 of a SDMX "Dataflow" and can be the artefact on which the VTL transformations are 2190 2191 executed (i.e., the Transformations are defined on Dataflows and are applied to 2192 Dataflow instances that can be Datasets). 2193

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

⁷ 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 <u>https://sdmx.org</u>.

⁸ In this chapter, in order to distinguish VTL and SDMX model artefacts, the VTL ones are written in the Arial font while the SDMX ones in Courier New



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

2202 **11.2 References to SDMX artefacts from VTL statements**

2203 **11.2.1 Introduction**

The VTL can manipulate SDMX artefacts (or objects) by referencing them through predefined conventional names (aliases).

- 2206
- The alias of an SDMX artefact can be its URN (Universal Resource Name), an
 abbreviation of its URN or another user-defined name.

In any case, the aliases used in the VTL Transformations have to be mapped to the
SDMX artefacts through the VtlMappingScheme and VtlMapping classes (see the
section of the SDMX IM relevant to the VTL). A VtlMapping allows specifying the
aliases to be used in the VTL Transformations, Rulesets⁹ or User Defined Operators¹⁰
to reference SDMX artefacts. A VtlMappingScheme is a container for zero or more
VtlMapping.

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

2222

2229

2233

2234

2235

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

2225 11.2.2 References through the URN

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

The SDMX URN¹¹ is the concatenation of the following parts, separated by special symbols like dot, equal, asterisk, comma, and parenthesis:

- SDMXprefix
 - SDMX-IM-package-name
 - class-name
 - agency-id

⁹ See also the section "VTL-DL Rulesets" in the VTL Reference Manual.

¹⁰ The VTLMappings are used also for User Defined Operators (UDO). Although UDOs 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 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.

¹¹ 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)").



2236 2237 2238 2239 2240	 maintainedobject-id maintainedobject-version container-object-id ¹² object-id The generic structure of the URN is the following:
2241 2242 2243 2244	SDMXprefix.SDMX-IM-package-name.class-name=agency-id:maintainedobject-id (maintainedobject-version).*container-object-id.object-id
2245 2246	The SDMXprefix is "urn:sdmx:org", always the same for all SDMX artefacts.
2240 2247 2248 2249 2250	The SDMX-IM-package-name is the concatenation of the string "sdmx.infomodel." with the package-name, which the artefact belongs to. For example, for referencing a Dataflow the SDMX-IM-package-name is "sdmx.infomodel.datastructure", because the class Dataflow belongs to the package "datastructure".
2251 2252 2253 2254 2255 2256	The class-name is the name of the SDMX object class, which the SDMX object belongs to (e.g., for referencing a Dataflow the class-name is "Dataflow"). The VTL can reference SDMX artefacts that belong to the classes Dataflow, Dimension, TimeDimension, Measure, DataAttribute, Concept, Codelist.
2250 2257 2258 2259 2260	The agency-id is the acronym of the agency that owns the definition of the artefact, for example for the Eurostat artefacts the agency-id is "ESTAT"). The agency-id can be composite (for example AgencyA.Dept1.Unit2).
2261 2262 2263 2264	The maintainedobject-id is the name of the maintained object which the artefact belongs to, and in case the artefact itself is maintainable ¹³ , coincides with the name of the artefact. Therefore the maintainedobject-id depends on the class of the artefact:
2265 2266 2267 2268 2269 2270 2271 2272 2273 2274 2275 2276	 if the artefact is a Dataflow, which is a maintainable class, the maintainedobject-id is the Dataflow name (dataflow-id); if the artefact is a Dimension, Measure, TimeDimension or DataAttribute, which are not maintainable and belong to the DataStructure maintainable class, the maintainedobject-id is the name of the DataStructure (dataStructure-id) which the artefact belongs to; if the artefact is a Concept, which is not maintainable and belongs to the ConceptScheme maintainable class, the maintainedobject-id is the name of the ConceptScheme (conceptScheme-id) which the artefact belongs to; if the artefact is a Codelist, which is a maintainable class, the maintainedobject-id is the codelist name (codelist-id).
2277 2278 2279	The maintainedobject-version is the version of the maintained object which the artefact belongs to (for example, possible versions are 1.0.0, 2.1.0, 3.1.2).
2280 2281 2282	The container-object-id does not apply to the classes that can be referenced in VTL Transformations, therefore is not present in their URN

¹² The container-object-id can repeat and may not be present.
¹³ i.e., the artefact belongs to a maintainable class



2283 The object-id is the name of the non-maintainable artefact (when the artefact is maintainable its name is already specified as the maintainedobject-id, see above), in 2284 particular it has to be specified: 2285

- 2286 2287
- 2288 2289

2290

2291

- if the artefact is a Dimension, TimeDimension, or Measure DataAttribute (the object-id is the name of one of the artefacts above, which are data structure components)
- if the artefact is a Concept (the object-id is the name of the Concept) •

For example, by using the URN, the VTL Transformation that sums two SDMX 2292 Dataflows DF1 and DF2 and assigns the result to a third persistent Dataflow DFR, 2293 assuming that DF1, DF2 and DFR are the maintainedobject-id of the three 2294 2295 Dataflows, that their version is 1.0.0 and their Agency is AG, would be written as¹⁴: 2296

```
2297
        'urn:sdmx:org.sdmx.infomodel.datastructure.Dataflow=AG:DFR(1.0.0)' <-</pre>
2298
        'urn:sdmx:org.sdmx.infomodel.datastructure.Dataflow=AG:DF1(1.0.0)' +
2299
        'urn:sdmx:org.sdmx.infomodel.datastructure.Dataflow=AG:DF2(1.0.0)'
```

2300 11.2.3 Abbreviation of the URN

2301 The complete formulation of the URN described above is exhaustive but verbose, even 2302 for very simple statements. In order to reduce the verbosity through a simplified identifier and make the work of transformation definers easier, proper abbreviations of 2303 2304 the URN are possible. Using this approach, the referenced artefacts remain intelligible 2305 in the VTL code by a human reader.

2307 The URN can be abbreviated by omitting the parts that are not essential for the identification of the artefact or that can be deduced from other available information, 2308 including the context in which the invocation is made. The possible abbreviations are 2309 2310 described below.

- The SDMXprefix can be omitted for all the SDMX objects, because it is a prefixed string (urn:sdmx:org), always the same for SDMX objects.
- The SDMX-IM-package-name can be omitted as well because it can be deduced from the class-name that follows it (the table of the SDMX-IM packages and classes that allows this deduction is in the SDMX 2.1 Standards - Section 5 -Registry Specifications, paragraph 6.2.3). In particular, considering the object classes of the artefacts that VTL can reference, the package is:

Dimension,

2319

2306

2311 2312

2313

2314

2315

2316 2317

2318

- 2320 2321
- o "datastructure" for the classes Dataflow, TimeDimension, Measure, DataAttribute,
- o "conceptscheme" for the class Concept,
- "codelist" for the class Codelist. 0
- 2323 The class-name can be omitted as it can be deduced from the VTL invocation. 2324 In particular, starting from the VTL class of the invoked artefact (e.g. dataset, component, identifier, measure, attribute, variable, valuedomain), which is 2325 known given the syntax of the invoking VTL operator¹⁵, the SDMX class can be 2326

¹⁴ 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.

¹⁵ For the syntax of the VTL operators see the VTL Reference Manual



2327 2328	deduced from the mapping rules between VTL and SDMX (see the section "Mapping between VTL and SDMX" hereinafter) ¹⁶ .
2329	 If the agency-id is not specified, it is assumed by default equal to the agency-
2330	id of the TransformationScheme. UserDefinedOperatorScheme Or
2331	BulesetScheme from which the artefact is invoked. For example, the agency-
2332	id can be omitted if it is the same as the invoking TransformationScheme
2222	and cannot be omitted if the artefact comes from another agency 17 . Take also
2334	into account that according to the VTL consistency rules the agency of the
2335	result of a Transformation must be the same as its
2336	TransformationScheme therefore the agency-id can be omitted for all the
2337	results (left part of Transformation statements)
2338	• As for the maintained chief id this is essential in some cases while in other
2330	• As for the maintained by Ject-id, this is essential in some cases while in other cases it can be omitted.
2340	\circ if the referenced artefact is a Dataflow which is a maintainable class
2341	the maintained object id is the dataflow id and obviously cannot be
2342	omitted.
2343	o if the referenced artefact is a Dimension. TimeDimension. Measure.
2344	DataAttribute which are not maintainable and belong to the
2345	DataStructure maintainable class the maintained object-id is the
2346	dataStructure-id and can be omitted, given that these components are
2347	always invoked within the invocation of a Dataflow whose
2348	dataStructure-id can be deduced from the SDMX structural definitions.
2349	o if the referenced artefact is a Concept, which is not maintainable and
2350	belong to the ConceptScheme maintainable class, the maintained
2351	object is the conceptScheme-id and cannot be omitted;
2352	o if the referenced artefact is a Codelist, which is a maintainable
2353	class, the maintainedobject-id is the codelist-id and obviously
2354	cannot be omitted.
2355	• When the maintainedobject-id is omitted, the maintainedobject-version is
2356	omitted too. When the maintainedobject-id is not omitted and the
2357	maintainedobject-version is omitted, the version 1.0.0 is assumed by default.
2358	• As said, the container-object-id does not apply to the classes that can be
2359	referenced in VTL Transformations, therefore is not present in their URN
2360	• The object-id does not exist for the artefacts belonging to the Dataflow,
2361	and Codelist classes, while it exists and cannot be omitted for the
2362	artefacts belonging to the classes Dimension, TimeDimension,
2363	Measure, DataAttribute and Concept, as for them the object-id is
2364	the main identifier of the artefact

¹⁶ 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, TimeDimension, Measure or DataAttribute) can be deduced from the data structure of the SDMX Dataflow, which the component belongs to.

¹⁷ 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)



The simplified object identifier is obtained by omitting all the first part of the URN, 2365 including the special characters, till the first part not omitted. 2366 2367 For example, the full formulation that uses the complete URN shown at the end of the 2368 2369 previous paragraph: 2370 2371 'urn:sdmx:org.sdmx.infomodel.datastructure.Dataflow=AG:DFR(1.0.0)' := 2372 'urn:sdmx:org.sdmx.infomodel.datastructure.Dataflow=AG:DF1(1.0.0)' + 2373 'urn:sdmx:org.sdmx.infomodel.datastructure.Dataflow=AG:DF2(1.0.0)' 2374 2375 by omitting all the non-essential parts would become simply: 2376 2377 DFR := DF1 + DF22378 2379 The references to the Codelists can be simplified similarly. For example, given the non-abbreviated reference to the Codelist AG:CL_FREQ(1.0.0), which is¹⁸: 2380 2381 2382 'urn:sdmx:org.sdmx.infomodel.codelist.Codelist=AG:CL FREQ(1.0.0)' 2383 2384 if the Codelist is referenced from a RulesetScheme belonging to the agency AG, 2385 omitting all the optional parts, the abbreviated reference would become simply¹⁹: 2386 2387 CL FREQ 2388 2389 As for the references to the components, it can be enough to specify the component-2390 Id, given that the dataStructure-Id can be omitted. An example of non-abbreviated reference, if the data structure is DST1 and the component is SECTOR, is the 2391 2392 following: 2393 2394 'urn:sdmx:org.sdmx.infomodel.datastructure.DataStructure=AG:DST1(1.0).SEC 2395 TOR' 2396 2397 corresponding fully abbreviated reference. from The if made а 2398 TransformationScheme belonging to AG, would become simply: 2399 2400 SECTOR 2401 For example, the Transformation for renaming the component SECTOR of the 2402 Dataflow **DF1** into SEC can be written as²⁰: 2403 2404 2405 'DFR(1.0.0)' := 'DF1(1.0.0)' [rename SECTOR to SEC] 2406 2407 In the references to the Concepts, which can exist for example in the definition of the 2408 VTL Rulesets, at least the conceptScheme-id and the concept-id must be 2409 specified. 2410

¹⁸ Single quotes are needed because this reference is not a VTL regular name.

¹⁹ Single quotes are not needed in this case because CL_FREQ is a VTL regular name.

 $^{^{20}}$ The result DFR(1.0.0) is be equal to DF1(1.0.0) save that the component SECTOR is called SEC



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

2413

2414 'urn:sdmx:org.sdmx.infomodel.conceptscheme.Concept=AG:CS1(1.0.0).SECTOR'
2415

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

2419 CS1(1.0.0).SECTOR

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

2424 2425 2426

2420

'DFR(1.0.0)' := between ('DF1(1.0.0)', 0, 25000)

The artefact (Component, Concept, Codelist ...) which the Values are referred to can be deduced from the context in which the reference is made, taking also into account the VTL syntax. In the Transformation above, for example, the values 0 and 2500 are compared to the values of the measures of DF1(1.0.0).

2431 **11.2.4 User-defined alias**

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

2434

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

2441 **11.2.5 References to SDMX artefacts from VTL Rulesets**

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

2447

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

2450

In the signature, given the mapping between VTL and SDMX better described in the
 following paragraphs, a reference to a VTL Value Domain becomes a reference to a
 SDMX Codelist, while a reference to a VTL Represented Variable becomes a
 reference to a SDMX Concept, assuming for it a definite representation²¹.

²¹ Rulesets of this kind cannot be reused when the referenced Concept has a different representation.


In general, for referencing SDMX Codelists and Concepts, the conventions described in the previous paragraphs apply. In the Ruleset syntax, the elements that reference SDMX artefacts are called "valueDomain" and "variable" for the Datapoint Rulesets and "ruleValueDomain", "ruleVariable", "condValueDomain" "condVariable" for the Hierarchical Rulesets). The syntax of the Ruleset signature allows also to define aliases of the elements above, these aliases are valid only within the specific Ruleset definition statement and cannot be mapped to SDMX.²²

2463

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

2467 **11.3 Mapping between SDMX and VTL artefacts**

2468 **11.3.1 When the mapping occurs**

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

2471 It should be remembered that VTL programs (i.e. Transformation Schemes) are 2472 represented in SDMX through the TransformationScheme maintainable class 2473 which is composed of Transformations (nameable artefacts). Each 2474 Transformation assigns the outcome of the evaluation of a VTL expression to a 2475 result: the input operands of the expression and the result can be SDMX artefacts.

Every time a SDMX object is referenced in a VTL Transformation as an input operand, there is the need to generate a VTL definition of the object, so that the VTL operations can take place. This can be made starting from the SDMX definition and applying a SDMX-VTL mapping method in the direction from SDMX to VTL. The possible mapping methods from SDMX to VTL are described in the following paragraphs and are conceived to allow the automatic deduction of the VTL definition of the object from the knowledge of the SDMX definition.

In the opposite direction, every time an object calculated by means of VTL must be
treated as a SDMX object (for example for exchanging it through SDMX), there is the
need of a SDMX definition of the object, so that the SDMX operations can take place.
The SDMX definition is needed for the VTL objects for which a SDMX use is
envisaged²³.

2488

2489 The mapping methods from VTL to SDMX are described in the following paragraphs as well, however they do not allow the complete SDMX definition to be automatically 2490 2491 deduced from the VTL definition, more than all because the former typically contains 2492 additional information in respect to the latter. For example, the definition of a SDMX 2493 DSD includes also some mandatory information not available in VTL (like the concept scheme to which the SDMX components refer, the assignmentStatus and attributeRelationship for the DataAttributes and so on). Therefore the mapping 2494 2495 2496 methods from VTL to SDMX provide only a general guidance for generating SDMX 2497 definitions properly starting from the information available in VTL, independently of how 2498 the SDMX definition it is actually generated (manually, automatically or part and part).

²² See also the section "VTL-DL Rulesets" in the VTL Reference Manual.

²³ If a calculated artefact is persistent, it needs a persistent definition, i.e. a SDMX definition in a SDMX environment. In addition, 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).



2499 **11.3.2 General mapping of VTL and SDMX data structures**

This section makes reference to the VTL "Model for data and their structure"²⁴ and the correspondent SDMX "Data Structure Definition"²⁵.

The main type of artefact that the VTL can manipulate is the VTL Data Set, which in general is mapped to the SDMX Dataflow. This means that a VTL Transformation, in the SDMX context, expresses the algorithm for calculating a derived Dataflow starting from some already existing Dataflows (either collected or derived).²⁶

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

2514 A VTL Data Set is structured by one and just one Data Structure and a VTL Data Structure can structure any number of Data Sets. Correspondingly, in the SDMX 2515 2516 context а **SDMX** Dataflow is structured by one and iust one 2517 DataStructureDefinition and one DataStructureDefinition can structure 2518 any number of Dataflows.

2520 A VTL Data Set has a Data Structure made of Components, which in turn can be Identifiers, Measures and Attributes. Similarly, a SDMX DataflowDefinition has 2521 2522 DataStructureDefinition made components that of can be а 2523 DimensionComponents, Measure and DataAttributes. In turn, a SDMX 2524 DimensionComponent can be а Dimension **or a** TimeDimension. 2525 Correspondingly, in the SDMX implementation of the VTL, the VTL Identifiers can be (optionally) distinguished in similar sub-classes (Simple Identifier, Time Identifier) even 2526 2527 if such a distinction is not evidenced in the VTL IM.

2528

2513

2519

2529 The possible mapping options are described in more detail in the following sections.

2530 **11.3.3 Mapping from SDMX to VTL data structures**

2531 **11.3.3.1 Basic Mapping**

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

2535 When transforming **from SDMX to VTL**, this method consists in leaving the 2536 components unchanged and maintaining their names and roles, according to the 2537 following table:

SDMX	VTL
Dimension	(Simple) Identifier
TimeDimension	(Time) Identifier

²⁴ See the VTL 2.0 User Manual

²⁵ See the SDMX Standards Section 2 – Information Model

 $^{^{26}}$ Besides the mapping between one SDMX <code>Dataflow</code> and one VTL Data Set, it is also possible to map distinct parts of a SDMX <code>Dataflow</code> to different VTL Data Set, as explained in a following paragraph.



Measure	Measure
DataAttribute	Attribute

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

2544

2545 With the Basic mapping, one SDMX observation²⁷ generates one VTL data point.

2546 **11.3.3.2 Pivot Mapping**

An alternative mapping method from SDMX to VTL is the **Pivot** mapping, which makes sense and is different from the Basic method only for the SDMX data structures that contain a Dimension that plays the role of measure dimension (like in SDMX 2.1) and just one Measure. Through this method, these structures can be mapped to multimeasure VTL data structures.

2553 In SDMX 2.1 the MeasureDimension was a subclass of DimensionComponent like Dimension and TimeDimension. In the current SDMX version, this subclass does 2554 not exist anymore, however a Dimension can have the role of measure dimension 2555 2556 (i.e. a Dimension that contributes to the identification of the measures). In SDMX 2.1 2557 a DataStructure could have zero or one MeasureDimensions, in the current 2558 version of the standard, from zero to many Dimension may have the role of measure 2559 dimension. Hereinafter a Dimension that plays the role of measure dimension is referenced for simplicity as "MeasureDimension", i.e. maintaining the capital letters 2560 and the courier font even if the MeasureDimension is not anymore a class in the 2561 2562 SDMX Information Model of the current SDMX version. For the sake of simplicity, the 2563 description below considers just one Dimension having the role **of** 2564 MeasureDimension (i.e., the more simple and common case). Nevertheless, it maintains its validity also if in the DataStructure there are more dimension with the 2565 2566 role of MeasureDimensions: in this case what is said about the 2567 MeasureDimension must be applied to the combination of all the 2568 MeasureDimensions considered as a joint variable²⁸.

Among other things, the Pivot method provides also backward compatibility with the SDMX 2.1 data structures that contained a MeasureDimension.

If applied to SDMX structures that do not contain any MeasureDimension, this
 method behaves like the Basic mapping (see the previous paragraph).

2575

2569

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

²⁷ Here an SDMX observation is meant to correspond to one combination of values of the DimensionComponents.

⁸ E.g., if in the data structure there exist 3 Dimensions C,D,E having the role of MeasureDimension, they should be considered as a joint MeasureDimension Z=(C,D,E); therefore when the description says "each possible value Cj of the MeasureDimension ..." it means "each possible combination of values (Cj, Dk, Ew) of the joint MeasureDimension Z=(C,D,E)".



2578	
• A SDMX simple dimension becomes a VTL (simple) identifier and a SD	MX
2580 TimeDimension becomes a VTL (time) identifier;	
• Each possible Code Cj of the SDMX MeasureDimension is mapped to a	√TL
2582 Measure, having the same name as the SDMX Code (i.e. Cj); the VTL Measure	sure
2583 Cj is a new VTL component even if the SDMX data structure has not suc	ch a
2584 Component;	
• The SDMX MeasureDimension is not mapped to VTL (it disappears in	the
2586 VTL Data Structure);	
• The SDMX Measure is not mapped to VTL as well (it disappears in the	√TL
2588 Data Structure);	
• An SDMX DataAttribute is mapped in different ways according to	its
2590 AttributeRelationship:	
2591 o If, according to the SDMX AttributeRelationship, the value	s of
2592 the DataAttribute do not depend on the values of	the
2593 MeasureDimension, the SDMX DataAttribute becomes a	√TL
2594 Attribute having the same name. This happens if	the
2595 AttributeRelationship is not specified (i.e. the DataAttrib	ute
2596 does not depend on any DimensionComponent and therefore i	s at
2597 data set level), or if it refers to a set (or a group) of dimensions w	nich
2598 does not include the MeasureDimension;	
2599 o Otherwise, if, according to the SDMX AttributeRelationship,	the
2600 values of the DataAttribute depend on the MeasureDimensi	.on,
2601 the SDMX DataAttribute is mapped to one VTL Attribute for e	ach
2602 possible Code of the SDMX MeasureDimension. By default,	the
2603 names of the VTL Attributes are obtained by concatenating the name	e of
2604 the SDMX DataAttribute and the names of the correspondent C	ode
2605 of the MeasureDimension separated by underscore. For example	e, if
2606 the SDMX DataAttribute is named DA and the possible Code	s of
2607 the SDMX MeasureDimension are named C1, C2,, Cn, then	the
2608 corresponding VTL Attributes will be named DA_C1, DA_C2,	,
2609 DA_Cn (if different names are desired, they can be achieved afterwa	ards
by renaming the Attributes through VTL operators).	
2611 • Like in the Basic mapping, the resulting VIL Attributes are conside	red
2012 as dependent on all the VIL identifiers (i.e. "at data point / observa	tion
2013 IEVEL, DECAUSE VIL DOES NOT NAVE THE SDIVIX NOTION OF ATTRIC	ute

2616 The summary mapping table of the "pivot" mapping from SDMX to VTL for the SDMX 2617 data structures that contain a MeasureDimension is the following:

SDMX	VTL	
Dimension	(Simple) Identifier	
TimeDimension	(Time) Identifier	
MeasureDimension &	One Measure for each Code of the	
one Measure	SDMX MeasureDimension	
DataAttribute not depending on the	Attribute	
MeasureDimension		
DataAttribute depending on the One Attribute for each Code of		
MeasureDimension	SDMX MeasureDimension	



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

At observation / data point level, calling Cj (j=1, ... n) the jth Code of the MeasureDimension:

2625 2626

2627 2628

2629

2636 2637

2618

- The set of SDMX observations having the same values for all the Dimensions except than the MeasureDimension become one multi-measure VTL Data Point, having one Measure for each Code Cj of the SDMX MeasureDimension;
- The values of the SDMX simple Dimensions, TimeDimension and DataAttributes not depending on the MeasureDimension (these components by definition have always the same values for all the observations of the set above) become the values of the corresponding VTL (simple) Identifiers, (time) Identifier and Attributes.
 The value of the Measure of the SDMX observation belonging to the set above
 - The value of the Measure of the SDMX observation belonging to the set above and having MeasureDimension=Cj becomes the value of the VTL Measure Cj
- For the SDMX DataAttributes depending on the MeasureDimension, the value of the DataAttribute DA of the SDMX observation belonging to the set above and having MeasureDimension=Cj becomes the value of the VTL Attribute DA_Cj
- 2642 11.3.3.3 From SDMX DataAttributes to VTL Measures
- 2643 In some cases, it may happen that the DataAttributes of the SDMX • 2644 DataStructure need to be managed as Measures in VTL. Therefore, a 2645 variant of both the methods above consists in transforming all the SDMX 2646 DataAttributes in VTL Measures. When DataAttributes are converted 2647 to Measures, the two methods above are called Basic A2M and Pivot A2M (the suffix "A2M" stands for Attributes to Measures). Obviously, the resulting 2648 VTL data structure is, in general, multi-measure and does not contain 2649 Attributes. 2650

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

2654

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

2658 11.3.4 Mapping from VTL to SDMX data structures

2659 **11.3.4.1 Basic Mapping**

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



Mapping table:

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

2668 2669

2669 2670

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

2671

2675

If the distinction between simple identifier and time identifier is not maintained in the
 VTL environment, the classification between Dimension and TimeDimension exists
 only in SDMX, as declared in the relevant DataStructureDefinition.

Regarding the Attributes, because VTL considers all of them "at observation level", the
corresponding SDMX DataAttributes should be put "at observation level" as well,
unless some different information about their AttributeRelationship is otherwise
available.

Note that the basic mappings in the two directions (from SDMX to VTL and vice-versa) are (almost completely) reversible. In fact, if a SDMX structure is mapped to a VTL structure and then the latter is mapped back to SDMX, the resulting data structure is like the original one (apart for the AttributeRelationship, that can be different if the original SDMX structure contains DataAttributes that are not at observation level). In reverse order, if a VTL structure is mapped to SDMX and then the latter is mapped back to VTL, the original data structure is obtained.

2688

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

2693 **11.3.4.2 Unpivot Mapping**

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

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

The multi-measures VTL structures are converted to SDMX Dataflows having an added MeasureDimension, which disambiguates the VTL multiple Measures, and a new Measure in place of the VTL ones, containing the values of the VTL Measures.

- 2703
- 2704 The **unpivot** mapping behaves like follows:
- like in the basic mapping, a VTL (simple) identifier becomes a SDMX
 Dimension and a VTL (time) identifier becomes a SDMX TimeDimension
 (as said, a measure identifier cannot exist in multi-measure VTL structures);



٠	a MeasureDimension component called "measure_name" is added to the
	SDMX DataStructure;

- a Measure component called "obs_value" is added to the SDMX DataStructure;
- each VTL Measure is mapped to a Code of the SDMX MeasureDimension
 having the same name as the VTL Measure (therefore all the VTL Measure
 Components do not originate Components in the SDMX DataStructure);
 - SDMX VTL Attribute becomes а DataAttribute having • а AttributeRelationship referred to all SDMX the DimensionComponents including the TimeDimension and except the MeasureDimension.
- 2718 2719

2716

2717

2708 2709 2710

2711

- 2720
- 2721

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

VTL	SDMX
(Simple) Identifier	Dimension
(Time) Identifier	TimeDimension
All Measure Components	MeasureDimension (having one Code
	for each VTL measure component) &
	ONE Measure
Attribute	DataAttribute depending on all
	SDMX Dimensions including the
	TimeDimension and except the
	MeasureDimension

2722	
2723	
2724	At observation / data point level:
2725	• a multi-measure VTL Data Point becomes a set of SDMX observations, one
2726	for each VTL Measure;
2727	• the values of the VTL Identifiers become the values of the corresponding
2728	SDMX DimensionComponents, for all the observations of the set above;
2729	• the name of the j th VTL Measure (e.g. "Cj") becomes the Code of the SDMX
2730	MeasureDimension of the jth observation of the set;
2731	• the value of the j th VTL Measure becomes the value of the SDMX Measure
2732	of the j th observation of the set;
2733	• the values of the VTL Attributes become the values of the corresponding
2734	SDMX DataAttributes (in principle for all the observations of the set
2735	above).
2736	If desired, this method can be applied also to mono-measure VTL structures, provided
2737	that none of the VTL Components has already the role of Measure Identifier. Like in
2738	the general case, a MeasureDimension component called "measure_name" is
2739	added to the SDMX DataStructure, in this case it has just one possible Code,
2740	corresponding to the name of the unique VTL Measure. The original VTL Measure
2741	would not become a Component in the SDMX data structure. The value of the VTL
2742	Measure would be assigned to the unique SDMX Measure called "obs_value".
2743	In any case, the resulting SDMX definitions must be compliant with the SDMX
2744	consistency rules. For example, the possible Codes of the SDMX
2745	MeasureDimension need to be listed in a SDMX Codelist, with proper id, agency



2746 and version; moreover, the SDMX DSD must have the assignmentStatus, which 2747 does not exist in VTL, the attributeRelationship for the DataAttributes and 2748 so on.

2749 **11.3.4.3 From VTL Measures to SDMX Data Attributes**

More than all for the multi-measure VTL structures (having more than one Measure Component), it may happen that the Measures of the VTL Data Structure need to be managed as DataAttributes in SDMX. Therefore a third mapping method consists in transforming some VTL measures in a corresponding SDMX Measures and all the other VTL Measures in SDMX DataAttributes. This method is called M2A ("M2A" stands for "Measures to DataAttributes").

All VTL Measures maintain their names in SDMX. The VTL Measure Components that
 become SDMX DataAttributes are the ones declared as DataAttributes in the
 target SDMX data structure definition.

2760

2756

2761 The mapping table is the following:

2762

VTL	SDMX
(Simple) Identifier	Dimension
(Time) Identifier	TimeDimension
Some Measures	Measure
Other Measures	DataAttribute
Attribute	DataAttribute

2763

Even in this case, the resulting SDMX definitions must be compliant with the SDMX
consistency rules. For example, the SDMX DSD must have the assignmentStatus,
which does not exist in VTL, the attributeRelationship for the
DataAttributes and so on.

2768 **11.3.5 Declaration of the mapping methods between data structures**

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

A customized mapping can be defined through the VtlMappingScheme and 2773 VtlDataflowMapping classes (see the section of the SDMX IM relevant to the VTL). 2774 2775 A VtlDataflowMapping allows specifying the mapping methods to be used for a 2776 specific dataflow, both in the direction from SDMX to VTL (toVtlMappingMethod) 2777 and from VTL to SDMX (fromVtlMappingMethod); in fact а 2778 VtlDataflowMapping associates the structured URN that identifies a SDMX Dataflow to its VTL alias and its mapping methods. 2779

2780

It is possible to specify the toVtlMappingMethod and fromVtlMappingMethod also for the conventional dataflow called "generic_dataflow": in this case the specified mapping methods are intended to become the default ones, overriding the "Basic" methods. In turn, the toVtlMappingMethod and fromVtlMappingMethod declared for a specific Dataflow are intended to override the default ones for such a Dataflow.



The VtlMappingScheme is a container for zero or more VtlDataflowMapping (it
 may contain also mappings towards artefacts other than dataflows).

2789 **11.3.6 Mapping dataflow subsets to distinct VTL Data Sets**

Until now it has been assumed to map one SMDX Dataflow to one VTL Data Set and 2790 vice-versa. This mapping one-to-one is not mandatory according to VTL because a 2791 2792 VTL Data Set is meant to be a set of observations (data points) on a logical plane, 2793 having the same logical data structure and the same general meaning, independently of the possible physical representation or storage (see VTL 2.0 User Manual page 24), 2794 2795 therefore a SDMX Dataflow can be seen either as a unique set of data observations 2796 (corresponding to one VTL Data Set) or as the union of many sets of data observations 2797 (each one corresponding to a distinct VTL Data Set).

As a matter of fact, in some cases it can be useful to define VTL operations involving definite parts of a SDMX Dataflow instead than the whole.²⁹

Therefore, in order to make the coding of VTL operations simpler when applied on parts of SDMX Dataflows, it is allowed to map distinct parts of a SDMX Dataflow to distinct VTL Data Sets according to the following rules and conventions. This kind of mapping is possible both from SDMX to VTL and from VTL to SDMX, as better explained below.³⁰

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

For example, assuming that the SDMX Dataflow DF1(1.0.0) has the Dimensions INDICATOR, TIME_PERIOD and COUNTRY, and that the user declares the Dimensions INDICATOR and COUNTRY as basis for the mapping (i.e. the mapping dimensions): the observations that have the same values for INDICATOR and COUNTRY would be mapped to the same VTL dataset (and vice-versa). In practice, this kind mapping is obtained like follows:

- 2814
- 2815
- 2816
- 2817
- 2818
- For a given SDMX Dataflow, the user (VTL definer) declares the DimensionComponents on which the mapping will be based, in a given order.³¹ Following the example above, imagine that the user declares the Dimensions INDICATOR and COUNTRY.

²⁹ 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 DimensionComponents of the Dataflow except the TimeDimension. 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.

³⁰ 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 <code>Dataflows</code> by means of VTL operators (e.g. "sub", "filter", "calc", "union" ...), maintaining a mapping one-to-one between SDMX <code>Dataflows</code> and VTL Data Sets.

³¹ 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 Data Sets is not used.



2819	•	The VTL Data Set is given a name using a special notation also called "ordered
2820		concatenation" and composed of the following parts:

 The reference to the SDMX Dataflow (expressed according to the rules described in the previous paragraphs, i.e. URN, abbreviated URN or another alias); for example DF(1.0.0);

- 2823or another alias); for example2824oa slash ("/") as a separator; 32
- 2825 o The reference to a specific part of the SDMX Dataflow above, 2826 expressed as the concatenation of the values that the SDMX 2827 DimensionComponents declared above must have, separated by 2828 dots (".") and written in the order in which these defined³³. 2829 For DimensionComponents are example POPULATION.USA would mean that such a VTL Data Set is mapped 2830 2831 to the SDMX observations for which the dimension INDICATOR is equal to POPULATION and the dimension COUNTRY is equal to USA. 2832

In the VTL Transformations, this kind of dataset name must be referenced between
single quotes because the slash ("/") is not a regular character according to the VTL
rules.

- 2836 Therefore, the generic name of this kind of VTL datasets would be:
- 2837 2838

2839

2843

2845

'DF(1.0.0)/INDICATORvalue.COUNTRYvalue'

Where DF(1.0.0) is the Dataflow and *INDICATORvalue* and *COUNTRYvalue* are placeholders for one value of the INDICATOR and COUNTRY dimensions. Instead the specific name of one of these VTL datasets would be:

2844 'DF(1.0.0)/POPULATION.USA'

In particular, this is the VTL dataset that contains all the observations of the Dataflow
DF(1.0.0) for which *INDICATOR* = POPULATION and *COUNTRY* = USA.

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

2850

As already said, the mapping from SDMX to VTL happens when the SDMX dataflows are operand of VTL Transformations, instead the mapping from VTL to SDMX happens when the VTL Data Sets that is result of Transformations³⁴ need to be treated as SDMX objects. This kind of mapping can be applied independently in the two directions and the Dimensions on which the mapping is based can be different in the two directions: these Dimensions are defined in the ToVtlSpaceKey and in the FromVtlSpaceKey classes respectively.

³² As a consequence of this formalism, a slash in the name of the VTL Data Set assumes the specific meaning of separator between the name of the Dataflow and the values of some of its Dimensions.

³³ 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.

³⁴ 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.



First, let us see what happens in the <u>mapping direction from SDMX to VTL</u>, i.e. when parts of a SDMX Dataflow (e.g. DF1(1.0.0)) need to be mapped to distinct VTL Data Sets that are operand of some VTL Transformations.

As already said, each VTL Data Set is assumed to contain all the observations of the SDMX Dataflow having INDICATOR=*INDICATORvalue* and COUNTRY= *COUNTRYvalue*. For example, the VTL dataset 'DF1(1.0.0)/POPULATION.USA' would contain all the observations of DF1(1.0.0) having INDICATOR = POPULATION and COUNTRY = USA.

- In order to obtain the data structure of these VTL Data Sets from the SDMX one, it is assumed that the SDMX DimensionComponents on which the mapping is based are dropped, i.e. not maintained in the VTL data structure; this is possible because their values are fixed for each one of the invoked VTL Data Sets³⁵. After that, the mapping method from SDMX to VTL specified for the Dataflow DF1(1.0.0) is applied (i.e. basic, pivot ...).
- 2873 In the example above. for all the datasets of the kind 'DF1(1.0.0)/INDICATORvalue.COUNTRYvalue', the dimensions INDICATOR and 2874 2875 COUNTRY would be dropped so that the data structure of all the resulting VTL Data 2876 Sets would have the identifier TIME PERIOD only.
- 2877 It should be noted that the desired VTL Data Sets (i.e. of the kind 'DF1(1.0.0)/
 2878 *INDICATORvalue.COUNTRYvalue*') can be obtained also by applying the VTL
 2879 operator "sub" (subspace) to the Dataflow DF1(1.0.0), like in the following VTL
 2880 expression:

2881 'DF1(1.0.0)/POPULATION.USA' := 2882 DF1(1.0.0) [sub INDICATOR="POPULATION", COUNTRY="USA"]; 2883 2884 'DF1(1.0.0)/POPULATION.CANADA' := 2885 DF1(1.0.0) [sub INDICATOR="POPULATION", COUNTRY="CANADA"]; 2886 2887 2888 In fact the VTL operator "sub" has exactly the same behaviour. Therefore, mapping 2899 different parts of a SDMX Dataflow to different VTL Data Sets in the direction from 2800 SDMX to VTL through the ordered consistencies potation is equivalent to a proper use

SDMX to VTL through the ordered concatenation notation is equivalent to a proper use
 of the operator "sub" on such a Dataflow. ³⁶

In the direction from SDMX to VTL it is allowed to omit the value of one or more
 DimensionComponents on which the mapping is based, but maintaining all the
 separating dots (therefore it may happen to find two or more consecutive dots and dots

2872

 $^{^{35}}$ If these <code>DimensionComponents</code> would not be dropped, the various VTL Data Sets resulting from this kind of mapping would have non-matching values for the Identifiers corresponding to the mapping Dimensions (e.g. POPULATION and COUNTRY). As a consequence, taking into account that the typical binary VTL operations at dataset level (+, -, *, / and so on) are executed on the observations having matching values for the identifiers, 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).

³⁶ 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. 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.



in the beginning or in the end). The absence of value means that for the corresponding 2895 Dimension all the values are kept and the Dimension is not dropped. 2896 For example, 'DF(1.0.0)/POPULATION.' (note the dot in the end of the name) is the 2897 VTL dataset that contains all the observations of the Dataflow DF(1.0.0) for which 2898 INDICATOR = POPULATION and COUNTRY = any value. 2899 2900 2901 This is equivalent to the application of the VTL "sub" operator only to the identifier 2902 **INDICATOR:** 2903 2904 'DF1(1.0.0)/POPULATION.' := 2905 DF1(1.0.0) [sub INDICATOR="POPULATION"]; 2906 Therefore the VTL Data Set 'DF1(1.0.0)/POPULATION.' would have the identifiers 2907 COUNTRY and TIME PERIOD. 2908 Heterogeneous invocations of the same Dataflow are allowed, i.e. omitting different 2909 Dimensions in different invocations. 2910 2911 Let us now analyse the mapping direction from VTL to SDMX. In this situation, distinct parts of a SDMX Dataflow are calculated as distinct VTL 2912 2913 datasets, under the constraint that they must have the same VTL data structure. 2914 For example, let us assume that the VTL programmer wants to calculate the SDMX 2915 Dataflow DF2(1.0.0) having the Dimensions TIME PERIOD, INDICATOR, and COUNTRY and that such a programmer finds it convenient to calculate separately the 2916 2917 parts of DF2(1.0.0) that have different combinations of values for INDICATOR and 2918 COUNTRY: 2919 each part is calculated as a VTL derived Data Set, result of a dedicated VTL 2920 Transformation: 37 2921 the data structure of all these VTL Data Sets has the TIME_PERIOD identifier 2922 and does not have the INDICATOR and COUNTRY identifiers.³⁸ Under these hypothesis, such derived VTL Data Sets can be mapped to DF2(1.0.0) by 2923 2924 declaring the DimensionComponents INDICATOR and COUNTRY as mapping dimensions³⁹. 2925 2926 2927 The corresponding VTL Transformations, assuming that the result needs to be persistent, would be of this kind: 40 2928 2929 'DF2(1.0.0)/INDICATORvalue.COUNTRYvalue' <- expression 2930 2931 Some examples follow, for some specific values of INDICATOR and COUNTRY: 2932 2933 'DF2(1.0.0)/GDPPERCAPITA.USA' expression11; < -2934 'DF2(1.0.0)/GDPPERCAPITA.CANADA' < expression12;

 $^{^{37}}$ If the whole DF2(1.0) is calculated by means of just one VTL Transformation, then the mapping between the SDMX <code>Dataflow</code> and the corresponding VTL dataset is one-to-one and this kind of mapping (one SDMX <code>Dataflow</code> to many VTL datasets) does not apply.

³⁸ This is possible as each VTL dataset corresponds to one particular combination of values of INDICATOR and COUNTRY.

³⁹ The mapping dimensions are defined as FromVtlSpaceKeys of the FromVtlSuperSpace of the VtlDataflowMapping relevant to DF2(1.0).

⁴⁰ the symbol of the VTL persistent assignment is used (<-)



2935 2936 2937 2938 2939	 'DF2(1.0.0)/POPGROWTH.USA' 'DF2(1.0.0)/POPGROWTH.CANAD 	<- expre ۲۰ <- express	ession21; ion22;	
2940 2941 2942 2943 2944	As said, it is assumed that these VTL deri the only identifier. In the mapping from VT and COUNTRY are added to the VTL data with the following values respectively:	ved Data Sets have to L to SMDX, the Dime structure on order to	the TIME_PERIOD as ensions INDICATOR obtain the SDMX one,	
2945 2946	VTL dataset	INDICATOR value	COUNTRY value	
2947 2948 2949	'DF2(1.0.0)/GDPPERCAPITA.USA' 'DF2(1.0.0)/GDPPERCAPITA.CANADA'	GDPPERCAPITA GDPPERCAPITA	USA CANADA	
2950 2951 2952 2953	'DF2(1.0.0)/POPGROWTH.USA' 'DF2(1.0.0)/POPGROWTH.CANADA' 	POPGROWTH POPGROWTH	USA CANADA	
2954 2955 2956 2957 2958 2959 2960 2961	It should be noted that the application of this many-to-one mapping from VTL to SDMX is equivalent to an appropriate sequence of VTL Transformations. These use the VTL operator "calc" to add the proper VTL identifiers (in the example, INDICATOR and COUNTRY) and to assign to them the proper values and the operator "union" in order to obtain the final VTL dataset (in the example DF2(1.0.0)), that can be mapped one-to-one to the homonymous SDMX Dataflow. Following the same example, these VTL Transformations would be:			
2962 2963 2964 2965	DF2bis_GDPPERCAPITA_USA := 'DF2 [calc ident identif	(1.0.0)/GDPPERCAPIT ifier INDICATOR := fier COUNTRY := "U	TA.USA' "GDPPERCAPITA", SA"];	
2965 2966 2967 2968 2969	<pre>DF2bis_GDPPERCAPITA_CANADA := 'DF2(2 [calc iden identif </pre>	1.0.0)/GDPPERCAPITA tifier INDICATOR:=' Fier COUNTRY:="CANA	.CANADA' 'GDPPERCAPITA", DA"];	
2970 2971 2972	DF2bis_POPGROWTH_USA := 'DF2(1.0 [calc ident identition	0.0)/POPGROWTH.USA' :ifier INDICATOR := fier COUNTRY := "U	"POPGROWTH", SA"];	
2973 2974 2975	<pre>DF2bis_POPGROWTH_CANADA' := 'DF2(1.0 [calc ident identified</pre>	.0)/POPGROWTH.CANAE ifier INDICATOR := fier COUNTRY := "C	DA' "POPGROWTH", ANADA"];	
2976				
2977 2978 2979 2980 2981 2982 2983	DF2(1.0) <- UNION DF2b DF2b DF2b DF2b);	(DF2bis_GDPPERCAPI is_GDPPERCAPITA_CAN is_POPGROWTH_USA', is_POPGROWTH_CANAD/	TA_USA', NADA',	
2984 2985 2986 2987	In other words, starting from the datasets example 'DF2(1.0)/GDPPERCAPITA.USA calculating other (non-persistent) DF2bis_GDPPERCAPITA_USA and so on	s explicitly calculated (`and so on), the VTL datasets () by adding the ident	d through VTL (in the first step consists in (in the example ifiers INDICATOR and	



2988 COUNTRY with the desired values (*INDICATORvalue* and *COUNTRYvalue*). Finally, 2989 all these non-persistent Data Sets are united and give the final result DF2(1.0)⁴¹, which 2990 can be mapped one-to-one to the homonymous SDMX Dataflow having the 2991 dimension components TIME_PERIOD, INDICATOR and COUNTRY.

Therefore, mapping different VTL datasets having the same data structure to different parts of a SDMX Dataflow, i.e. in the direction from VTL to SDMX, through the ordered concatenation notation is equivalent to a proper use of the operators "calc" and "union" on such datasets. ⁴²

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

3000

3001 **11.3.7 Mapping variables and value domains between VTL and SDMX**

3002 With reference to the VTL "model for Variables and Value domains", the following 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 DimensionComponent Or a Measure or a DataAttribute) belonging to one
Dennegented Veriable	specific Dataflow ⁴³
Represented variable	Representation a definite
Value Domain	Representation (see the Structure Pattern in the Base Package)
Enumerated Value Domain / Code List	Codelist
Code	Code (for enumerated
	DimensionComponent, Measure,
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

⁴¹ The result is persistent in this example but it can be also non persistent if needed.

⁴² 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.

⁴³ Through SDMX Constraints, it is possible to specify the values that a Component of a Dataflow can assume.



	to a valid value (for non-enumerated
	Representations)
Value Domain Subset / Set	This abstraction does not exist in SDMX
Enumerated Value Domain Subset /	This abstraction does not exist in SDMX
Enumerated Set	
Described Value Domain Subset /	This abstraction does not exist in SDMX
Described Set	
Set list	This abstraction does not exist in SDMX

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

Another difference consists in the fact that all Value Domains are considered as identifiable objects in VTL either if enumerated or not, while in SDMX the Codelist (corresponding to a VTL enumerated Value Domain) is identifiable, while the SDMX non-enumerated Representation (corresponding to a VTL non-enumerated Value Domain) is not identifiable. As a consequence, the definition of the VTL Rulesets, which in VTL can refer either to enumerated or non-enumerated value domains, in SDMX can refer only to enumerated Value Domains (i.e. to SDMX Codelists).

As for the mapping between VTL variables and SDMX Concepts, it should be noted that these artefacts do not coincide perfectly. In fact, the VTL variables are represented variables, defined always on the same Value Domain ("Representation" in SDMX) independently of the data set / data structure in which they appear⁴⁴, while the SDMX Concepts can have different Representations in different DataStructures.⁴⁵ This means that one SDMX Concept can correspond to many VTL Variables, one for each representation the Concept has.

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

3031DS_c := DS_a + DS_b(where DS_a, DS_b, DS_c are VTL Data Sets)3032if the matching components in DS_a and DS_b (e.g. ref_date, geo_area, sector ...)3033refer to the same general representation. In simpler words, DS_a and DS_b must use3034the same values/codes (for ref_date, geo_area, sector ...), otherwise the relevant3035values would not match and the result of the operation would be wrong.

As mentioned, the property above is not enforced by construction in SDMX, and different representations of the same Concept can be not compatible one another (for example, it may happen that geo_area is represented by ISO-alpha-3 codes in DS_a and by ISO alpha-2 codes in DS_b). Therefore, it will be up to the definer of VTL

⁴⁴ 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.

⁴⁵ 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.



Transformations to ensure that the VTL expressions are consistent with the actual representations of the correspondent SDMX Concepts.

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

3049

3050 11.4 Mapping between SDMX and VTL Data Types

3051 **11.4.1 VTL Data types**

According to the VTL User Guide the possible operations in VTL depend on the data types of the artefacts. For example, numbers can be multiplied but text strings cannot. In the VTL Transformations, the compliance between the operators and the data types of their operands is statically checked, i.e., violations result in compile-time errors.

3056 3057

The VTL data types are sub-divided in scalar types (like integers, strings, etc.), which are the types of the scalar values, and compound types (like Data Sets, Components, Rulesets, etc.), which are the types of the compound structures. See below the diagram of the VTL data types, taken from the VTL User Manual:





3062

Figure 22 – VTL Data Types

3063

3064 The VTL scalar types are in turn subdivided in basic scalar types, which are elementary (not defined in term of other data types) and Value Domain and Set scalar types, which 3065 3066 are defined in terms of the basic scalar types.

The VTL basic scalar types are listed below and follow a hierarchical structure in terms 3067

3068 of supersets/subsets (e.g. "scalar" is the superset of all the basic scalar types):





Figure 23 – VTL Basic Scalar Types

3071 11.4.2 VTL basic scalar types and SDMX data types

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

3074 The internal representation is the format used within a VTL system to represent (and 3075 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 3076 3077 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 3078 3079 days since a predefined date (e.g. from 01/01/4713 BC up to 31/12/5874897 AD like 3080 in Postgres) while two possible external representations are the formats YYYY-MM-3081 GG and MM-GG-YYYY (e.g. respectively 2010-12-31 and 12-31-2010).

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).

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.

3091 SDMX data types are conceived instead to support the data exchange, therefore they 3092 do have a format, which is known by the users and correspond, in VTL terms, to 3093 external representations. Therefore, for each VTL basic scalar type there can be more 3094 SDMX data types (the latter are explained in the section "General Notes for 3095 Implementers" of this document and are actually much more numerous than the 3096 former).

3097

3098The following paragraphs describe the mapping between the SDMX data types and3099the VTL basic scalar types. This mapping shall be presented in the two directions of3100possible conversion, i.e. from SDMX to VTL and vice-versa.

3101

The conversion from SDMX to VTL happens when an SDMX artefact acts as inputs of a VTL Transformation. As already said, in fact, at compile time the VTL needs to know the VTL type of the operands in order to check their compliance with the VTL operators



- and at runtime it must convert the values from their external (SDMX) representationsto the corresponding internal (VTL) ones.
- 3107
- 3108 The opposite conversion, i.e. from VTL to SDMX, happens when a VTL result, i.e. a
- VTL Data Set output of a Transformation, must become a SDMX artefact (or part of it).
 The values of the VTL result must be converted into the desired (SDMX) external
- 3111 representations (data types) of the SDMX artefact.

3112 **11.4.3 Mapping SDMX data types to VTL basic scalar types**

The following table describes the default mapping for converting from the SDMX data 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: {true, false})	boolean



Statistical Data an	d Metadata	eXchange
---------------------	------------	----------

URI	string
(corresponds to the XML Schema xs:anyURI;	5
absolute or relative Uniform Resource Identifier	
Reference)	
Count	integer
(an integer following a sequential pattern,	
increasing by 1 for each occurrence)	
InclusiveValueRange	number
(decimal number within a closed interval, whose	
bounds are specified in the SDMX representation	
by the facets minValue and maxValue)	
ExclusiveValueRange	number
(decimal number within an open interval, whose	
bounds are specified in the SDMX representation	
by the facets minValue and maxValue)	
Incremental	number
(decimal number the increased by a specific	
interval (defined by the interval facet), which is	
typically enforced outside of the XML validation)	
ObservationalTimePeriod	time
(superset of StandardTimePeriod and	
ImeRange)	
Standard I imePeriod	time
(superset of Basic limePeriod and	
Reporting LimePeriod)	
Basic I imePeriod	date
(superset of Gregorian TimePeriod and Date Time)	data
Gregorian TimePeriod	date
(superset of Gregorian Year, Gregorian YearMonth,	
GragorianVoar	data
	uale
GregorianVearMonth / GregorianMonth	date
(YYYY-MM)	uale
GregorianDay	date
(YYYY-MM-DD)	date
ReportingTimePeriod	time period
(superset of RepostingYear, ReportingSemester,	
ReportingTrimester. ReportingQuarter.	
ReportingMonth. ReportingWeek. ReportingDay)	
ReportingYear	time period
(YYYY-A1 – 1 vear period)	
ReportingSemester	time period
(YYYY-Ss – 6 month period)	
ReportingTrimester	time period
$(\dot{Y}\dot{Y}Y-T\dot{t} - 4 \text{ month period})$	
ReportingQuarter	time_period
(YYY-Qq – 3 month period)	—
ReportingMonth	time_period
(YYYY-Mmm – 1 month period)	
ReportingWeek	time_period



Statistical Data and Metadata eXchange

(YYYY-Www – 7 day period; following ISO 8601 definition of a week in a year)	
ReportingDay (YYYY-Dddd – 1 day period)	time_period
DateTime (YYYY-MM-DDThh:mm:ss)	date
TimeRange (YYYY-MM-DD(Thh:mm:ss)?/ <duration>)</duration>	time
Month (MM; speicifies 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	Metadata type – not applicable
KeyValues	Metadata type – not applicable
IdentifiableReference	Metadata type – not applicable
DataSetReference	Metadata type – not applicable
AttachmentConstraintReference	Metadata type – not applicable

3115

Figure 14 - Mappings from SDMX data types to VTL Basic Scalar Types

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

3121 11.4.4 Mapping VTL basic scalar types to SDMX data types

- The following table describes the default conversion from the VTL basic scalar types
- 3123 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)</date></date>
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"

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

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

3128 3129 The custom output formats can be specified by means of the VTL formatting mask described in the section "Type Conversion and Formatting Mask" of the VTL Reference 3130 Manual. Such a section describes the masks for the VTL basic scalar types "number", 3131 3132 "integer", "date", "time", "time_period" and "duration" and gives examples. As for the types "string" and "boolean" the VTL conventions are extended with some other special 3133 3134 characters as described in the following table. actors for the formatt . . .

VIL 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 durat	ion	
С	century	
Υ	year	
S	semester	
Q	quarter	
Μ	month	
W	week	
D	day	
h	hour digit (by default on 24 hours)	
Μ	minute	
S	second	
D	decimal of second	
Р	period indicator (representation in one digit for the duration)	
Р	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)
Dav	First character uppercase, then lowercase textual representation of
Day	the devusing (or Mondey)
	the day using (e.g. monday)
String	
Х	any string character
Z	any string character from "A" to "z"
9	any string character from "0" to "9"
Boolean	
В	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
Ν	fixed number of digits used in the preceding textual representation
	of the month or the day

The default conversion, either standard or customized, can be used to deduce automatically the representation of the components of the result of a VTL Transformation. In alternative, the representation of the resulting SDMX Dataflow can be given explicitly by providing its DataStructureDefinition. In other words, the representation specified in the DSD, if available, overrides any default conversion⁴⁶.

3142 **11.4.5 Null Values**

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

On the other side, the VTL programs can produce in output NULL values for Measures 3146 and Attributes (Null values are not allowed in the Identifiers). In the conversion from 3147 VTL to SDMX, it is assumed that a NULL in VTL becomes a missing value in SDMX. 3148 In the conversion from VTL to SDMX, the default assumption can be overridden, 3149 separately for each VTL basic scalar type, by specifying which the value that 3150 represents the NULL in SDMX is. This can be specified in the attribute "nullValue" 3151 of the CustomType artefact (see also the section Transformations and Expressions of 3152 3153 the SDMX information model). A CustomType belongs to a CustomTypeScheme, which can be referenced by one or more TransformationScheme (i.e. VTL 3154 3155 programs). The overriding assumption is applied for all the SDMX Dataflows 3156 calculated in the TransformationScheme.

⁴⁶ The representation given in the DSD should obviously be compatible with the VTL data type.



3157 **11.4.6 Format of the literals used in VTL Transformations**

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

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

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

3171 A different format can be specified in the attribute "vtlLiteralFormat" of the 3172 CustomType artefact (see also the section Transformations and Expressions of the 3173 SDMX information model).

Like in the case of the conversion of NULLs described in the previous paragraph, the overriding assumption is applied, for a certain VTL basic scalar type, if a value is found for the vtlLiteralFormat attribute of the CustomType of such VTL basic scalar

3177 type. The overriding assumption is applied for all the literals of a related VTL3178 TransformationScheme.

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



3181 **12 Structure Mapping**

3182 **12.1 Introduction**

The purpose of SDMX structure mapping is to transform datasets from one dimensionality to another. In practice, this means that the input and output datasets conform to different Data Structure Definition.

- 3186 Structure mapping does not alter the observation values and is not intended to perform 3187 any aggregations or calculations.
- 3188 An input series maps to:
- 3189 a. Exactly one output series; or
- b. Multiple output series with different Series Keys, but the same observation
 values; or
- 3192 c. Zero output series where no source rule matches the input Component values. 3193
- 3194 Typical use cases include:
- Transforming received data into a common internal structure;
- Transforming reported data into the data collector's preferred structure;
- Transforming unidimensional datasets⁴⁷ to multi-dimensional; and
- Transforming internal datasets with a complex structure to a simpler structure with fewer dimensions suitable for dissemination.

3200 **12.21-1** *structure maps*

1-1 (pronounced 'one to one') mappings support the simple use case where the value
of a Component in the source structure is translated to a different value in the target,
usually where different classification schemes are used for the same Concept.

In the example below, ISO 2-character country codes are mapped to their ISO 3character equivalent.

3207

Country	Alpha-2 code	Alpha-3 code
Afghanistan	AF	AFG
Albania	AL	ALB
Algeria	DZ	DZA
American Samoa	AS	ASM
Andorra	AD	AND
etc		

3208

Different source values can also map to the same target value, for example when deriving regions from country codes.

3211

⁴⁷ Unidimensional datasets are those with a single 'indicator' or 'series code' dimension.



Statistical Data and	Metadata eXchange
----------------------	-------------------

Source Component:	Target Component:
REF_AREA	REGION
FR	EUR
DE	EUR
IT	EUR
ES	EUR
BE	EUR

3213 12.3N-n structure maps

N-n (pronounced 'N to N') mappings describe rules where a specified combination of values in multiple source Components map to specified values in one or more target Components. For example, when mapping a partial Series Key from a highly multidimensional cube (like Balance of Payments) to a single 'Indicator' Dimension in a target Data Structure.

3219

3220 Example:

	1 -	l
Rule	Source	Target
1	If FREQUENCY=A; and ADJUSTMENT=N; and MATURITY=L.	Set INDICATOR=A_N_L
2	If FREQUENCY=M; and ADJUSTMENT=S_A1; and MATURITY=TY12.	Set INDICATOR=MON_SAX_12

3221 3222

N-n rules can also set values for multiple source Components.

Rule	Source	Target	
1	lf	Set	
	FREQUENCY=A; and	INDICATOR=A_N_L,	
	ADJUSTMENT=N; and	STATUS=QXR15,	
	MATURITY=L.	NOTE="Unadjusted".	
2	lf	Set	
	FREQUENCY=M; and	INDICATOR=MON_SAX_12,	
	ADJUSTMENT=S_A1; and	STATUS=MPM12,	
	MATURITY=TY12.	NOTE="Seasonally Adjusted"	

3223



3224 **12.4 Ambiguous mapping rules**

A structure map is ambiguous if the rules result in a dataset containing multiple series with the same Series Key.

3227

A simple example mapping a source dataset with a single dimension to one with multiple dimensions is shown below:

3230

Source	Target	Output Series Key
SERIES_CODE=XMAN_Z_21	Dimensions INDICATOR=XM FREQ=A ADJUSTMENT=N Attributes UNIT_MEASURE=_Z	XM:A:N
SERIES_CODE=XMAN_Z_34	Dimensions INDICATOR=XM FREQ=A ADJUSTMENT=N Attributes UNIT_MEASURE=_Z COMP_ORG=34	XM:A:N

3231

The above behaviour can be okay if the series XMAN_Z_21 contains observations for different periods of time then the series XMAN_Z_34. If however both series contain observations for the same point in time, the output for this mapping will be two observations with the same series key, for the same period in time.

3236 **12.5 Representation maps**

Representation Maps replace the SDMX 2.1 Codelist Maps and are used describe
explicit mappings between source and target Component values.

3240 The source and target of a Representation Map can reference any of the following:

- 3241
- and larger of a Representation Map Ca a. Codelist
- 3241
- b. Free Text (restricted by type, e.g String, Integer, Boolean)
- c. Valuelist
- 3243 3244

A Representation Map mapping ISO 2-character to ISO 3-character Codelists would take the following form:

CL_ISO_ALPHA2	CL_ISO_ALPHA3
AF	AFG
AL	ALB
DZ	DZA
AS	ASM
AD	AND
etc	

3247

3248 A Representation Map mapping free text country names to an ISO 2-character Codelist

3249 could be similarly described:



Text	CL_ISO_ALPHA2
"Germany"	DE
"France"	FR
"United Kingdom"	GB
"Great Britain"	GB
"Ireland"	IE
"Eire"	IE
etc	

Valuelists, introduced in SDMX 3.0, are equivalent to Codelists but allow the
 maintenance of non-SDMX identifiers. Importantly, their IDs do not need to conform to
 IDType, but as a consequence are not Identifiable.

When used in Representation Maps, Valuelists allow Non-SDMX identifiers containing characters like £, \$, % to be mapped to Code IDs, or Codes mapped to non-SDMX identifiers.

In common with Codelists, each item in a Valuelist has a multilingual name giving it a human-readable label and an optional description. For example:

Value	Locale	Name
\$	en	United States Dollar
%	En	Percentage
	fr	Pourcentage

3259

- 3260 Other characteristics of Representation Maps:
- Support the mapping of multiple source Component values to multiple Target
 Component values as described in section 12.3 on n-to-n mappings; this covers
 also the case of mapping an Attribute with an array representation to map
 combinations of values to a single target value;
- Allow source or target mappings for an Item to be optional allowing rules such as 'A maps to nothing' or 'nothing maps to A'; and
- Support for mapping rules where regular expressions or substrings are used to match source Component values. Refer to section 12.6 for more on this topic.

3269 **12.6 Regular expression and substring rules**

- It is common for classifications to contain meanings within the identifier, for example
 the code Id 'XULADS' may refer to a particular seasonality because it starts with the
 letters XU.
- With SDMX 2.1 each code that starts with XU had to be individually mapped to the same seasonality, and additional mappings added when new Codes were added to the Codelists. This led to many hundreds or thousands of mappings which can be more efficiently summarised in a single conceptual rule:
- 3277 If starts with 'XU' map to 'Y'



These rules are described using either regular expressions, or substrings for simpler use cases.

3280 12.6.1 Regular expressions

- 3281 Regular expression mapping rules are defined in the Representation Map.
- 3282 Below is an example set of regular expression rules for a particular component.

Regex	Description	Output
A	Rule match if input = 'A'	OUT_A
^[A-G]	Rule match if the input starts with letters A to G	OUT_B
A B	Rule match if input is either 'A' or 'B'	OUT_C

3283

- Like all mapping rules, the output is either a Code, a Value or free text depending on the representation of the Component in the target Data Structure Definition.
- 3286 If the regular expression contains capture groups, these can be used in the definition 3287 of the output value, by specifying n as an output value where **n** is the number of the 3288 capture group starting from 1. For example

3289

Regex	Target output	Example Input	Example Output
([0-9]{4})[0- 9]([0-9]{1})	\1-Q\2	200933	2009-Q3

3290

As regular expression rules can be used as a general catch-all if nothing else matches,
the ordering of the rules is important. Rules should be tested starting with the highest
priority, moving down the list until a match is found.

3294 The following example shows this:

Priority	Regex	Description	Output
1	A	Rule match if input = 'A'	OUT_A
2	В	Rule match if input = 'B'	OUT_B
3	[A-Z]	Any character A-Z	OUT_C

3295

- The input 'A' matches both the first and the last rule, but the first takes precedence having the higher priority. The output is OUT_A.
- The input 'G' matches on the last rule which is used as a catch-all or default in this example.

3300 **12.6.2 Substrings**

3301 Substrings provide an alternative to regular expressions where the required section of 3302 an input value can be described using the number of the starting character, and the 3303 length of the substring in characters. The first character is at position 1.



3304 For instance:

Input String	Start	Length	Output
ABC_DEF_XYZ	5	3	DEF
XULADS	1	2	XU

3305

3306 Sub-strings can therefore be used for the conceptual rule *If starts with 'XU' map to Y* 3307 as shown in the following example:

Start	Length	Source	Target
1	2	XU	Y

3308 **12.7 Mapping non-SDMX time formats to SDMX formats**

- 3309 Structure mapping allows non-SDMX compliant time values in source datasets to be 3310 mapped to an SDMX compliant time format.
- 3311 Two types of time input are defined:
- a. Pattern based dates a string which can be described using a notation like
 dd/mm/yyyy or is represented as the number of periods since a point in time, for
 example: 2010M001 (first month in 2010), or 2014D123 (123rd day in 2014); and
- b.Numerical based datetime a number specifying the elapsed periods since a
 fixed point in time, for example Unix Time is measured by the number of
 milliseconds since 1970.

The output of a time-based mapping is derived from the output Frequency, which is either explicitly stated in the mapping or defined as the value output by a specific Dimension or Attribute in the output mapping. If the output frequency is unknown or if the SDMX format is not desired, then additional rules can be provided to specify the output date format for the given frequency Id. The default rules are:

Frequency	Format	Example
A	YYYY	2010
D	YYYY-MM-DD	2010-01-01
1	YYYY-MM-DD- Thh:mm:ss	2010-01T20:22:00
Μ	YYYY-MM	2010-01
Q	YYYY-Qn	2010-Q1
S	YYYY-Sn	2010-S1
Т	YYYY-Tn	2010-T1



W	YYYY-Wn	YYYY-W53
---	---------	----------

3327

3328 3329

3330

3331

3332

3333 3334

In the case where the input frequency is lower than the output frequency, the mapping
defaults to end of period, but can be explicitly set to start, end or mid-period.

There are two important points to note:

- 1. The output frequency determines the output date format, but the default output can be redefined using a Frequency Format mapping to force explicit rules on how the output time period is formatted.
- 2. To support the use case of changing frequency the structure map can optionally provide a start of year attribute, which defines the year start date in MM-DD format. For example: YearStart=04-01.

3335 **12.7.1 Pattern based dates**

Date and time formats are specified by date and time pattern strings based on Java's Simple Date Format. Within date and time pattern strings, unquoted letters from 'A' to 'Z' and from 'a' to 'z' are interpreted as pattern letters representing the components of a date or time string. Text can be quoted using single quotes (') to avoid interpretation. """ represents a single quote. All other characters are not interpreted; they're simply copied into the output string during formatting or matched against the input string during parsing.

3343 Due to the fact that dates may differ per locale, an optional property, defining the locale 3344 of the pattern, is provided. This would assist processing of source dates, according to 3345 the given locale⁴⁸. An indicative list of examples is presented in the following table:

English (en)	Australia (AU)	en-AU
English (en)	Canada (CA)	en-CA
English (en)	United Kingdom (GB)	en-GB
English (en)	United States (US)	en-US
Estonian (et)	Estonia (EE)	et-EE
Finnish (fi)	Finland (FI)	fi-Fl
French (fr)	Belgium (BE)	fr-BE
French (fr)	Canada (CA)	fr-CA
French (fr)	France (FR)	fr-FR
French (fr)	Luxembourg (LU)	fr-LU
French (fr)	Switzerland (CH)	fr-CH
German (de)	Austria (AT)	de-AT
German (de)	Germany (DE)	de-DE

⁴⁸ A list of commonly used locales can be found in the Java supported locales: https://www.oracle.com/java/technologies/javase/jdk8-jre8-suported-locales.html



German (de)	Luxembourg (LU)	de-LU
German (de)	Switzerland (CH)	de-CH
Greek (el)	Cyprus (CY)	el-CY <u>(*)</u>
Greek (el)	Greece (GR)	el-GR
Hebrew (iw)	Israel (IL)	iw-IL
Hindi (hi)	India (IN)	hi-IN
Hungarian (hu)	Hungary (HU)	hu-HU
Icelandic (is)	Iceland (IS)	is-IS
Indonesian (in)	Indonesia (ID)	in-ID <u>(*)</u>
Irish (ga)	Ireland (IE)	ga-IE <u>(*)</u>
Italian (it)	Italy (IT)	it-IT

3347 Examples

- 3348 22/06/1981 would be described as dd/MM/YYYY, with locale en-GB
- 3349 2008-mars-12 would be described as YYYY-MMM-DD, with locale fr-FR
- 3350 22 July 1981 would be described as dd MMMM YYYY, with locale en-US
- 3351 22 Jul 1981 would be described as dd MMM YYYY
- 3352 2010 D62 would be described as YYYYDnn (day 62 of the year 2010)
- The following pattern letters are defined (all other characters from 'A' to 'Z' and from 'a' to 'z' are reserved):

Letter	Date or Time Component	Presentation	Examples
G	Era designator	Text	AD
уу	Year short (upper case is Year of Week ⁴⁹)	Year	96
уууу	Year Full (upper case is Year of Week)	Year	1996
MM	Month number in year starting with 1	Month	07
MMM	Month name short	Month	Jul
MMMM	Month name full	Month	July
ww	Week in year	Number	27
W	Week in month	Number	2
DD	Day in year	Number	189
dd	Day in month	Number	10
F	Day of week in month	Number	2
E	Day name in week	Text	Tuesday; Tue

⁴⁹ yyyy represents the calendar year while YYYY represents the year of the week, which is only relevant for 53 week years



U	Day number of week (1 = Monday,, 7 = Sunday)	Number	1
HH	Hour in day (0-23)	Number	0
kk	Hour in day (1-24)	Number	24
KK	Hour in am/pm (0-11)	Number	0
hh	Hour in am/pm (1-12)	Number	12
mm	Minute in hour	Number	30
SS	Second in minute	Number	55
S	Millisecond	Number	978
n	Number of periods, used after a SDMX Frequency Identifier such as M, Q, D (month, quarter, day)	Number	12

3356 The model is illustrated below:



3357

3358Figure 24 showing the component map mapping the SOURCE_DATE Dimension to the3359TIME_PERIOD dimension with the additional information on the component map to3360describe the time format





3363

3368

3371

3372

Figure 25 showing an input date format, whose output frequency is derived from the output value of the FREQ Dimension

3364 12.7.2 Numerical based datetime

Where the source datetime input is purely numerical, the mapping rules are defined by the **Base** as a valid SDMX Time Period, and the **Period** which must take one of the following enumerated values:

- day
- second
- 3369 3370 •
 - millisecondmicrosecond
 - nanosecond

Numerical datetime systems	Base	Period
Epoch Time (UNIX)	1970	millisecond
Milliseconds since 01 Jan 1970		
Windows System Time	1601	millisecond
Milliseconds since 01 Jan 1601		

3373

The example above illustrates numerical based datetime mapping rules for two commonly used time standards.

3376 The model is illustrated below:

		Type=numerical Base=1970 Period=millisecond
Compone	ent Map	Output Fred=A
SOURCE_DATE	TIME_PERIOD	

3377

3378Figure 26 showing the component map mapping the SOURCE_DATE Dimension to the3379TIME_PERIOD Dimension with the additional information on the component map to

3380 describe the numerical datetime system in use



3381 **12.7.3 Mapping more complex time inputs**

3382 VTL should be used for more complex time inputs that cannot be interpreted using thepattern based on numerical methods.

3384 **12.8Using TIME_PERIOD in mapping rules**

The source TIME_PERIOD Dimension can be used in conjunction with other input Dimensions to create discrete mapping rules where the output is conditional on the time period value.

3388 The main use case is setting the value of Observation Attributes in the target dataset.

Rule	Source	Target
1	lf	Set
	INDICATOR=XULADS; and	OBS_CONF=F
	TIME_PERIOD=2007.	
2	lf	Set
	INDICATOR=XULADS; and	OBS_CONF=F
	TIME_PERIOD=2008.	
3	lf	Set
	INDICATOR=XULADS; and	OBS_CONF=F
	TIME_PERIOD=2009.	
4	lf	Set
	INDICATOR=XULADS; and	OBS_CONF= C
	TIME_PERIOD=2010.	

3389 In the example above, OBS_CONF is an Observation Attribute.

3390 **12.9 Time span mapping rules using validity periods**

3391 Creating discrete mapping rules for each TIME_PERIOD is impractical where rules
 3392 need to cover a specific span of time regardless of frequency, and for high-frequency
 3393 data.

3394 Instead, an optional validity period can be set for each mapping.

By specifying validity periods, the example from Section 12.8 can be re-written using two rules as follows:

Rule	Source	Target
1	If INDICATOR=XULADS. Validity Period start period=2007 end period=2009	Set OBS_CONF=F
2	If INDICATOR=XULADS. Validity Period start period=2010	Set OBS_CONF=F

3397

In Rule 1, start period resolves to the start of the 2007 period (2007-01-01T00:00:00), and the end period resolves to the very end of 2009 (2009-12-31T23:59:59). The rule



3400 will hold true regardless of the input data frequency. Any observations reporting data 3401 for the Indicator XULADS that fall into that time range will have an OBS CONF value 3402 of F.

3403 In Rule 2, no end period is specified so remains in effect from the start of the period (2010-01-01T00:00:00) until the end of time. Any observations reporting data for the 3404 3405 Indicator XULADS that fall into that time range will have an OBS CONF value of C.

12.10 Mapping examples 3406

3407 12.10.1 Many to one mapping (N-1)

Source	Мар То
FREQ="A"	FREQ="A"
ADJUSTMENT="N"	REF AREA="PL"
REF_AREA="PL"	COUNTERPART AREA="W0"
COUNTERPART_AREA="W0"	INDICATOR="IND ABC"
REF SECTOR="S1"	_
COUNTERPART SECTOR="S1"	
ACCOUNTING ENTRY="B"	
STO="B5G"	

3408

- 3409 The bold Dimensions map from source to target verbatim. The mapping simply 3410 specifies:
- 3411
- FREQ => FREQ REF AREA=> REF AREA 3412
- 3413 COUNTERPART AREA=> COUNTERPART AREA 3414
- 3415 No Representation Mapping is required. The source value simply copies across unmodified. 3416
- 3417

3418 The remaining Dimensions all map to the Indicator Dimension. This is an example of 3419 many Dimensions mapping to one Dimension. In this case a Representation 3420 Mapping is required, and the mapping first describes the input 'partial key' and how 3421 this maps to the target indicator: 3422

3423 N:S1:S1:B:B5G => IND ABC

3424 3425 Where the key sequence is based on the order specified in the mapping (i.e. ADJUSTMENT, REF SECTOR, etc will result in the first value N being taken from 3426 3427 ADJUSTMENT as this was the first item in the source Dimension list.

3428

3429 Note: The key order is NOT based on the Dimension order of the DSD, as the mapping 3430 needs to be resilient to the DSD changing.

3431

3432 12.10.2 Mapping other data types to Code Id

3433 In the case where the incoming data type is not a string and not a code identifier i.e. 3434 the source Dimension is of type Integer and the target is Codelist. This is supported by 3435 the RepresentationMap. The RepresentationMap source can reference a Codelist, 3436 Valuelist, or be free text, the free text can include regular expressions.

3437 The following representation mapping can be used to explicitly map each age to an 3438 output code.


Source Input	Desired Output
Free Text	Code Id
0	Α
1	Α
2	Α
3	В
4	В

3439

3440 If this mapping takes advantage of regular expressions it can be expressed in two 3441 rules:

Regular Expression	Desired Output
[0-2]	А
[3-4]	В

3442

3443 12.10.3 Observation Attributes for Time Period

This use case is where a specific observation for a specific time period has an attribute value.

Input INDICATOR	Input TIME_PERIOD	Output OBS_CONF
XULADS	2008	С
XULADS	2009	С
XULADS	2010	С

3446

3447	Or using a validity period on the Representation Mapping:						
	Input INDICATOR	Valid From/ Valid To	Output OBS_CONF				
	XULADS	2008/2010	С				

3448

3449 12.10.4 Time mapping

This use case is to create a time period from an input that does not respect SDMX Time Formats.

3452The Component Mapping from SYS_TIME to TIME_PERIOD specifies itself as a time3453mapping with the following details:

Source Value	Source Mapping	Target Frequency	Output
18/07/1981	dd/MM/yyyy	А	1981

3454

When the target frequency is based on another target Dimension value, in this example
 the value of the FREQ Dimension in the target DSD.
 Source Value

		Dimensio	n	Output
18/07/1981	dd/MM/yyyy	FREQ		1981-07-18 (when FREQ=D)
When the source	is a numerical forn	nat		
Source Value	Start Period	Interval	Target FREQ	Output

3460



3461 When the source frequency is lower than the target frequency additional information 3462 can be provided for resolve to start of period, end of period, or mid period, as shown 3463 in the following example:

Source Value	Source Mapping	Target Frequency	Output
		Dimension	-
1981	уууу	D – End of Period	1981-12-31
When the start of	vear is April 1 st the Stru	icture Man has YearStart	-04-01
When the start of Source Value	year is April 1 st the Stru Source Mapping	ucture Map has YearStart Target Frequency	=04-01: Output
When the start of Source Value	year is April 1 st the Stru Source Mapping	ucture Map has YearStart Target Frequency Dimension	=04-01: Output

3466



13 ANNEX Semantic Versioning 3467

13.1 Introduction to Semantic Versioning 3468

In the world of versioned data modelling exists a dreaded place called "dependency 3469 hell." The bigger your data model through organisational, national or international 3470 harmonisation grows and the more artefacts you integrate into your modelling, the 3471 3472 more likely you are to find yourself, one day, in this pit of despair.

3473

In systems with many dependencies, releasing new artefact versions can quickly 3474 3475 become a nightmare. If the dependency specifications are too tight, you are in danger of version lock (the inability to upgrade an artefact without having to release new 3476 3477 versions of every dependent artefact). If dependencies are specified too loosely, you will inevitably be bitten by version promiscuity (assuming compatibility with more future 3478 3479 versions than is reasonable). Dependency hell is where you are when version lock 3480 and/or version promiscuity prevent you from easily and safely moving your data 3481 modelling forward.

3482 3483 As a very successful solution to the similar problem in software development, 3484 "Semantic Versioning" semver.org proposes a simple set of rules and requirements 3485 that dictate how version numbers are assigned and incremented. These rules make 3486 also perfect sense in the world of versioned data modelling and help to solve the "dependency hell" encountered with previous versions of SDMX. SDMX 3.0 applies 3487 3488 thus the Semantic Versioning rules on all versioned SDMX artefacts. Once you release 3489 a versioned SDMX artefact, you communicate changes to it with specific increments

3490 to your version number.

3491

3492 This SDMX 3.0(.0) specification inherits the original semver.org 2.0.0 wording 3493 (license: Creative Commons - CC BY 3.0) and applies it to versioned SDMX structural artefacts. Under this scheme, version numbers and the way they change 3494 3495 convey meaning about the underlying data structures and what has been modified from 3496 one version to the next. 3497

13.2 Semantic Versioning Specification for SDMX 3.0(.0) 3498

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT". 3499 "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this 3500 3501 document are to be interpreted as described in RFC 2119. 3502

In the following, "versioned" artefacts are understood to be semantically versioned 3503 SDMX structural artefacts. 3504

- 3505
- 3506
- All versioned SDMX artefacts MUST specify a version number.

3507 The version number of immutable versioned SDMX artefacts MUST take the 3508 form X.Y.Z where X, Y, and Z are non-negative integers and MUST NOT 3509 contain leading zeroes. X is the major version, Y is the minor version, and Z is the patch version. Each element MUST increase numerically. For instance: 3510 3511 1.9.0 -> 1.10.0 -> 1.11.0.



- Once an SDMX artefact with an X.Y.Z version has been publicly released, the contents of that version MUST NOT be modified. That artefact version is considered stable. Any modifications MUST be released as a new version.
- Major version zero (0.y.z) is for initial modelling. Anything MAY change at any time. The public artefact SHOULD NOT be considered stable.
- Version 1.0.0 defines the first stable artefact. The way in which the version number is incremented after this release is dependent on how this public artefact changes.
- Patch version Z (x.y.Z | x > 0) MUST be incremented if only backwards compatible property changes are introduced. A property change is defined as an internal change that does not affect the relationship to other artefacts. These are changes in the artefact's or artefact element's names, descriptions and annotations that MUST NOT alter their meaning.
- Minor version Y (x.Y.z | x > 0) MUST be incremented if a new, backwards compatible element is introduced to a stable artefact. These are additional items in ItemScheme artefacts. It MAY be incremented if substantial new information is introduced within the artefact's properties. It MAY include patch level changes. Patch version MUST be reset to 0 when minor version is incremented.
- Major version X (X.y.z | X > 0) MUST be incremented if any backwards incompatible changes are introduced to a stable artefact. These often relate to deletions of items in ItemSchemes or to backwards incompatibility introduced due to changes in references to other artefacts. A major version change MAY also include minor and patch level changes. Patch and minor version MUST be reset to 0 when major version is incremented.
- 3537 A mutable version, e.g. used for public drafts or as pre-release, MUST be denoted by appending an Extension that consists of a hyphen and a series of 3538 3539 dot separated identifiers immediately following the patch version (x.y.z-EXT). 3540 Identifiers MUST comprise only ASCII alphanumerics and hyphen [0-9A-Za-z-]. Identifiers MUST NOT be empty. Numeric identifiers MUST NOT include 3541 3542 leading zeroes. However, to foster harmonisation and general comprehension 3543 it is generally recommended to use the standard extension "-draft". Extended versions have a lower precedence than the associated stable version. An 3544 extended version indicates that the version is unstable and it might not satisfy 3545 3546 the intended compatibility requirements as denoted by its associated stable 3547 version. When making changes to an SDMX artefact with an extended version number then one is not required to increment the version if those changes are 3548 kept within the allowed scope of the version increment from the previous 3549 3550 version (if that existed), otherwise also here the before mentioned version increment rules for X.Y.Z apply. Concretely, a version X.0.0-EXT will allow for 3551 any changes, a version X.Y.0-EXT will allow only for minor changes and a 3552 3553 version X.Y.Z-EXT will allow only for any patch changes, as defined above. Extension examples: 1.0.0-draft, 1.0.0-draft.1, 1.0.0-0.3.7, 1.0.0-x.7.z.92. 3554
- Precedence refers to how versions are compared to each other when ordered.
 Precedence MUST be calculated by separating the version into major, minor, patch and extension identifiers in that order. Precedence is determined by the first difference when comparing each of these identifiers from left to right as follows: Major, minor, and patch versions are always compared numerically.



Example: $1.0.0 < 2.0.0 < 2.1.0 < 2.1.1$. When major, minor, and patch are equal,
an extended version has lower precedence than a stable version. Example:
1.0.0-draft < 1.0.0. Precedence for two extended versions with the same major,
minor, and patch version MUST be determined by comparing each dot
separated identifier from left to right until a difference is found as follows:
identifiers consisting of only digits are compared numerically and identifiers with
letters or hyphens are compared lexically in ASCII sort order. Numeric
identifiers always have lower precedence than non-numeric identifiers. A larger
set of extension fields has a higher precedence than a smaller set, if all of the
preceding identifiers are equal. Example: 1.0.0-draft < 1.0.0-draft.1 < 1.0.0-
draft.prerelease < 1.0.0-prerelease < 1.0.0-prerelease.2 < 1.0.0-prerelease.11
< 1.0.0-rc.1 < 1.0.0.

- The reasons for version changes MAY be documented in brief form in an artefact's annotation of type "CHANGELOG".

3575 13.3 Backus–Naur Form Grammar for Valid SDMX 3.0(.0) 3576 Semantic Versions

```
3577
        <valid semver> ::= <version core>
3578
                          | <version core> "-" <extension>
3579
3580
        <version core> ::= <major> "." <minor> "." <patch>
3581
3582
        <major> ::= <numeric identifier>
3583
3584
        <minor> ::= <numeric identifier>
3585
3586
        <patch> ::= <numeric identifier>
3587
3588
        <extension> ::= <dot-separated extension identifiers>
3589
3590
        <dot-separated extension identifiers> ::= <extension identifier>
3591
                                                     | <extension identifier> "." <dot-
3592
        separated extension identifiers>
3593
3594
        <extension identifier> ::= <alphanumeric identifier>
3595
                                     | <numeric identifier>
3596
3597
        <alphanumeric identifier> ::= <non-digit>
3598
                                      | <non-digit> <identifier characters>
                                      | <identifier characters> <non-digit>
| <identifier characters> <non-digit> <identifier</pre>
3599
3600
3601
        characters>
3602
3603
        <numeric identifier> ::= "0"
3604
                                | <positive digit>
3605
                                 | <positive digit> <digits>
3606
3607
        <identifier characters> ::= <identifier character>
3608
                                    | <identifier character> <identifier characters>
3609
3610
        <identifier character> ::= <digit>
3611
                                  | <non-digit>
3612
3613
        <non-digit> ::= <letter>
                      1 "-"
3614
3615
3616
        <digits> ::= <digit>
```



| <digit> <digits>

<digit> ::= "0"

| <positive digit>

2	<positive< th=""><th>a dig</th><th>git></th><th>:</th><th>:= "1</th><th>1"</th><th> "2</th><th>2"</th><th> "3</th><th>3"</th><th> "4</th><th>1"</th><th> "5</th><th>5"</th><th> "6</th><th>5"</th><th>1 "7</th><th>7 ''</th><th> "8</th><th>3 ''</th><th> "9</th><th>)"</th></positive<>	a dig	git>	:	:= "1	1"	"2	2"	"3	3"	"4	1"	"5	5"	"6	5"	1 "7	7 ''	"8	3 ''	"9) "
ł	<letter></letter>	::=	"A"	Ι	"в"	Ι	"C"	T	"D"	Ι	"E"	Т	"F"	T	"G"	Т	"Н"	Τ	"I"	T	"J"	
5		1	"K"	Т	"L"	Т	"M"		"N"	Т	"0"	Т	"P"	Т	"Q"	Т	"R"	Т	"S"	T	"T"	
5		1	ייטיי	Т	"V"	T	"W"	T	"X"	Т	"Y"	T	"Z"	T	"a"	Т	"b"	T	"c"	1	"d"	
7		1	"e"	Т	"£"	Т	"q"	Т	"h"	Т	"i"	Т	"j"	Т	"k"	Т	"1"	Т	"m"	1	"n"	
3		1	"o"	Т	"p"	Т	"q"	Т	"r"	Т	"s"	Т	"t"	Т	"u"	Т	"v"	Т	"w"	1	"x"	
)		Ì	"y"	Ì	- "z"	-	-	-						-				-				

3631	13.4 Dependency Management in SDMX 3.0(.0):
3632 3633	MAJOR, MINOR or PATCH version parts in SDMX 3.0 artefact references CAN be wildcarded using "+" as extension:
3634	• X+.Y.Z means the currently latest available version >= X.Y.Z
3635 3636	 Example: "2+.3.1" means the currently latest available version >= "2.3.1" (even if not backwards compatible)
3637	 Typical use case: references in SDMX Categorisations
3638 3639	 X.Y+.Z means the currently latest available backwards compatible version >= X.Y.Z
3640 3641 3642	 Example: "2.3+.1" means the currently latest available version >= "2.3.1" and < "3.0.0" (all backwards compatible versions >= "2.3.1")
3643	 Typical use case: references in SDMX DSD
3644 3645	 X.Y.Z+ means the currently latest available forwards and backwards compatible version >= X.Y.Z
3646 3647 3648	 Example: "2.3.1+" means the currently latest available version >= "2.3.1" and < "2.4.0" (all forwards and backwards compatible versions >= "2.3.1")
3649 3650 3651 3652	 Non-versioned and 2-digit version SDMX structural artefacts CAN reference any other non-versioned or versioned (whether SemVer or not) SDMX structural artefacts.
3653 3654	 Semantically versioned artefacts MUST only reference other semantically versioned artefacts.
3655 3656	 Wildcarded references in a stable artefact implicitly target only future stable versions of the referenced artefacts within the defined wildcard scope.
3657 3658 3659 3660	 Example: The reference to "AGENCY_ID:CODELIST_ID(2.3+.1)" in an artefact "AGENCY_ID:DSD_ID(2.2.1)" resolves to artefact "AGENCY_ID:CODELIST_ID(2.4.3)" if that was currently the latest available stable version.



- Wildcarded references in a version-extended artefact implicitly target future stable and version-extended versions of the referenced artefacts within the defined wildcard scope.
 - Example: The reference to "AGENCY_ID:CODELIST_ID(2.3+.1)" in an artefact "AGENCY_ID:DSD_ID(2.2.1-draft)" resolves to artefact "AGENCY_ID:CODELIST_ID(2.5.0-draft)" if that was currently the latest available version.
- References to specific version-extended artefacts MAY be used, but those cannot be combined with a wildcard.
- 3670 3671

3672 3673

3664

3665

3666 3667

> • Example: The reference to "AGENCY_ID:CODELIST_ID(2.5.0draft)" in an artefact "AGENCY_ID:DSD_ID(2.2.1)" resolves to artefact "AGENCY_ID:CODELIST_ID(2.5.0-draft)", which might be subject to continued backwards compatible changes.

3674 Because both, wildcarded references and references to version-extended artefacts, 3675 allow for changes in the referenced artefacts, care needs to be taken when choosing 3676 the appropriate references in order to achieve the required limitation in the allowed 3677 scope of changes.

3678

367913.5 Upgrade and conversions of artefacts defined with
previous SDMX standard versions to Semantic Versioning

Because SDMX standardises the interactions between statistical systems, which cannot all be upgraded at the same time, the new versioning rules cannot be applied to existing artefacts in EDIFACT, SDMX 1.0, 2.0 or 2.1. SemVer can only be applied to structural artefacts that are newly modelled with the SDMX 3.0 Information Model. Migrating to SemVer means migrating to the SDMX 3.0 Information Model, to its new API version and new versions of its exchange message formats.

3687

3688 To migrate SDMX structural artefacts created previously to SDMX 3.0.0: 3689

3690 If the artefacts do not need versioning, then the new artefacts based on the SDMX 3.0
3691 Information Model SHOULD be made version-less, e.g., a previous artefact with the
3692 non-final version 1.0 and that doesn't need versioning becomes non-versioned. This
3693 will be the case for all AgencyScheme artefacts.

3694 3695 If artefa

3695 If artefact versioning is required and SDMX 3.0.0 Semantic Versioning is available
3696 within the tools and processes used, then it is recommended to switch to Semantic
3697 Versioning with the following steps:

- 36981. Complement the missing version parts with 0s to make the version number3699SemVer-compliant using the MAJOR.MINOR.PATCH-EXTENSION syntax:
- 3700Example: Version 2 becomes version 2.0.0 and version 3.1 becomes version37013.1.0.
- 370214. Replace the "isFinal=false" property by the version extensions "-
draft" (or alternatively "-unstable" or "-nonfinal" depending
on the use case).



3705Example: Version 1.3 with isFinal=true becomes version 1.3.0 and version37061.3 with isFinal=false becomes version 1.3.0-draft.

3707 If artefact versioning is required but semantic versioning cannot be applied, then
3708 version strings used in previous versions of the Standard (e.g., version=1.2) may
3709 continue to be used.
3710

Note: Like for other not fully backwards compatible SDMX 3.0 features, also some cases of semantically versioned SDMX 3.0 artefacts cannot be converted back to earlier SDMX versions. This is the case when one or more extensions have been created in parallel to the corresponding stable version. In this case, only the stable version SHOULD be converted to a final version (e.g., 3.2.1 becomes 3.2.1 final, and 3.2.1-draft cannot be converted back).

3717

3718 **13.6 FAQ for Semantic Versioning**

3719 My organisation is new to SDMX and starts to implement 3.0 or starts to 3720 implement a new process fully based on SDMX 3.0. Which versioning scheme 3721 should be used? 3722

3723 If all counterparts involved in the process and all tools used for its implementation are3724 SDMX 3.0-ready, then it is recommended to:

- in general, use semantic versioning;
- exceptionally, do not use versioning for artefacts that do not require it, e.g. artefacts that never change, that are only used internally or for which communication on changes with external parties or systems is not required.
- 3729

3730 How should I deal with revisions in the 0.y.z initial modelling phase?

3731

3736

3741

3743

The simplest thing to do is start your initial modelling release at 0.1.0 and then increment the minor version for each subsequent release.

3735 How do I know when to release 1.0.0?

If your data model is being used in production, it should probably already be 1.0.0. If
you have a stable artefact on which users have come to depend, you should be 1.0.0.
If you're worrying a lot about backwards compatibility, you should probably already be
1.0.0.

3742 Doesn't this discourage rapid modelling and fast iteration?

Major version zero is all about rapid modelling. If you're changing the artefact every day you should either still be in version 0.y.z or on the next (minor or) major version for a separate modelling.

3747

3748If even the tiniest backwards incompatible changes to the public artefact require3749a major version bump, won't I end up at version 42.0.0 very rapidly?

3750

This is a question of responsible modelling and foresight. Incompatible changes should not be introduced lightly to a data model that has a lot of dependencies. The cost that



3753 must be incurred to upgrade can be significant. Having to bump major versions to 3754 release incompatible changes means you will think through the impact of your 3755 changes, and evaluate the cost/benefit ratio involved. 3756

3757 Documenting the version changes in an artefact's annotation of type "CHANGELOG" is too much work! 3758

It is your responsibility as a professional modeller to properly document the artefacts 3760 3761 that are intended for use by others. Managing data model complexity is a hugely 3762 important part of keeping a project efficient, and that's hard to do if nobody knows how 3763 to use your data model, or what artefacts are safe to reuse. In the long run, SDMX 3.0 3764 Semantic Versioning can keep everyone and everything running smoothly.

However, refrain from overdoing. Nobody can and will read too long lists of changes. 3765 3766 Thus, keep it to the absolute essence, and mainly use it for short announcements. You 3767 can even skip the changelog if the change is impact-less. For all complete reports, a new API feature could be more useful to automatically generate a log of differences 3768 between two versions. 3769 3770

3771 What do I do if I accidentally release a backwards incompatible change as a 3772 minor version?

3774 As soon as you realise that you've broken the SDMX 3.0 Semantic Versioning 3775 specification, fix the problem and release a new minor version that corrects the 3776 problem and restores backwards compatibility. Even under this circumstance, it is unacceptable to modify versioned releases. If it's appropriate, document the offending 3777 3778 version and inform your users of the problem so that they are aware of the offending 3779 version. 3780

What if I inadvertently alter the public artefact in a way that is not compliant with 3781 3782 the version number change (i.e. the modification incorrectly introduces a major 3783 breaking change in a patch release)? 3784

3785 Use your best judgement. If you have a huge audience that will be drastically impacted by changing the behaviour back to what the public artefact intended, then it may be 3786 3787 best to perform a major version release, even though the property change could strictly 3788 be considered a patch release. Remember, SDMX 3.0.0 Semantic Versioning is all 3789 about conveying meaning by how the version number changes. If these changes are important to your users, use the version number to inform them. 3790

- 3792
- 3793

3791

3759

3773

How should I handle deprecating elements?

3794 Deprecating existing elements is a normal part of data modelling and is often required 3795 to make forward progress or follow history (changing classifications, evolving reference 3796 areas). When you deprecate part of your stable artefact, you should issue a new minor 3797 version with the deprecation in place (e.g. add the new country code but still keep the old country code) and with a "CHANGELOG" annotation announcing the deprecation 3798 3799 (e.g. the intention to remove the old country code in a future version). Before you 3800 completely remove the functionality in a new major release there should be at least 3801 one minor release that contains the deprecation so that users can smoothly transition 3802 to the new artefact.

3803

3804 Does SDMX 3.0.0 Semantic Versioning have a size limit on the version string?



3805 3806 No, but use good judgement. A 255 character version string is probably overkill, for 3807 example. In addition, specific SDMX implementations may impose their own limits on the size of the string. Remember, it is generally recommended to use the standard 3808 extension "-draft". 3809 3810 Is "v1.2.3" a semantic version? 3811 3812 3813 No, "v1.2.3" is not a semantic version. The semantic version is "1.2.3". 3814 3815 Is there a suggested regular expression (RegEx) to check an SDMX 3.0.0 3816 Semantic Versioning string? 3817 3818 There are two: 3819 3820 One with named groups for those systems that support them (PCRE [Perl Compatible 3821 Regular Expressions, i.e. Perl, PHP and R], Python and Go). 3822 3823 original Reduced version (without SemVer "build metadata") from: 3824 https://regex101.com/r/Ly7O1x/3/ 3825 3826 ^(?P<major>0|[1-9]\d*)\.(?P<minor>0|[1-9]\d*)\.(?P<patch>0|[1-9]\d*) (?:-(?P<extension>(?:0|[1-9]\d*|\d*[a-zA-Z-][0-9a-zA-Z-] 3827 3828]*)(?:\.(?:0|[1-9]\d*|\d*[a-zA-Z-][0-9a-zA-Z-]*))*))?\$ 3829 And one with numbered capture groups instead (so cg1 = major, cg2 = minor, cg3 = 3830 patch and cq4 = extension) that is compatible with ECMA Script (JavaScript), PCRE 3831 3832 (Perl Compatible Regular Expressions, i.e. Perl, PHP and R), Python and Go. 3833 3834 Reduced version (without original SemVer "build metadata") from: 3835 https://regex101.com/r/vkijKf/1/ 3836 3837 (0|[1-9]d*).(0|[1-9]d*).(0|[1-9]d*).(2:-((2:0)[1-0]d*)9]\d*|\d*[a-zA-Z-][0-9a-zA-Z-]*)(?:\.(?:0|[1-9]\d*|\d*[a-zA-Z-3838 3839][0-9a-zA-Z-]*))*))?\$ 3840 3841 Must I adopt semantic versioning rules when switching to SDMX 3.0? 3842 3843 No. If backwards compatibility with pre-existing tools and processes is required, then 3844 it is possible to continue using the previous versioning scheme (with up to two version 3845 parts MAJOR.MINOR). Semantic versioning is indicated only for those use cases 3846 where a proper artefact versioning is required. If versioning does not apply to some or 3847 all of your artefacts, then rather migrate to non-versioned SDMX 3.0 artefacts. 3848 3849 May I mix artefacts that follow semantic versioning with artefacts that don't? 3850 3851 Artefacts that are not (semantically) versioned may reference artefacts that are 3852 semantically versioned, but those are fully safe to use only when not extended. However, the reverse is not true: non-semantically-versioned artefacts do not offer 3853 3854 change guarantees, and, therefore, should not be referenced by semantically versioned artefacts. 3855



I have plenty of artefacts. I'm happy with my current versioning policy and I don't want to use SemVer! Can I still migrate to SDMX 3.0, and if so, what do I need to do?

3861 Yes, of course, you can. The introduction of semantic versioning is done in a way which is largely backward compatible with previous versions of the standard, so you can keep 3862 3863 your existing 2-digit version numbers (1.0, 1.1, 2.0, etc.) if that is required by your current tools and processes. However, if not using SemVer then pre-SDMX 3.0 final 3864 3865 artefacts will be migrated as non-final and mutable in SDMX 3.0. There are also many 3866 good reasons to move to SemVer, and the migration is encouraged. Be assured that 3867 there will be tools out there that will assist you doing this in an efficient and convenient 3868 way.

3869

I have plenty of artefacts versioned 'X.Y'. I want to make some of them immutable, and enjoy the benefits provided by semantic versioning. Some other artefacts however must remain mutable (i.e. non final). However, in both cases, I'd like adopt the semantic versioning. What do I need to do?

For artefacts that will be made immutable and are therefore safe to use, simply append a '.0' to the current version (use X.Y.0) when migrating to Semantic Versioning. E.g., if the version of your artefact is currently 1.10, then migrate to 1.10.0.

3878

3890

3895

For artefacts that remain mutable, and therefore do not bring the guarantees of semantic versioning, if you want to benefit from the advantages of semantic versioning, then simply append '.0-notfinal' to the version string. So, if the version of your artefact is currently 1.10, use 1.10.0-notfinal instead. Indeed, other extensions can be used depending on your use case.

I have adopted SDMX 3.0 with the semantic versioning conventions for the version strings of all my artefacts, regardless of whether these are stable (e.g. 1.0.0) or unstable (e.g. 1.0.0-notfinal, 1.0.0-draft, etc.). However, I still receive artefacts from organizations that have not yet adopted SemVer conventions for the version strings. How should I treat these?

- The only artefacts that are safe to use, are those that are semantically versioned. Starting with SDMX 3.0, these artefacts MUST use the SEMVER version string to indicate this fact and the version string of these artefacts MUST be expressed as X.Y.Z (e.g. 2.1.0). Extended versions bring some limited guarantees for changes.
- All other artefacts are in principle unsafe. They might be safe in practice but the SDMX
 standard does not bring any guarantees in that respect, and these artefacts may
 change in unpredictable ways.
- 3900 In practice, the migration approach will often mirror the way in which organisations 3901 have migrated between earlier SDMX versions. Rarely, the new data models used 3902 mixed SDMX standard versions in their dependencies, and if they did then standard 3903 conversions were put in place. A typical method is to first migrate the re-used artefacts from the previous SDMX version to SDMX 3.0 and while doing so to apply the 3904 3905 appropriate new semantic version string. From that point onwards, you can enjoy the 3906 advantages of the new SDMX versioning features for all those artefacts that require 3907 appropriate versioning.