



Context Information Management (CIM); NGSI-LD API

Disclaimer

The present document has been produced and approved by the cross-cutting Context Information Management (CIM) ETSI Industry Specification Group (ISG) and represents the views of those members who participated in this ISG.
It does not necessarily represent the views of the entire ETSI membership.

Reference

RGS/CIM-009v191

Keywords

API, architecture, digital twins, GAP, information model, interoperability, NGSI-LD, smart agriculture, smart city, smart industry, smart manufacturing, smart water, WoT

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - APE 7112B
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° w061004871

Important notice

The present document can be downloaded from the
[ETSI Search & Browse Standards](#) application.

The present document may be made available in electronic versions and/or in print. The content of any electronic and/or print versions of the present document shall not be modified without the prior written authorization of ETSI. In case of any existing or perceived difference in contents between such versions and/or in print, the prevailing version of an ETSI deliverable is the one made publicly available in PDF format on [ETSI deliver](#) repository.

Users should be aware that the present document may be revised or have its status changed,
this information is available in the [Milestones listing](#).

If you find errors in the present document, please send your comments to
the relevant service listed under [Committee Support Staff](#).

If you find a security vulnerability in the present document, please report it through our
[Coordinated Vulnerability Disclosure \(CVD\)](#) program.

Notice of disclaimer & limitation of liability

The information provided in the present deliverable is directed solely to professionals who have the appropriate degree of experience to understand and interpret its content in accordance with generally accepted engineering or other professional standard and applicable regulations.

No recommendation as to products and services or vendors is made or should be implied.

No representation or warranty is made that this deliverable is technically accurate or sufficient or conforms to any law and/or governmental rule and/or regulation and further, no representation or warranty is made of merchantability or fitness for any particular purpose or against infringement of intellectual property rights.

In no event shall ETSI be held liable for loss of profits or any other incidental or consequential damages.

Any software contained in this deliverable is provided "AS IS" with no warranties, express or implied, including but not limited to, the warranties of merchantability, fitness for a particular purpose and non-infringement of intellectual property rights and ETSI shall not be held liable in any event for any damages whatsoever (including, without limitation, damages for loss of profits, business interruption, loss of information, or any other pecuniary loss) arising out of or related to the use of or inability to use the software.

Copyright Notification

No part may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm except as authorized by written permission of ETSI.

The content of the PDF version shall not be modified without the written authorization of ETSI.

The copyright and the foregoing restriction extend to reproduction in all media.

© ETSI 2025.
All rights reserved.

Contents

Intellectual Property Rights	19
Foreword.....	19
Modal verbs terminology.....	19
Executive summary	19
Introduction	20
1 Scope	21
2 References	21
2.1 Normative references	21
2.2 Informative references.....	22
3 Definition of terms, symbols and abbreviations.....	23
3.1 Terms.....	23
3.2 Symbols.....	26
3.3 Abbreviations	26
4 Context Information Management Framework.....	28
4.1 Introduction	28
4.2 NGSI-LD Information Model.....	28
4.2.1 Introduction.....	28
4.2.2 NGSI-LD Meta Model.....	29
4.2.3 Cross Domain Ontology	30
4.2.4 NGSI-LD domain-specific models and instantiation	31
4.2.5 UML representation.....	31
4.3 NGSI-LD Architectural Considerations	32
4.3.1 Introduction.....	32
4.3.2 Centralized architecture	33
4.3.3 Distributed architecture.....	33
4.3.4 Federated architecture.....	34
4.3.5 NGSI-LD API Structure and Implementation Options	35
4.3.6 Distributed Operations.....	40
4.3.6.1 Introduction.....	40
4.3.6.2 Additive Registrations.....	41
4.3.6.3 Proxied Registrations	41
4.3.6.4 Limiting Cascading Distributed Operations.....	42
4.3.6.5 Extra information to provide when contacting Context Source	42
4.3.6.6 Additional pre- and post-processing of extra information when contacting Context Source.....	42
4.3.6.7 Querying and Retrieving Distributed Entities as Unitary Operations	43
4.3.6.8 Backwards compatibility of Context Source payloads.....	44
4.3.7 Snapshots	46
4.4 Core and user NGSI-LD @context	46
4.5 NGSI-LD Data Representation.....	47
4.5.0 Introduction.....	47
4.5.1 NGSI-LD Entity Representation.....	47
4.5.2 NGSI-LD Property Representations	48
4.5.2.1 Introduction.....	48
4.5.2.2 Normalized NGSI-LD Property	48
4.5.2.3 Concise NGSI-LD Property	50
4.5.3 NGSI-LD Relationship Representations	51
4.5.3.1 Introduction.....	51
4.5.3.2 Normalized NGSI-LD Relationship.....	52
4.5.3.3 Concise NGSI-LD Relationship.....	53
4.5.4 Simplified Representation.....	54
4.5.5 Multi-Attribute Support	60
4.5.5.1 Introduction.....	60
4.5.5.2 Processing of Conflicting Transient Entities.....	60

4.5.5.3	Processing of Conflicting Attributes	60
4.5.6	Temporal Representation of an Entity	61
4.5.7	Temporal Representation of a Property	61
4.5.8	Temporal Representation of a Relationship	61
4.5.9	Simplified temporal representation of an Entity	61
4.5.10	Entity Type List Representation	65
4.5.11	Detailed Entity Type List Representation	65
4.5.12	Entity Type Information Representation	66
4.5.13	Attribute List Representation	66
4.5.14	Detailed Attribute List Representation	66
4.5.15	Attribute Information Representation	67
4.5.16	GeoJSON Representation of Entities	67
4.5.16.0	Foreword	67
4.5.16.1	Top-level "geometry" field selection algorithm	67
4.5.16.2	GeoJSON Representation of an individual Entity	67
4.5.16.3	GeoJSON Representation of Multiple Entities	68
4.5.17	Simplified GeoJSON Representation of Entities	68
4.5.17.0	Foreword	68
4.5.17.1	Simplified GeoJSON Representation of an individual Entity	68
4.5.17.2	Simplified GeoJSON Representation of multiple Entities	69
4.5.18	NGSI-LD LanguageProperty Representations	69
4.5.18.1	Introduction	69
4.5.18.2	Normalized NGSI-LD LanguageProperty	70
4.5.18.3	Concise NGSI-LD LanguageProperty	70
4.5.19	Aggregated temporal representation of an Entity	71
4.5.19.0	Foreword	71
4.5.19.1	Supported behaviours for aggregation functions	72
4.5.20	NGSI-LD VocabProperty Representations	73
4.5.20.1	Introduction	73
4.5.20.2	Normalized NGSI-LD VocabProperty	74
4.5.20.3	Concise NGSI-LD VocabProperty	74
4.5.21	NGSI-LD ListProperty Representations	75
4.5.21.1	Introduction	75
4.5.21.2	Normalized NGSI-LD ListProperty	75
4.5.21.3	Concise NGSI-LD ListProperty	75
4.5.22	NGSI-LD ListRelationship Representations	76
4.5.22.1	Introduction	76
4.5.22.2	Normalized NGSI-LD ListRelationship	76
4.5.22.3	Concise NGSI-LD ListRelationship	77
4.5.23	NGSI-LD Linked Entity Retrieval	77
4.5.23.1	Introduction	77
4.5.23.2	Inline Linked Entity Representation	78
4.5.23.3	Flattened Linked Entity Representation	78
4.5.24	NGSI-LD JsonProperty Representations	78
4.5.24.1	Introduction	78
4.5.24.2	Normalized NGSI-LD JsonProperty	78
4.5.24.3	Concise NGSI-LD JsonProperty	79
4.5.25	NGSI-LD EntityMap Representation	79
4.6	Data Representation Restrictions	80
4.6.1	Supported text encodings	80
4.6.2	Supported names	80
4.6.3	Supported data types for Values	80
4.6.4	Supported Content	81
4.6.5	Supported data types for LanguageMaps	82
4.6.6	Ordering of Entities in arrays having more than one instance of the same Entity	82
4.7	Geospatial Properties	82
4.7.1	GeoJSON Geometries	82
4.7.2	Representation of GeoJSON Geometries in JSON-LD	83
4.7.3	Concise NGSI-LD GeoProperty	83
4.8	Temporal Properties	83
4.9	NGSI-LD Query Language	84
4.10	NGSI-LD Geoquery Language	92

4.11	NGSI-LD Temporal Query Language	94
4.12	NGSI-LD Pagination	95
4.13	Counting the Number of Results	96
4.14	Supporting Multiple Tenants	96
4.15	NGSI-LD Language Filter	96
4.16	Supporting Multiple Entity Types	97
4.17	NGSI-LD Entity Type Selection Language	97
4.18	NGSI-LD Scopes	98
4.19	NGSI-LD Scope Query Language	98
4.20	NGSI-LD Distributed Operation names	99
4.21	NGSI-LD Attribute Projection Language	101
4.22	Transient Storage of Entities and Attributes	102
4.23	Entity Ordering	102
4.23.1	Introduction	102
4.23.2	Datatype Comparison Order	102
4.23.3	NGSI-LD Entity Ordering Language	103
5	API Operation Definition	104
5.1	Introduction	104
5.2	Data Types	104
5.2.1	Introduction	104
5.2.2	Common members	104
5.2.3	@context	105
5.2.4	Entity	105
5.2.5	Property	106
5.2.6	Relationship	108
5.2.7	GeoProperty	110
5.2.8	EntityInfo	111
5.2.9	CSourceRegistration	111
5.2.10	RegistrationInfo	115
5.2.11	TimeInterval	115
5.2.12	Subscription	116
5.2.13	GeoQuery	118
5.2.14	NotificationParams	119
5.2.14.1	NotificationParams data type definition	119
5.2.14.2	Output only members	120
5.2.15	Endpoint	121
5.2.16	BatchOperationResult	121
5.2.17	BatchEntityError	122
5.2.18	UpdateResult	122
5.2.19	NotUpdatedDetails	122
5.2.20	EntityTemporal	122
5.2.21	TemporalQuery	123
5.2.22	KeyValuePair	124
5.2.23	Query	124
5.2.24	EntityTypeList	126
5.2.25	EntityType	126
5.2.26	EntityTypeInfo	127
5.2.27	AttributeList	127
5.2.28	Attribute	127
5.2.29	Feature	128
5.2.30	FeatureCollection	128
5.2.31	FeatureProperties	129
5.2.32	LanguageProperty	129
5.2.33	EntitySelector	131
5.2.34	RegistrationManagementInfo	132
5.2.35	VocabProperty	132
5.2.36	ListProperty	134
5.2.37	ListRelationship	136
5.2.38	JsonProperty	138
5.2.39	EntityMap	140
5.2.40	Context Source Identity	141

5.2.41	Snapshot.....	141
5.2.42	ExecutionResultDetails.....	143
5.2.43	OrderingParams	144
5.2.44	AggregationParams.....	144
5.3	Notification data types.....	145
5.3.1	Notification	145
5.3.2	CSourceNotification	146
5.3.3	TriggerReasonEnumeration	146
5.3.4	SnapshotNotification	146
5.4	NGSI-LD Fragments	147
5.5	Common Behaviours.....	148
5.5.1	Introduction.....	148
5.5.2	Error types	148
5.5.3	Error response payload body	149
5.5.4	General NGSI-LD validation.....	149
5.5.5	Default @context assignment	150
5.5.6	Operation execution and generic error handling	150
5.5.7	Term to URI expansion or compaction.....	150
5.5.8	Partial Update Patch Behaviour	151
5.5.9	Pagination Behaviour.....	153
5.5.9.1	General Pagination Behaviour	153
5.5.9.2	Pagination option using limit and offset.....	154
5.5.9.3	Pagination with Entity maps	154
5.5.10	Multi-Tenant Behaviour	154
5.5.11	More than one instance of the same Entity in an Entity array	155
5.5.11.0	Foreword	155
5.5.11.1	Batch Entity Creation case	155
5.5.11.2	Batch Entity Creation or Update (Upsert) case	155
5.5.11.3	Batch Entity Update case	155
5.5.11.4	Batch Entity Delete case	155
5.5.11.5	Batch Entity Merge case	156
5.5.12	Merge Patch Behaviour	156
5.5.13	Limiting operations to local scope	157
5.5.14	Distributed Transactional Behaviour	158
5.5.15	Snapshot Behaviour	158
5.6	Context Information Provision	159
5.6.1	Create Entity	159
5.6.1.1	Description	159
5.6.1.2	Use case diagram	159
5.6.1.3	Input data	159
5.6.1.4	Behaviour	159
5.6.1.5	Output data.....	160
5.6.2	Update Attributes.....	160
5.6.2.1	Description	160
5.6.2.2	Use case diagram	160
5.6.2.3	Input data	161
5.6.2.4	Behaviour	161
5.6.2.5	Output data.....	162
5.6.3	Append Attributes.....	162
5.6.3.1	Description	162
5.6.3.2	Use case diagram	162
5.6.3.3	Input data	163
5.6.3.4	Behaviour	163
5.6.3.5	Output data.....	164
5.6.4	Partial Attribute update.....	164
5.6.4.1	Description	164
5.6.4.2	Use case diagram	164
5.6.4.3	Input data	164
5.6.4.4	Behaviour	165
5.6.4.5	Output data.....	165
5.6.5	Delete Attribute	165
5.6.5.1	Description	165

5.6.5.2	Use case diagram	166
5.6.5.3	Input data	166
5.6.5.4	Behaviour	166
5.6.5.5	Output data	167
5.6.6	Delete Entity	167
5.6.6.1	Description	167
5.6.6.2	Use case diagram	167
5.6.6.3	Input data	168
5.6.6.4	Behaviour	168
5.6.6.5	Output data	168
5.6.7	Batch Entity Creation	168
5.6.7.1	Description	168
5.6.7.2	Use case diagram	168
5.6.7.3	Input data	169
5.6.7.4	Behaviour	169
5.6.7.5	Output data	170
5.6.8	Batch Entity Creation or Update (Upsert)	170
5.6.8.1	Description	170
5.6.8.2	Use case diagram	170
5.6.8.3	Input data	171
5.6.8.4	Behaviour	171
5.6.8.5	Output data	173
5.6.9	Batch Entity Update	173
5.6.9.1	Description	173
5.6.9.2	Use case diagram	173
5.6.9.3	Input data	173
5.6.9.4	Behaviour	174
5.6.9.5	Output data	175
5.6.10	Batch Entity Delete	175
5.6.10.1	Description	175
5.6.10.2	Use case diagram	175
5.6.10.3	Input data	175
5.6.10.4	Behaviour	176
5.6.10.5	Output data	177
5.6.11	Create or Update (Upsert) Temporal Evolution of an Entity	177
5.6.11.1	Description	177
5.6.11.2	Use case diagram	177
5.6.11.3	Input data	177
5.6.11.4	Behaviour	177
5.6.11.5	Output data	178
5.6.12	Add Attributes to Temporal Evolution of an Entity	178
5.6.12.1	Description	178
5.6.12.2	Use case diagram	178
5.6.12.3	Input data	179
5.6.12.4	Behaviour	179
5.6.12.5	Output data	180
5.6.13	Delete Attribute from Temporal Evolution of an Entity	180
5.6.13.1	Description	180
5.6.13.2	Use case diagram	180
5.6.13.3	Input data	180
5.6.13.4	Behaviour	181
5.6.13.5	Output data	181
5.6.14	Modify Attribute instance in Temporal Evolution of an Entity	181
5.6.14.1	Description	181
5.6.14.2	Use case diagram	182
5.6.14.3	Input data	182
5.6.14.4	Behaviour	182
5.6.14.5	Output data	183
5.6.15	Delete Attribute instance from Temporal Evolution of an Entity	183
5.6.15.1	Description	183
5.6.15.2	Use case diagram	183
5.6.15.3	Input data	184

5.6.15.4	Behaviour	184
5.6.15.5	Output data	185
5.6.16	Delete Temporal Evolution of an Entity	185
5.6.16.1	Description	185
5.6.16.2	Use case diagram	185
5.6.16.3	Input data	185
5.6.16.4	Behaviour	185
5.6.16.5	Output data	186
5.6.17	Merge Entity	186
5.6.17.1	Description	186
5.6.17.2	Use case diagram	186
5.6.17.3	Input data	187
5.6.17.4	Behaviour	187
5.6.17.5	Output data	188
5.6.18	Replace Entity	188
5.6.18.1	Description	188
5.6.18.2	Use case diagram	188
5.6.18.3	Input data	189
5.6.18.4	Behaviour	189
5.6.18.5	Output data	190
5.6.19	Replace Attribute	190
5.6.19.1	Description	190
5.6.19.2	Use case diagram	190
5.6.19.3	Input data	190
5.6.19.4	Behaviour	190
5.6.19.5	Output data	191
5.6.20	Batch Entity Merge	191
5.6.20.1	Description	191
5.6.20.2	Use case diagram	191
5.6.20.3	Input data	192
5.6.20.4	Behaviour	192
5.6.20.5	Output data	193
5.6.21	Purge Entities	193
5.6.21.1	Description	193
5.6.21.2	Use case diagram	193
5.6.21.3	Input data	194
5.6.21.4	Behaviour	194
5.6.21.5	Output data	195
5.7	Context Information Consumption	196
5.7.1	Retrieve Entity	196
5.7.1.1	Description	196
5.7.1.2	Use case diagram	196
5.7.1.3	Input data	196
5.7.1.4	Behaviour	197
5.7.1.5	Output data	198
5.7.2	Query Entities	199
5.7.2.1	Description	199
5.7.2.2	Use case diagram	199
5.7.2.3	Input data	199
5.7.2.4	Behaviour	201
5.7.2.5	Output data	203
5.7.3	Retrieve Temporal Evolution of an Entity	204
5.7.3.1	Description	204
5.7.3.2	Use case diagram	204
5.7.3.3	Input data	204
5.7.3.4	Behaviour	205
5.7.3.5	Output data	206
5.7.4	Query Temporal Evolution of Entities	206
5.7.4.1	Description	206
5.7.4.2	Use case diagram	206
5.7.4.3	Input data	207
5.7.4.4	Behaviour	208

5.7.4.5	Output Data	211
5.7.5	Retrieve Available Entity Types	211
5.7.5.1	Description	211
5.7.5.2	Use case diagram	211
5.7.5.3	Input data	212
5.7.5.4	Behaviour	212
5.7.5.5	Output data	212
5.7.6	Retrieve Details of Available Entity Types	212
5.7.6.1	Description	212
5.7.6.2	Use case diagram	212
5.7.6.3	Input data	213
5.7.6.4	Behaviour	213
5.7.6.5	Output data	213
5.7.7	Retrieve Available Entity Type Information	213
5.7.7.1	Description	213
5.7.7.2	Use case diagram	213
5.7.7.3	Input data	214
5.7.7.4	Behaviour	214
5.7.7.5	Output data	214
5.7.8	Retrieve Available Attributes	214
5.7.8.1	Description	214
5.7.8.2	Use case diagram	214
5.7.8.3	Input data	215
5.7.8.4	Behaviour	215
5.7.8.5	Output data	215
5.7.9	Retrieve Details of Available Attributes	215
5.7.9.1	Description	215
5.7.9.2	Use case diagram	215
5.7.9.3	Input data	216
5.7.9.4	Behaviour	216
5.7.9.5	Output data	216
5.7.10	Retrieve Available Attribute Information	216
5.7.10.1	Description	216
5.7.10.2	Use case diagram	216
5.7.10.3	Input data	217
5.7.10.4	Behaviour	217
5.7.10.5	Output data	217
5.7.11	Architecture-related aspects of retrieval of Entity Types and Attributes	217
5.8	Context Information Subscription	218
5.8.1	Create Subscription	218
5.8.1.1	Description	218
5.8.1.2	Use case diagram	218
5.8.1.3	Input data	218
5.8.1.4	Behaviour	218
5.8.1.5	Output data	220
5.8.2	Update Subscription	220
5.8.2.1	Description	220
5.8.2.2	Use case diagram	220
5.8.2.3	Input data	220
5.8.2.4	Behaviour	221
5.8.2.5	Output data	221
5.8.3	Retrieve Subscription	221
5.8.3.1	Description	221
5.8.3.2	Use case diagram	221
5.8.3.3	Input data	222
5.8.3.4	Behaviour	222
5.8.3.5	Output data	222
5.8.4	Query Subscriptions	222
5.8.4.1	Description	222
5.8.4.2	Use case diagram	222
5.8.4.3	Input data	223
5.8.4.4	Behaviour	223

5.8.4.5	Output data	223
5.8.5	Delete Subscription	223
5.8.5.1	Description	223
5.8.5.2	Use case diagram	223
5.8.5.3	Input data	224
5.8.5.4	Behaviour	224
5.8.5.5	Output data	224
5.8.6	Notification behaviour	224
5.9	Context Source Registration	226
5.9.1	Introduction	226
5.9.2	Register Context Source	227
5.9.2.1	Description	227
5.9.2.2	Use case diagram	227
5.9.2.3	Input data	227
5.9.2.4	Behaviour	227
5.9.2.5	Output data	228
5.9.3	Update Context Source Registration	228
5.9.3.1	Description	228
5.9.3.2	Use case diagram	228
5.9.3.3	Input data	229
5.9.3.4	Behaviour	229
5.9.3.5	Output data	230
5.9.4	Delete Context Source Registration	230
5.9.4.1	Description	230
5.9.4.2	Use case diagram	230
5.9.4.3	Input data	230
5.9.4.4	Behaviour	230
5.9.4.5	Output data	231
5.10	Context Source Discovery	231
5.10.1	Retrieve Context Source Registration	231
5.10.1.1	Description	231
5.10.1.2	Use case diagram	231
5.10.1.3	Input data	231
5.10.1.4	Behaviour	231
5.10.1.5	Output data	232
5.10.2	Query Context Source Registrations	232
5.10.2.1	Description	232
5.10.2.2	Use case diagram	232
5.10.2.3	Input data	232
5.10.2.4	Behaviour	233
5.10.2.5	Output data	234
5.11	Context Source Registration Subscription	234
5.11.1	Introduction	234
5.11.2	Create Context Source Registration Subscription	234
5.11.2.1	Description	234
5.11.2.2	Use case diagram	234
5.11.2.3	Input data	235
5.11.2.4	Behaviour	235
5.11.2.5	Output data	236
5.11.3	Update Context Source Registration Subscription	236
5.11.3.1	Description	236
5.11.3.2	Use case diagram	236
5.11.3.3	Input data	237
5.11.3.4	Behaviour	237
5.11.3.5	Output data	237
5.11.4	Retrieve Context Source Registration Subscription	237
5.11.4.1	Description	237
5.11.4.2	Use case diagram	238
5.11.4.3	Input data	238
5.11.4.4	Behaviour	238
5.11.4.5	Output data	238
5.11.5	Query Context Source Registration Subscriptions	238

5.11.5.1	Description	238
5.11.5.2	Use case diagram	238
5.11.5.3	Input data	239
5.11.5.4	Behaviour	239
5.11.5.5	Output data	239
5.11.6	Delete Context Source Registration Subscription	239
5.11.6.1	Description	239
5.11.6.2	Use case diagram	239
5.11.6.3	Input data	240
5.11.6.4	Behaviour	240
5.11.6.5	Output data	240
5.11.7	Notification behaviour	240
5.12	Matching Context Source Registrations	241
5.13	Storing, Managing and Serving @contexts	242
5.13.1	Introduction	242
5.13.2	Add @context	243
5.13.2.1	Description	243
5.13.2.2	Use case diagram	243
5.13.2.3	Input data	244
5.13.2.4	Behaviour	244
5.13.2.5	Output data	244
5.13.3	List @contexts	244
5.13.3.1	Description	244
5.13.3.2	Use case diagram	244
5.13.3.3	Input data	245
5.13.3.4	Behaviour	245
5.13.3.5	Output data	245
5.13.4	Serve @context	246
5.13.4.1	Description	246
5.13.4.2	Use case diagram	246
5.13.4.3	Input data	246
5.13.4.4	Behaviour	246
5.13.4.5	Output data	247
5.13.5	Delete and Reload @context	247
5.13.5.1	Description	247
5.13.5.2	Use case diagram	247
5.13.5.3	Input data	247
5.13.5.4	Behaviour	247
5.13.5.5	Output data	248
5.14	Context Source Entity Mapping	248
5.14.1	Retrieve EntityMap	248
5.14.1.1	Description	248
5.14.1.2	Use case diagram	248
5.14.1.3	Input data	248
5.14.1.4	Behaviour	249
5.14.1.5	Output data	249
5.14.2	Update EntityMap	249
5.14.2.1	Description	249
5.14.2.2	Use case diagram	249
5.14.2.3	Input data	249
5.14.2.4	Behaviour	249
5.14.2.5	Output data	250
5.14.3	Delete EntityMap	250
5.14.3.1	Description	250
5.14.3.2	Use case diagram	250
5.14.3.3	Input data	250
5.14.3.4	Behaviour	250
5.14.3.5	Output data	250
5.14.4	Create EntityMap for Query Entities	251
5.14.4.1	Description	251
5.14.4.2	Use case diagram	251
5.14.4.3	Input data	251

5.14.4.4	Behaviour	252
5.14.4.5	Output data	254
5.14.5	Create EntityMap for Query Temporal Evolution of Entities	254
5.14.5.1	Description	254
5.14.5.2	Use case diagram	254
5.14.5.3	Input data	254
5.14.5.4	Behaviour	255
5.14.5.5	Output Data	257
5.15	Context Source Identity Information	257
5.15.1	Retrieve Context Source Identity Information	257
5.15.1.1	Description	257
5.15.1.2	Use case diagram	257
5.15.1.3	Input data	258
5.15.1.4	Behaviour	258
5.15.1.5	Output data	258
5.16	Snapshot Functionality	258
5.16.1	Create Snapshot	258
5.16.1.1	Description	258
5.16.1.2	Use case diagram	258
5.16.1.3	Input data	259
5.16.1.4	Behaviour	259
5.16.1.5	Output data	260
5.16.2	Clone Snapshot	260
5.16.2.1	Description	260
5.16.2.2	Use case diagram	260
5.16.2.3	Input data	260
5.16.2.4	Behaviour	260
5.16.2.5	Output data	261
5.16.3	Retrieve Snapshot Status	261
5.16.3.1	Description	261
5.16.3.2	Use case diagram	261
5.16.3.3	Input data	262
5.16.3.4	Behaviour	262
5.16.3.5	Output data	262
5.16.4	Update Snapshot Status	262
5.16.4.1	Description	262
5.16.4.2	Use case diagram	262
5.16.4.3	Input data	262
5.16.4.4	Behaviour	262
5.16.4.5	Output data	263
5.16.5	Delete Snapshot	263
5.16.5.1	Description	263
5.16.5.2	Use case diagram	263
5.16.5.3	Input data	263
5.16.5.4	Behaviour	263
5.16.5.5	Output data	264
5.16.6	Snapshot status notification behaviour	264
5.16.7	Purge Snapshots	264
5.16.7.1	Description	264
5.16.7.2	Use case diagram	264
5.16.7.3	Input data	264
5.16.7.4	Behaviour	265
5.16.7.5	Output data	265
6	API HTTP Binding	265
6.1	Introduction	265
6.2	Global Definitions and Resource Structure	265
6.3	Common Behaviours	269
6.3.1	Introduction	269
6.3.2	Error Types	269
6.3.3	Reporting errors	269
6.3.4	HTTP request preconditions	270

6.3.5	JSON-LD @context resolution	270
6.3.6	HTTP response common requirements	271
6.3.7	Representation of Entities	272
6.3.8	Notification behaviour	273
6.3.9	Csource Notification behaviour	274
6.3.10	Pagination behaviour	274
6.3.11	Including system Attributes	276
6.3.12	Simplified or aggregated temporal representation of Entities	276
6.3.13	Counting number of results	277
6.3.14	Tenant specification	277
6.3.15	GeoJSON representation of spatially bound entities	277
6.3.16	Expiration time for cached @contexts	277
6.3.17	Distributed Operations Caching and Timeout Behaviour	278
6.3.18	Limiting Distributed Operations	279
6.3.19	Extra information to provide when contacting Context Source	279
6.3.20	Invalid parameters	279
6.3.21	Optional profile information regarding versioning and datatype conformance	280
6.3.22	Snapshot specification	280
6.4	Resource: entities/	280
6.4.1	Description	280
6.4.2	Resource definition	280
6.4.3	Resource methods	280
6.4.3.1	POST	280
6.4.3.2	GET	281
6.4.3.3	DELETE	285
6.5	Resource: entities/{entityId}	287
6.5.1	Description	287
6.5.2	Resource definition	288
6.5.3	Resource methods	288
6.5.3.1	GET	288
6.5.3.2	DELETE	290
6.5.3.3	PUT	291
6.5.3.4	PATCH	293
6.6	Resource: entities/{entityId}/attrs/	294
6.6.1	Description	294
6.6.2	Resource definition	294
6.6.3	Resource methods	295
6.6.3.1	POST	295
6.6.3.2	PATCH	296
6.7	Resource: entities/{entityId}/attrs/{attrId}	297
6.7.1	Description	297
6.7.2	Resource definition	297
6.7.3	Resource methods	298
6.7.3.1	PATCH	298
6.7.3.2	DELETE	299
6.7.3.3	PUT	300
6.8	Resource: csourceRegistrations/	302
6.8.1	Description	302
6.8.2	Resource definition	302
6.8.3	Resource methods	302
6.8.3.1	POST	302
6.8.3.2	GET	303
6.9	Resource: csourceRegistrations/{registrationId}	305
6.9.1	Description	305
6.9.2	Resource definition	305
6.9.3	Resource methods	306
6.9.3.1	GET	306
6.9.3.2	PATCH	306
6.9.3.3	DELETE	307
6.10	Resource: subscriptions/	308
6.10.1	Description	308
6.10.2	Resource definition	308

6.10.3	Resource methods.....	308
6.10.3.1	POST.....	308
6.10.3.2	GET.....	309
6.11	Resource: subscriptions/{subscriptionId}	310
6.11.1	Description.....	310
6.11.2	Resource definition	310
6.11.3	Resource methods.....	310
6.11.3.1	GET.....	310
6.11.3.2	PATCH	311
6.11.3.3	DELETE	311
6.12	Resource: csourceSubscriptions/.....	312
6.12.1	Description.....	312
6.12.2	Resource definition	312
6.12.3	Resource methods.....	312
6.12.3.1	POST.....	312
6.12.3.2	GET.....	313
6.13	Resource: csourceSubscriptions/{subscriptionId}.....	314
6.13.1	Description.....	314
6.13.2	Resource definition	314
6.13.3	Resource methods.....	314
6.13.3.1	GET.....	314
6.13.3.2	PATCH	315
6.13.3.3	DELETE	316
6.14	Resource: entityOperations/create.....	316
6.14.1	Description.....	316
6.14.2	Resource definition	317
6.14.3	Resource methods.....	317
6.14.3.1	POST.....	317
6.15	Resource: entityOperations/upsert.....	318
6.15.1	Description.....	318
6.15.2	Resource definition	318
6.15.3	Resource methods.....	319
6.15.3.1	POST.....	319
6.16	Resource: entityOperations/update.....	320
6.16.1	Description.....	320
6.16.2	Resource definition	321
6.16.3	Resource methods.....	321
6.16.3.1	POST.....	321
6.17	Resource: entityOperations/delete.....	322
6.17.1	Description.....	322
6.17.2	Resource definition	322
6.17.3	Resource methods.....	323
6.17.3.1	POST.....	323
6.18	Resource: temporal/entities/	324
6.18.1	Description.....	324
6.18.2	Resource definition	324
6.18.3	Resource methods.....	325
6.18.3.1	POST.....	325
6.18.3.2	GET.....	325
6.19	Resource: temporal/entities/{entityId}	329
6.19.1	Description.....	329
6.19.2	Resource definition	329
6.19.3	Resource methods.....	329
6.19.3.1	GET.....	329
6.19.3.2	DELETE	331
6.20	Resource: temporal/entities/{entityId}/attrs/	332
6.20.1	Description.....	332
6.20.2	Resource definition	332
6.20.3	Resource methods.....	332
6.20.3.1	POST.....	332
6.21	Resource: temporal/entities/{entityId}/attrs/{attrId}.....	333
6.21.1	Description.....	333

6.21.2	Resource definition	333
6.21.3	Resource methods	334
6.21.3.1	DELETE	334
6.22	Resource: temporal/entities/{entityId}/attrs/{attrId}/ {instanceId}	334
6.22.1	Description	334
6.22.2	Resource definition	335
6.22.3	Resource methods	335
6.22.3.1	PATCH	335
6.22.3.2	DELETE	336
6.23	Resource: entityOperations/query	336
6.23.1	Description	336
6.23.2	Resource definition	336
6.23.3	Resource methods	337
6.23.3.1	POST	337
6.24	Resource: temporal/entityOperations/query	338
6.24.1	Description	338
6.24.2	Resource definition	338
6.24.3	Resource methods	339
6.24.3.1	POST	339
6.25	Resource: types/	339
6.25.1	Description	339
6.25.2	Resource definition	340
6.25.3	Resource methods	340
6.25.3.1	GET	340
6.26	Resource: types/{type}	341
6.26.1	Description	341
6.26.2	Resource definition	341
6.26.3	Resource methods	341
6.26.3.1	GET	341
6.27	Resource: attributes/	342
6.27.1	Description	342
6.27.2	Resource definition	342
6.27.3	Resource methods	342
6.27.3.1	GET	342
6.28	Resource: attributes/{attrId}	343
6.28.1	Description	343
6.28.2	Resource definition	343
6.28.3	Resource methods	343
6.28.3.1	GET	343
6.29	Resource: jsonldContexts/	344
6.29.1	Description	344
6.29.2	Resource definition	344
6.29.3	Resource methods	344
6.29.3.1	POST	344
6.29.3.2	GET	345
6.30	Resource: jsonldContexts/{contextId}	346
6.30.1	Description	346
6.30.2	Resource definition	346
6.30.3	Resource methods	346
6.30.3.1	GET	346
6.30.3.2	DELETE	347
6.31	Resource: entityOperations/merge	348
6.31.1	Description	348
6.31.2	Resource definition	348
6.31.3	Resource methods	348
6.31.3.1	POST	348
6.32	Resource: entityMaps/{entityMapId}	350
6.32.1	Description	350
6.32.2	Resource definition	350
6.32.3	Resource methods	350
6.32.3.1	GET	350
6.32.3.2	PATCH	351

6.32.3.3	DELETE	351
6.33	Resource: info/sourceIdentity.....	352
6.33.1	Description.....	352
6.33.2	Resource definition	352
6.33.3	Resource methods.....	352
6.33.3.1	GET.....	352
6.34	Resource: entityMaps	353
6.34.1	Description.....	353
6.34.2	Resource definition	353
6.34.3	Resource methods.....	353
6.34.3.1	GET.....	353
6.34.3.2	POST.....	354
6.35	Resource: temporal/entityMaps.....	355
6.35.1	Description.....	355
6.35.2	Resource definition	355
6.35.3	Resource methods.....	355
6.35.3.1	GET.....	355
6.35.3.2	POST.....	356
6.36	Resource: snapshots	357
6.36.1	Description.....	357
6.36.2	Resource definition	357
6.36.3	Resource methods.....	357
6.36.3.1	POST.....	357
6.36.3.2	DELETE	358
6.37	Resource: snapshots/{snapshotId}	359
6.37.1	Description.....	359
6.37.2	Resource definition	359
6.37.3	Resource methods.....	359
6.37.3.1	GET.....	359
6.37.3.2	PATCH	360
6.37.3.3	DELETE	360
6.38	Resource: snapshots/{snapshotId}/clone.....	361
6.38.1	Description.....	361
6.38.2	Resource definition	361
6.38.3	Resource methods.....	361
6.38.3.1	POST.....	361
7	API MQTT Notification Binding	362
7.1	Introduction	362
7.2	Notification behaviour.....	362
Annex A (normative):	NGSI-LD identifier considerations	364
A.1	Introduction	364
A.2	Entity identifiers	364
A.3	NGSI-LD namespace	364
Annex B (normative):	Core NGSI-LD @context definition.....	365
Annex C (informative):	Examples of using the API	371
C.1	Introduction	371
C.2	Entity Representation	371
C.2.1	Property Graph	371
C.2.2	Vehicle Entity.....	372
C.2.3	Parking Entity.....	385
C.2.4	@context	391
C.3	Context Source Registration.....	392
C.4	Context Subscription	392

C.5	HTTP REST API Examples	393
C.5.1	Introduction	393
C.5.2	Create Entity of Type Vehicle	393
C.5.2.1	HTTP Request	393
C.5.2.2	HTTP Response	393
C.5.3	Query Entities.....	393
C.5.3.1	Introduction.....	393
C.5.3.2	HTTP Request	394
C.5.3.3	HTTP Response	394
C.5.4	Query Entities (Pagination)	394
C.5.4.1	Introduction.....	394
C.5.4.2	HTTP Request	394
C.5.4.3	HTTP Response	394
C.5.5	Temporal Query	395
C.5.5.1	Introduction.....	395
C.5.5.2	HTTP Request #1	395
C.5.5.3	HTTP Response #1	395
C.5.5.4	HTTP Request #2	396
C.5.5.5	HTTP Response #2	396
C.5.6	Temporal Query (Simplified Representation)	396
C.5.6.1	Introduction.....	396
C.5.6.2	HTTP Request	396
C.5.6.3	HTTP Response	397
C.5.7	Retrieve Available Entity Types	397
C.5.7.1	Introduction.....	397
C.5.7.2	HTTP Request	397
C.5.7.3	HTTP Response	397
C.5.8	Retrieve Details of Available Entity Types	398
C.5.8.1	Introduction.....	398
C.5.8.2	HTTP Request	398
C.5.8.3	HTTP Response	398
C.5.9	Retrieve Available Entity Type Information	399
C.5.9.1	Introduction.....	399
C.5.9.2	HTTP Request	399
C.5.9.3	HTTP Response	399
C.5.10	Retrieve Available Attributes	400
C.5.10.1	Introduction.....	400
C.5.10.2	HTTP Request	400
C.5.10.3	HTTP Response	400
C.5.11	Retrieve Details of Available Attributes	400
C.5.11.1	Introduction.....	400
C.5.11.2	HTTP Request	400
C.5.11.3	HTTP Response	401
C.5.12	Retrieve Available Attribute Information.....	401
C.5.12.1	Introduction.....	401
C.5.12.2	HTTP Request	401
C.5.12.3	HTTP Response	402
C.5.13	Query Entities (Natural Language Filtering).....	402
C.5.13.1	Introduction.....	402
C.5.13.2	HTTP Request	402
C.5.13.3	HTTP Response	402
C.5.14	Temporal Query (Aggregated Representation)	403
C.5.14.1	Introduction.....	403
C.5.14.2	HTTP Request	403
C.5.14.3	HTTP Response	403
C.5.15	Scope Queries.....	404
C.5.15.1	Introduction.....	404
C.5.15.2	HTTP Request	404
C.5.15.3	HTTP Response	404
C.5.16	Temporal Scope Queries	405
C.5.16.1	Introduction.....	405
C.5.16.2	HTTP Request	405

C.5.16.3	HTTP Response	405
C.6	Date Representation	407
C.7	@context utilization clarifications	408
C.8	Link header utilization clarifications.....	409
C.9	@context processing clarifications.....	411
C.10	ValueType datatype utilization clarifications.....	412
C.11	Entity with digital signature for a Property	412
Annex D (informative):	Transformation Algorithms.....	414
D.1	Introduction	414
D.2	Algorithm for transforming an NGSI-LD Entity into a JSON-LD document (ALG1).....	414
D.3	Algorithm for transforming an NGSI-LD Property into JSON-LD (ALG1.1)	415
D.4	Algorithm for transforming an NGSI-LD Relationship into JSON-LD (ALG1.2).....	416
Annex E (informative):	RDF-compatible specification of NGSI-LD meta-model.....	417
Annex F (informative):	Conventions and syntax guidelines.....	418
Annex G (informative):	Localization and Internationalization Support.....	419
G.0	Foreword	419
G.1	Introduction	419
G.1.0	Foreword	419
G.1.1	Associating an Entity with a Natural Language	419
G.1.2	Associating a Property with a Natural Language	419
G.1.3	Associating as equivalent entity	420
G.2	Natural Language Collation Support.....	420
G.2.0	Foreword	420
G.2.1	Maintain collations as metadata	421
G.2.2	Route language sensitive queries via a proxy.....	421
G.3	Localization of Dates, Currency formats, etc.....	421
G.3.0	Foreword	421
G.3.1	Localizing Dates.....	421
Annex H (informative):	Suggested actuation workflows.....	423
H.1	Actuators and feedback to the consumer.....	423
H.2	Architecture for actuation.....	423
H.3	Structure of Commands and additional Properties	424
H.3.0	Introduction	424
H.3.1	Property for listing available commands	425
H.3.2	Properties for command endpoints	425
H.4	Communication model	427
H.4.1	Possible communication models	427
H.4.2	Subscription/notification model	427
H.4.3	Forwarding model	428
H.5	Implementation of the subscription-based actuation workflow	429
H.6	Implementation of the registration-based actuation workflow	430
Annex I (informative):	Change history	433
History		436

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to normative deliverables may have been declared to ETSI. The declarations pertaining to these essential IPRs, if any, are publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the [ETSI IPR online database](#).

Pursuant to the ETSI Directives including the ETSI IPR Policy, no investigation regarding the essentiality of IPRs, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

DECT™, **PLUGTESTS™**, **UMTS™** and the ETSI logo are trademarks of ETSI registered for the benefit of its Members. **3GPP™**, **LTE™** and **5G™** logo are trademarks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners. **oneM2M™** logo is a trademark of ETSI registered for the benefit of its Members and of the oneM2M Partners. **GSM®** and the GSM logo are trademarks registered and owned by the GSM Association.

Foreword

This Group Specification (GS) has been produced by ETSI Industry Specification Group (ISG) cross-cutting Context Information Management (CIM).

Modal verbs terminology

In the present document **"shall"**, **"shall not"**, **"should"**, **"should not"**, **"may"**, **"need not"**, **"will"**, **"will not"**, **"can"** and **"cannot"** are to be interpreted as described in clause 3.2 of the [ETSI Drafting Rules](#) (Verbal forms for the expression of provisions).

"must" and **"must not"** are **NOT** allowed in ETSI deliverables except when used in direct citation.

Executive summary

The present document formally describes the Context Information Management API (NGSI-LD) Specification. The Context Information Management API allows users to provide, consume and subscribe to context information in multiple scenarios and involving multiple stakeholders. Context information is modelled as attributes (properties and relationships) of context entities, also referred to as "digital twins", representing real-world assets. It enables close to real-time access to information coming from many different sources (not only IoT data sources).

Introduction

The present document defines the NGSI-LD API Specification. This Context Information Management API allows users to provide, consume and subscribe to context information in multiple scenarios and involving multiple stakeholders. Context information is modelled as attributes of context entities, also referred to as "digital twins", representing real-world assets (e.g. a bus in a city or a luggage claim ticket). Because of that, the NGSI-LD API is often used to bring standardized access to digital twin data.

The ongoing status of the NGSI-LD API can be found in [i.17].

The ETSI ISG CIM has decided to give the name "NGSI-LD" to the Context Information Management API. The rationale is to reinforce the fact that the present document leverages on the former OMA NGSI 9 and 10 interfaces [i.3] and FIWARE® NGSIv2 [i.9] to incorporate the latest advances from Linked Data.

Most of the NGSI-LD API and the ETSI ISG CIM information model work referenced here was created with the support of the following European Union Horizon 2020 research projects: No. 732851 (FI-NEXT), No. 723156 (WISE-IoT), No. 732240 (SynchroniCity) and No. 731993 (AutoPilot), No. 814918 (Fed4IoT), No. 779852 (IoTCrawler), No. 731884 (IoF2020), including many contributions from members of the FIWARE® Community.

1 Scope

The present document defines a standardized API for Context Information Management (NGSI-LD API) enabling close to real-time (right-time) access to context/digital twin information coming from many different sources (not only IoT data sources). The present document defines how such an API enables applications to perform updates on context, register context providers which can be queried to get updates on context, query information on current and historic context information and subscribe to receive notifications of context changes. The criteria for choice of the API characteristics are based on requirements resulting from the Use Cases ETSI GR CIM 002 [i.1] and other work items ETSI GR CIM 007 [i.2] and ETSI GS CIM 006 [i.8] on security and on the information model. The present document supersedes prior versions, including ETSI GS CIM 004 [i.16].

2 References

2.1 Normative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found in the [ETSI docbox](#).

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long-term validity.

The following referenced documents are necessary for the application of the present document.

- [1] W3C® Recommendation 25 February 2014: "[RDF Schema 1.1](#)".
- [2] W3C® Recommendation 16 July 2020: "[JSON-LD 1.1 - A JSON-based Serialization for Linked Data](#)".
- [3] [IETF RFC 7231](#): "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content".
- [4] [IETF RFC 7232](#): "Hypertext Transfer Protocol (HTTP/1.1): Conditional Requests".
- [5] [IETF RFC 3986](#): "Uniform Resource Identifier (URI): Generic Syntax".
- [6] [IETF RFC 8259](#): "The JavaScript Object Notation (JSON) Data Interchange Format".
- [7] [IETF RFC 8288](#): "Web Linking".
- [8] [IETF RFC 7946](#): "The GeoJSON Format".
- [9] [IETF RFC 8141](#): "Uniform Resource Names (URNs)".
- [10] [IETF RFC 7807](#): "Problem Details for HTTP APIs".
- [11] [IEEE 1003.2™-1992](#): "IEEE Standard for Information Technology - Portable Operating System Interfaces (POSIX™) - Part 2: Shell and Utilities".
- [12] [IETF RFC 5234](#): "Augmented BNF for Syntax Specifications: ABNF".
- [13] [Unicode® Technical Standard #10](#): "Unicode Collation Algorithm".
- [14] [Open Geospatial Consortium Inc. OGC 06-103r4](#): "OpenGIS® Implementation Standard for Geographic information - Simple feature access - Part 1: Common architecture".
- [15] [Codes for Units of Measure used in International Trade Revision 9](#).
- [16] [IETF RFC 7396](#): "JSON Merge Patch".

- [17] [ISO 8601: 2019](#): "Date and time — Representations for information interchange".
- [18] [IETF RFC 2818](#): "HTTP Over TLS".
- [19] [IETF RFC 5246](#): "The Transport Layer Security (TLS) Protocol Version 1.2".
- [20] [IANA Registry of Link Relation Types](#).
- [21] [ECMA-262 Specification](#): "ECMAScript® 2022 language specification".
- [22] The Unicode Consortium: "[Unicode® 15.0.0](#)".
- [23] [IETF RFC 3987](#): "Internationalized Resource Identifiers (IRIs)".
- [24] OASIS Standard: "[MQTT Version 3.1.1 Plus Errata 01](#)". Edited by Andrew Banks and Rahul Gupta. 10 December 2015.
- [25] OASIS Standard: "[MQTT Version 5.0](#)". Edited by Andrew Banks, Ed Briggs, Ken Borgendale and Rahul Gupta. 07 March 2019.
- [26] [IETF RFC 7240](#): "Prefer Header for HTTP".
- [27] [IETF RFC 7230](#): "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing".
- [28] [IETF RFC 5646](#): "Tags for Identifying Languages".
- [29] [IETF RFC 3282](#): "Content Language Headers".
- [30] [IETF RFC 7234](#): "Hypertext Transfer Protocol (HTTP/1.1): Caching".
- [31] [IETF RFC 7233](#): "Hypertext Transfer Protocol (HTTP/1.1): Range Requests".
- [32] IANA: "[Hypertext Transfer Protocol \(HTTP\) Warn Codes](#)".
- [33] [IETF RFC 6906](#): "The 'profile' Link Relation Type".
- [34] W3C® Recommendation 25 February 2014: "[RDF 1.1 Concepts and Abstract Syntax](#)".
- [35] [ETSI GS CIM 019](#): "cross-cutting Context Information Management (CIM); handling of provenance information in NGSI-LD".
- [36] [IETF RFC 6067](#): "BCP 47 Extension U".
- [37] [ETSI GS CIM 047](#): "Context Information Management (CIM); OpenAPI Specification for NGSI-LD API".

2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long-term validity.

The following referenced documents may be useful in implementing an ETSI deliverable or add to the reader's understanding, but are not required for conformance to the present document.

- [i.1] [ETSI GR CIM 002 \(V1.1.1\)](#): "Context Information Management (CIM); Use Cases (UC)".
- [i.2] [ETSI GR CIM 007](#): "Context Information Management (CIM); Security and Privacy".
- [i.3] [OMA-TS-NGSI Context Management-V1 0-20120529-A](#): "NGSI Context Management".
- [i.4] ETSI TS 103 264 (V3.1.1): "SmartM2M; Smart Applications; Reference Ontology and oneM2M Mapping".

- [i.5] [NGSI-LD Wrapper](#): "Experimental proxy for adaptation between FIWARE® and NGSI-LD".
- [i.6] Graph Databases: "New Opportunities for Connected Data". O'Reilly 2nd Edition. Webber, Robinson, et al. ISBN:1491930896 9781491930892.
- [i.7] [JSON-LD Playground](#): "Experimentation tool for JSON-LD".
- [i.8] [ETSI GS CIM 006](#): "Context Information Management (CIM); Information Model (MOD0)".
- [i.9] [FIWARE®-NGSI v2 Specification](#).
- [i.10] [IETF RFC 6902](#): "JavaScript Object Notation (JSON) Patch".
- [i.11] [JSON Schema Validation: A Vocabulary for Structural Validation of JSON](#).
- [i.12] [OpenAPI™ Specification](#).
- [i.13] [NGSI-LD JSON Schemas](#).
- [i.14] [NGSI-LD OpenAPI™ Specification](#).
- [i.15] [NGSI-LD Examples](#).
- [i.16] ETSI GS CIM 004 (V1.1.2): "Context Information Management (CIM); Application Programming Interface (API)".
- [i.17] ETSI ISG CIM: "[NGSI-LD Status](#)".
- [i.18] [Regulation \(EU\) 2016/679](#) of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation).
- [i.19] [MQTT URI Scheme](#).
- [i.20] [GeoJSON-LD 1.0](#). Base context for processing GeoJSON according to the JSON-LD processing model.
- [i.21] [ETSI GR CIM 008](#): "Context Information Management (CIM); NGSI-LD Primer".
- [i.22] [IoT Agent Library](#).

3 Definition of terms, symbols and abbreviations

3.1 Terms

For the purposes of the present document, the following terms apply:

NOTE 1: The letters "NGSI-LD" were added to most terms to confirm that they are distinct from other terms of similar/same name in use in other organizations, however, in the present document the letters "NGSI-LD" are generally omitted for brevity.

NOTE 2: The use of URI in the context of the present document also includes the use of International Resource Identifiers (IRIs) as defined in IETF RFC 3987 [23], which extends the use of characters to Unicode characters [22] beyond the ASCII character set, enabling the support of languages other than English.

NGSI-LD Attribute: reference to both an NGSI-LD Property and to an NGSI-LD Relationship

NGSI-LD Attribute Instance (in case of temporal representation of NGSI-LD Entities): reference to an NGSI-LD Attribute, at a specific moment in time of its temporal evolution, usually identified by its instanceId

NGSI-LD Central Broker: NGSI-LD Context Broker that only uses a local storage when serving NGSI-LD requests, without involving any external Context Sources

NGSI-LD Context Broker: architectural component that implements all the NGSI-LD interfaces

NGSI-LD Context Consumer: agent that uses the query and subscription functionality of NGSI-LD to retrieve context information

NGSI-LD Context Producer: agent that uses the NGSI-LD context provision and/or registration functionality to provide or announce the availability of its context information to an NGSI-LD Context Broker

NGSI-LD Context Registry: software functional element where Context Sources register the information that they can provide

NOTE: It is used by Distribution Brokers and Federation Brokers to find the appropriate Context Sources which can provide the information required for serving an NGSI-LD request.

NGSI-LD Context Source: source of context information which implements the NGSI-LD consumption and subscription (and possibly provision) interfaces defined by the present document

NOTE: It is usually registered with an NGSI-LD Registry so that it can announce what kind of information it can provide, when requested, to Context Consumers and Brokers.

NGSI-LD Context Source Registration: description of the information that can be provided by a Context Source, which is used when registering the Context Source with the Context Registry

NGSI-LD Core API: core part of the NGSI-LD API that has to be implemented by all Brokers, including operations for providing or managing Entities and Attributes, operations for consuming Entities and checking which Entity Types and Attributes Entities are available in the system and operations for subscribing to Entities, receiving notifications and managing subscriptions

NGSI-LD Distribution Broker: NGSI-LD Context Broker that uses both local context information and registration information from an NGSI-LD Context Registry, to access matching context information from a set of distributed Context Sources

NGSI-LD Element: any JSON element that is defined by the NGSI-LD API

NGSI-LD Entity: informational representative of something that is supposed to exist in the real world, physically or conceptually

NOTE: In the NGSI-LD API, any instance of such an entity is **uniquely identified by a URI**, and characterized by reference to one or more **NGSI-LD Entity Type(s)**.

NGSI-LD Entity Map: mapping of NGSI-LD Entity ids to Context Source Registrations used in maintaining atomicity of transactions performed by Distribution Brokers and Federation Brokers

NGSI-LD Entity Type: categorization of an NGSI-LD Entity as belonging to a class of similar entities, or sharing a set of characteristic properties

NOTE: In the NGSI-LD API, an NGSI-LD Entity Type is **uniquely identified by a URI**.

EXAMPLE 1: "Vehicle" is an NGSI-LD Entity Type and is identified with a proper URI.

EXAMPLE 2: Bob's private car whose plate number is "ABCD1234" is an NGSI-LD Entity whose NGSI-LD Entity Type name is "Vehicle".

EXAMPLE 3: Alice's motorhome has a unique URI as id, but can be assigned multiple NGSI-LD Entity types, e.g. "Vehicle" and "Home".

NGSI-LD External Linked Entity: Linked Entity that is identified through a **dereferenceable URI** which does not exist within the current NGSI-LD system

NOTE: It can exist within another NGSI-LD system or within a non-NGSI-LD system.

EXAMPLE: An NGSI-LD Entity, whose Entity Type name is "Book", can be externally linked, through the "wasWrittenBy" relationship, to a resource identified by the URI "http://dbpedia.org/resource/Mark_Twain".

NGSI-LD Federation Broker: Distribution Broker that federates information from multiple underlying NGSI-LD Context Brokers and across domains

NGSI-LD GeoProperty: subclass of NGSI-LD Property which is a description instance which associates a main characteristic, i.e. an **NGSI-LD Value**, to either an NGSI-LD Entity, an NGSI-LD Relationship or another NGSI-LD Property, that uses the special *hasValue* property to define its target value and holds a geographic location in GeoJSON format

NGSI-LD Internal Linked Entity: Linked Entity that exists within the current NGSI-LD system

EXAMPLE: An NGSI-LD Entity, whose Entity Type name is "Vehicle", can be internally linked, through the "isParkedAt" relationship, to another NGSI-LD Entity, of Type name "Parking", identified by the URI "urn:ngsi-ld:Parking:Downtown1".

NGSI-LD JSONLDContext API: part of the NGSI-LD API that is used to host, serve and manage JSON-LD @contexts

NGSI-LD JsonProperty: subclass of NGSI-LD Property which is a description instance which associates a raw JSON literal value as a defined main characteristic to an NGSI-LD Entity, an NGSI-LD Relationship or another NGSI-LD Property and that uses the special *hasJson* (a subproperty of *hasValue*) property to define its target value

NOTE: The target value contains data which is not available for interpretation.

NGSI-LD LanguageProperty: subclass of NGSI-LD Property which is a description instance which associates a set of strings in different natural languages as a defined main characteristic, i.e. an **NGSI-LD Map**, to an NGSI-LD Entity, an NGSI-LD Relationship or another NGSI-LD Property and that uses the special *hasLanguageMap* (a subproperty of *hasValue*) property to define its target value

NGSI-LD Linked Entity: NGSI-LD Entity referenced from another NGSI-LD Entity (the linking NGSI-LD Entity) via an NGSI-LD Relationship

NGSI-LD Linking Entity: NGSI-LD Entity which is the subject of a Relationship to another NGSI-LD Entity (the linked NGSI-LD Entity) or an external resource (identified by a URI)

NGSI-LD ListProperty: description instance which associates an ordered array of main characteristics, i.e. **NGSI-LD Values**, to either an NGSI-LD Entity, an NGSI-LD Relationship or another NGSI-LD Property and that uses the special *hasValueList* property to define its target value

NGSI-LD ListRelationship: description of an ordered array of directed links between a subject which is either an NGSI-LD Entity, an NGSI-LD Property or another NGSI-LD Relationship on one hand, and a series of objects, which are NGSI-LD Entities, on the other hand, and which uses the special *hasObjectList* property to define its target objects

EXAMPLE: "A bus route services the following bus stops" can be represented by an NGSI-LD *ListRelationship* whose name is "route" which holds an array of directed links towards a series of NGSI-LD Entities of type (Type name) "BusStop".

NGSI-LD Map: JSON-LD language map in the form of key-value pairs holding the string representation of a main characteristic in a series of natural languages

EXAMPLE: "Bob's vehicle is currently parked on a street which is known as 'Grand Place' in French and 'Grote Markt' in Dutch" can be represented by an NGSI-LD LanguageProperty whose name is "street" which holds an NGSI-LD Map of two key-value pairs containing both the French ("fr") and Dutch ("nl") exonyms of the street name.

NGSI-LD Null: "urn:ngsi-ld:null" or {"@none": "urn:ngsi-ld:null"} used as an encoding for null values

NGSI-LD Property: description instance which associates a main characteristic, i.e. an **NGSI-LD Value**, to either an NGSI-LD Entity, an NGSI-LD Relationship or another NGSI-LD Property and that uses the special *hasValue* property to define its target value

EXAMPLE: "Bob's vehicle's speed is 40 km/h" can be represented by an NGSI-LD Property, whose name is "speed", and which characterizes an NGSI-LD Entity, whose NGSI-LD Type is "Vehicle". Such a name can be expanded to a fully qualified name in the form of a URI, for instance "http://example.org/Vehicle" or "http://example.org/speed".

NGSI-LD Query: collection of criteria used to select a sub-set of NGSI-LD Elements, matching the criteria

NGSI-LD Registry API: part of the NGSI-LD API that is implemented by the Context Registry, including operations for registering Context Sources and managing Context Source Registrations (CSRs), operations for retrieving and discovering CSRs, and operations for subscribing to CSRs and receiving notifications

NGSI-LD Relationship: description of a directed link between a subject which is either an NGSI-LD Entity, an NGSI-LD Property or another NGSI-LD Relationship on one hand, and an object, or unordered array of objects, each of which is an NGSI-LD Entity, on the other hand, and which uses the special *hasObject* property to define its target object

EXAMPLE 1: An NGSI-LD Entity of type "Vehicle" can be the subject of an NGSI-LD Relationship whose object is an NGSI-LD Entity of type "Parking".

EXAMPLE 2: An NGSI-LD Entity of type "Vehicle" can be the subject of an NGSI-LD Relationship whose object is an array of NGSI-LD Entities of type "Passenger".

NGSI-LD Scope: hierarchical structuring of Entities that enables restricting query results and notifications

NGSI-LD Snapshot: set of NGSI-LD Entities, retrieved through one or more NGSI-LD Queries, which are locally persisted, providing a relatively consistent view of the situation when the queries were executed, and on which Context Consumers can execute further local NGSI-LD Operations

NGSI-LD Temporal API: part of the NGSI-LD API pertaining to the Temporal Evolution of Entities, including operations for providing and managing the Temporal Evolution of Entities and Attributes, and operations for consuming the Temporal Evolution of Entities

NGSI-LD Temporal Evolution of an Entity: sequence of values attributed to an NGSI-LD Entity over time, i.e. its history or future predictions

NGSI-LD Tenant: user or group of users that utilize a single instance of a system implementing the NGSI-LD API (NGSI-LD Context Source or NGSI-LD Broker) in isolation from other users or groups of users of the same instance, so that any information related to one Tenant (e.g. Entities, Subscriptions, Context Source Registrations) are only visible to users of the same Tenant, but not to users of a different Tenant

NGSI-LD Value: JSON value (i.e. a string, a number, *true* or *false*, an object, an array), or JSON-LD typed value (i.e. a string as the lexical form of the value together with a type, defined by an XSD base type or more generally an IRI), or JSON-LD structured value (i.e. a set, a list, a language-tagged string)

EXAMPLE: Bob's private car 'speed' NGSI-LD Value is the number 100 (kilometres per hour).

NGSI-LD VocabProperty: subclass of NGSI-LD Property which is a description instance which associates a string value which can be coerced to a URI as a defined main characteristic to an NGSI-LD Entity, an NGSI-LD Relationship or another NGSI-LD Property and that uses the special *hasVocab* (a subproperty of *hasValue*) property to define its target value

EXAMPLE: "Bob's car is a non-commercial vehicle" can be represented by an NGSI-LD VocabProperty whose name is "category" which holds the string value "non-commercial". If the associated JSON-LD context defines the term "non-commercial" as "http://example.com/non-commercial", then the returned value shall be expanded using JSON-LD type coercion into the URI the "http://example.com/non-commercial".

3.2 Symbols

Void.

3.3 Abbreviations

For the purposes of the present document, the following abbreviations apply:

ABNF	Augmented Backus-Naur Form
ALG1	ALGORITHM for transforming an NGSI-LD Entity into a JSON-LD document

AM	Ante Meridiem
API	Application Programming Interface
ASCII	American Standard Code for Information Interchange
BNF	Backus Naur Form
CH	Switzerland
CSR	Context Source Registration
ECMA	European Computer Manufacturers Association
EU	European Union
FI	Future Internet
FQN	Fully Qualified Name
GB	Great Britain
GDPR	General Data Protection Regulation
GeoJSON	Geographic JavaScript Object Notation
GeoJSON-LD	Geographic JavaScript Object Notation - Linked Data
GIS	Geographic Information System
GPS	Global Positioning System
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IANA	Internet Assigned Numbers Authority
ID	IDentifier
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IoT	Internet of Things
IRI	Internationalized Resource Identifier
ISG	Industry Specification Group
ISO	International Organization for Standardization
JSON	JavaScript Object Notation
JSON-LD	JSON Linked Data
LD	Linked Data
LWM2M	LightWeight Machine to Machine
M2M	Machine to Machine
MIME	Multi-purpose Internet Mail Extensions
MQTT	Message Queuing Telemetry Transport
N/A	Not Applicable
NGSI	Next Generation Service Interfaces
NGSILD	Next Generation Service Interfaces Linked Data (same as NGSI-LD)
NID	Namespace IDentifier
NSS	Namespace Specific String
OAS	Open API Specification
OMA	Open Mobile Alliance
oneM2M	oneM2M Partnership Project
PM	Post Meridiem
POSIX	Portable Operating System Interface
QoS	Quality of Service
RDF	Resource Description Format
REST	REpresentational State Transfer
RFC	Request For Comments
SAREF	Smart Applications REference ontology
TCP	Transport Control Protocol
TLS	Transport Layer Security
TS	Technical Specification
UCA	Unicode Collation Algorithm
UL	Ultra Light
UML	Unified Modelling Language
URI	Uniform Resource Identifier
URL	Universal Resource Locator
URN	Uniform Resource Name
UTC	Coordinated Universal Time
UTF	Unicode (or Universal Coded Character Set) Transformation Format
XSD	XML Schema Definition

4 Context Information Management Framework

4.1 Introduction

The present clause describes the technical design principles behind the context information management framework supported by NGSI-LD. As stated in clause 3.1, the letters "NGSI-LD" which are part of most terms, to confirm that they are distinct from other terms of similar/same name in use in other organizations, are generally omitted in the present document for brevity. In the present document, a number of rather obvious typographic conventions and syntax guidelines are followed, and the reader is referred to annex F for details.

4.2 NGSI-LD Information Model

4.2.1 Introduction

The NGSI-LD Information Model prescribes the structure of context information that shall be supported by an NGSI-LD system. It specifies the data representation mechanisms that shall be used by the NGSI-LD API itself. In addition, it specifies the structure of the Context Information Management vocabularies to be used in conjunction with the API.

The NGSI-LD Information Model is defined at two levels (see Figure 4.2.1-1): the foundation classes which correspond to the Core Meta-model and the Cross-Domain Ontology. The former amounts to a formal specification of the "property graph" model [i.6]. The latter is a set of generic, transversal classes which are aimed at avoiding conflicting or redundant definitions of the same classes in each of the domain-specific ontologies. Below these two levels, domain-specific ontologies or vocabularies can be devised. For instance, the SAREF Ontology ETSI TS 103 264 [i.4] can be mapped to the NGSI-LD Information Model, so that smart home applications will benefit from this Context Information Management API specification.

The version of the cross-domain model proposed by the present document is a minimal one, aimed at defining the classes used in this release of the API specification. It has been extended by other work items like ETSI GS CIM 006 [i.8], with classes defining extra concepts such as mobile vs. stationary entities, instantaneous vs. static properties, etc.




Figure 4.2.1-1: Overview of the NGSI-LD Information Model Structure

4.2.2 NGSI-LD Meta Model

Figure 4.2.2-1 provides a graphical representation of the NGSI-LD Meta-Model in terms of classes and their relationships. To provide additional clarity an informal (non-normative) mapping to the Property Graph Model is also presented.




Figure 4.2.2-1: NGSI-LD Core Meta-Model

Implementations shall support the NGSI-LD Meta-model as follows:

- An **NGSI-LD Entity** is a subclass of `rdfs:Resource` [1].
- An **NGSI-LD Relationship** is a subclass of `rdfs:Resource` [1].
- An **NGSI-LD Property** is a subclass of `rdfs:Resource` [1].
- An **NGSI-LD Value** shall be either a `rdfs:Literal` or a node object (in JSON-LD syntax) to represent complex data structures [1].
- An **NGSI-LD Property** shall have a **value**, stated through *hasValue*, which is of type `rdf:Property` [1]. An **NGSI-LD Relationship** shall have an **object** stated through *hasObject* which is of type `rdf:Property` [1].

4.2.3 Cross Domain Ontology




Figure 4.2.3-1: NGSi-LD Core Meta-Model plus the Cross-Domain Ontology

Figure 4.2.3-1 describes the concepts introduced by the NGSi-LD Cross-Domain Ontology, which shall be supported by implementations as follows:

- **Geo Properties:** Are intended to convey geospatial information and implementations shall support them as defined in clause 4.7.
- **Temporal Properties:** Are non-reified Properties (represented only by their Value) that convey temporal information for capturing the time series evolution of other Properties; implementations shall support them as defined in clause 4.8.
- **Language Properties:** Are intended to convey different versions of the same textual values, whenever a version for each language (for instance: English, Spanish) is needed.
- **"unitCode" Property:** Is a Property intended to provide the units of measurement of an NGSi-LD Value. Implementations shall support it as defined in clause 4.5.2.
- **"scope" Property:** Is a Property that enables putting Entities into a hierarchical structure. Implementations shall support it as defined in clause 4.18.
- **LanguageMaps:** Are a special type of NGSi-LD Value intended to convey the different values of Language Properties, stated through an *hasLanguageMap*, which is of type *rdf:Property* [1] and is itself a subproperty of *hasValue*.

- **Geometry Values:** Are a special type of NGSI-LD Value intended to convey geometries corresponding to geospatial properties. Implementations shall support them as defined in clause 4.7.
- **Time Values:** Are a special type of NGSI-LD Value intended to convey time instants or intervals representations. Implementations shall support them as defined in clause 4.6.3.

Clause 4.4 defines the Core JSON-LD *@context* which includes the URIs which correspond to the concepts introduced above.

4.2.4 NGSI-LD domain-specific models and instantiation

This clause is informative and is intended to illustrate the relationship between the NGSI-LD Information Model and NGSI-LD Domain-specific models.

Figure 4.2.4-1 shows an example of an NGSI-LD domain-specific model. Domain-specific models introduce the specific entity types required for a particular domain. Figure 4.2.4-1 shows the types "Car", "Parking", "Street", "Gate". Entity types can have further subtypes, e.g. "OffStreetParking" as subtype of "Parking".




Figure 4.2.4-1: Cross-Domain Ontology and instantiation

In addition, two different NGSI-LD Properties are introduced (*hasState*, *reliability*).

The *adjacentTo* Relationship links entities of type "Parking" with entities of type "Street".

4.2.5 UML representation

This clause is informative and is intended to show how the NGSI-LD information model could be described using UML diagrams. The aim of this diagram is to help those readers less familiar with ontology representations or RDF [1] to understand the NGSI-LD Information Model.

In Figure 4.2.5-1 NGSI-LD Entity, Relationship, Property and Value are represented as UML classes. UML associations are used to interrelate these classes while keeping the structure and semantics defined by the NGSI-LD Information Model.




Figure 4.2.5-1: NGSI-LD information model as UML

4.3 NGSI-LD Architectural Considerations

4.3.1 Introduction

The NGSI-LD API is intended to be primarily an API and does not define a specific architecture. It is envisioned that the NGSI-LD API can be used in different architectural settings and the architectural assumptions of the API are kept to a minimum.

As it is not possible to elaborate all possible architectures in which the NGSI-LD API could be used, three prototypical architectures are presented. The NGSI-LD API shall enable efficient support for all of them, i.e. the design decisions for the NGSI-LD API take these prototypical architectures into consideration. A real system architecture utilizing the NGSI-LD API can map to one, take elements from multiple or combine all of the prototypical architectures.

The NGSI-LD API implicitly defines two sets of Entities:

- the "current state";
- the "temporal evolution" (both the past and possibly future predictions).

The NGSI-LD API is structured into a **Core API** and an optional **Temporal API**. The **Core API** manages the current state of Entities. The **Temporal API** is optional and manages the **Temporal Evolution of Entities**. Brokers that intend to implement the **Temporal API** should consider updating the **Temporal Evolution of an Entity** whenever the "current state" is modified via the **Core API**.

4.3.2 Centralized architecture

Figure 4.3.2-1 shows a centralized architecture. In the centre is a **Central Broker** that stores all the context information. There are **Context Producers** that use update operations to update the context information in the **Central Broker** and there are **Context Consumers** that request context information from the **Central Broker**, either using synchronous one-time query or asynchronous subscribe/notify operations. The **Central Broker** answers all requests from its storage. Figure 4.3.2-1 shows one component that acts as both **Context Producer** and **Context Consumer**. The general assumption is that components can have multiple roles, so such components are not explicitly shown in clauses 4.3.3 and 4.3.4.




Figure 4.3.2-1: Centralized architecture

4.3.3 Distributed architecture

Figure 4.3.3-1 shows a distributed architecture. The underlying idea here is that all information is stored by the **Context Sources**. **Context Sources** implement the query and subscription part of the NGSI-LD API as a **Context Broker** does. They register themselves with the **Context Registry**, providing information about what context information they can provide, but not the context information itself, e.g. a certain **Context Source** registers that it can provide the indoor temperature for Building A and Building B or that it can provide the speed of cars in a geographic region covering the centre of a city.




Figure 4.3.3-1: Distributed architecture

Context Consumers can query or subscribe to the Distribution Broker. On each request, the Distribution Broker discovers or does a discovery subscription to the Context Registry for relevant Context Sources, i.e. those that may provide context information relevant to the respective request from the Context Consumer. The Distribution Broker then queries or subscribes to each relevant Context Source, if possible it aggregates the context information retrieved from the Context Sources and provides them to the Context Consumer. In this mode of operation, it is not visible to the Context Consumer, whether the Context Broker is a Central Broker or a Distribution Broker. Alternatively, the architecture allows that Context Consumers can discover Context Sources through the Context Registry themselves and then directly request from Context Sources. This is shown in Figure 4.3.3-1 with the fine dashed arrows.

4.3.4 Federated architecture

The federated architecture shown in Figure 4.3.4-1 is used in cases where existing domains are to be federated. For example, different departments in a city operate their own Context Broker-based NGSI-LD infrastructure, but applications should be able to easily access all available information using just one point of access. The architecture works in the same way as the distributed architecture described in clause 4.3.3, except that instead of simple Context Sources, whole domains are registered with the respective Context Broker as point of access. Typically, the domains will be registered to the federation Context Registry on a more coarse-grained level, providing scopes, in particular geographic scopes, that can then be matched to the scopes provided in the requests. For example, instead of registering individual entities like buildings, the domain would be registered with having information about entities of type building within a geographic area. Applications then query or subscribe for entities within a geographic scope, e.g. buildings in a certain area of the city. The Federation Broker discovers the domain Context Brokers that can provide relevant information, forwards the request to these Context Brokers and aggregates the results, so the application gets the result in the same way as in the centralized and distributed cases.




Figure 4.3.4-1: Federated architecture

A domain itself can use a centralized or distributed architecture or could even utilize a federated architecture that federates sub-domains.

As in the distributed case, it is also possible that applications discover relevant domains through the federation-level Context Registry and directly contact the Context Brokers in the individual domains.

4.3.5 NGSI-LD API Structure and Implementation Options

As stated in clause 4.3.1, the NGSI-LD API is structured into a Core API and an optional Temporal API. To support the distributed and federated architectures described in clauses 4.3.3 and 4.3.4 respectively, distributed versions of the operations are needed. They are listed separately under Distributed API and require the operations of the Registry API. The Registry API consists of the operations to be implemented by the Context Registry. Furthermore, the JSONLDContext API provides functionality for storing, managing, and serving JSON-LD *@contexts*. The APIs are structured according to their functionalities, which is also reflected in how the operations are structured in clause 5. Table 4.3.5-1 introduces the API structure, the respective functionalities and lists the operations for each functionality, pointing to the clauses in which they are defined. The distributed versions of the operations are separately shown in Table 4.3.5-1 under Distributed API, but there is a single clause for each of the operations describing both the centralized and distributed behaviour. In addition, the Distributed API has support operations only needed in distributed and federated architectures.

Table 4.3.5-1: NGSI-LD API structure

API	Functionality	Operations
Core API	Context Information Provision - operations for providing or managing Entities and Attributes	5.6.1 Create Entity 5.6.2 Update Attributes 5.6.3 Append Attributes 5.6.4 Partial Attribute Update 5.6.5 Delete Attribute 5.6.6 Delete Entity 5.6.7 Batch Entity Creation 5.6.8 Batch Entity Upsert 5.6.9 Batch Entity Update 5.6.10 Batch Entity Delete 5.6.17 Merge Entity 5.6.18 Replace Entity 5.6.19 Replace Attribute 5.6.20 Batch Entity Merge
	Context Information Consumption - operations for consuming Entities and checking for which Entity Types and Attributes Entities are available in the system	5.7.1 Retrieve Entity 5.7.2 Query Entities 5.7.5 Retrieve Available Entity Types 5.7.6 Retrieve Details of Available Entity Types 5.7.7 Retrieve Available Entity Type Information 5.7.8 Retrieve Available Attributes 5.7.9 Retrieve Details of Available Attributes 5.7.10 Retrieve Available Attribute Information
	Context Information Subscription - operations for subscribing to Entities, receiving notifications and managing subscriptions	5.8.1 Create Subscription 5.8.2 Update Subscription 5.8.3 Retrieve Subscription 5.8.4 Query Subscription 5.8.5 Delete Subscription 5.8.6 Notification
Temporal API	Temporal Context Information Provision - operations for providing or managing the Temporal Evolution of Entities and Attributes	5.6.11 Upsert Temporal Evolution of an Entity 5.6.12 Add Attributes to Temporal Evolution of an Entity 5.6.13 Delete Attributes from Temporal Evolution of an Entity 5.6.14 Partial Update Attribute instance 5.6.15 Delete Attribute Instance 5.6.16 Delete Temporal Evolution of an Entity
	Temporal Context Information Consumption - operations for consuming the Temporal Evolution of Entities	5.7.3 Retrieve Temporal Evolution of an Entity 5.7.4 Query Temporal Evolution of Entities

API	Functionality	Operations
Distributed API	Distributed Context Information Provision - operations for providing or managing Entities and Attributes	5.6.1 Create Entity (distributed) 5.6.2 Update Attributes (distributed) 5.6.3 Append Attributes (distributed) 5.6.4 Partial Attribute Update (distributed) 5.6.5 Delete Attribute (distributed) 5.6.6 Delete Entity (distributed) 5.6.7 Batch Entity Creation (distributed) 5.6.8 Batch Entity Upsert (distributed) 5.6.9 Batch Entity Update (distributed) 5.6.10 Batch Entity Delete (distributed) 5.6.17 Merge Entity (distributed) 5.6.18 Replace Entity (distributed) 5.6.19 Replace Attribute (distributed) 5.6.20 Batch Entity Merge (distributed)
	Distributed Context Information Consumption - operations for consuming Entities and checking for which Entity Types and Attributes Entities are available in the system	5.7.1 Retrieve Entity (distributed) 5.7.2 Query Entities (distributed) 5.7.5 Retrieve Available Entity Types (distributed) 5.7.6 Retrieve Details of Available Entity Types (distributed) 5.7.7 Retrieve Available Entity Type Information (distributed) 5.7.8 Retrieve Available Attributes (distributed) 5.7.9 Retrieve Details of Available Attributes (distributed) 5.7.10 Retrieve Available Attribute Information (distributed)
	Distributed Context Information Subscription - operations for subscribing to Entities, receiving notifications and managing subscriptions	5.8.1 Create Subscription (distributed) 5.8.2 Update Subscription (distributed) 5.8.3 Retrieve Subscription (distributed) 5.8.4 Query Subscription (distributed) 5.8.5 Delete Subscription (distributed) 5.8.6 Notification (distributed)
	Distributed Temporal Context Information Provision - operations for providing or managing the Temporal Evolution of Entities and Attributes	5.6.11 Upsert Temporal Evolution of an Entity (distributed) 5.6.12 Add Attributes to Temporal Evolution of an Entity (distributed) 5.6.13 Delete Attributes from Temporal Evolution of an Entity (distributed) 5.6.14 Partial Update Attribute instance (distributed) 5.6.15 Delete Attribute Instance (distributed) 5.6.16 Delete Temporal Evolution of an Entity (distributed)
	Distributed Temporal Context Information Consumption - operations for consuming the Temporal Evolution of Entities	5.7.3 Retrieve Temporal Evolution of an Entity (distributed) 5.7.4 Query Temporal Evolution of Entities (distributed)
	Support operations for distributed operations	5.14.1 Retrieve EntityMap 5.14.2 Update EntityMap 5.14.3 Delete EntityMap 5.14.4 Create EntityMap for Query Entities 5.14.5 Create EntityMap for Query Temporal Evolution of Entities 5.15.1 Retrieve Context Source Identity Information
Registry API	Context Source Registration - operations for registering Context Sources and managing Context Source Registrations (CSRs)	5.9.2 Register Context Source 5.9.3 Update CSR 5.9.4 Delete CSR
	Context Source Discovery - operations for retrieving and discovering CSRs	5.7.1 Retrieve CSR 5.7.2 Query CSRs
	Context Source Registration Subscription - operations for subscribing to CSRs, receiving notifications and managing CSRs	5.11.2 Create CSR Subscription 5.11.3 Update CSR Subscription 5.11.4 Retrieve CSR Subscription 5.11.5 Query CSR Subscription 5.11.6 Delete CSR Subscription 5.11.7 CSR Notification

API	Functionality	Operations
Snapshot API	Operations for creating and managing Snapshots.	5.16.1 Create Snapshot 5.16.2 Clone Snapshot 5.16.3 Retrieve Snapshot Status 5.16.4 Update Snapshot Status 5.16.5 Delete Snapshot 5.16.6 Snapshot Status Notification
JSONLDContext API	Storing, managing and serving <i>@contexts</i>	5.13.2 Add <i>@context</i> 5.13.3 List <i>@contexts</i> 5.13.4 Serve <i>@context</i> 5.13.5 Delete and Reload <i>@context</i>

All Context Brokers shall implement the Core API. Context Brokers supporting distributed and federated deployments shall also implement the Distributed API. Temporal API and Registry API can be implemented by a Broker or by a separate temporal component and Context Registry respectively. Table 4.3.5-2 shows the possible implementation configurations. A temporal component implementing the Temporal API can also be used completely independently of a Context Broker. The Snapshot API and the JSONLDContext API are optional. The managing and serving of *@contexts* can also be handled by an independent, stand-alone component.

Table 4.3.5-2: Main implementation configurations

Description	Temporal API	Registry API
Central Broker without temporal support	none	none
Central Broker with integrated temporal component	local	none
Central Broker with separate temporal component	separate	none
Context Broker supporting distributed and federated deployments without temporal support and with integrated Context Registry	none	local
Context Broker supporting distributed and federated deployments with integrated temporal component and integrated Context Registry	local	local
Context Broker supporting distributed and federated deployments with separate temporal component and integrated Context Registry	separate	local
Context Broker supporting distributed and federated deployments without temporal support and separate Context Registry	none	separate
Context Broker supporting distributed and federated deployments with integrated temporal component and separate Context Registry	local	separate
Context Broker supporting distributed and federated deployments with separate temporal component and separate Context Registry	separate	separate

Table 4.3.5-3 shows which operations are implemented and used by the other architectural roles as introduced in clause 4.3.2, clause 4.3.3 and clause 4.3.4. In addition, there are separate roles for the Temporal API, i.e. Temporal Context Producer, Temporal Context Source and Temporal Context Consumer. For completeness, the roles of Context Repository and Temporal Context Repository have been introduced, implementing the Context Information Provision and Temporal Context Information Provision functionalities, respectively. In practice, components implementing the latter roles will also implement functionalities for consuming or processing the stored information. Actual components can have multiple roles at the same time, e.g. a Context Broker can implement all roles at the same time. Context Consumers typically only interact with Context Brokers, but in alternative setups, as shown in Figure 4.3.3-1, they can also directly interact with the Context Registry and then directly contact Context Sources.

Table 4.3.5-3: Operations implemented by the various NGSI-LD Roles

NGSI-LD Role	Implements	Uses
Context Consumer	5.8.6 Notification - if supporting asynchronous interactions In case of direct interactions with Context Registry: 5.11.7 CSR Notification - if supporting asynchronous interactions	5.7.1 Retrieve Entity 5.7.2 Query Entities 5.7.5 Retrieve Available Entity Types 5.7.6 Retrieve Details of Available Entity Types 5.7.7 Retrieve Available Entity Type Information 5.7.8 Retrieve Available Attributes 5.7.9 Retrieve Details of Available Attributes 5.7.10 Retrieve Available Attribute Information 5.8.1 Create Subscription 5.8.2 Update Subscription 5.8.3 Retrieve Subscription 5.8.4 Query Subscription 5.8.5 Delete Subscription 5.14.1 Retrieve EntityMap 5.14.2 Update EntityMap 5.14.3 Delete EntityMap 5.14.4 Create EntityMap for Query Entities 5.14.5 Create EntityMap for Query Temporal Evolution of Entities 5.15.1 Retrieve Context Source Identity Information 5.16.1 Create Snapshot 5.16.2 Clone Snapshot 5.16.3 Retrieve Snapshot Status 5.16.4 Update Snapshot Status 5.16.5 Delete Snapshot 5.16.6 Snapshot Status Notification In case of direct interactions with Context Registry: 5.7.1 Retrieve CSR 5.7.2 Query CSRs if supporting asynchronous interactions: 5.11.2 Create CSR Subscription 5.11.3 Update CSR Subscription 5.11.4 Retrieve CSR Subscription 5.11.5 Query CSR Subscription 5.11.6 Delete CSR Subscription
Context Producer	none	5.6.1 Create Entity 5.6.2 Update Attributes 5.6.3 Append Attributes 5.6.4 Partial Attribute Update 5.6.5 Delete Attribute 5.6.6 Delete Entity 5.6.7 Batch Entity Creation 5.6.8 Batch Entity Upsert 5.6.9 Batch Entity Update 5.6.10 Batch Entity Delete 5.6.17 Merge Entity 5.6.18 Replace Entity 5.6.19 Replace Attribute 5.6.20 Batch Entity Merge

NGSI-LD Role	Implements	Uses
Context Source	5.7.1 Retrieve Entity 5.7.2 Query Entities 5.7.5 Retrieve Available Entity Types 5.7.6 Retrieve Details of Available Entity Types 5.7.7 Retrieve Available Entity Type Information 5.7.8 Retrieve Available Attributes 5.7.9 Retrieve Details of Available Attributes 5.7.10 Retrieve Available Attribute Information 5.8.1 Create Subscription 5.8.2 Update Subscription 5.8.3 Retrieve Subscription 5.8.4 Query Subscription 5.8.5 Delete Subscription 5.14.1 Retrieve EntityMap 5.14.2 Update EntityMap 5.14.3 Delete EntityMap 5.14.4 Create EntityMap for Query Entities 5.14.5 Create EntityMap for Query Temporal Evolution of Entities 5.15.1 Retrieve Context Source Identity Information	5.8.6 Notification 5.9.2 Register Context Source 5.9.3 Update CSR 5.9.4 Delete CSR
Context Repository	5.6.1 Create Entity 5.6.2 Update Attributes 5.6.3 Append Attributes 5.6.4 Partial Attribute Update 5.6.5 Delete Attribute 5.6.6 Delete Entity 5.6.7 Batch Entity Creation 5.6.8 Batch Entity Upsert 5.6.9 Batch Entity Update 5.6.10 Batch Entity Delete 5.6.17 Merge Entity 5.6.18 Replace Entity 5.6.19 Replace Attribute 5.6.20 Batch Entity Merge 5.16.1 Create Snapshot 5.16.2 Clone Snapshot 5.16.3 Retrieve Snapshot Status 5.16.4 Update Snapshot Status 5.16.5 Delete Snapshot 5.16.6 Snapshot Status Notification	none
Temporal Context Consumer	In case of direct interactions with Context Registry: 5.11.7 CSR Notification - if supporting asynchronous interactions	5.7.3 Retrieve Temporal Evolution of an Entity 5.7.4 Query Temporal Evolution of Entities In case of direct interactions with Context Registry: 5.7.1 Retrieve CSR 5.7.2 Query CSRs if supporting asynchronous interactions: 5.11.2 Create CSR Subscription 5.11.3 Update CSR Subscription 5.11.4 Retrieve CSR Subscription 5.11.5 Query CSR Subscription 5.11.6 Delete CSR Subscription
Temporal Context Producer	None	5.6.11 Upsert Temporal Evolution of an Entity 5.6.12 Add Attributes to Temporal Evolution of an Entity 5.6.13 Delete Attributes from Temporal Evolution of an Entity 5.6.14 Partial Update Attribute instance 5.6.15 Delete Attribute Instance 5.6.16 Delete Temporal Evolution of an Entity
Temporal Context Source	5.7.3 Retrieve Temporal Evolution of an Entity 5.7.4 Query Temporal Evolution of Entities	5.9.2 Register Context Source 5.9.3 Update CSR 5.9.4 Delete CSR

NGSI-LD Role	Implements	Uses
Temporal Context Repository	5.6.11 Upsert Temporal Evolution of an Entity 5.6.12 Add Attributes to Temporal Evolution of an Entity 5.6.13 Delete Attributes from Temporal Evolution of an Entity 5.6.14 Partial Update Attribute instance 5.6.15 Delete Attribute Instance 5.6.16 Delete Temporal Evolution of an Entity	none
Context Registry	5.9.2 Register Context Source 5.9.3 Update CSR 5.9.4 Delete CSR 5.7.1 Retrieve CSR 5.7.2 Query CSRs 5.11.2 Create CSR Subscription 5.11.3 Update CSR Subscription 5.11.4 Retrieve CSR Subscription 5.11.5 Query CSR Subscription 5.11.6 Delete CSR Subscription	5.11.7 CSR Notification

4.3.6 Distributed Operations

4.3.6.1 Introduction

One fundamental concept underpinning all of the prototypical architectures described above (clauses 4.3.2, 4.3.3 and 4.3.4) is the idea that Entity data does not need to be centralized within a single **Context Broker**. When reading context information, a **Context Broker** can be used as a single point of access to retrieve Entity data found distributed across multiple associated **Context Brokers** each receiving a **context consumption** request. Similarly, when modifying an Entity, a single request to a **Context Broker** may result in the operation being distributed and different parts of that Entity being updated across multiple **Context Brokers** each receiving a **context provision** request.

As long as there is only a centralized **Context Broker**, i.e. there are no **Context Sources** registered, all NGSI-LD requests, with few exceptions such as Update Attributes (see clause 5.6.2) and the batch operations (see clauses 5.6.7, 5.6.8, 5.6.9, 5.6.10 and 5.6.20), can either be successfully executed completely, or result in an error. In the distributed case, all requests can be partially successful. For the centralized case described above, only specific operations, such as Update Attributes and the batch operations, can be partially successful.

It is the responsibility of the **Context Broker** to respect the registration parameters when issuing distributed requests. For instance, if a registration states that only Entities of a given type are offered, the distributed request does not contain additional types. Such a strict requirement is justified because **Context Sources** (the receivers of the distributed request) are not in a position to determine whether a request has been triggered by a registration, rendering them unable to ensure that the registration parameters are respected in the first place. This applies for any kind of context data a **Context Broker** can exchange such as Entity IDs, entity types, attribute names, geofenced areas, etc. Ultimately, all constraints specified in the registration shall be respected.

When a **Context Source** is registered, an operation mode is selected. This defines the basis for distributed operations and also defines whether or not the **Context Broker** is permitted to hold context data about the Entities and Attributes locally itself.

If two registered **Context Sources** are providing context data for the same Attribute, the Attribute instances can be distinguished by *datasetId*. The mechanism for determining which data shall be returned is defined in clause 4.5.5.

It is possible to restrict a registered **Context Source** to operate on a specific Entity type or list of Entity types. In order for **Context Broker** hierarchies to support and restrict the distribution of such limited operations, the Entity type selector (see clause 4.17) can be added as a filter on forwarded requests even where its presence initially seems redundant.

Furthermore, registered **Context Sources** may indicate that they are only willing to respond to a limited subset of API operations. **Context Brokers** shall respect this, to avoid unnecessarily sending distributed operation requests which are always guaranteed to fail. For example, a **Context Source** may consistently refuse certain API operations since it does not support them. Alternatively, some **Context Source** endpoints (such as updates) may be protected for use by authorized users only, and not accessible to a **Context Broker** without those rights. Limited access is likely to be the case in extended data sharing scenarios, where a registered **Context Source**, and the data held within it, may belong to an external third party.

For the endpoints served, all registered **Context Sources** shall support the normalized representation of Entities as default. Support of additional representation formats is optional and will depend on the implementation. System generated attributes such as *modifiedAt* and *createdAt* (see clause 4.8) should be supported by registered **Context Sources**, at a minimum no error shall be returned if they are not available when requested.

4.3.6.2 Additive Registrations

For additive registrations, the **Context Broker** is permitted to hold context data about the Entities and Attributes locally itself, and also obtain data from external sources. Context provisioning operations are serviced both locally by the **Context Broker** itself, and also distributed on to the registered sources.

An **inclusive Context Source Registration** specifies that the **Context Broker** considers all registered **Context Sources** as equals and will distribute operations to those **Context Sources** even if relevant context data is available directly within the **Context Broker** itself (in which case, all results will be integrated in the final response). Data from every **Context Source** registered by an inclusive **Context Source Registration** is requested with an equal priority. This is the default mode of operation.

An **auxiliary Context Source Registration** never overrides data held directly within a **Context Broker**. Auxiliary distributed operations are limited to context information consumption operations (see clause 5.7). Context data from auxiliary context sources is only included if it is supplementary to the context data otherwise available to the **Context Broker**. Auxiliary **Context Source Registrations** are always accepted as there can never be a conflict.

4.3.6.3 Proxied Registrations

For proxied registrations, the **Context Broker** itself is not permitted to hold context data about the registered Entities and Attributes locally (thus all registered context data is obtained from the external registered sources). Unregistered Attributes of an Entity are permitted to be held locally; when context provisioning operations are received, registered Attributes are distributed on to the registered sources and never serviced directly by the **Context Broker** itself.

An **exclusive Context Source Registration** specifies that all of the registered context data is held in a single location external to the **Context Broker**. The **Context Broker** itself holds no data locally about the registered Attributes and no overlapping proxied **Context Source Registrations** shall be supported for the same combination of registered Attributes on the Entity. An **exclusive** registration shall always relate to specific Attributes found on a single Entity. Thus, the registration shall define **both**:

- An entity id (i.e. an id pattern or Entity type defining a group of entities is not supported for **exclusive** registrations).
- Attributes.

Once an **exclusive Context Source Registration** has been created, no further exclusive or redirect **Context Source Registrations** can be created for that same combination of Entity ID and Attributes.

A **redirect Context Source Registration** also specifies that the registered context data is held in a location external to the **Context Broker**. It is possible to register (any combination of):

- A whole Entity by id or id pattern (i.e. without specifying individual Attributes in the registration; in this case, all Attributes are held externally).
- Entities by Entity type only (with or without specifying individual Attributes).

- Attributes only.

Potentially multiple distinct **redirect** registrations can apply at the same time. The **Context Broker** itself holds no data locally in conflict to the registration. In the case that multiple overlapping **redirect** registrations are defined, operations are distributed to all registered **Context Sources**.

4.3.6.4 Limiting Cascading Distributed Operations

When creating a registration, it is unknown whether the requested data is held at the distributed endpoint, or it is in turn distributed via further registrations. It is necessary to include a binding-specific mechanism to request operations only on the registered endpoint itself to avoid cascades of an excessive lengths, duplicates or loops.

Furthermore, it is not known if any distributed endpoints of a registered **Context Source** are in turn reliant on previously encountered **Context Sources** thus causing an infinite loop. Therefore, when processing a distributed operation, a specific field listing all previously encountered **Context Sources** (e.g. a **Via** header in the response in case of HTTP binding (IETF RFC 7230 [27])) shall be passed as part of the request and this field can be used to exclude duplicated sources from matching as context source registrations.

In the case of multi-tenancy (see clause 4.14) each **Tenant** found within each registered **Context Source** shall be considered separately.

4.3.6.5 Extra information to provide when contacting Context Source

If the optional array (of *KeyValuePair* type, as defined by clause 5.2.22) *contextSourceInfo* of the *CSourceRegistration* is present, it contains, whatever extra information the **Context Broker** shall convey when contacting the **Context Source**. This can be information the **Context Broker** needs to successfully communicate with the **Context Source** (e.g. Authorization material), or for the **Context Source** to correctly interpret the received content (e.g. the Link URL to fetch an *@context*). The method for conveying this information is binding-specific, e.g. using headers in the case of HTTP.

Instead of providing the actual value, the special value "urn:ngsi-ld:request" can be used to indicate that the respective value is to be taken from the request that triggered the given request, if present.

EXAMPLE: If the key value pair "user": "urn:ngsi-ld:request" is part of *contextSourceInfo* of the *CSourceRegistration*, the **Context Broker** checks if "user" was conveyed in the triggering request. If this is the case, e.g. "user": "abcd", "user": "abcd" is also conveyed when contacting the **Context Source**.

As **Tenant** information, if applicable, is directly specified in the *CSourceRegistration*, it shall not be part of *contextSourceInfo*. Binding-specific information that is used for setting up the connection or is specific for an interaction, e.g. Content-length in HTTP, cannot be overridden by *contextSourceInfo*. If present, such information shall be ignored.

4.3.6.6 Additional pre- and post-processing of extra information when contacting Context Source

The following key-values have a specific well-defined meaning when defined as elements within the optional array *contextSourceInfo* of the *CSourceRegistration*.

If the key "accept" is defined:

- the value shall be a MIME type acceptable to the **Context Broker** (one of: "application/json", "application/ld+json").
- the response from the distributed endpoint shall be returned in this defined format and if necessary, the **Context Broker** shall be responsible for converting this to the desired content type when aggregating responses to the initial request.

If the key "contentType" is defined:

- the value shall be a MIME type acceptable to the Context Broker (one of: "application/json", "application/ld+json").
- the Context Broker shall provide the request and the associated @context as required by the MIME type when distributing the request to the context source endpoint, regardless of how it was provided in the initial request.

If the key "jsonldContext" is defined:

- the value shall correspond to a URL reference as defined by the JSON-LD specification [2], section 3.1.
- the Context Broker shall apply a compaction operation as defined by the JSON-LD specification [2], section 4.1.5 over both payload and query parameters using the JSON-LD Context supplied in the value of the "jsonldContext" key-value pair, prior to distributing the request to the context source endpoint and forwarding with this JSON-LD context using an appropriate binding. Additionally, if a payload is defined in the initial request to the Context Broker, the "Content-Type" of the forwarded request shall be "application/json" and the Context Broker shall remove any @context members from the payload prior to distributing the request to the context source endpoint.

If the key "ngsildConformance" is defined:

- The value shall define in the form *major.minor*, for example 1.5.
- The Context Broker shall apply a backwards compatibility operation over the payload (as defined by clause 4.3.6.8) prior to distributing the request to the context source endpoint such that the forwarded payload conforms to the specified version of the NGSI-LD specification.

4.3.6.7 Querying and Retrieving Distributed Entities as Unitary Operations

Context Broker architectures assume that Entity data does not need to be centralized within a single Context Broker, however, when querying context information, Entity data retrieval can be considered as a unitary operation, masking the fact that each registered Context Broker is receiving a separate distributed **Context Consumption** request.

To process each **Context Consumption** request efficiently, and to support consistent pagination, it is necessary for the Context Broker to initially make a broad request to each registered **Context Source** whose registration is matching the request.

In the case of a **query Entities operation** (clause 5.7.2) or **query Temporal Evolution of Entities operation** (clause 5.7.4), a list of Entity identifiers is returned and stored together with the registration information in an Entity map. Only the Entities whose identifiers are contained in the Entity map are considered when rendering the result pages. Filtering based on queries, geoqueries, scope queries or attribute filters, as provided in the original query request, is applied to the Entities before adding the Entity to a result page, i.e. identifiers of Entities not matching the filters at the time of checking are removed from the Entity map.

In the case of a **retrieve Entity operation** (clause 5.7.1) or **retrieve Temporal Evolution of an Entity operation** (clause 5.7.3), an Entity map can be used to make subsequent retrievals of the same Entity more efficient, as the Entity map provides the information about which Context Source(s) store relevant Entity information, and other Context Sources do not have to be considered.

This Entity mapping is an internal operation, not usually exposed to the end user, however it is necessary to explicitly define a consistent mechanism for Entity map creation, caching and retrieval.

A specific field pointing to the location of a cached EntityMap (e.g. a custom header in the response in case of HTTP binding, see Table 6.4.3.2-2) shall be returned within the response of a query, whenever this is requested by the client. Similarly, the reuse of a previously created EntityMap can be requested by passing the same specific field into a request.

Since an **exclusive Context Source Registration** already specifies that all context data is held in a single location, its relevance to a distributed query can be inferred within the Registered Context Source without the use of an EntityMap.

When executing **Context Consumption** or **Subscription** operations, a significant optimization in performance can be achieved if it is known a priori whether individual Entities are themselves distributed among **Context Brokers** and **Context Sources**, or if each **Context Broker** and **Context Source** always stores complete Entities. In the latter case all parameters used for filtering such as queries, geoqueries, scope queries or attribute filters can be forwarded and applied locally, whereas in the former case, the Entity first has to be assembled by the **Context Broker** and only then the filtering can be applied. Since being able to apply filters locally is significantly more efficient, a parameter can indicate that, for the given request, only Entities are to be expected that are stored in their entirety on each **Context Broker** and each **Context Sources**, and there are no Entities that are themselves stored in a distributed fashion. Such Entities are also referred to as split Entities. **Context Broker** implementations should enable configuring a default for this parameter, so deployments where no split Entities are to be expected can filter locally and thus be more efficient.

In the case of split Entities, EntityMaps initially only store "candidate Entities" as no filters could be applied, because only a part of the Entity was available. In the process of pagination, the filters will be (re-)checked. Any Entity not (or no longer) fulfilling the filter shall be removed from the EntityMap.

4.3.6.8 Backwards compatibility of Context Source payloads

When retrieving Entity data found distributed across multiple associated **Context Brokers** each **Context Source** is sent a **context consumption** request. A **Context Broker** shall assume that every **Context Source** will return valid NGSI-LD Entity data in a format that it understands and it shall reject data that is invalid. However, since the definition of a valid NGSI-LD Entity has broadened with each version of the NGSI-LD specification, it is possible that a registered **Context Source** could respond with valid NGSI-LD Entity data which does not fit the narrower confines of a previous NGSI-LD specification.

Therefore, when making a **context consumption** request, a **Context Broker** may wish to indicate that it is only capable of interpreting responses which conform to a specific NGSI-LD specification, in which case the **Context Source** shall endeavour to amend its payload accordingly. Table 4.3.6.8-1 describes a minimal level of support for a NGSI-LD Entity data as specified by each version of the NGSI-LD specification, and the expected fallback behaviour required if a **context consumption** request is made to receive data conformant to an earlier version of the NGSI-LD specification. Table 4.3.6.8-2 describes conformance fallbacks for an NGSI-LD Property (and its subclasses) and Table 4.3.6.8-2 describes conformance fallbacks for an NGSI-LD Relationship (and its subclasses).

The version of the NGSI-LD specification requested and the conformant version returned is defined in the form *major.minor*, for example 1.5.

Table 4.3.6.8-1: NGSI-LD Entity data type attribute support

Name	Data Type	Definition	Version Introduced	Conformant Data Fallback
id	String	Valid URI	1.0	
type	String or String[] (see note 1)		1.0	
expiresAt	String	<i>DateTime</i> as mandated by clause 4.22	1.9	Remove attribute from payload
location	GeoProperty	Default geospatial Property of an entity. See clause 4.7	1.0	
observationSpace	GeoProperty	See clause 4.7	1.0	
operationSpace	GeoProperty	See clause 4.7	1.0	
scope	String or String[]	See clause 4.18	1.4	Remove attribute from payload
<Property name>	Property or Property[] (see note 2)	Property as mandated by clause 4.5.2.	1.0	
	GeoProperty or GeoProperty[] (see note 2)	GeoProperty as mandated by clause 4.5.2.	1.0	
	LanguageProperty or LanguageProperty[] (see note 2)	LanguageProperty as mandated by clause 4.5.18.	1.4	Reformat attribute as Property
	JsonProperty or JsonProperty[] (see note 2)	JsonProperty as mandated by clause 4.5.24.	1.8	Reformat attribute as Property

Name	Data Type	Definition	Version Introduced	Conformant Data Fallback
	VocabProperty or VocabProperty[] (see note 2)	VocabProperty as mandated by clause 4.5.20.	1.8	Reformat attribute as Property
	ListProperty or ListProperty[] (see note 1)	ListProperty as mandated by clause 4.5.21.	1.8	Reformat attribute as Property
<Relationship name>	Relationship or Relationship[] (see note 3)	Relationship as mandated by clause 4.5.3.	1.0	
	ListRelationship or ListRelationship[] (see note 3)	ListRelationship as mandated by clause 4.5.22.	1.8	Reformat attribute as Relationship
NOTE 1: From 1.3 onwards, an Entity <i>type</i> can be assigned multiple values. For 1.0 backwards compatibility only return a single element with preference to the first instance. NOTE 2: From 1.3 onwards, multiple instances of a <i>Property</i> (or subclass of <i>Property</i>) identified by the same Property name can be separated by <i>datasetId</i>. For 1.0 backwards compatibility only return a single element of Property Array with preference to the default instance. NOTE 3: From 1.3 onwards, multiple instances of a <i>Relationship</i> (or subclass of <i>Relationship</i>) identified by the same Relationship name can be separated by <i>datasetId</i>. For 1.3 onwards. For 1.0 backwards compatibility only return a single element of Relationship Array with preference to the default instance.				

Table 4.3.6.8-2: NGSI-LD Property data type attribute support

Name	Data Type	Definition	Version Introduced	Conformant Data Fallback
type	String		1.0	
value	Any JSON value as defined by IETF RFC 8259 [6]	See NGSI-LD Value definition in clause 3.1	1.0	
datasetId	String	Valid URI as mandated by clause 4.5.5	1.3	Remove attribute from payload
expiresAt	String	<i>DateTime</i> as mandated by clause 4.22	1.9	Remove attribute from payload
observedAt	String	<i>DateTime</i> as mandated by clause 4.8	1.3	Remove attribute from payload
unitCode	String	As mandated by [15]	1.3	Remove attribute from payload
valueType	String	See clause 4.5.2	1.9	Remove attribute from payload

Table 4.3.6.8-3: NGSI-LD Relationship data type attribute support

Name	Data Type	Definition	Version Introduced	Conformant Data Fallback
type	String	Valid URI	1.0	
object	String or String[]		1.0	
datasetId	String	Valid URI as mandated by clause 4.5.5	1.3	Remove attribute from payload
expiresAt	String	<i>DateTime</i> as mandated by clause 4.22	1.9	Remove attribute from payload
objectType	String or String[]	See clause 4.5.23	1.8	Remove attribute from payload
observedAt	String	<i>DateTime</i> as mandated by clause 4.8	1.3	Remove attribute from payload

When responding to a **context consumption** request to supply data conforming to a specific NGSI-LD specification, Context Sources should indicate the version of the specification the returned payload **actually conforms to**. In general, Context Sources will not be expected to be flexible enough to supply payloads conformant to all past and future versions of the specification, but the requesting Context Broker may use supplied version information when collating data from multiple Context Sources. and to validate and amend received payloads.

4.3.7 Snapshots

Context information can be dynamic, e.g. when a query result contains many Entities, and the user has to paginate through the result set, the values encountered later may be from a much later point in time than earlier results, making them incomparable. To get more consistent results, the Snapshot concept is introduced, which can be optionally implemented by Context Brokers. It enables "freezing" the result by retrieving all results immediately and persisting them locally. *Context Consumers* can then query and analyse the Snapshot in a much more consistent way.

This is especially relevant for distributed cases, where information initially was distributed across multiple *Context Brokers* and *Context Sources*.

Snapshots offer the full functionality of NGSI-LD available locally and can be used as a basis for simulations that enable making predictions or "what-if" analyses, which are often needed for digital twins.

4.4 Core and user NGSI-LD @context

NGSI-LD serialization is based on JSON-LD [2], a JSON-based format to serialize Linked Data. The *@context* in JSON-LD is used to expand terms, provided as short-hand strings, to concepts, specified as URIs, and vice versa, to compact URIs into terms. The Core NGSI-LD (JSON-LD) *@context* is defined as a JSON-LD *@context* which contains:

- The core terms needed to uniquely represent the key concepts defined by the NGSI-LD Information Model, as mandated by clause 4.2.
- The terms needed to uniquely represent all the members that define the API-related Data Types, as mandated by clauses 5.2 and 5.3.
- A fallback *@vocab* rule to expand or compact user-defined terms to a default URI, in case there is no other possible expansion or compaction as per the current *@context*.
- The core NGSI-LD *@context* defines the term "id", which is mapped to *@id*, and the term "type", which is mapped to *@type*. Since *@id* and *@type* are what is typically used in JSON-LD, they may also be used in NGSI-LD requests instead of "id" and "type" respectively, wherever this is applicable. In NGSI-LD responses, only "id" and "type" shall be used.

NGSI-LD compliant implementations shall support such Core *@context*, which shall be implicitly present when processing or generating context information. Furthermore, the Core *@context* is protected and shall remain immutable and invariant during expansion or compaction of terms. Therefore, and as per the JSON-LD processing rules [2], when processing NGSI-LD content, implementations shall consider the Core *@context* as if it were in the **last** position of the *@context* array. Nonetheless, for the sake of compatibility and cleanness, data providers should generate JSON-LD content that conveys the Core *@context* in the last position.

For the avoidance of doubt, when rendering NGSI-LD Elements, the Core *@context* **shall always be treated** as if it had been originally placed **in the last position**, so that, if needed, upstream JSON-LD processors can properly expand as NGSI-LD or override the resulting JSON-LD documents provided by API implementations.

The NGSI-LD Core *@context* is publicly available at <https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld> and shall contain all the terms as mandated by annex B.

In addition to the terms defined by the Core NGSI-LD *@context* (mandatory as per annex B), a user *@context* should be provided and it should contain the following terms:

- One term associated to the Entity Type, mapping the Entity Type name with its Type Identifier (URI).
- One term associated to the name of each *Property* or any of its subclasses mapping the Property name with its Property Identifier (URI). If the Property's range is a data type different than a native JSON type, then it shall be conveyed explicitly under this term by using a nested JSON object in the form:
 - "@type": <Datatype's URI>.
 - "@id": <Property's URI>.

- One term associated to the name of each *Relationship* mapping the Relationship name with the Relationship Identifier (URI) in the form:
 - "@type": "@id".
 - "@id": <Relationship's URI>.

Whereas the Core NGSI-LD *@context* contains JSON-LD Scoped Contexts, the user *@context* shall not contain JSON-LD Scoped Contexts (see [2], clause 4.1.8), as described in clause 5.5.7, because this would enable overwriting terms defined in the Core NGSI-LD *@context*.

Depending on the binding, the *@context* may not just be provided embedded with the rest of the JSON content, but there could be other options. For example, in the HTTP binding, the *@context* can be made available through a Link header (see clause 6.3.5).

4.5 NGSI-LD Data Representation

4.5.0 Introduction

All NGSI-LD elements are represented in JSON-LD [2]. For the use with the API, the compacted JSON-LD representation is used, i.e. short terms are used, which are expanded by the component implementing the NGSI-LD API using a JSON-LD *@context*, typically provided as part of the request. As described in clause 4.4, the NGSI-LD Core *@context* is always considered to be part of the *@context* to be used.

The use of JSON-LD for NGSI-LD elements has some implications for the use of *null* values, as JSON-LD interprets setting elements to *null* as elements to be removed when performing JSON-LD expansion. Thus, *null* cannot be used as a value in NGSI-LD.

To nevertheless allow deletions as part of NGSI-LD operations that update NGSI-LD data, which is typically handled by setting the respective JSON key to *null* (e.g. as in IETF RFC 7396 [16]), the URI "urn:ngsi-ld:null" is used as a replacement for *null* in all places where URI strings are valid JSON values. If an array is required, an array with one single NGSI-LD Null is used. For *languageMap*, the JSON object {"@none": "urn:ngsi-ld:null"} is to be used as explained in clause 4.5.18. These encodings of *null* are referred to as NGSI-LD Null.

For representing deleted elements in notifications and in the temporal representation, the URI "urn:ngsi-ld:null" is used as a Property *value* or Relationship *object* and the JSON object {"@none": "urn:ngsi-ld:null"} for the "languageMap" of a Language Property, respectively.

As *null* cannot be used as a value in JSON-LD, there is still the possibility of using a JSON null literal represented as {"@type": "@json", "@value": null} in JSON-LD instead. JSON literals are not to be expanded in JSON-LD and thus the respective element is not removed during JSON-LD expansion.

4.5.1 NGSI-LD Entity Representation

An NGSI-LD Entity shall be represented by an object encoded using JSON-LD [2]. The rules described below state the encoding that shall be supported by implementations. Annex D provides a computational description of this process in terms of an algorithm.

The JSON-LD object contains the following members:

Mandatory

- "id" whose value shall be a URI that identifies the Entity.
- "type" whose value shall be equal to the Entity Type name or an unordered JSON array with multiple Entity Type names in case of an Entity that has multiple Entity Types.

Optional

- "expiresAt": a string as mandated by clause 4.22.

- "scope" whose value shall be a Scope as defined in clause 4.18 or an unordered JSON array with multiple Scopes in case of an Entity that has multiple Scopes.
- "@context" a JSON-LD *@context* as described in clause 4.4.
- One member for each *Property* as per the rules stated in clause 4.5.2. In case of multiple *Property* instances with the same Property name as described in clause 4.5.5, all instances are provided as an unordered JSON array, and *datasetId* (see clause 4.5.5) shall be used to distinguish them.
- One member for each *Relationship* as per the rules stated in clause 4.5.3. In case of multiple *Relationship* instances with the same Relationship name as described in clause 4.5.5, all instances are provided as an unordered JSON array, and *datasetId* (see clause 4.5.5) shall be used to distinguish them.

NOTE 1: In the following, the term Attribute is used when referring in the text to both a Property and a Relationship (see definition of NGSI-LD Attribute in clause 3.1).

NOTE 2: When GeoJSON representation is selected, the layout of the Entities changes, see clause 4.5.16 for details.

Terms defined in the Core Context as non-reified Properties (such as "datasetId", "instanceId", etc.) shall not be used as Attribute names.

Attributes shall not contain any embedded *@context*, as described in clause 5.5.7.

4.5.2 NGSI-LD Property Representations

4.5.2.1 Introduction

An NGSI-LD Property, its value and sub-attributes can be represented in two equally valid lossless formats. The **normalized representation** is a JSON-LD document that is complete with respect to mandatory members. The **concise representation** is a terser alternative, which makes various implicit assumptions against the payloads and removes redundancy from them.

Both normalized and concise representation of Properties shall be supported by implementations and can be selected by Context Consumers through specific request parameters. An example of this representation can be found in annex C, clause C.2.2.

4.5.2.2 Normalized NGSI-LD Property

An NGSI-LD Property in normalized representation shall be represented by a member whose key is the Property name (a term) and whose value is a JSON-LD object (or JSON-LD array with such JSON-LD objects, if there are multiple instances with the same Property name, as described in clause 4.5.5), which includes the following members:

Mandatory

- "type": the fixed value "Property".
- "value": the Property Value (see definition of NGSI-LD Value in clause 3.1). The datatype URI can be placed in the optional "valueType" member.

An NGSI-LD Null (explained in clause 4.5.0 and defined in clause 3.1) can be used as the right-hand side of the "value" during partial update patch and merge patch (see clauses 5.5.8 and 5.5.12) to indicate a deletion of the *Property*, as well as in notifications and in temporal evolutions (for encoding a deleted *Property*).

It should be noted that in the JSON-LD serialization, a **number** data type does not allow for leading zeros, and octal and hexadecimal formats are not used. In the context broker's internal representation, a **number** is equivalent to either an integer or floating-point number, depending on if the number has a non-zero fractional part. The degree of precision of a floating-point number held within a context broker will depend upon the implementation, and insignificant digits may be lost during the deserialization and serialization process.

Optional

- "datasetId": a URI as mandated by clause 4.5.5.
- "expiresAt": a string as mandated by clause 4.22.
- "ngsildproof": a Property with the non-reified subproperties "entityIdSealed" and "entityTypeSealed" as specified in [35]. The value of its "value" element shall be an object containing the W3C® Data integrity "proof" structure [35]. See clause C.11 for an example.
- "observedAt": a string as mandated by clause 4.8.
- "unitCode": a string representing the measurement unit corresponding to the Property value. It shall be encoded using the UNECE/CEFACT Common Codes for Units of Measurement [15].
- "valueType": a string value which shall be type coerced into a datatype URI. It is a non-reified alternative to the use of a native JSON-LD @type for the Property value. When it is a JSON primitive, it should align this with the RDF datatype [34].
- For each of the Properties this Property is associated with, a member whose key (a term) is the Property name and value is the result of serializing a *Property* (or any of its subclasses) in **normalized representation** (see clause 4.5.2.2).
- For each of the Relationships this Property is associated with, a member whose key (a term) is the Relationship name and value is the result of serializing a *Relationship* in **normalized representation** (see clause 4.5.3.2).

System Generated

- "createdAt": a string as mandated by clause 4.8.
- "deletedAt": a string as mandated by clause 4.8.
- "instanceId": a URI uniquely identifying a Property instance, as mandated by clause 4.5.7.
- "modifiedAt": a string as mandated by clause 4.8.

Output Only

- "previousValue": only provided only in the case of notifications and only if the *showChanges* option is explicitly requested. It represents the previous Property Value, before the triggering change. The representation is the same as that of "value".

Furthermore, an NGSI-LD Property in the normalized representation shall never include the following members:

Prohibited

- "entity": shall never be present, as it is used during inline **Linked Entity retrieval** to define the target Entity of a *Relationship's* object.
- "entityIdSealed" and "entityTypeSealed": shall never be present, unless the Property name is "ngsildproof".
- "entityList": shall never be present, as it is used during inline **Linked Entity retrieval** to define the target Entities of a *ListRelationship's* object.
- "json" and "previousJson": shall never be present, as they define a *JsonProperty* value.
- "languageMap" and "previousLanguageMap": shall never be present, as they define a *LanguageProperty* value.
- "object" and "previousObject": shall never be present, as they define a *Relationship's* object URI.

- "objectList" and "previousObjectList": shall never be present, as they define an ordered array of *Relationship* object URIs.
- "valueList" and "previousValueList": shall never be present, as they define an ordered array of *Property* values.
- "vocab" and "previousVocab": shall never be present, as they define a *VocabProperty* value.

4.5.2.3 Concise NGSI-LD Property

An NGSI-LD Property **without sub-attributes** shall be represented in a concise but lossless representation by a member whose key is the Property name (a term) and whose value is the Property Value (see definition of NGSI-LD Value in clause 3.1). In this case the concise representation is equivalent to simplified representation (see clause 4.5.4). If the provided value is a JSON object and contains a "type" field, the whole Attribute shall be treated as a normalized representation (see clause 4.5.2.2).

An NGSI-LD Property which **includes sub-attributes** shall be represented in a concise but lossless representation by a member whose key is the Property name (a term) and whose value is a JSON-LD object (or JSON-LD array with such JSON-LD objects if there are multiple instances with the same Property name as described in clause 4.5.5) which includes the following members:

Mandatory

- "value": the Property Value (see definition of NGSI-LD Value in clause 3.1). The datatype URI can be placed in the optional "valueType" member.

An NGSI-LD Null (explained in clause 4.5.0 and defined in clause 3.1) can be used as the right-hand side of the "value" during partial update patch and merge patch (see clauses 5.5.8 and 5.5.12) to indicate a deletion of the *Property*, as well as in notifications and in temporal evolutions (for encoding a deleted *Property*).

It should be noted that in the JSON-LD serialization, a **number** data type does not allow for leading zeros, and octal and hexadecimal formats are not used. In the context broker's internal representation, a **number** is equivalent to either an integer or floating-point number, depending on if the number has a non-zero fractional part. The degree of precision of a floating-point number held within a context broker will depend upon the implementation, and insignificant digits may be lost during the deserialization and serialization process.

Optional

- "datasetId": a URI as mandated by clause 4.5.5.
- "expiresAt": a string as mandated by clause 4.22.
- "ngsildproof": a Property with the non-reified subproperties "entityIdSealed" and "entityTypeSealed" as specified in [35]. The value of its "value" element shall be an object containing the W3C® Data integrity "proof" structure [35]. See clause C.11 for an example.
- "observedAt": a string as mandated by clause 4.8.
- "type": If missing, "Property" can be inferred by the presence of the "value" attribute. An exception to this inference rule occurs for geospatial Property Values, where the "GeoProperty" sub-type shall be inferred instead, if the Property Value resolves to a supported GeoJSON geometry (see clause 4.7).
- "unitCode": a string representing the measurement unit corresponding to the Property value. It shall be encoded using the UNECE/CEFACT Common Codes for Units of Measurement [15].
- "valueType": a string value which shall be type coerced into a datatype URI. It is a non-reified alternative to the use of a native JSON-LD @type for the Property value. When it is a JSON primitive, it should align this with the RDF datatype [34].
- For each of the Properties this Property is associated with, a member whose key (a term) is the Property name and value is the result of serializing a *Property* (or any of its subclasses) in **concise representation** (see clause 4.5.2.3).

- For each of the Relationships this Property is associated with, a member whose key (a term) is the Relationship name and value is the result of serializing a *Relationship* (or any of its subclasses) in **concise representation** (see clause 4.5.3.3).

System Generated

- "createdAt": a string as mandated by clause 4.8.
- "deletedAt": a string as mandated by clause 4.8.
- "instanceId": a URI uniquely identifying a Property instance as mandated by clause 4.5.7.
- "modifiedAt": a string as mandated by clause 4.8.

Output Only

- "previousValue": only provided if the *showChanges* option is explicitly requested. It represents the previous Property Value, before the triggering change. The representation is the same as that of "value".

Furthermore, an NGSI-LD Property in the concise representation shall never include the following members:

Prohibited

- "entity": shall never be present, as it is used during inline Linked Entity **retrieval** to define the target Entity of a *Relationship*'s object.
- "entityIdSealed" and "entityTypeSealed" shall never be present, unless the Property name is "ngsildproof".
- "entityList": shall never be present, as it is used during inline Linked Entity **retrieval** to define the target Entities of a *ListRelationship*'s object.
- "languageMap" and "previousLanguageMap": shall never be present, as they define a *LanguageProperty* value.
- "json" and "previousJson": shall never be present, as they define a *JsonProperty* value.
- "object" and "previousObject": shall never be present, as they define a *Relationship*'s object URI.
- "objectList" and "previousObjectList": shall never be present, as they define an ordered array of *Relationship* object URIs.
- "vocab" and "previousVocab": shall never be present, as they define a *VocabProperty* value.
- "valueList" and "previousValueList": shall never be present, as they define an ordered array of *Property* values.

4.5.3 NGSI-LD Relationship Representations

4.5.3.1 Introduction

An NGSI-LD Relationship, its value and sub-attributes can be represented in two equally valid lossless formats. The **normalized representation** is a JSON-LD document that is complete with respect to mandatory members. The **concise representation** is a terser alternative, which makes various implicit assumptions against the payloads and removes redundancy from them.

Both normalized and concise representation of Relationships shall be supported by implementations and can be selected by Context Consumers through specific request parameters. An example of this representation can be found in annex C, clause C.2.2.

4.5.3.2 Normalized NGSI-LD Relationship

An NGSI-LD Relationship in normalized representation shall be represented by a member whose key is the Relationship name (a term) and whose value is a JSON-LD object (or JSON-LD array with such JSON-LD objects, if there are multiple instances with the same Relationship name, as described in clause 4.5.5) with the following terms:

Mandatory

- "object": the Relationship's object represented by a URI or array of URIs.
An NGSI-LD Null (explained in clause 4.5.0 and defined in clause 3.1) can be used as the right-hand side of the "object" during partial update patch and merge patch (see clauses 5.5.8 and 5.5.12) to indicate a deletion of the *Relationship*, as well as in notifications and in temporal evolutions (for encoding a deleted *Relationship*).
- "type": the fixed value "Relationship".

Optional

- "datasetId": a URI as mandated by clause 4.5.5.
- "expiresAt": a string as mandated by clause 4.22.
- "ngsildproof": a Property, as mandated by clause 4.5.2, with the non-reified subproperties "entityIdSealed" and "entityTypeSealed" as specified in [35]. The value of its "value" element shall be an object containing the W3C® Data integrity "proof" structure [35]. See clause C.11 for an example.
- "objectType": a string as mandated by clause 4.5.23.
- "observedAt": a string as mandated by clause 4.8.
- For each Relationship this Relationship is associated with, a member whose key is the Relationship name (a term) and whose value is the result of serializing a Relationship (or any of its subclasses) as per the rules of representation of a *Relationship* in **normalized representation** (see clause 4.5.3.2).
- For each Property this Relationship is associated with, a member whose key is the Property name (a term) and whose value is the result of serializing a Property (or any of its subclasses) as per the rules of representation of a *Property* in **normalized representation** (see clause 4.5.2.2).

System Generated

- "createdAt": a string as mandated by clause 4.8.
- "deletedAt": a string as mandated by clause 4.8.
- "instanceId": a URI uniquely identifying a Relationship instance as mandated by clause 4.5.8.
- "modifiedAt": a string as mandated by clause 4.8.

Output Only

- "entity": only provided in case of **Linked Entity retrieval**, and only if the inline *join* option is explicitly requested (see clause 4.5.23.2), where it used to define the target **Linked Entity** of a Relationship's object in **normalized representation**.
- "previousObject": only provided if the *showChanges* option is explicitly requested. It represents the previous Relationship "object", before the triggering change. The representation is the same as that of "object".

Furthermore, an NGSI-LD Relationship in the normalized representation shall never include the following members:

Prohibited

- "entityIdSealed" and "entityTypeSealed" shall never be present.

- "entityList" shall never be present, as it is used during inline **Linked Entity retrieval** to define the target Entities of a *ListRelationship*'s object.
- "languageMap" and "previousLanguageMap": shall never be present, as they define a *LanguageProperty* value.
- "json" and "previousJson": shall never be present, as they define a *JsonProperty* value.
- "objectList" and "previousObjectList": shall never be present, as they define an ordered array of *Relationship* object URIs.
- "unitCode": shall never be present, as *Relationships* are unitless.
- "value" and "previousValue": shall never be present, as they define a *Property* value.
- "valueList" and "previousValueList": shall never be present, as they define an ordered array of *Property* values.
- "vocab" and "previousVocab": shall never be present, as they define a *VocabProperty* value.

4.5.3.3 Concise NGSI-LD Relationship

An NGSI-LD Relationship shall be represented in a concise but lossless representation by a member whose key is the Relationship name (a term) and whose value is a JSON-LD object (or JSON-LD array with such JSON-LD objects if there are multiple instances with the same Relationship name as described in clause 4.5.5) with the following terms:

Mandatory

- "object": the Relationship's object represented by a URI or array of URIs.
An NGSI-LD Null (explained in clause 4.5.0 and defined in clause 3.1) can be used as the right-hand side of the "object" during partial update patch and merge patch (see clauses 5.5.8 and 5.5.12) to indicate a deletion of the *Relationship*, as well as in notifications and in temporal evolutions (for encoding a deleted *Relationship*).

Optional

- "datasetId": a URI as mandated by clause 4.5.5.
- "expiresAt": a string as mandated by clause 4.22.
- "ngsildproof": a Property, as mandated by clause 4.5.2, with the non-reified subproperties "entityIdSealed" and "entityTypeSealed" as specified in [35]. The value of its "value" element shall be an object containing the W3C® Data integrity "proof" structure [35]. See clause C.11 for an example.
- "observedAt": a string as mandated by clause 4.8.
- "objectType": a string as mandated by clause 4.5.23.
- "type": If missing, "Relationship" can be inferred by the presence of the "object" attribute.
- For each Relationship this Relationship is associated with, a member whose key is the Relationship name (a term) and whose value is the result of serializing a Relationship (or any of its subclasses) as per the rules of representation of a *Relationship* in **concise representation**. (see clause 4.5.3.3).
- For each Property this Relationship is associated with, a member whose key is the Property name (a term) and whose value is the result of serializing a Property (or any of its subclasses) as per the rules of representation of a *Property* in **concise representation**. (see clause 4.5.2.3).

System Generated

- "createdAt": a string as mandated by clause 4.8.

- "deletedAt": a string as mandated by clause 4.8.
- "instanceId": a URI uniquely identifying a Relationship instance as mandated by clause 4.5.8.
- "modifiedAt": a string as mandated by clause 4.8.

Output Only

- "entity": only provided in case of **Linked Entity retrieval**, and only if the inline *join* option is explicitly requested (see clause 4.5.23.2), where it used to define the target **Linked Entity** of a Relationship's object in **concise representation**.
- "previousObject": only provided if the *showChanges* option is explicitly requested. It represents the previous Relationship "object", before the triggering change. The representation is the same as that of "object".

Furthermore, an NGSI-LD Relationship in the concise representation shall never include the following members:

Prohibited

- "entityIdSealed" and "entityTypeSealed" shall never be present.
- "entityList": shall never be present, as it is used during inline **Linked Entity retrieval** (see clause 4.5.23.2) to define the target Entities of a *ListRelationship*'s object.
- "languageMap" and "previousLanguageMap": shall never be present, as they define a *LanguageProperty* value.
- "json" and "previousJson": shall never be present, as they define a *JsonProperty* value.
- "objectList" and "previousObjectList": shall never be present, as they define an ordered array of *Relationship* object URIs.
- "unitCode": shall never be present, as Relationships are unitless.
- "valueList" and "previousValueList": shall never be present, as they define an ordered array of *Property* values.
- "value" and "previousValue": shall never be present, as they define a *Property* value.
- "vocab" and "previousVocab": shall never be present, as they define a *VocabProperty* value.

Notwithstanding the definition above, during partial update patch and merge patch (see clauses 5.5.8 and 5.5.12), an NGSI-LD Relationship with a value of NGSI-LD Null and without a *datasetId* should be represented in a concise representation by a member whose key is the Relationship name (a term) and whose value is "urn:ngsi-ld:null".

4.5.4 Simplified Representation

The NGSI-LD specification defines an abbreviated, lossy representation of Entities, which allows consuming only entity data (the target object of each Relationship or the value of each Property) corresponding to the Properties or Relationships whose subject is the Entity itself i.e. the own Attributes of the Entity. The simplified representation of Entities shall be supported by implementations and can be selected by **Context Consumers** through specific request parameters. An example of this representation can be found in annex C, clause C.2.2.

The simplified representation of an Entity shall be a JSON-LD object containing the following members:

Mandatory

- "id" whose value shall be a URI that identifies the Entity.
- "type" whose value shall be equal to the Entity Type name or an unordered JSON array with multiple Entity Type names in case of an Entity that has multiple Entity Types.

Optional

- "@context", a JSON-LD @context as described in clause 4.4.
- For each *Property* a member whose key is the Property name (a term) and whose value is the Property Value.

EXAMPLE 1: "name": "David Robert Jones"

- In the multi-attribute case (see clause 4.5.5), the simplified representation of a Property (or any of its subtypes) changes. Each *Property* consists of a key-value pair, the key being the Property name (a term) and the value being a JSON Object, which contains a single Attribute with a key called "dataset", and its value in turn is a JSON Object holding a series of key-value pairs, one for each *datasetId*, where the value corresponds to the simplified representation of the property value. The default *datasetId* (where present) is represented by the JSON-LD keyword "@none".

EXAMPLE 2:

```
"name": {
  "dataset": {
    "@none": "David Robert Jones",
    "urn:ngsi-ld:datasetId:001": "David Bowie",
    "urn:ngsi-ld:datasetId:002": "Ziggy Stardust"
  }
}
```

- For each *GeoProperty*, a member whose key is the Property name (a term) and whose value is the Property Value.

EXAMPLE 3: "location": {"type": "Point", "coordinates": [13.3986, 52.5547]}

- For each *LanguageProperty*, a member whose key is the Property name (a term) and whose value is a JSON Object containing a single Attribute with a key called "languageMap" where the value shall correspond to a *LanguageProperty languageMap*.

EXAMPLE 4: "says": {"languageMap": {"en": "yes", "fr": "oui"}}

- For each *JsonProperty*, a member whose key is the Property name (a term) and whose value is a JSON Object containing a single Attribute with a key called "json" where the value shall correspond to the raw JSON data that cannot be held in JSON-LD format. Within the example below, the *id* attribute is never expanded to @id and therefore does not have to be defined as a URI, and the *value* attribute is never expanded to ngsi-ld:hasValue.

EXAMPLE 5:

```
"parkingTickets": {
  "json": {
    "id": "85a6cc52-0589-45f9",
    "value": "Overstay 60 minutes"
  }
}
```

- For each *VocabProperty*, a member whose key is the Property name (a term) and whose value is a JSON Object containing a single Attribute with a key called "vocab" where the value shall correspond to a *VocabProperty vocab*.

EXAMPLE 6: "gender": {"vocab": "Male"}

- For each *Relationship* a term whose key is the Relationship name (a term) and whose value is the Relationship's Object (represented as a URI or array of URIs).

EXAMPLE 7: "providedBy": "urn:ngsi-ld:Device:31415"

EXAMPLE 8:

```
"devices": [
  "urn:ngsi-ld:Device:14142",
  "urn:ngsi-ld:Device:13562",
  "urn:ngsi-ld:Device:37309"
]
```

- In the inline Linked Entity **retrieval** case (see clause 4.5.23.2), the simplified representation of a *Relationship* changes. Any *Relationship* which targets an Entity stored locally or includes an *objectType* Attribute is returned as a JSON object holding key-value pairs corresponding to the data from the Relationship's *object* URI in simplified format.

EXAMPLE 9:

```
"providedBy": {
  "id": "urn:ngsi-ld:Device:31415",
  "type": "Device",
  "brandName": "Anemometer",
  "manufacturerName": "Cyberdyne Systems",
  "windSpeed": 60
}
```

EXAMPLE 10:

```
"devices": [
  {
    "id": "urn:ngsi-ld:Device:14142",
    "type": "Device",
    "brandName": "Anemometer",
    "manufacturerName": "Cyberdyne Systems",
    "windSpeed": 60
  },
  {
    "id": "urn:ngsi-ld:Device:13562",
    "type": "Device",
    "brandName": "Hygromometer",
    "manufacturerName": "Acme Corporation",
    "humidity": 64
  },
  {
    "id": "urn:ngsi-ld:Device:37309",
    "type": "Device",
    "brandName": "Barometer",
    "manufacturerName": "Trask Industries",
    "pressure": 760
  }
]
```

- In the flattened Linked Entity **retrieval** case (see clause 4.5.23.3), the simplified representation of a *Relationship* changes to return multiple Entities. Any *Relationships* which target Entities stored locally or include an *objectType* Attribute are returned as part of the array as an additional JSON object holding key-value pairs corresponding to the data from the *Relationship's object* URI in simplified format.

EXAMPLE 11:

```
[
  {
    ...etc
    "providedBy": "urn:ngsi-ld:Device:31415"
  },
  {
    "id": "urn:ngsi-ld:Device:31415",
```



```

    "type": "Device",
    "brandName": "Anemometer",
    "manufacturerName": "Cyberdyne Systems",
    "windSpeed": 60
  }
]

```

EXAMPLE 12:

```

[
  {
    ... etc
    "providedBy": [
      "urn:ngsi-ld:Device:14142",
      "urn:ngsi-ld:Device:13562",
      "urn:ngsi-ld:Device:37309"
    ]
  },
  {
    "id": "urn:ngsi-ld:Device:14142",
    "type": "Device",
    "brandName": "Anemometer",
    "manufacturerName": "Cyberdyne Systems",
    "windSpeed": 60
  },
  {
    "id": "urn:ngsi-ld:Device:13562",
    "type": "Device",
    "brandName": "Hygromometer",
    "manufacturerName": "Acme Corporation",
    "humidity": 64
  },
  {
    "id": "urn:ngsi-ld:Device:37309",
    "type": "Device",
    "brandName": "Barometer",
    "manufacturerName": "Trask Industries",
    "pressure": 760
  }
]

```

- In the multi-attribute case (see clause 4.5.5), the simplified representation of a *Relationship* changes. Each *Relationship* consists of a key-value pair, the key being the Relationship name (a term) and the value being a JSON Object containing a single Attribute with a key called "dataset" and its value in turn is a JSON Object holding a series of key-value pairs, one for each *datasetId*, where the value corresponds to the object of the *Relationship*. The default *datasetId* (where present) is represented by the JSON-LD keyword "@none".

EXAMPLE 13:

```

"providedBy": {
  "dataset": {
    "@none": "urn:ngsi-ld:Device:31415",
    "urn:ngsi-ld:datasetId:001": "urn:ngsi-ld:Device:27182",
    "urn:ngsi-ld:datasetId:002": "urn:ngsi-ld:Device:14142"
  }
}

```

- For each *ListProperty* a member whose key is the Property name (a term) and whose value is an ordered array holding the Property Values.

EXAMPLE 14: "periods": ["First", "Second", "Third", "Fourth"]

- In the multi-attribute case (see clause 4.5.5), the simplified representation of a *ListProperty* changes. Each *ListProperty* consists of a key-value pair, the key being the Property name (a term) and the value being a JSON Object containing a single Attribute with a key called "dataset" and its value in turn is a JSON Object holding a series of key-value pairs, one for each *datasetId*, where the value corresponds to the simplified representation of the *ListProperty valueList*. The default *datasetId* (where present) is represented by the JSON-LD keyword "@none".

EXAMPLE 15:

```
"periods": {
  "dataset": {
    "@none": ["First", "Second", "Third", "Fourth"],
    "urn:ngsi-ld:datasetId:001": ["1st", "2nd", "3rd", "4th"],
    "urn:ngsi-ld:datasetId:002": ["Primary", "Secondary", "Tertiary",
    "Quaternary"]
  }
}
```

- For each *ListRelationship* a term whose key is the Relationship name (a term) and whose value is an ordered array holding the *ListRelationship's* Objects (represented as URIs).

EXAMPLE 16:

```
"route": [
  "urn:ngsi-ld:BusStop:0101",
  "urn:ngsi-ld:BusStop:0102",
  "urn:ngsi-ld:BusStop:9912"
]
```

- In the inline Linked Entity **retrieval** case (see clause 4.5.23.2), the simplified representation of a *ListRelationship* changes. Any *ListRelationship* which targets Entities stored locally or includes an *objectType* Attribute is returned as an ordered array of JSON objects holding key-value pairs corresponding to the data from the *ListRelationship's* target *objectList* URIs in simplified format.

EXAMPLE 17:

```
"route": [
  {
    "id": "urn:ngsi-ld: BusStop:0101",
    "type": "BusStop",
    "stopName": "High Street",
    "location": {"type": Point, coordinates [54.112,0.334]}}
  },
  {
    "id": "urn:ngsi-ld: BusStop:0102",
    "type": "BusStop",
    "stopName": "Station Road",
    "location": {"type": Point, coordinates [54.101,0.302]}}
  },
  {
    "id": "urn:ngsi-ld: BusStop:9912",
    "type": "BusStop",
    "stopName": "Mornington Crescent",
    "location": {"type": Point, coordinates [54.142,0.332]}}
  }
]
```

- In the flattened Linked Entity **retrieval** case (see clause 4.5.23.3), the simplified representation of a *ListRelationship* changes to return multiple Entities. Any *ListRelationships* which target Entities stored locally or include an *objectType* Attribute are returned as additional JSON objects holding key-value pairs corresponding to the data from the *ListRelationship's* target *objectList* URIs in simplified format.

EXAMPLE 18:

```
[
  {
    ...etc
    "route": [
      "urn:ngsi-ld:BusStop:0101",
      "urn:ngsi-ld:BusStop:0102",
      "urn:ngsi-ld:BusStop:9912"
    ]
  },
  {
    "id": "urn:ngsi-ld: BusStop:0101"
    "type": "BusStop",
    "stopName": "High Street",
    "location": {"type": "Point", "coordinates": [54.112,0.334]}}
  {
    "id": "urn:ngsi-ld: BusStop:0102",
    "type": "BusStop",
    "stopName": "Station Road",
    "location": {"type": "Point", "coordinates": [54.101,0.302]}}
  },
  {
    "id": "urn:ngsi-ld:BusStop:9912"
    "type": "BusStop",
    "stopName": "Mornington Crescent",
    "location": {"type": "Point", "coordinates": [54.142,0.332]}}
  }
]
```

- In the multi-attribute case (see clause 4.5.5), the simplified representation of a *ListRelationship* changes. Each *ListRelationship* consists of a key-value pair, the key being the Relationship name (a term) and the value being a JSON Object containing a single Attribute with a key called "dataset" and its value in turn is a JSON Object holding a series of key-value pairs, one for each *datasetId*, where the value corresponds to the array of objects of the relationship list. The default *datasetId* (where present) is represented by the JSON-LD keyword "@none".

EXAMPLE 19:

```
"route": {
  "dataset": {
    "@none": [
      "urn:ngsi-ld:BusStop:0101",
      "urn:ngsi-ld:BusStop:0102",
      "urn:ngsi-ld:BusStop:9912"
    ],
    "urn:ngsi-ld:datasetId:001": [
      "urn:ngsi-ld:BusStop:0101",
      "urn:ngsi-ld:BusStop:0102",
      "urn:ngsi-ld:BusStop:0022"
    ],
    "urn:ngsi-ld:datasetId:002": [
      "urn:ngsi-ld:BusStop:0101",
      "urn:ngsi-ld:BusStop:0102",
      "urn:ngsi-ld:BusStop:0022",
      "urn:ngsi-ld:BusStop:9912"
    ]
  }
}
```

NOTE: When the simplified GeoJSON representation is selected, the layout of the Entities changes, see clause 4.5.17 for details.

4.5.5 Multi-Attribute Support

4.5.5.1 Introduction

For each Entity, there can be Attributes that simultaneously have more than one instance. In the case of Properties, there may be more than one source at a time that provides a Property instance, e.g. based on independent sensor measurements with different quality characteristics. For instance, take a speedometer and a GPS both providing the current speed of a car. In the case of Relationships, there may be non-functional Relationships requiring separate metadata attached to each object, e.g. for a room, there may be multiple "contains" Relationships to all sorts of objects currently in the room that have been put there by different people (a separate "placedBy" Relationship-of-the-Relationship) and which are dynamically changing over time.

To be able to explicitly manage such multi-attributes, the optional *datasetId* property is used, which is of datatype URI, or equal to the JSON-LD keyword "@none". It is introduced for Properties and Relationships in clauses 4.5.2 and 4.5.3 respectively. If a *datasetId* is provided when creating, updating, appending or deleting Attributes, only instances with the same *datasetId* are affected, leaving instances with another *datasetId* or an instance without a *datasetId* untouched. If no *datasetId* is provided, or "datasetId": "@none" is supplied, it is considered as the default Attribute instance. Thus, the creation, updating, appending or deleting of Attributes without providing a *datasetId* only affects the default Attribute instance. There can only be one default Attribute instance for an Attribute with a given Attribute name in any request or response. An example can be found in annex C, clause C.2.2.

When requesting Entity information, if there are multiple instances of matching Attributes these are returned as arrays of Attributes, instead of a single Attribute element. The *datasetId* of the default Attribute instance is never explicitly included in responses, i.e. a default Attribute instance does not have a *datasetId*.

The *datasetId* can be used to create different "views" of Entities. All Attribute instances sharing the same *datasetId* can be seen as one "view". If one or more *datasetIds* are specified in the request, only the Attribute instances that match one of the *datasetIds* will be returned. The *datasetId* of the default Attribute instance can be specified as "@none". In case that no Attribute instances match the provided *datasetIds*, then the Attribute shall not be returned with the Entity. If no *datasetId* is provided, then all available Attribute instances will be returned.

There is no multi-attribute support for non-reified Attributes, in particular this applies to the Temporal Properties *createdAt*, *modifiedAt*, *deletedAt*, *expiresAt* and *observedAt*, and also the *unitCode* Property.

4.5.5.2 Processing of Conflicting Transient Entities

In case of conflicting information when an Entity is received from a registered Context Source and marked with an *expiresAt DateTime*, but this *expiresAt* is not duplicated across all versions of the Entity received across all registered Context Sources, the following mechanism shall be used to differentiate.

For each version of the Entity received where the *expiresAt* non-reified attribute is present at the Entity level:

- If no *expiresAt* attribute is present as a property of an Attribute, a non-reified *expiresAt* attribute is added as a property of that Attribute corresponding to the *expiresAt* found on the Entity.
- If the *expiresAt* attribute is present as a property of an Attribute, and the value on the Attribute is further in the future than the *expiresAt* value found on the Entity, the value of the *expiresAt* on the Attribute is overwritten to corresponding to the earlier *DateTime*.

4.5.5.3 Processing of Conflicting Attributes

In case of conflicting information for an Attribute, where a *datasetId* is duplicated, but there are differences in the other attribute data, if an *expiresAt DateTime* is present on the Attribute and the date lies in the past, it shall be discarded, thereafter the one with the most recent *observedAt DateTime*, if present, and otherwise the one with the most recent *modifiedAt DateTime* shall be provided. If no other mechanism for determining the most current Attribute instance is found, the NGSI-LD system shall choose the Attribute instance at random and the result is indeterminate.

When conflicting information is received for the non-reified *expiresAt* attribute at an Entity level:

- If an *expiresAt* attribute is missing from at least one version of the Entity received from registered Context Sources, the *expiresAt* attribute is removed from the Entity.

- Otherwise, the *expiresAt* attribute of the Entity is set to the *DateTime* corresponding to the *expiresAt DateTime* which is furthest in the future.

4.5.6 Temporal Representation of an Entity

The temporal representation of an Entity is the way to represent the **Temporal Evolution of an Entity**: the Entity shall be as mandated by clause 4.5.1, but for each Property and Relationship their temporal representation shall be provided as mandated by clauses 4.5.7 and 4.5.8 and the Scope (if present) shall be represented as a temporal representation of a Property (clause 4.5.7) that can only have the non-reified Temporal Properties *createdAt*, *modifiedAt*, *deletedAt* and *observedAt* as sub-Properties. This is required to represent the Scope of a **Temporal Evolution of an Entity**. In case the Temporal Evolution of the Scope is updated as the result of a change from the **Core API**, the *observedAt* sub-Property should be set as a copy of the *modifiedAt* sub-Property.

4.5.7 Temporal Representation of a Property

The temporal evolution of a Property (for instance, its historical evolution or future predictions) is composed of the sequence of instances of the referred Property during a period of time within its lifetime.

The temporal representation of a Property shall be provided as an Array of JSON-LD objects (or a single JSON-LD object in case only one is present), each one representing an instance of the Property (as mandated by clause 4.5.2) at a particular point in time, which is recorded as a Temporal Property of the instance (typically *observedAt*). See example in annex C, clause C.5.6. In case the *Property* is deleted, an instance of the *Property* is recorded with its *value* set to the URI "urn:ngsi-ld:null" and the *deletedAt* Temporal Property set.

Systems should maintain an *instanceId* for each such *Property* instance. Without such an *instanceId*, it is not possible to selectively modify or delete temporal information via the NGSI-LD API. The consequences of this may be severe in the case of modification or deletion requests for legal reasons, e.g. GDPR [i.18]. When implementing the NGSI-LD API on storage systems that do NOT allow modification or deletion, similar problems may be encountered.

4.5.8 Temporal Representation of a Relationship

The temporal evolution of a Relationship (for instance, its historical evolution or future predictions) is composed of the sequence of instances of the referred Relationship during a period of time within its lifetime.

The temporal representation of an NGSI-LD Relationship shall be provided as an Array of JSON-LD objects (or a single JSON-LD object in case only one is present), each one representing an instance of the Relationship (as mandated by clause 4.5.3) at a particular point in time, which is recorded as a Temporal Property of the instance (typically *observedAt*). See example in annex C, clause C.5.5. In case the *Relationship* is deleted, an instance of the *Relationship* is recorded with its *object* set to the URI "urn:ngsi-ld:null" and the *deletedAt* Temporal Property set.

Systems should maintain an *instanceId* for each such Relationship instance. Without such an *instanceId*, it is not possible to selectively modify or delete temporal information via the NGSI-LD API. The consequences of this may be severe in the case of modification or deletion requests for legal reasons, e.g. GDPR [i.18]. When implementing the NGSI-LD API on storage systems that do NOT allow modification or deletion, similar problems may be encountered.

4.5.9 Simplified temporal representation of an Entity

The NGSI-LD specification defines an alternative, abbreviated temporal representation of **Temporal Evolution of Entities**, which allows consuming temporal Entity data in a more straightforward manner. The simplified temporal representation of Entities shall be supported by implementations and can be selected by **Context Consumers** through specific request parameters. An example can be found in annex C, clause C.5.6.

The simplified temporal representation of an Entity shall be a JSON-LD object containing the following members:

Mandatory

- "id" whose value shall be a URI that identifies the Entity.
- "type" whose value shall be equal to the Entity Type name or an unordered JSON array with multiple Entity Type names in case of an Entity that has multiple Entity Types.

Optional

- "@context", a JSON-LD *@context* as described in clause 4.4.
- For each *Property*, a member whose key is the Property name (a term), the member value shall be a JSON-LD object labelled with the type "Property". Such JSON-LD object shall only contain a member whose key shall be "values". The value of the referred "values" member shall be a JSON-LD Array that shall contain as many array elements as *Property* instances (i.e. data points of the concerned *Property*) being represented. Each array element shall be another Array containing exactly two array elements: the first element shall be a Property value and the second element shall correspond to the associated Temporal Property (for instance *observedAt*).

EXAMPLE 1:

```
"name": {
  "type": "Property",
  "values": [
    [
      "Joe Bloggs",
      "2022-08-09T18:25:02Z"
    ],
    [
      "Bill Smith",
      "2022-08-10T18:25:02Z"
    ]
  ]
}
```

- For each *GeoProperty*, a member whose key is the Property name (a term), the member value shall be a JSON-LD object labelled with the type "GeoProperty". Such JSON-LD object shall only contain a member whose key shall be "values". The value of the referred "values" member shall be a JSON-LD Array that shall contain as many array elements as *GeoProperty* instances (i.e. data points of the concerned *GeoProperty*) being represented. Each array element shall be another Array containing exactly two array elements: the first element shall be a *GeoProperty* value and the second element shall correspond to the associated Temporal Property (for instance *observedAt*).
- For each *LanguageProperty*, a member whose key is the Property name (a term), the member value shall be a JSON-LD object labelled with the type "LanguageProperty". Such JSON-LD object shall only contain a member whose key shall be "languageMaps". The value of the referred "languageMaps" member shall be a JSON-LD Array that shall contain as many array elements as *LanguageProperty* instances (i.e. data points of the concerned *LanguageProperty*) being represented. Each array element shall be another Array containing exactly two array elements: the first element shall be a JSON Object containing a single Attribute with a key called "languageMap" where the value shall correspond to a *LanguageProperty languageMap* and the second element shall correspond to the associated Temporal Property (for instance *observedAt*).

EXAMPLE 2:

```
"says": {
  "type": "LanguageProperty",
  "languageMaps": [
    [
      {"languageMap": {"en": "yes", "fr": "oui"}},
      "2022-08-09T18:25:02Z"
    ],
    [
      {"languageMap": {"en": "no", "fr": "non"}},
      "2022-08-10T18:25:02Z"
    ]
  ]
}
```

- For each *ListProperty*, a member whose key is the Property name (a term), the member value shall be a JSON-LD object labelled with the type "ListProperty". Such JSON-LD object shall only contain a member whose key shall be "valueLists". The value of the referred "valueLists" member shall be a JSON-LD Array that shall contain as many array elements as *ListProperty* instances (i.e. data points of the concerned *ListProperty*) being represented. Each array element shall be another array containing exactly two elements: the first element shall be an ordered array of Property values, and the second element shall correspond to the associated Temporal Property (for instance *observedAt*).

EXAMPLE 3:

```
"period": {
  "type": "ListProperty",
  "valueLists": [
    [
      ["First", "Second", "Third", "Fourth"],
      "2022-08-09T18:25:02Z"
    ],
    [
      ["1st", "2nd", "3rd", "4th"],
      "2022-08-10T18:25:02Z"
    ]
  ]
}
```

- For each *JsonProperty*, a member whose key is the Property name (a term), the member value shall be a JSON-LD object labelled with the type "JsonProperty". Such JSON-LD object shall only contain a member whose key shall be "jsons". The value of the referred "jsons" member shall be a JSON-LD Array that shall contain as many array elements as *JsonProperty* instances (i.e. data points of the concerned *JsonProperty*) being represented. Each array element shall be another Array containing exactly two array elements: the first element shall be a JSON Object containing a single Attribute with a key called "json", where the value shall correspond to raw JSON data that cannot be held in JSON-LD format and the second element shall correspond to the associated Temporal Property (for instance *observedAt*).

EXAMPLE 4:

```
"parkingTickets": {
  "type": "JsonProperty",
  "jsons": [
    [
      {
        "json": [
          {
            "id": "85a6cc52-0589-45f9",
            "value": "Overstay 60 minutes"
          }
        ]
      },
      "2022-08-09T18:25:02Z"
    ],
    [
      {
        "json": [
          {
            "id": "85a6cc52-0589-45f9",
            "value": "Overstay 60 minutes"
          },
          {
            "id": "x5c56s-0589-45f9",
            "value": "Overstay 45 minutes"
          }
        ]
      },
      "2022-08-10T18:25:02Z"
    ]
  ]
}
```

```
    ]
  }
}
```

- For each *VocabProperty*, a member whose key is the Property name (a term), the member value shall be a JSON-LD object labelled with the type "VocabProperty". Such JSON-LD object shall only contain a member whose key shall be "vocabs". The value of the referred "vocabs" member shall be a JSON-LD Array that shall contain as many array elements as *VocabProperty* instances (i.e. data points of the concerned *VocabProperty*) being represented. Each array element shall be another Array containing exactly two array elements: the first element shall be a JSON Object containing a single Attribute with a key called "vocab", where the value shall correspond to a *VocabProperty* *vocab* and the second element shall correspond to the associated Temporal Property (for instance *observedAt*).

EXAMPLE 5:

```
"gender": {
  "type": "VocabProperty",
  "vocabs": [
    [
      {"vocab": "Male"},
      "2022-08-09T18:25:02Z"
    ],
    [
      {"vocab": "Female"},
      "2022-08-10T18:25:02Z"
    ]
  ]
}
```

- For each *Relationship*, a term whose key is the Relationship name (a term). The member value shall be a JSON-LD object labelled with the type "Relationship". Such JSON-LD object shall only contain a member whose key shall be "objects". The value of the referred "objects" member shall be a JSON-LD Array that shall contain as many array elements as *Relationship* instances (i.e. data points of the concerned *Relationship*) being represented. Each array element shall be another array containing exactly two elements: the first element shall be a *Relationship* object (a URI or array of URIs) and the second element shall correspond to the associated Temporal Property (for instance *observedAt*).

EXAMPLE 6:

```
"spouse": {
  "type": "Relationship",
  "objects": [
    [
      "urn:ngsi-ld:Person:123455",
      "2022-08-09T18:25:02Z"
    ],
    [
      "urn:ngsi-ld:Person:999999",
      "2022-08-10T18:25:02Z"
    ]
  ]
}
```

EXAMPLE 7:

```
"activeDevices": {
  "type": "Relationship",
  "objects": [
    [
      ["urn:ngsi-ld:Device:14142", "urn:ngsi-ld:Device:13562"],
      "2022-08-09T18:25:02Z"
    ],
  ],
}
```



```
[
  ["urn:ngsi-ld:Device:14142", "urn:ngsi-ld:Device:13562", "urn:ngsi-ld:Device:37309"],
  "2022-08-10T18:25:02Z"
]
}
```

- For each *ListRelationship*, a term whose key is the Relationship name (a term), the member value shall be a JSON-LD object labelled with the type "ListRelationship". Such JSON-LD object shall only contain a member whose key shall be "objectLists". The value of the referred "objectLists" member shall be a JSON-LD Array that shall contain as many array elements as *ListRelationship* instances (i.e. data points of the concerned *ListRelationship*) being represented. Each array element shall be another array containing exactly two elements: the first element shall be an ordered array of Relationship objects (URIs) and the second element shall correspond to the associated Temporal Property (for instance *observedAt*).

EXAMPLE 7:

```
"membersPresent": {
  "type": "ListRelationship",
  "objectLists": [
    [
      ["urn:ngsi-ld:Person:Alice", "urn:ngsi-ld:Person:Bob"],
      "2022-08-09T18:25:02Z"
    ],
    [
      ["urn:ngsi-ld:Person:Alice", "urn:ngsi-ld:Person:Eve", "urn:ngsi-ld:Person:Mallory"],
      "2022-08-10T18:25:02Z"
    ]
  ]
}
```

4.5.10 Entity Type List Representation

The entity type list representation is used to consume information about entity types. The entity type list representation shall be a JSON-LD object containing the following members:

Mandatory

- "id" whose value shall be a URI that identifies the entity type list.
- "type": the fixed value "EntityTypeList".
- "typeList": JSON-LD array containing the entity type names.

Optional

- "@context" a JSON-LD *@context* as described in clause 4.4.

4.5.11 Detailed Entity Type List Representation

The detailed entity type list representation is used to consume detailed information about entity types including the names of attributes that instances of each entity type can have. The detailed entity type list representation shall be an array of JSON-LD objects containing the following members:

Mandatory

- "attributeNames": JSON-LD array containing the names of attributes that instances of the entity type can have.
- "id" whose value shall be the URI that identifies the entity type.

- "type": the fixed value "EntityType".
- "typeName": name of entity type, short name if contained in *@context*.

Optional

- "@context": a JSON-LD *@context* as described in clause 4.4.

4.5.12 Entity Type Information Representation

The entity type information representation is used to consume detailed information about an entity type. The entity type information representation shall be a JSON-LD object containing the following members:

Mandatory

- "id" whose value shall be the URI that identifies the entity type.
- "type": the fixed value "EntityTypeInfo".
- "typeName": the URI that identifies the entity type (short name in case of availability in *@context*).

Optional

- "@context": a JSON-LD *@context* as described in clause 4.4.

4.5.13 Attribute List Representation

The attribute list representation is used to consume information about attributes. The attribute list representation shall be a JSON-LD object containing the following members:

Mandatory

- "attributeList": JSON-LD array containing the attribute names.
- "id": whose value shall be a URI that identifies the attribute list.
- "type": the fixed value "AttributeList".

Optional

- "@context": a JSON-LD *@context* as described in clause 4.4.

4.5.14 Detailed Attribute List Representation

The detailed attribute list representation is used to consume detailed information about attributes including the names of entity types that have instances with attributes, which have the respective attribute name. The detailed attribute list representation shall be an array of JSON-LD objects containing the following members:

Mandatory

- "attributeName": the URI that identifies the attribute (short name in case of availability in *@context*).
- "id": whose value shall be the URI that identifies the attribute.
- "type": the fixed value "Attribute".

Optional

- "@context": a JSON-LD *@context* as described in clause 4.4.
- "typeNames": an array of the names of entity types that have instances with attributes, which have the respective attribute name.

4.5.15 Attribute Information Representation

The attribute information representation is used to consume detailed information about an attribute. The attribute information representation shall be a JSON-LD object containing the following members:

Mandatory

- "attributeName": the URI that identifies the attribute (short name in case of availability in *@context*).
- "id": whose value shall be the URI that identifies the attribute.
- "type": the fixed value "Attribute".

Optional

- "@context": a JSON-LD *@context* as described in clause 4.4.
- "attributeCount": number of instances of this attribute.
- "attributeTypes": an array of attribute types (e.g. Property, Relationship, GeoProperty) for which instances with the attribute name exist.
- "typeNames": an array of the names of entity types that have instances with attributes, which have the respective attribute name.

4.5.16 GeoJSON Representation of Entities

4.5.16.0 Foreword

The NGSI-LD specification defines an alternative representation of Entities, to make NGSI-LD responses compatible with GIS (Geographic Information System) applications which support the GeoJSON format [8] and/or GeoJSON-LD [i.20].

Every NGSI-LD Entity can be represented as a GeoJSON *Feature* object, where a *Feature* object represents any spatially bounded thing as defined by its geometry.

4.5.16.1 Top-level "geometry" field selection algorithm

A parameter of the request (named *geometryProperty*) may be used to indicate the name of the *GeoProperty* to be selected. If this parameter is not present, then the default name of "location" shall be used.

If the selected *GeoProperty* has multiple instances as described in clause 4.5.5, either a *datasetId* shall be specified, in order to define which instance of the value is to be selected, or a default attribute instance exists, which is then selected, if no *datasetId* was specified.

If an entity lacks the *GeoProperty* as specified or the value does not hold a valid GeoJSON *geometry* object then the *geometry* shall be undefined and returned with a value of *null* - which is syntactically valid GeoJSON.

4.5.16.2 GeoJSON Representation of an individual Entity

The GeoJSON representation of a spatially bounded Entity is defined as a single GeoJSON *Feature* object including the following members:

Mandatory

- "geometry": The value of the selected *GeoProperty* (a GeoJSON geometry object) used to define the spatial location of the Entity. Note that no sub-Attributes of the selected *GeoProperty* are present in the representation.
- "id": the Entity id.

- "properties": A JSON object containing the following members:
 - "type": the Entity Type name of the Entity or an unordered JSON array with the Entity Type names of the Entity.
 - One member for each *Property* (including the selected *GeoProperty*) as per the rules stated in clause 4.5.2. In case of multiple *Property* instances with the same Property name as described in clause 4.5.5, all instances are provided as an unordered JSON array.
 - One member for each *Relationship* as per the rules stated in clause 4.5.3. In case of multiple *Relationship* instances with the same Relationship name as described in clause 4.5.5, all instances are provided as an unordered JSON array.
- "type": the fixed value "Feature".

Optional

- A JSON-LD *@context* as described in clause 4.4 if requested as part of the payload body.

This representation shall be fully compliant with *Feature* as defined within IETF RFC 7946 [8].

An example can be found in annex C, clause C.2.3.

4.5.16.3 GeoJSON Representation of Multiple Entities

The GeoJSON representation of a list of spatially bounded Entities is defined as a single GeoJSON *FeatureCollection* object containing an array of GeoJSON *Feature* objects as follows:

Mandatory

- "type": the fixed value "FeatureCollection".
- "features": a JSON array of GeoJSON *Feature* objects as defined in clause 4.5.16.2. Note that separate *@context* elements for each *Feature* will not be present in the payload body.

Optional

- A JSON-LD *@context* as described in clause 4.4 if requested as part of the payload body.

This representation shall be fully compliant with *FeatureCollection* as defined within IETF RFC 7946 [8].

An example can be found in annex C, clause C.2.3.

4.5.17 Simplified GeoJSON Representation of Entities

4.5.17.0 Foreword

When both simplified (see clause 4.5.4) and GeoJSON representation is requested, the following simplified GeoJSON representation compatible with GIS systems shall be returned.

4.5.17.1 Simplified GeoJSON Representation of an individual Entity

The simplified GeoJSON representation of a spatially bounded Entity is defined as a single GeoJSON *Feature* object as follows:

Mandatory

- "geometry": The value of the selected *GeoProperty* (a GeoJSON geometry object) used to define the spatial location of the Entity.
- "id": the Entity id.
- "type": the fixed value "Feature".

- "properties": An array containing the following attributes:
 - "type": **Mandatory** - the Entity Type name of the Entity or an unordered JSON array with the Entity Type names of the Entity.
 - For each *Property* (including the selected *GeoProperty*) a member whose key is the Property name (a term) and whose value is the Property Value. In the multi-attribute case, each Property consists of a key-value pair, the key being the Property name (a term) and the value being a JSON Object, which contains a single Attribute with a key called "dataset", and its value in turn is a JSON Object holding a series of key-value pairs, one for each *datasetId*, where the value corresponds to the simplified representation of the property value. The default *datasetId* (where present) is represented by the JSON-LD keyword "@none".
 - For each *Relationship* a term whose key is the Relationship name (a term) and whose value is the Relationship's Object (represented as a URI). In the multi-attribute case, each Relationship consists of a key-value pair, the key being the Relationship name (a term) and the value being a JSON Object containing a single Attribute with a key called "dataset" and its value in turn is a JSON Object holding a series of key-value pairs, one for each *datasetId* where the value corresponds to the object of the relationship. The default *datasetId* (where present) is represented by the JSON-LD keyword "@none".

Optional

- A JSON-LD @context as described in clause 4.4 if requested as part of the payload body.

The selection of the geometry field is defined in clause 4.5.16.1.

This representation shall be fully compliant with *Feature* as defined within IETF RFC 7946 [8].

An example can be found in annex C, clause C.2.3.

4.5.17.2 Simplified GeoJSON Representation of multiple Entities

The simplified GeoJSON representation of a list of spatially bounded Entities is defined as a single GeoJSON *FeatureCollection* object containing an array of GeoJSON *Feature* objects as follows:

Mandatory

- "features": a JSON array of simplified GeoJSON *Feature* objects as defined in clause 4.5.17.1. Note that separate @context elements for each *Feature* will not be present in the payload body.
- "type": the fixed value "FeatureCollection".

Optional

- A JSON-LD @context as described in clause 4.4 if requested as part of the payload body.

This representation shall be fully compliant with *FeatureCollection* as defined within IETF RFC 7946 [8].

4.5.18 NGSI-LD LanguageProperty Representations

4.5.18.1 Introduction

NGSI-LD defines a specialized type of Property named *LanguageProperty*, defined by the NGSI-LD @context described by the present document in clause 4.4.

When dealing with NGSI-LD Entities, implementations shall interpret the JSON-LD nodes of type *LanguageProperty* as per clause 4.5.18.2 (when in normalized representation) or clause 4.5.18.3 (when in concise representation).

Both normalized and concise representation of *LanguageProperties* shall be supported by implementations and can be selected by Context Consumers through specific request parameters. An example of this representation can be found in annex C, clause C.2.2.

4.5.18.2 Normalized NGSI-LD LanguageProperty

An NGSI-LD *LanguageProperty* shall be represented in normalized representation by a member whose key is the Property name (a term), whose value is the same as the JSON-LD object in **NGSI-LD Property Representation** defined in clause 4.5.2.2, with the following differences:

Mandatory

- "languageMap": a JSON object consisting of a set of a non-empty language tags as defined by IETF RFC 5646 [28] or the language tag "@none" which represents a default language, with each language tag mapping to a single string or array of strings. It represents a more specialized *value*.

An NGSI-LD Null used during partial update patch and merge patch (see clauses 5.5.8 and 5.5.12) shall be encoded as the JSON object {"@none": "urn:ngsi-ld:null"}. The same representation is also used to indicate a deletion in notifications and in temporal evolutions for encoding a deleted *LanguageProperty*.

- "type": the fixed value "LanguageProperty".

Output Only

- "previousLanguageMap": only provided only in the case of notifications and only if the *showChanges* option is explicitly requested. It represents the previous *LanguageProperty languageMap*, before the triggering change. The representation is the same as that of "languageMap".

Furthermore, an NGSI-LD *LanguageProperty* in the normalized representation shall never include the following members:

Prohibited

- "unitCode": shall never be present, as language maps are always strings and hence unitless.
- "value" and "previousValue": shall never be present, as *value* is a generalization of *languageMap*.

4.5.18.3 Concise NGSI-LD LanguageProperty

An NGSI-LD *LanguageProperty* shall be represented in concise but lossless representation by a member whose key is the Property name (a term), whose value is the same as the JSON-LD object in **NGSI-LD Property Representation** defined in clause 4.5.2.3, with the following differences:

Mandatory

- "languageMap": a JSON object consisting of a set of a non-empty language tags as defined by IETF RFC 5646 [28] or the language tag "@none" which represents a default language, with each language tag mapping to a single string or array of strings. It represents a more specialized *value*.

An NGSI-LD Null used during partial update patch and merge patch (see clauses 5.5.8 and 5.5.12) shall be encoded as the JSON object {"@none": "urn:ngsi-ld:null"}. The same representation is also used to indicate a deletion in notifications and the temporal evolutions for encoding a deleted *LanguageProperty*.

Optional

- "type": If missing, "LanguageProperty" can be inferred by the presence of the "languageMap" attribute.

Output Only

- "previousLanguageMap": only provided only in the case of notifications and only if the *showChanges* option is explicitly requested. It represents the previous *LanguageProperty languageMap*, before the triggering change. The representation is the same as that of "languageMap".

Furthermore, an NGSI-LD *LanguageProperty* in the concise representation shall never include the following members:

Prohibited

- "unitCode": shall never be present, as language maps are always strings and hence unitless.
- "value" and "previousValue": shall never be present, as it is a generalization of "languageMap".

Notwithstanding the definition above, during partial update patch and merge patch (see clauses 5.5.8 and 5.5.12), an NGSI-LD LanguageProperty with a value of NGSI-LD Null without a *datasetId* should be represented in a concise representation by a member whose key is the LanguageProperty name (a term) and whose value is "urn:ngsi-ld:null".

4.5.19 Aggregated temporal representation of an Entity

4.5.19.0 Foreword

The NGSI-LD specification defines an alternative temporal representation of Entities, called aggregated temporal representation, which allows consuming temporal Entity data after applying an aggregation function on the values of the Attribute instances. The aggregated temporal representation of Entities shall be supported by implementations supporting the temporal representation of Entities and can be selected by Context Consumers through specific request parameters. An example can be found in annex C, clause C.5.14.

The aggregation function is applied according to the following principles:

- An aggregation method specifies the function used to aggregate the values (e.g. sum, mean, etc.). A Context Consumer can ask for many aggregation methods in one request.
- The duration of an aggregation period specifies the duration of each period to be used when applying the aggregation function on the values of a Temporal Entity.

The aggregated temporal representation of an Entity shall include the following:

- A JSON-LD object containing the following members:
 - *id*, *type* and *@context* as described in clause 4.5.1.
 - For each *Property* a member whose key is the Property name (a term). The member value shall be a JSON-LD object labelled with the type "Property". Such JSON-LD object shall contain one member per aggregation method requested by the Context Consumer. Each member uses the aggregation method name as a key. The value of each member shall be a JSON-LD Array that shall contain as many array elements as there are periods in the time range of the query. Each array element shall be another Array containing exactly three array elements in the following order:
 - 1) the value obtained after applying the aggregation method over the period;
 - 2) the start *DateTime* of the corresponding period;
 - 3) the end *DateTime* of the corresponding period.
 - For each *Relationship* a term whose key is the Relationship name (a term). The member value shall be a JSON-LD object labelled with the type "Relationship". Such JSON-LD object shall contain one member per aggregation method requested by the Context Consumer. Each member uses the aggregation method name as a key. The value of each member shall be a JSON-LD Array that shall contain as many array elements as there are periods in the time range of the query. Each array element shall be another array containing exactly three array elements in the following order:
 - 1) the value obtained after applying the aggregation method over the period;
 - 2) the start *DateTime* of the corresponding period;
 - 3) the end *DateTime* of the corresponding period.

An example of this aggregated temporal representation can be found in annex C, clause C.5.14.

4.5.19.1 Supported behaviours for aggregation functions

In order to support such aggregation functions, two parameters are defined:

- *aggrMethods*, to express the aggregation methods to apply.
- *aggrPeriodDuration* to express the duration of the period to consider in each step of the aggregation.

The duration is expressed using the ISO 8601 [17] Duration Representation and in particular using the following format and conventions:

- The duration shall be a string in the format P[n]Y[n]M[n]DT[n]H[n]M[n]S or P[n]W, where [n] is replaced by the value for each of the date and time elements that follow the [n], P is the duration designator and T is the time designator. For example, "P3Y6M4DT12H30M5S" represents a duration of "three years, six months, four days, twelve hours, thirty minutes, and five seconds".
- Date and time elements including their designator may be omitted if their value is zero.
- Lower-order elements may be omitted for reduced precision.
- A duration of 0 second (e.g. expressed as "PT0S" or "P0D") is valid and is interpreted as a duration spanning the whole time range specified by the temporal query.
- Alternative representations based on combined date and time representations are not allowed.

The values supported by the *aggrMethods* parameter are the following:

- *aggrMethods* = "avg" / "distinctCount" / "max" / "min" / "stddev" / "sum" / "sumsq" / "totalCount"

The semantics of the different aggregation methods defined above is as follows, and shall be supported by compliant implementations:

Table 4.5.19.1-1: Semantics of aggregation methods for Properties on JSON native data types

Aggregation Method	JSON String	JSON Number	JSON Object	JSON Array	JSON Boolean (see note)
avg	N/A	Calculate the average of the values inside the period	N/A	Calculate the average number of the sizes of the arrays inside the period	Calculate the average of the values inside the period
distinctCount	Calculate the count of distinct values inside the period				
max	Calculate the last value in lexicographical order inside the period	Calculate the maximum value inside the period	N/A	Calculate the maximum size of the arrays inside the period	Calculate the maximum value inside the period
min	Calculate the first value in lexicographical order inside the period	Calculate the minimum value inside the period	N/A	Calculate the minimum size of the arrays inside the period	Calculate the minimum value inside the period
stddev	N/A	Calculate the standard deviation of the values inside the period	N/A	N/A	Calculate the standard deviation of the values inside the period
sum	N/A	Calculate the sum of the values inside the period	N/A	Calculate the sum of the sizes of the arrays inside the period	Calculate the sum of the values inside the period

Aggregation Method	JSON String	JSON Number	JSON Object	JSON Array	JSON Boolean (see note)
sumsq	N/A	Calculate the sum of the square of the values inside the period	N/A	N/A	Calculate the sum of the square of the values inside the period
totalCount	Calculate the number of times the value has been updated inside the period				
NOTE: For the purpose of aggregation, <i>true</i> is considered as a value of 1, <i>false</i> is considered as a value of 0.					

Table 4.5.19.1-2: Semantics of aggregation methods for Properties on other supported data types

Aggregation Method	DateTime	Date	Time	URI
avg	N/A	N/A	Calculate the average time inside the period (e.g. to apply on an event that occurs at non fixed times, like the time a car enters a given parking)	N/A
distinctCount	Calculate the count of distinct values inside the period			
max	Calculate the maximum value inside the period	Calculate the maximum value inside the period	Calculate the maximum value inside the period	N/A
min	Calculate the minimum value inside the period	Calculate the minimum value inside the period	Calculate the minimum value inside the period	N/A
stddev	N/A	N/A	N/A	N/A
sum	N/A	N/A	N/A	N/A
sumsq	N/A	N/A	N/A	N/A
totalCount	Calculate the number of times the value has been updated inside the period			

Table 4.5.19.1-3: Semantics of aggregation methods for Relationships

Aggregation Method	Relationship
avg	N/A
distinctCount	Calculate the count of distinct relationships targets inside the period
max	N/A
min	N/A
stddev	N/A
sum	N/A
sumsq	N/A
totalCount	Calculate the number of times the relationship has been updated inside the period

4.5.20 NGSI-LD VocabProperty Representations

4.5.20.1 Introduction

NGSI-LD defines a specialized type of Property named `VocabProperty`, defined by the NGSI-LD *@context* described by the present document in clause 4.4.

When dealing with NGSI-LD Entities, implementations shall interpret the JSON-LD nodes of type `VocabProperty` as per clause 4.5.20.2 (when in normalized representation) or clause 4.5.20.3 (when in concise representation).

Both normalized and concise representation of `VocabProperties` shall be supported by implementations and can be selected by `Context Consumers` through specific request parameters. An example of this representation can be found in annex C, clause C.2.2.

4.5.20.2 Normalized NGSI-LD VocabProperty

An NGSI-LD `VocabProperty` shall be represented in normalized representation by a member whose key is the Property name (a term), whose value is the same as the JSON-LD object in **NGSI-LD Property Representation** defined in clause 4.5.2.2, with the following differences:

Mandatory

- `"type"`: the fixed value `"VocabProperty"`.
- `"vocab"`: a JSON object consisting of a single string or array of strings which can be type coerced into an IRI or array of IRIs. It represents a more specialized *value*.

Output Only

- `"previousVocab"`: only provided in case of notifications and only if the *showChanges* option is explicitly requested. It represents the previous `VocabProperty vocab`, before the triggering change. The representation is the same as that of `"vocab"`.

Furthermore, an NGSI-LD `VocabProperty` in the normalized representation shall never include the following members:

Prohibited

- `"unitCode"`: shall never be present, as vocabs are always strings and hence unitless.
- `"value"` and `"previousValue"`: shall never be present, as *value* is a generalization of *vocab*.

4.5.20.3 Concise NGSI-LD VocabProperty

An NGSI-LD `VocabProperty` shall be represented in concise but lossless representation by a member whose key is the Property name (a term), whose value is the same as the JSON-LD object in **NGSI-LD Property Representation** defined in clause 4.5.2.3, with the following differences.

Mandatory

- `"vocab"`: a JSON object consisting of a single string or array of strings which can be type coerced into an IRI or array of IRIs. It represents a more specialized *value*.

Optional

- `"type"`: If missing, `"VocabProperty"` can be inferred by the presence of the `"vocab"` attribute.

Output Only

- `"previousVocab"`: only provided only in the case of notifications and only if the *showChanges* option is explicitly requested. It represents the previous `VocabProperty vocab`, before the triggering change. The representation is the same as that of `"vocab"`.

Furthermore, an NGSI-LD `VocabProperty` in the concise representation shall never include the following members:

Prohibited

- `"unitCode"`: shall never be present, as vocabs are always strings and hence unitless.
- `"value"` and `"previousValue"`: shall never be present, as it is a generalization of *vocab*.

4.5.21 NGSI-LD ListProperty Representations

4.5.21.1 Introduction

NGSI-LD defines a specialized type of Property named *ListProperty*, defined by the NGSI-LD *@context* described by the present document in clause 4.4.

When dealing with NGSI-LD Entities, implementations shall interpret the JSON-LD nodes of type *ListProperty* as per clause 4.5.21.2 (when in normalized representation) or clause 4.5.21.3 (when in concise representation).

Both normalized and concise representation of *ListProperties* shall be supported by implementations and can be selected by Context Consumers through specific request parameters. An example of this representation can be found in annex C, clause C.2.2.

4.5.21.2 Normalized NGSI-LD ListProperty

An NGSI-LD *ListProperty* shall be represented in normalized representation by a member whose key is the Property name (a term), whose value is the same as the JSON-LD object in **NGSI-LD Property Representation** defined in clause 4.5.2.2, with the following differences:

Mandatory

- "type": the fixed value "*ListProperty*".
- "valueList": a JSON representation ordered array of Property Values (see definition of NGSI-LD Value in clause 3.1). It represents a more specialized *value*.

An array consisting of a single NGSI-LD Null (explained in clause 4.5.0 and defined in clause 3.1) can be used as the right-hand side of the "valueList" during partial update patch and merge patch (see clauses 5.5.8 and 5.5.12) to indicate a deletion of the *ListProperty*, as well as in notifications and in temporal evolutions (for encoding a deleted *ListProperty*).

Output Only

- "previousValueList": only provided only in the case of notifications and only if the *showChanges* option is explicitly requested. It represents the previous *ListProperty valueList*, before the triggering change. The representation is the same as that of "valueList".

Furthermore, an NGSI-LD *ListProperty* in the normalized representation shall never include the following members:

Prohibited

- "value" and "previousValue": shall never be present, as *value* is a generalization of *valueList*.

4.5.21.3 Concise NGSI-LD ListProperty

An NGSI-LD *ListProperty* shall be represented in concise but lossless representation by a member whose key is the *ListProperty* name (a term), whose value is the same as the JSON-LD object in **NGSI-LD Property Representation** defined in clause 4.5.2.3, with the following differences:

Mandatory

- "valueList": a JSON representation of a ordered array of Property Values (see definition of NGSI-LD Value in clause 3.1). It represents a more specialized *value*.

An array consisting of a single NGSI-LD Null (explained in clause 4.5.0 and defined in clause 3.1) can be used as the right-hand side of the "valueList" during partial update patch and merge patch (see clauses 5.5.8 and 5.5.12) to indicate a deletion of the *ListProperty*, as well as in notifications and in temporal evolutions (for encoding a deleted *ListProperty*).

Optional

- "type": If missing, "*ListProperty*" can be inferred by the presence of the "valueList" attribute.

Output Only

- "previousValueList": only provided only in the case of notifications and only if the *showChanges* option is explicitly requested. It represents the previous *ListProperty* "valueList", before the triggering change. The representation is the same as that of "valueList".

Furthermore, an NGSI-LD *ListProperty* in the concise representation shall never include the following members:

Prohibited

- "value" and "previousValue": shall never be present, as *value* is a generalization of *valueList*.

4.5.22 NGSI-LD ListRelationship Representations

4.5.22.1 Introduction

NGSI-LD defines a specialized type of Relationship named *ListRelationship*, defined by the NGSI-LD *@context* described by the present document in clause 4.4.

When dealing with NGSI-LD Entities, implementations shall interpret the JSON-LD nodes of type *ListRelationship* as per clause 4.5.22.2 (when in normalized representation) or clause 4.5.22.3 (when in concise representation).

Both normalized and concise representation of *ListRelationships* shall be supported by implementations and can be selected by Context Consumers through specific request parameters. An example of this representation can be found in annex C, clause C.2.2.

4.5.22.2 Normalized NGSI-LD ListRelationship

An NGSI-LD *ListRelationship* shall be represented in normalized representation by a member whose key is the Relationship name (a term), whose value is the same as the JSON-LD object in **NGSI-LD Relationship Representation** defined in clause 4.5.3.2, with the following differences:

Mandatory

- "objectList": a JSON representation of an ordered array of *Relationship* Objects each consisting of a JSON object containing a single Attribute with a key called "object" and its value holding the *Relationship*'s object represented by a URI.

An array consisting of a single NGSI-LD Null (explained in 4.5.0 and defined in clause 3.1) can be used as the right-hand side of the "objectList" during partial update patch and merge patch (see clauses 5.5.8 and 5.5.12) to indicate a deletion of the *ListRelationship*, as well as in notifications and in temporal evolutions (for encoding a deleted *ListRelationship*).

- "type": the fixed value "*ListRelationship*".

Output Only

- "entityList": only provided in case of **Linked Entity retrieval**, and only if the inline *join* option is explicitly requested (see clause 4.5.23.2), where it is used to define the target **Linked Entities** of a *ListRelationship*'s *objectList* in **normalized representation**.
- "previousObjectList": only provided in the case of notifications and only if the *showChanges* option is explicitly requested. It represents the previous *ListRelationship* *objectList*, before the triggering change. The representation is the same as that of "objectList".

Furthermore, an NGSI-LD *ListRelationship* in the normalized representation shall never include the following members:

Prohibited

- "entity": shall never be present as *entity* is a generalization of *entityList*.

- "object" and "previousObject": shall never be present as *object* is a generalization of *objectList*.

4.5.22.3 Concise NGSI-LD ListRelationship

An NGSI-LD *ListRelationship* shall be represented in concise but lossless representation by a member whose key is the Relationship name (a term), whose value is the same as the JSON-LD object in **NGSI-LD Relationship Representation** defined in clause 4.5.3.3, with the following differences:

Mandatory

- "objectList": this represents a more specialized *object* and is represented by an ordered array of either:
 - a JSON objects containing a single Attribute with a key called "object" and its value holding the *Relationship*'s object represented by a URI.
 - Strings representing URIs only, where expansion to *Relationship objects* can be inferred by the presence of the "objectList" attribute.

An array consisting of a single NGSI-LD Null (explained in 4.5.0 and defined in clause 3.1) can be used as the right-hand side of the "objectList" during partial update patch and merge patch (see clauses 5.5.8 and 5.5.12) to indicate a deletion of the *ListRelationship*, as well as in notifications and in temporal evolutions (for encoding a deleted *ListRelationship*).

Optional

- "type": If missing, "ListRelationship" can be inferred by the presence of the "objectList" attribute.

Output Only

- "entityList": only provided if the inline *join* option is explicitly requested (see clause 4.5.23.2), where it is used to define the target *Linked Entities* of a *ListRelationship*'s "objectList" in **concise representation**.
- "previousObjectList": only provided only in the case of notifications and only if the *showChanges* option is explicitly requested. It represents the previous *ListRelationship objectList*, before the triggering change. *Optional*. The representation is the same as that of "objectList".

Furthermore, an NGSI-LD *ListRelationship* in the concise representation shall never include the following members:

Prohibited

- "entity": shall never be present as *entity* is a generalization of *entityList*.
- "object" and "previousObject": shall never be present as *object* is a generalization of *objectList*.

4.5.23 NGSI-LD Linked Entity Retrieval

4.5.23.1 Introduction

Since Entities are uniquely identifiable by a URI, it is possible to traverse across the Entity graph directly from a *Linking Entity* to a *Linked Entity*. It is therefore sometimes convenient to be able to query or retrieve data via a single Context Broker request and to receive a response including both *Linking Entities* and dependent *Linked Entities* directly.

The concept of Entity graph retrieval is a common concept amongst graph databases and it allows for more structured queries (see clause 4.9) and the complete serialization of an Entity and its dependents.

When retrieving *Linked Entities*, it is necessary to limit retrieval to avoid cascades of an excessive length, duplicates or loops. Only *Relationships* targeting a locally stored Entity or *Relationships* annotated with an *objectType* whose *object* is an *Internal Linked Entity* are considered to be retrievable in this manner.

4.5.23.2 Inline Linked Entity Representation

With the inline representation, the Context Broker response shall only consist of Linking Entities - either a single Linking Entity, or an array consisting of Linking Entities. The additional Entity data from Linked Entities is returned via a Sub-Attribute added to annotated Relationships. This inline representation is generated for output only.

An example of this representation can be found in annex C, clause C.2.2.

4.5.23.3 Flattened Linked Entity Representation

With the flattened representation, the Context Broker response shall always consist of an array of Entities. This array will consist of both Linking Entities and Linked Entities (where the retrieved Linked Entities defined by an annotated Relationship), are appended to the array. This flattened representation allows for batch operations (see clauses 5.6.7, 5.6.8, 5.6.9 and 5.6.10) be applied directly using the response from the Context Broker.

An example of this representation can be found in annex C, clause C.2.2.

4.5.24 NGSI-LD JsonProperty Representations

4.5.24.1 Introduction

NGSI-LD defines a specialized type of Property named JsonProperty, defined by the NGSI-LD @context described by the present document in clause 4.4.

When dealing with NGSI-LD Entities, implementations shall interpret the JSON-LD nodes of type JsonProperty as per clause 4.5.24.2 (when in normalized representation) or clause 4.5.24.3 (when in concise representation).

Both normalized and concise representation of JsonProperties shall be supported by implementations and can be selected by Context Consumers through specific request parameters. An example of this representation can be found in annex C, clause C.2.2.

4.5.24.2 Normalized NGSI-LD JsonProperty

An NGSI-LD JsonProperty shall be represented in normalized representation by a member whose key is the Property name (a term), whose value is the same as the JSON-LD object in **NGSI-LD Property Representation** defined in clause 4.5.2.2, with the following differences:

Mandatory

- "type": the fixed value "JsonProperty".
- "json": a raw JSON object (or array of objects) which contains data which is not available for JSON-LD interpretation. This means that the attributes within the JSON object are **never subject to JSON-LD term expansion or compaction**. It represents a more specialized value.

Output Only

- "previousJson": only provided in case of notifications and only if the showChanges option is explicitly requested. It represents the previous JsonProperty json, before the triggering change. The representation is the same as that of "json".

Furthermore, an NGSI-LD JsonProperty in the normalized representation shall never include the following members:

Prohibited

- "unitCode": shall never be present, as raw JSON objects are unitless.
- "value" and "previousValue": shall never be present, as value is a generalization of json.

4.5.24.3 Concise NGSI-LD JsonProperty

An NGSI-LD `JsonProperty` shall be represented in concise but lossless representation by a member whose key is the Property name (a term), whose value is the same as the JSON-LD object in **NGSI-LD Property Representation** defined in clause 4.5.2.3, with the following differences:

Mandatory

- `"json"`: a raw JSON object which contains data which is not available for JSON-LD interpretation. This means that the attributes within the JSON object are **never subject to JSON-LD term expansion or compaction**. It represents a more specialized *value*.

Optional

- `"type"`: If missing, `"JsonProperty"` can be inferred by the presence of the `"json"` attribute.

Output Only

- `"previousJson"`: only provided in case of notifications and only if the *showChanges* option is explicitly requested. It represents the previous `JsonProperty json`, before the triggering change. The representation is the same as that of `"json"`.

Furthermore, an NGSI-LD `JsonProperty` in the concise representation shall never include the following members:

Prohibited

- `"unitCode"`: shall never be present, as raw JSON objects are unitless.
- `"value"` and `"previousValue"`: shall never be present, as it is a generalization of *json*.

4.5.25 NGSI-LD EntityMap Representation

The `EntityMap` representation is used by `Context Brokers` to ensure unity when querying across distributed operations. It is an active mapping of `Context Source Registrations` to Entities which are relevant to an ongoing Context Information Consumption request (see clause 4.3.6.7). The `EntityMap` representation shall be a JSON-LD object containing the following members:

Mandatory

- `"id"`: whose value shall be the URI that identifies the attribute.
- `"type"`: the fixed value `"EntityMap"`.
- `"expiredBy"`: a `DateTime` string (as defined in clause 4.6.3) for encoding a timestamp indicating the time at which the `EntityMap` can no longer be used.

Optional

- `"@context"`: a JSON-LD *@context* as described in clause 4.4.

Output Only

- `"entityMap"`: a JSON-LD index map holding a unique list of entity ids (URIs) each of which lists the registrations that were fired by the previous request and successfully returned data.
- `"linkedMaps"`: a JSON-LD index map holding a unique list of `Context Source Registrations` ids (URIs) each of which lists the *entityMap* used by a `Context Source`.

4.6 Data Representation Restrictions

4.6.1 Supported text encodings

NGSI-LD implementations shall support the **UTF-8** text encoding format. To avoid interoperability problems, applications shall provide JSON content encoded using UTF-8 and NGSI-LD systems shall also expose such JSON content using UTF-8.

4.6.2 Supported names

Even though the JSON serialization format allows inclusion of any character in the Unicode space, NGSI-LD restricts Entity Type names, Property names and Relationship names to the following ABNF grammar:

```
nameChar = unicodeNumber / unicodeLetter
nameChar =/ %x5F ; _
name = unicodeLetter *nameChar
```

- *unicodeNumber* is any Unicode character that has *Number* as a Category [22]. With Unicode-capable regular expression (RegEx) parsers, such a character may be matched by `\p{N}`.
- *unicodeLetter* is any Unicode character that has *Letter* as a Category [22]. With Unicode-capable regular expression (RegEx) parsers, such a character may be matched by `\p{L}`.

In order to avoid name clashing, names can be prefixed as specified by the following BNF grammar:

```
prefix = unicodeLetter *nameChar
name =/ prefix %x3A unicodeLetter *nameChar ; prefix:name
```

When receiving a JSON-LD object with a name (Type, Property, Relationship) including characters different than those expressed above, implementations should raise an error of type *BadRequestData*.

4.6.3 Supported data types for Values

Compliant NGSI-LD implementations shall support the following data types for representing Values:

- All the JSON native data types as mandated by IETF RFC 8259 [6], section 3.
- All the GeoJSON *Geometries* [8] with the exception of *GeometryCollection*.
- **DateTime** string for encoding a timestamp, i.e. a calendar date together with a time of day, expressed in **UTC**, using the ISO 8601 [17] Complete Representation and in particular using the 'Extended Format', as described below:
 - The timestamp shall be a string containing *Year, Month, Day, Hours, Minutes, Seconds and time zone* components using the format `YYYY-MM-DDThh:mm:ssZ` as defined in ISO 8601 [17]. In this representation, the character "-" is used to separate the calendar date components, the character "T" is used to indicate the start of the time-of-day portion, the character ":" is used to separate the time-of-day components, and the trailing character "Z" is used to convey the time zone.
 - All the referred components shall appear in the string; reduced representations are not permitted.
 - The Seconds component may optionally contain a decimal fraction. In this case the string shall contain two integer digits, followed by a decimal point and then one or more fractional digits, up to a maximum of six. For example, `YYYY-MM-DDThh:mm:ss.ssssssZ`. In requests, also a comma instead of a decimal point may be used as separator for compatibility reasons.

NOTE 1: In previous versions of NGSI-LD, only the comma was supported as ISO 8601 [17] states that it is the preferred option. However, in practice the decimal point is more commonly used.

- The trailing timestamp component shall contain the time zone related information and shall always be equal to the character "Z". Therefore, all timestamps shall be expressed in **UTC**.

- **Date** string for encoding a calendar date. It uses ISO 8601 [17] Complete Representation using the 'Extended Format', as described below:
 - It shall be a string containing *Year*, *Month*, *Day* components using the format *YYYY-MM-DD* as defined in ISO 8601 [17]. In this representation, the character "-" is used to separate the calendar date components.
 - All the referred components shall appear in the string; reduced representations are not permitted.
 - **Time** string for encoding a local time expressed in **UTC**. It uses ISO 8601 [17] Complete Representation using the 'Extended Format', as described below:
 - It shall be a string containing *Hours*, *Minutes* and *Seconds* components using the format *hh:mm:ssZ* as defined in ISO 8601 [17]. In this representation, the character ":" is used to separate the local time components.
 - All the referred components shall appear in the string; reduced representations are not permitted.
 - The *Seconds* component may optionally contain a decimal fraction. In this case the string shall contain two integer digits, followed by a decimal point and then one or more fractional digits, up to a maximum of six. For example, *hh:mm:ss.ssssssz*. In requests, also a comma instead of a decimal point may be used as separator for compatibility reasons.
- NOTE 2: In previous versions of NGSI-LD, only the comma was supported as ISO 8601 [17] states that it is the preferred option. However, in practice the decimal point is more commonly used.
- The string shall not contain expressions of the difference between local time and UTC. All representations shall be interpreted as being expressed in **UTC**.
 - URI as mandated by ISO 8601 [17], Appendix A, production rule named 'URI'.

Implementations may support additional data types different to those enumerated above, for instance:

- JSON-LD typed value (i.e. a string as the lexical form of the value together with a type, defined by an XSD base type or more generally an IRI).
- JSON-LD structured value (e.g. a set, a list).

4.6.4 Supported Content

In principle, context information providers can publish any kind of data serialized in JSON and encoded in UTF-8. Nonetheless, to avoid security problems caused by script injection attacks or other attack vectors, implementations should consider that the incoming data from a client may contain the following characters:

- %x3C; <
- %x3E; >
- %x22; "
- %x27; '
- %x3D; =
- %x3B; ;
- %x28; (
- %x29;)

When receiving entities (context information) encoded in JSON format and containing values that include the above characters, implementations should decide how to resolve the possible security problems that may be generated by the data. In all cases, implementations shall preserve the representation of the content of the values provided by the context information providers and return the original content when replying to **context consumption requests**.

If implementations decide to raise an error, the error shall be *BadRequestData*.

4.6.5 Supported data types for LanguageMaps

Compliant NGSI-LD implementations shall support the following data types for representing LanguageMaps:

- A JSON object consisting of a series of key-value pairs where the keys shall be JSON strings representing IETF RFC 5646 [28] language codes or the JSON-LD "@none" for representing default when no more specific language is found, and the values shall be JSON strings or arrays of JSON strings. Additionally, the *languageMap* encoding {"@none": "urn:ngsi-ld:null"} shall be used to represent an NGSI-LD Null during partial update patch and merge patch (see clauses 5.5.8 and 5.5.12) and for representing deleted Language Properties in notifications and in temporal evolutions.

4.6.6 Ordering of Entities in arrays having more than one instance of the same Entity

Some services (batch operations, clauses 5.6.7, 5.6.8, 5.6.9 and 5.6.10) operate on an array of entities, as input, and if this array contains more than one instance of the same entity, then these entity instances shall come in chronological order, i.e. the first entity instance in the array shall be older than the second, the second shall be older than the third, etc.

Without this assumption, there is no way for the request to be treated correctly, as the entity instances are often used for replacing or modifying the prior entity instance.

4.7 Geospatial Properties

4.7.1 GeoJSON Geometries

Geospatial Properties in NGSI-LD shall be represented using **GeoJSON** Geometries [8]. With the aim of highlighting and encoding those Properties which convey geospatial characteristics, NGSI-LD defines a special type of Property named *GeoProperty*, defined by the Core NGSI-LD *@context* described by the present document in clause 4.4.

When dealing with NGSI-LD Entities, implementations shall interpret JSON-LD nodes of type *GeoProperty* just as conventional Properties but with the additional requirement that the Value corresponding to such Property shall be a GeoJSON Geometry. All the Geometries defined by [8] are allowed except *GeometryCollection*. In addition, implementations should take the necessary steps to create the corresponding geo-indexes so that information can be properly returned when geo-queries are executed.

NGSI-LD defines the following Properties of type *GeoProperty*. Preferably these Properties should be used if they semantically fit, but if necessary, additional Properties of type *GeoProperty* can be defined by *Context Producers*:

- **location** is defined as the geospatial Property representing the geographic location of the Entity, e.g. the location of a building or the current location of a car.
- **observationSpace** is defined as the geospatial Property representing the geographic location that is being observed, e.g. by a sensor. For example, in the case of a camera, the location of the camera and the observation space are different and can be disjoint.
- **operationSpace** is defined as the geospatial Property representing the geographic location in which an Entity, e.g. an actuator is active. For example, a crane can have a certain operation space.

The defined Properties can also be used as part of *Context Source Registrations* (see clause 5.2.9). In this case they represent locations in which Entities with the respective geospatial Properties are contained. For example, a *Context Source* that monitors the location of cars in a city may be represented by a *Context Source Registration* whose Property *location* corresponds to the space of the city in which the location of cars is monitored.

4.7.2 Representation of GeoJSON Geometries in JSON-LD

There are certain types of GeoJSON geometries, for instance GeoJSON Polygon, whose coordinates are represented using nested array structures (through the *coordinates* member). Such representation may introduce serialization problems when transforming JSON-LD content into RDF graphs.

Also, when using whole GeoJSON geometries (consisting of *type* and *coordinates*) in an NGSI-LD document, its JSON syntax is only preserved in the regular JSON-LD representation (with separate *@context*), but not in an expanded representation. To handle resulting problems, optionally, whole GeoJSON geometries can be represented as a JSON string.

Implementations shall accept the referred encoded string value, if and only if, it can be parsed into a JSON Object, as mandated by IETF RFC 8259 [6], meeting the syntax and restrictions mandated by IETF RFC 7946 [8] when representing a valid Geometry of the type specified.

For the avoidance of doubt, regular encodings of GeoJSON geometries (as JSON Object) shall also be accepted by implementations, but Context Producers should consider the implications in terms of RDF compatibility.

GeoJSON coordinates shall be considered as consisting of values of a JSON-LD floating point **number** data type. The degree of precision of a latitude and longitude (and optional altitude) held within a context broker will depend upon the implementation. Implementors should note that the GeoJSON specification itself recommends 6 decimal places for latitude and longitude which equates to roughly 10cm of precision.

4.7.3 Concise NGSI-LD GeoProperty

Notwithstanding the restrictions defined in clause 4.5.2.3, an NGSI-LD GeoProperty without additional sub-attributes shall be represented in a concise but lossless representation by a member whose key is the Property name (a term) and whose value is the Property Value (see definition of terms in clause 3.1) which itself is also a supported GeoJSON geometry.

Mandatory

- "type": shall be a supported GeoJSON geometry type as defined in clause 4.7.1.
- "coordinates": shall be present, as defined by the relevant GeoJSON Geometry [8].

When parsing a geospatial value submitted in the concise representation, it shall be possible for the NGSI-LD system to infer the *GeoProperty* type. Error handling of the payload is left ambiguous if the NGSI-LD system is unable to distinguish a payload as either a Property or a *GeoProperty*.

Furthermore, an NGSI-LD GeoProperty which includes additional Properties or Relationships shall be treated in the same manner as an ordinary NGSI-LD Property (see clause 4.5.2.3) with the exception that if the Property *value* resolves to a supported GeoJSON geometry, the type "GeoProperty" shall be inferred.

4.8 Temporal Properties

NGSI-LD defines the following Properties of type *TemporalProperty* that shall be supported by implementations:

- **observedAt** is defined as the temporal Property at which a certain Property or Relationship became valid or was observed. For example, a temperature Value was measured by the sensor at this point in time.
- **createdAt** is defined as the temporal Property at which the Entity, Property or Relationship was entered into an NGSI-LD system.
- **modifiedAt** is defined as the temporal Property at which the Entity, Property or Relationship was last modified in an NGSI-LD system, e.g. in order to correct a previously entered incorrect value.
- **deletedAt** is defined as the temporal Property at which the Entity, Property or Relationship was deleted from an NGSI-LD system.
- **expiresAt** is defined as the temporal Property at which the Entity, Property, Relationship, CSourceRegistration, Subscription or Snapshot should be deleted from an NGSI-LD system.

- **lastUsedAt** is defined as the temporal Property at which a Snapshot has been most recently used, i.e. when the most recent operation has been executed on this Snapshot.

Temporal Properties in NGSI-LD shall be represented based on the *DateTime* data type as mandated by clause 4.6.3.

NOTE 1: For simplicity reasons, a *TemporalProperty* is represented only by its Value, i.e. no *Properties* of *TemporalProperty* nor *Relationships* of *TemporalProperty* can be conveyed. In more formal language, a *TemporalProperty* does not allow reification.

NOTE 2: It is important to remark that the term *TemporalProperty* has been reserved for the semantic tagging of non-reified structural timestamps (*observedAt*, *createdAt*, *modifiedAt*, *deletedAt*, *expiresAt*), which capture the temporal evolution of Attributes. **Only such structural timestamps can be used as timeproperty in Temporal Queries as mandated by clause 4.11.**

NOTE 3: User-defined Properties whose value is a time value (*Date*, *DateTime* or *Time*) are defined as *Property*, not as *TemporalProperty*, and are serialized in NGSI-LD as shown in annex C, clause C.6.

Whenever a *TemporalProperty* value is unknown by a registered Context Source, the Property shall be omitted rather than sending an error response.

4.9 NGSI-LD Query Language

The NGSI-LD Query Language shall be supported by implementations. It is intended to:

- filter out Entities by Attribute Values (target is the *value* member of a *Property*, see Table 5.2.5-1, or the *object* member of a *Relationship*, see Table 5.2.6-1);
- filter out Context Sources by the values of properties that describe them, defined when Context Sources are registered (target is the name of a Context Source Property member of the CSourceRegistration, see Table 5.2.9-1).
- filter out Snapshots by the values of the Snapshot data type members, i.e. when filtering Snapshots, only the names of the members defined for Snapshot in Table 5.2.43-1 are allowed values for AttrName.

In this clause, three string parameters are defined in order to fully specify an NGSI-LD Query:

- **q**, to express the desired query;
- **expandValues**, to define the list of attributes whose values should be expanded against the supplied *@context* using JSON-LD type coercion prior to executing the query in the Context Broker.

Optional

- **jsonKeys**, to define the list of attributes whose values uninterpretable as JSON-LD and should not be expanded against the supplied *@context* using JSON-LD type coercion prior to executing the query in the Context Broker. *Optional*

In case of HTTP binding, whenever the string acting as a filter is part of the HTTP binding's URI, then it shall be URI-encoded (percent-encoded, as described in IETF RFC 3986 [5]).

The grammar that encodes the syntax of the **q** parameter, expressed in ABNF format [12], is the NGSI-LD Query Language. It is described below (it has been validated using <https://github.com/ietf-tools/bap>), and it shall be supported by implementations:

```
Query = (QueryTerm / QueryTermAssoc) *(LogicalOp (QueryTerm / QueryTermAssoc))
QueryTermAssoc = %x28 QueryTerm *(LogicalOp QueryTerm) %x29 ; (QueryTerm)
QueryTerm = Attribute
QueryTerm =/ Attribute Operator ComparableValue
QueryTerm =/ Attribute equal CompEqualityValue
QueryTerm =/ Attribute unequal CompEqualityValue
QueryTerm =/ Attribute patternOp RegExp
QueryTerm =/ Attribute notPatternOp RegExp
Attribute = LinkedEntityRelation
LinkedEntityRelation = AttrName %x7B LinkedEntityPath %x7D ; AttrName{LinkedEntityPath}
LinkedEntityRelation =/ ValuePath
LinkedEntityPath = *1(EntityType 1*(%x2C EntityType) %x3A) AttrName %x7B LinkedEntityPath %x7D
```

```

; *1(EntityType 1*(,EntityType):)AttrName{LinkedEntityPath}
LinkedEntityPath = / ValuePath
ValuePath = DottedPath *1(%x5B DottedPath %x5D) ; DottedPath *1([DottedPath])
DottedPath = AttrName *(%x2E AttrName) ; AttrName *(.AttrName)
Operator = equal / unequal / greaterEq / greater / lessEq / less
ComparableValue = Number / quotedStr / dateTime / date / time
OtherValue = false / true
Value = ComparableValue / OtherValue
Range = ComparableValue dots ComparableValue
ValueList = Value 1*(%x2C Value) ; Value 1*(, Value)
CompEqualityValue = OtherValue / ValueList / Range / URI
equal = %x3D %x3D ; ==
unequal = %x21 %x3D ; !=
greater = %x3E ; >
greaterEq = %x3E %x3D ; >=
less = %x3C ; <
lessEq = %x3C %x3D ; <=
patternOp = %x7E %x3D ; ~=
notPatternOp = %x21 %x7E %x3D ; !~=
dots = %x2E %x2E ; ..
TermChar = unicodeNumber / unicodeLetter
TermChar = / %x5F ; _
AttrName = unicodeLetter *TermChar
EntityType = unicodeLetter *TermChar
quotedStr = String ; "char"
andOp = %x3B ; &
orOp = %x7C ; |
LogicalOp = andOp / orOp

```

- `unicodeNumber` is any Unicode character that has *Number* as a Category [22]. With Unicode-capable regular expression (RegEx) parsers, such a character may be matched by `\p{N}`.
- `unicodeLetter` is any Unicode character that has *Letter* as a Category [22]. With Unicode-capable regular expression (RegEx) parsers, such a character may be matched by `\p{L}`.
- `Number` shall be a number as mandated by the JSON Specification, following the ABNF Grammar, production rule named `number`, section 6 of IETF RFC 8259 [6].
- `String` shall be a text string as mandated by the JSON Specification, following the ABNF Grammar, production rule named `String`, section 7 of IETF RFC 8259 [6].
- `char` shall be a character as mandated by the JSON Specification, ABNF Grammar, production rule named `char`, section 7 of IETF RFC 8259 [6].
- `false` shall be conformant with the JSON ABNF Grammar, production rule named `false`, section 3 of IETF RFC 8259 [6]. It is intended to represent the Boolean value corresponding to *false*.
- `true` shall be conformant with the JSON ABNF Grammar, production rule named `true`, section 3 of IETF RFC 8259 [6]. It is intended to represent the Boolean value corresponding to *true*.
- `RegExp` shall be a regular expression as mandated by IEEE 1003.2™ [11].
- `dateTime` shall be a *DateTime* value as mandated by clause 4.6.3.
- `time` shall be a *Time* value as mandated by clause 4.6.3.
- `date` shall be a *Date* value as mandated by clause 4.6.3.
- `URI` shall be a URI as mandated by IETF RFC 3986 [5] or an IRI as mandated by IETF RFC 3987 [23], appendix A, production rule named *URI*.

A **Query Term** (production rule `QueryTerm`) defines a predicate which serves as a matching condition for Entities. The constituent parts of a Query Term are:

- an attribute path (production rule named `Attribute`).
- an optional pair composed by an operator (production rule named `Operator`) and a value (production rule named `Value`).

The attribute path (production rule *Attribute*) is a simple name *AttrName*, optionally followed by a dot-separated list of more *AttrName* (see later example 8), optionally followed by one trailing list of more dot-separated *AttrNames* enclosed in one pair of square brackets (see later example 9). The attribute path is always a composition of short hand names and not a fully qualified ones, because, when the query language is used, an *@context* properly defining all the terms (as per clause 5.5.7) shall be issued.

EXAMPLE 0: `?q=temperature` (checks for the existence of the attribute temperature).

EXAMPLE 1: `?q=temperature==20`.

EXAMPLE 2: `?q=brandName!="Mercedes"`.

EXAMPLE 3: `?q=isParked=="urn:ngsi-ld:OffStreetParking:Downtown1"`.

EXAMPLE 4: A query encoded as an HTTP Query String. Note that this is HTTP binding specific, to be used via GET method, as defined in clause 6.4.3.2. The NGSI-LD query language string is conveyed by means of parameter **q**.

`?q=speed>50;brandName!="Mercedes"`. Also note that (as stated above) URI-encoding (percent-encoding) is required if the query string contains reserved characters (see IETF RFC 3986 [5] and IETF RFC 3987 [23], for the exact list of them).

EXAMPLE 5: `?q=isMonitoredBy` (to query Entities that have the Attribute *isMonitoredBy*).

Query Terms may be combined through logical operators that shall be supported by implementations as follows:

- The production rule *andOp* defines a logical AND operator conveying that the requested entities are those which meet at the same time the conditions posed by all the Query Terms affected by such an operator.
- The production rule *orOp* defines a logical OR operator conveying that the requested entities are those which meet any of the conditions posed by the Query Terms affected by such an operator.
- When evaluating logical conditions, and in the absence of specific Query Term associations (see below), the logical AND operator shall take precedence over the logical OR operator.

Association of Query Terms shall be supported by implementations as per the grammar included by the present clause (production rule named *QueryTermAssoc*). An association of Query Terms is composed of the combination of different Query Terms linked by logical operators (AND, OR) and delimited by parenthesis. The evaluation of an association of Query Terms shall always take precedence over individual, non-associated Query Terms.

EXAMPLE 6: `?q=((speed>50|rpm>3000);brandName=="Mercedes")`

EXAMPLE 7: `?q=(temperature>=20;temperature<=25)|capacity<=10`

The following example 8 shows the syntax of an attribute path that is defined by the production rule *Attribute*, as a dot-separated list of names. Such a list is intended to address a *Property* or *Relationship* included by the matching entities subjacent graph, in accordance with the following rules:

- Every name in the list shall be expanded to a URI (Fully Qualified Name) as mandated by clause 5.5.7.
- The first name shall refer to a *Property* or *Relationship* (top level element) whose subject shall be a matching Entity. Strictly speaking, and as per the JSON-LD representation rules, such (fully qualified) name shall be equal to the (fully qualified) name of the concerned *Property* or *Relationship*.
- Each other name (if present) represents a (sub)Property or (sub)Relationship, starting with the top-level element as subject and continuing through the graph traversal. The element addressed by the last name in the list is defined as the target element. If only one name is present in the attribute path, then the target element is the top level one.

EXAMPLE 8: `?q=temperature.observedAt>=2017-12-24T12:00:00Z`

If the target element is a *Property*, the **target value** is defined as the Value associated to such Property. If a Property has multiple instances (identified by its respective *datasetId*), and no *datasetId* is explicitly addressed, the target value shall be any Value of such instances.

If the target element is a *LanguageProperty*, and no target language is specified, the **target value** is defined as a value from any of the key-value pairs held within the *languageMap* associated to such *LanguageProperty*.

If the target element is a *ListProperty*, the **target value** is defined as the *valueList* array associated to such a *ListProperty*.

If the target element is a *LanguageProperty* and a target language is specified, the **target value** is defined as the Value associated to the matching key-value pair held within the *languageMap* associated to such *LanguageProperty* where the key matches the target language.

If the target element is a *VocabProperty*, the **target value** shall be expanded according to the *@context*.

If the target element is a *Relationship*, the **target object** is defined as the *object* associated (represented as a URI or array of URIs) to such *Relationship*. If a *Relationship* has multiple instances (identified by its respective *datasetId*), and no *datasetId* is explicitly addressed, the target object shall be any object of such instances.

If the target element is a *ListRelationship*, the **target object** is defined as the array of *objects* associated (represented as URIs) to such *ListRelationship*.

When a Query Term only defines an attribute path (production rule named *Attribute*), the matching Entities shall be those which define the target element (*Property* or a *Relationship*), regardless of any target value or object.

Lastly, implementations shall support queries involving specific data subitems belonging to a Property Value (**seed target value**) represented by a JSON object structure (compound value). For that purpose, an attribute path may additionally contain a **trailing path** (enclosed in a single pair of square brackets that signal that the overall path is now entering the compound value) composed of a dot-concatenated list of JSON member names, and intended to address a specific data subitem (member) within the **seed target value**. When such a trailing path is present, implementations shall interpret and evaluate it (against the seed target value) as a *MemberExpression* of ECMA 262 [21], in dot notation, as clarified therein at section Property Accessors). If the evaluation of such *MemberExpression* does not result in a defined value, the target element shall be considered as non-existent for the purpose of query resolution.

EXAMPLE 9: `?q=address[city]=="Berlin"`. The trailing path is `[city]`. It is used to refer to a particular subitem within the value of the *address Property*, which is a complex JSON object representing a postal address. Refer to the following NGSI-LD Entity:

```
{
  "id": "urn:ngsi-ld:placedescription:123",
  "type": "PlaceDescription",
  "address": {
    "type": "Property",
    "value": {
      "city": "Berlin",
      "street": "Ulrich Strasse"
    }
  }
}
```

EXAMPLE 10: `?q=sensor.rawdata[airquality.particulate]==40`. The trailing path is `[airquality.particulate]`. The *particulate Property* of the compound JSON object is targeted. Refer to the following NGSI-LD Entity:

```
{
  "id": "urn:ngsi-ld:particulatemeasurement:345",
  "type": "ParticulateMeasurement",
  "sensor": {
    "type": "Property",
    "value": 40,
    "rawdata": {
      "type": "Property",
      "value": {
        "airquality": {
          "particulate": 40,
          "PM20": 85
        }
      }
    }
  }
}
```

```

    }
  }
}

```

EXAMPLE 11: `?q=parkingTickets[value]=="Overstay 60 minutes"&jsonKeys=parkingTickets`. The trailing path is `parkingTickets`. The `parkingTickets` *Property* of the JSON object is targeted, but the target value raw is JSON, and is not expanded to `ngsi-ld:hasValue` using the core `@context`. Refer to the following NGSI-LD Entity:

```

{
  "id": "urn:ngsi-ld:Car:6152s",
  "type": "Car",
  "parkingTickets": {
    "type": "JsonProperty",
    "json": {
      "id": "85a6cc52-0589-45f9",
      "value": "Overstay 60 minutes"
    }
  }
}

```

EXAMPLE 12: `?q=gender==Male&expandValues=gender`. The trailing path is `gender`. The `gender` *Property* of JSON object is targeted, but the target value is first expanded to a URI using the supplied `@context`. Refer to the following NGSI-LD Entity:

```

{
  "id": "urn:ngsi-ld:Person:678",
  "type": "Person",
  "gender": {
    "type": "VocabProperty",
    "vocab": "Male",
  }
},
"@context": {
  "Male": "http://example.org/Male",
}
}

```

The filter can also apply to a *Property* or *Relationship* of an NGSI-LD Entity targeted by a (recursively) followed *Relationship*, for example as part of a linked entity retrieval (clause 4.5.23).

EXAMPLE 13: `?q=sensor{humidity}==40`. The trailing path is `sensor{humidity}`. The query targets entities with a *sensor Relationship* and makes a sub-query on matching target objects which have the matching *humidity* Attribute. Refer to the following NGSI-LD Entities:

```

{
  "id": "urn:ngsi-ld:WeatherStation:123",
  "type": "WeatherStation",
  "sensor": {
    "type": "Relationship",
    "objectType": "Device",
    "object": "urn:ngsi-ld:Device:345"
  }
}

{
  "id": "urn:ngsi-ld:Device:345",
  "type": "Device",
  "humidity": {
    "type": "Property",

```



```

    "value": 40
  }
}

```

As not knowing the Entity Type targeted by a *Relationship* could make the query significantly more expensive, a hint for the required Entity Type can be provided, so only such NGSI-LD Entities need to be considered.

EXAMPLE 14: `?q=sensor{Device:humidity}==40`. The trailing path is `sensor{Device:humidity}`. The query targets entities with a `sensor.entityType = "Device"` within a *Relationship* and then makes a sub-query on matching target *objects* which have the matching *humidity* Attribute. The *entityType* hint results in a faster lookup. Refer to the following NGSI-LD Entities.

```

{
  "id": "urn:ngsi-ld:WeatherStation:123",
  "type": "WeatherStation",
  "sensor": {
    "type": "Relationship",
    "objectType": "Device",
    "object": "urn:ngsi-ld:Device:345"
  }
}

{
  "id": "urn:ngsi-ld:Device:345",
  "type": "Device",
  "humidity": {
    "type": "Property",
    "value": 40
  }
}

```

If the target element corresponds to a *Relationship* or *ListRelationship*, the combination of such target element with any operator different than *equal* or *unequal* shall result in **not matching**.

A **Query Term value** shall be any of the following (depending on the operator used):

- A literal value (string, number, date, etc.) (production rule named `value`).
- A range of values (production rule named `Range`), specified as a minimum and a maximum value.
- A regular expression (production rule named `RegExp`).
- A URI (production rule named `URI`).
- A comma-separated list of literal values (production rule named `valueList`).

When comparing dates or times, the order relation considered shall be a temporal one.

When it comes to comparing text strings, implementations:

- shall follow the recommendations defined by IETF RFC 8259 [6], section 8.3.
- should support the Unicode Collation Algorithm (UCA), as defined by [13].

URI comparison should be performed so that the number of false negatives is minimized, as recommended by IETF RFC 3986 [5], section 6.

The semantics of the different logical operators used by Query Terms are described as follows and shall be supported by compliant implementations:

- **Existence** (only attribute is specified). A matching entity shall contain the target element.

- **Equal** operator (production rule named `equal`). A matching Entity shall contain the target element and meet any of the following conditions:
 - The Query Term value, e.g. `color=="red"`:
 - Is identical or equivalent to the target value (e.g. matches `"red"`).
 - Is included in the target value, if the latter is an array (e.g. matches `["blue", "red", "green"]`).
 - If the Query Term value is a list of values (production rule named `valueList`), e.g. `color=="black", "red"`:
 - The target value is identical or equivalent to any of the list values (e.g. matches `"red"`).
 - The target value includes any of the Query Term values, if the target value is an array (e.g. matches `["red", "blue"]`).
 - If the Query Term value is a range (production rule named `Range`), e.g. `temperature==10..20`:
 - The target value is in the interval between the minimum and maximum of the range (both included) (e.g. matches `15`).
 - The Query Term value target element corresponds to a *LanguageProperty* and a natural language is specified e.g. `color[en]=="red"`:
 - a match is found as the value of the key-value pair corresponding to the specified natural language of the *languageMap* (e.g. matches `{"fr": "rouge", "en": "red", "de": "rot"}`) but not `{"fr": "red", "en": "black", "de": "blue"}`).
 - a match is found as a single element from the array of values of the key-value pair corresponding to the specified natural language of the *languageMap* (e.g. matches `{"fr": ["chat", "rouge"], "en": ["red", "cat"], "de": ["rote", "Katze"]}`) but not `{"fr": ["chat", "rouge"], "en": ["coal", "black"], "de": ["blaue", "Engel"]}`).
 - The Query Term value target element corresponds to a *LanguageProperty* and no natural language is specified e.g. `color[*]=="red"`:
 - any match is found in the values of the key-value pairs of the *languageMap* (e.g. matches `{"fr": "rouge", "en": "red", "de": "rote"}`).
 - a match is found as a single element of the array of values of the key-value pairs of the *languageMap* (e.g. matches `{"fr": "chat", "rouge"], "en": ["red", "cat"], "de": ["rote", "Katze"]}`).
 - The Query Term value is a URI and the target element corresponds to a *Relationship*, a *ListRelationship* or a *VocabProperty*, e.g. `color=="http://example/red"`:
 - Is identical to the target value (e.g. matches `"http://example.com/red"`).
 - Is included in the target value, if the latter is an array (e.g. matches `["http://example.com/blue", "http://example.com/red", "http://example.com/green"]`).
 - If the Query Term value target element corresponds to a *Relationship*, a *ListRelationship* or a *VocabProperty* and is a list of URIs (production rule named `ValueList`), e.g. `color=="http://example/black", "http://example/red"`:
 - The target value is identical to any of the list values (e.g. matches `"http://example.com/red"`).
 - The target value includes any of the Query Term values, if the target value is an array (e.g. matches `["http://example.com/red", "http://example.com/blue"]`).

- If there is no equality between the target value data type and the Query Term value data type, then it shall be considered as not matching.
- **Unequal** operator (production rule named `unequal`). A matching entity shall contain the target element and meet any of the following conditions:
 - The Query Term value, e.g. `color!="red"`:
 - Is neither identical nor equivalent to the target value (e.g. matches "black").
 - Is not included in the target value, if the latter is an array (e.g. matches ["blue", "black", "green"], but not ["blue", "red", "green"]).
 - If the Query Term value is a list of values (production rule named `valueList`), e.g. `color!="black", "red"`:
 - The target value is neither identical nor equivalent to any of the list values (e.g. matches "blue").
 - The target value does not include any of the list values, if the target value is an array (e.g. matches ["blue", "yellow", "green"], but not ["blue", "red", "green"]).
 - If the Query Term value is a range (production rule named `Range`), e.g. `temperature!=10..20`:
 - The target value is not in the interval between the minimum and the maximum (both included) (e.g. matches 9).
 - The Query Term value target element corresponds to a *LanguageProperty* and a natural language is specified e.g. `color[en]!="red"`:
 - No matching value is found as the value of the specified language key of a *languageMap* where a language filter is specified. (e.g. matches {"fr": "noir", "en": "black", "de": "schwarz"} but not {"fr": "rouge", "en": "red", "de": "rot"}).
 - No matching value is found as a single element from the array of values of the key-value pair corresponding to the specified natural language of the *languageMap* (e.g. matches {"fr": ["chat", "rouge"], "en": ["coal", "black"], "de": ["blaue", "Engel"]}) but not {"fr": ["rouge", "noir"], "en": ["red", "black"], "de": ["rot", "schwarz"]}).
 - The Query Term value target element corresponds to a *LanguageProperty* and no language filter is specified e.g. `color[*]!="red"`:
 - No matching value is found in any of the values of the key-value pairs of a *languageMap* (e.g. matches {"fr": "noir", "en": "black", "de": "schwarz"}, but not {"fr": "rouge", "en": "red", "de": "rot"}).
 - No matching value is found as a single element from the array of values of the key-value pair corresponding to the specified natural language of the *languageMap* (e.g. matches {"fr": ["chat", "rouge"], "en": ["coal", "black"], "de": ["blaue", "Engel"]}) but not {"fr": ["rouge", "noir"], "en": ["red", "black"], "de": ["rot", "schwarz"]}).
 - The Query Term value is a URI and the target element corresponds to a *Relationship*, a *ListRelationship* or a *VocabProperty*, e.g. `color!="http://example.com/red"`:
 - Is not identical to the target value (e.g. matches "http://example.com/black").
 - Is not included in the target value, if the latter is an array (e.g. matches ["http://example.com/blue", "http://example.com/black", "http://example.com/green"], but not ["http://example.com/blue", "http://example.com/red", "http://example.com/green"]).

- If the Query Term value target element corresponds to a *Relationship*, a *ListRelationship* or a *VocabProperty* and is a list of URIs e.g. `color!="http://example.com/black", "http://example.com/red"`:
 - The target value is not identical to any of the list values (e.g. matches `"http://example.com/blue"`).
 - The target value does not include any of the list values, if the target value is an array (e.g. matches `["http://example.com/blue", "http://example.com/yellow", "http://example.com/green"]`, but not `["http://example.com/blue", "http://example.com/red", "http://example.com/green"]`).
- If the data type of the target value and the data type of the Query Term value are different, then they shall be considered unequal.
- **Greater than** operator (production rule named `greater`). For an entity to match, it shall contain the target element and the target value has to be strictly greater than the Query Term value:
 - If there is no equality between the target value data type and the Query Term value data type then it shall be considered as not matching.
- **Less than** operator (production rule named `less`). For an entity to match, it shall contain the target element and the target value shall be strictly less than the value:
 - If there is no equality between the target value data type and the Query Term value data type then it shall be considered as not matching.
- **Greater or equal than** (production rule named `greaterEq`). A matching entity shall meet any of the *Greater than* or the *Equal* conditions for single values.
- **Less or equal than** (production rule named `lessEq`). A matching entity shall meet any of the *Less than* or the *Equal* conditions for single values.
- **Match pattern** (production rule named `patternOp`). A matching entity shall contain the target element and the target value shall be in the L(R) of the regular pattern specified by the Query Term:
 - If the target value data type is different than *String* then it shall be considered as not matching.
- **Do not match pattern** (production rule named `notPatternOp`). A matching entity shall contain the target element and the target value shall not be in the L(R) of the regular pattern specified by the Query Term:
 - If the target value data type is different than *String* then it shall be considered as not matching.

4.10 NGSI-LD Geoquery Language

The NGSI-LD Geoquery language shall be supported by implementations. It is intended to define predicates which allow testing whether a specific topological spatial relationship exists between a pair of geometries: a target geometry and a reference geometry. The target geometry represents a geospatial Property of an Entity, typically, the location of the Entity.

A total of four parameters are defined in order to fully specify an NGSI-LD Geoquery:

- **georel**, to express the desired geospatial relationship;
- **geometry**, to express the type of the reference geometry;
- **coordinates**, to express the reference geometry;
- **geoproperty**, to express the target geometry of an Entity. This parameter is optional, *location* is the default.

The following grammar defines the syntax for the geospatial relationships (parameter name **georel**):

```
andOp = %x3B                ; ;
equal = %x3D %x3D          ; ==
georel = nearRel / withinRel / containsRel / overlapsRel / intersectsRel / equalsRel / disjointRel
nearRel = nearOp andOp distance equal PositiveNumber ; near;max(min)Distance==x (in meters)
distance = "maxDistance" / "minDistance"
nearOp = "near"
withinRel = "within"
containsRel = "contains"
intersectsRel = "intersects"
equalsRel = "equals"
disjointRel = "disjoint"
overlapsRel = "overlaps"
```

PositiveNumber shall be a non-zero positive number as mandated by the JSON Specification. Thus, it shall follow the ABNF Grammar, production rule named Number, section 6 of IETF RFC 8259 [6], excluding the 'minus' symbol and excluding the number 0.

Reference geometries shall be specified by:

- A geometry type (parameter name **geometry**) as defined by the GeoJSON specification (IETF RFC 7946 [8], section 1.4), except *GeometryCollection*.
- A coordinates (parameter name **coordinates**) element which shall represent the coordinates of the reference geometry as mandated by IETF RFC 7946 [8], section 3.1.1.

Target geometry, i.e. the target Entity's *GeoProperty* to which the geoquery is to be applied, can be specified by an extra parameter named **geoproperty**. The *GeoProperty*'s name shall be specified as short hand name and not a fully qualified one, because, when the query language is used, an *@context* properly defining all the terms (as per clause 5.5.7) shall be issued. If no *geoproperty* is specified, the geoquery is applied to the default *Property location* (see clause 4.7.1).

Note that proper URL encoding shall be used by HTTP binding API clients when using these examples.

EXAMPLE 1: ?georel=near;maxDistance==2000
&geometry=Point
&coordinates=[8,40]
&geoproperty=observationSpace

EXAMPLE 2: ?georel=within
&geometry=Polygon
&coordinates=[[100.0,0.0],[101.0,0.0],[101.0,1.0],[100.0,1.0],
[100.0,0.0]]
&geoproperty=location

EXAMPLE 3: A query encoded as an HTTP Query String. Note that this is HTTP binding specific, to be used via GET method, as defined in clause 6.4.3.2.

```
?georel=near;maxDistance==2000
&geometry=Point
&coordinates=[8,40]
```

The semantics of the different geospatial relationships defined above is as follows, and shall be supported by compliant implementations:

- **near** statement (production rule named nearRel):
 - maxDistance modifier. For an entity to match it has to be within the buffer geometric object (as defined by [14]) given by the reference geometry, with distance (in meters) equal to the number conveyed (production rule named PositiveNumber).
 - minDistance modifier. For an entity to match it has to be disjoint with the buffer geometric object (as defined by [14]) given by the reference geometry, with distance (in meters) equal to the number conveyed (production rule named PositiveNumber).

- **equals** relationship (production rule named `equalsRel`). For an entity to match, the target geometry shall be equal, as specified by [14], to the reference geometry.
- **disjoint** relationship (production rule named `disjointRel`). For an entity to match, the target geometry shall be disjoint, as specified by [14], to the reference geometry.
- **intersects** relationship (production rule named `intersectsRel`). For an entity to match, the target geometry shall intersect, as specified by [14], with the reference geometry.
- **within** relationship (production rule named `withinRel`). For an entity to match, the target geometry shall be within, as specified by [14], the reference geometry.
- **contains** relationship (production rule named `containsRel`). For an entity to match, the target geometry shall contain, as specified by [14], the reference geometry.
- **overlaps** relationship (production rule named `overlapsRel`). For an entity to match, the target geometry shall overlap, as specified by [14], the reference geometry.

When resolving geo-queries, Entities which do not convey the target *GeoProperty* of the query shall be considered as non-matching.

4.11 NGSI-LD Temporal Query Language

The NGSI-LD Temporal Query language shall be supported by implementations. It is intended to define predicates which allow testing whether Temporal Properties of NGSI-LD Entities, Properties and Relationships, are within certain temporal constraints. In particular it can be used to request historic Property values and Relationships that were valid within the specified timeframe.

The following grammar defines the syntax that shall be supported:

```
timerel = beforeRel / afterRel / betweenRel
beforeRel = "before"
afterRel = "after"
betweenRel = "between"
```

The points in time for comparison are defined as follows:

- A **timeAt** parameter, which shall represent the comparison point for the *before* and *after* relation and the starting point for the *between* relation. It shall be represented as *DateTime* (mandated by clause 4.6.3).
- An **endTimeAt** parameter, which is only used for the *between* relation and shall represent the end point for comparison. It shall be represented as *DateTime* (mandated by clause 4.6.3).

The Temporal Property (see clause 4.8) to which the temporal query is to be applied can be specified by **timeproperty**. If no *timeproperty* is specified, the temporal query is applied to the default Temporal Property *observedAt*.

EXAMPLE 1: `?timerel=before
&timeAt=2017-12-13T14:20:00Z`

EXAMPLE 2: `?timerel=between
&timeAt=2017-12-13T14:20:00Z
&endTimeAt=2017-12-13T14:40:00Z
&timeproperty=modifiedAt`

EXAMPLE 3: Temporal query encoded as HTTP Query String, note that this is HTTP binding specific, to be used via GET method, as defined in clause 6.18.3.2.

```
?timerel=between
&timeproperty=observedAt
&timeAt=2017-12-13T14:20:00Z
```

The semantics of the different temporal relations defined above is as follows, and shall be supported by compliant implementations:

- **before** relationship (production rule named *beforeRel*). For a Temporal Property to match, the value of the specified Temporal Property (or *observedAt* as default) has to be before the time specified by *timeAt*. The specified value is used as an exclusive bound in the Temporal Query;
- **after** relationship (production rule named *afterRel*). For a Temporal Property to match, the value of the specified Temporal Property (or *observedAt* as default) has to be after the time specified by *timeAt*. The specified value is used as an inclusive bound in the Temporal Query;
- **between** relationship (production rule named *betweenRel*). For a Temporal Property to match, the value of the specified Temporal Property (or *observedAt* as default) has to be after the time specified by *timeAt* and before the time specified by *endTimeAt*. In the Temporal Query, the value specified for the lower bound of the range is inclusive and the value specified for the upper bound of the range is exclusive.

When resolving temporal queries, Entities which do not convey the target Temporal Property of the query shall be considered as non-matching.

4.12 NGSI-LD Pagination

NGSI-LD operations can potentially return a result set including a large number of NGSI-LD Elements, so that pagination of query results shall be supported by compliant implementations.

The list of operations that incur this behaviour is as follows:

- Query Entities (clause 5.7.2).
- Query Subscriptions (clause 5.8.4).
- Query Context Source Registrations (clause 5.10.2).
- Query Context Source Registration Subscriptions (clause 5.11.5).
- Query Temporal Evolution of Entities (clause 5.7.4).

Nonetheless, the NGSI-LD API is agnostic about specific pagination mechanisms and only defines the behaviour that shall be observed by NGSI-LD Systems.

For each operation above, NGSI-LD Systems shall:

- provide a mechanism to iterate through the NGSI-LD Elements of a result set without exhausting NGSI-LD Client or Broker resources;
- provide a mechanism to flag NGSI-LD Clients when there are remaining NGSI-LD Elements to be traversed as part of a result set;
- allow NGSI-LD Clients specifying a limit (page size), as a parameter of API operations, to the number of NGSI-LD Elements (at a maximum) retrieved by the implementation for each pagination iteration;
- define a **default limit** (default page size) to the number of NGSI-LD Elements retrieved per pagination iteration;
- allow NGSI-LD Clients iterating forwards and backwards through a result set.

NGSI-LD implementations should:

- avoid Denial of Service attacks or other potential security risks, by defining a hard limit to the size of generated response payload body while paginating. For instance, certain queries can be rejected by issuing an error of type *TooManyResults*.

NGSI-LD implementations may:

- warn NGSI-LD Clients when result sets become invalid due to dynamic changes in NGSI-LD Elements (additions, deletions) occurred while iterating over pages.

The concrete realization of the features described above might depend on each API binding. Nonetheless, NGSI-LD Systems shall implement pagination features as mandated by the present clause, for any API binding.

4.13 Counting the Number of Results

Given that NGSI-LD Query operations can potentially return a result set including a large number of NGSI-LD Elements and that pagination of query results shall be supported (see clause 4.12), compliant implementations shall also support a mechanism for relaying to the client the number of expected resulting elements, when a query is executed.

A specific field (e.g. a custom header in the response in case of HTTP binding, see clause 6.3.13) shall be returned within the response of a query, whenever this is requested by the client.

Mechanisms for limiting the number of returned NGSI-LD Elements are independent of the counting mechanism, so that, potentially, a client can issue a query that limits to zero the number of desired results but asks for the count to be present.

This is useful for client-side planning and fine-tuning of subsequent queries and their parameters.

4.14 Supporting Multiple Tenants

The concept of a **Tenant** is that a user or group of users utilizes a single instance of an NGSI-LD system (**Context Source** or **Context Broker**) in isolation from other users or groups of users of the same instance, which are considered to be different **Tenants**. Thus a multi-tenant NGSI-LD system is a system where a single software instance is used by different users or groups of users, the **Tenants**, where any information related to one **Tenant** (e.g. **Entities**, **Subscriptions**, **Context Source Registrations**) are only visible to users of the same **Tenant**, but not to users of a different **Tenant**. Typically, multi-tenancy is used together with an access control mechanism, enforcing the isolation of **Tenants**, however access control and other security-related aspects are out-of-scope of the present document.

The NGSI-LD API optionally enables multi-tenant systems. To support this, **Tenant** information can be optionally specified in NGSI-LD API operations. The operation then only applies to the targeted **Tenant**. As all information of one **Tenant** is isolated from other **Tenants**, the NGSI-LD API operations for managing, retrieving and subscribing to entity information, but also any context source related operations only apply to the information of the specified **Tenant** in isolation and never have any effect on the information of other **Tenants**.

As the support and use of **Tenants** is optional, any operation not explicitly specifying a **Tenant** targets a default **Tenant**, which always exists. NGSI-LD systems not supporting multiple **Tenants** should raise an error of type *NoMultiTenantSupport* if a **Tenant** is specified. To enable **Context Sources** to be part of tenant-based distributed or federated systems, **Tenant** information can optionally be specified in **Context Source Registrations**. When contacting the respective **Context Sources**, the **Tenant** information from the **Context Source Registration** has to be used. If no **Tenant** information is present in the **Context Source Registration**, no **Tenant** information is to be used and thus the default **Tenant** is targeted on the registered **Context Source**. This enables integrating **Context Sources** not supporting multi-tenancy in a distributed system with a **Tenant**-based **Context Broker** or integrating local **Tenants** in a federated system using a different **Tenant**.

4.15 NGSI-LD Language Filter

The NGSI-LD Language Filter shall be supported by implementations. It is intended to form a mechanism which allows just one matching string value of *LanguageProperties* of NGSI-LD Entities to be converted to an NGSI-LD Property, where the value will be a string for the specified natural language.

The following grammar defines the syntax that shall be supported by the filter:

`lang = langtag`

Where the langtag is defined according to the rules as mandated by IETF RFC 5646 [28], and IETF RFC 3282 [29]. If the Context Broker cannot serve any matching language, it shall default to any supported language. This behaviour can be triggered by specifying `lang=""` in the filter (see example 3).

In any case, the attribute in question shall be augmented with an additional non-reified subproperty *lang* indicating the actual language returned.

EXAMPLE 1: Specified natural language - return LanguageProperties as strings in English only.

`lang="en"`

EXAMPLE 2: Multiple natural languages with no ranked preference - return LanguageProperties as strings in either Swiss French or French.

`lang="fr-CH, fr"`

EXAMPLE 3: Wildcard - return LanguageProperties as a string in any supported language.

`lang=""`

EXAMPLE 4: Quality value ranking - return all LanguageProperties as a string in Swiss French or French with no ranked preference, fallback to English as a second choice and finally fallback to any other supported language.

`lang="fr-CH, fr;q=0.9, en;q=0.8, *,q=0.5"`

4.16 Supporting Multiple Entity Types

From NGSI-LD API version 1.5.1 onwards, multiple Entity Types for any Entity are supported. An Entity is uniquely identified by its id, so whenever information is provided for an Entity with a given id, it is considered part of the same Entity, regardless of the Entity Type(s) specified. To avoid unexpected behaviour, Entity Types can be implicitly added by all operations that update or append attributes. There is no operation to remove Entity Types from an Entity. The philosophy here is to assume that an Entity always had all Entity Types, but possibly not all Entity Types have previously been known in the system. The only option to remove an Entity Type is to delete the Entity and re-create it with the same id. Alternatively, a batch upsert with 'replace' update mode can be used, as described in clause 5.6.8.

4.17 NGSI-LD Entity Type Selection Language

The NGSI-LD Entity Type Selection Language shall be supported by implementations. It is intended to select only those Entities that have the specified Entity Type(s), possibly among others. Entity Types are specified as a disjunction of elements, where each element can either directly be an Entity Type or a conjunction of multiple Entity Types. The logical operators are the same as in the NGSI-LD Query Language specified in clause 4.9. As a disjunction of Entity Types can also be seen as a list, and to be compatible with previous versions of the NGSI-LD API, a comma can be used as an alternative representation of the *or* operator. For logical *and* grouping parenthesis are needed.

```
EntityTypes = OrEntityType *(orOp OrEntityType)           ; OrEntityType|OrEntityType
OrEntityType = %x28 EntityType *(andOp EntityType) %x29 ; (EntityType;EntityType)
OrEntityType = EntityType                                   ; EntityType
andOp = %x3B                                              ; ;
orOp = %x7C / %x2C                                       ; | ,
```

EntityType is either a valid name as specified in clause 4.6.2 or a URI.

EXAMPLE 1: Entities of type Building or House:

`Building|House`

Alternative Representation:

`Building,House`

EXAMPLE 2: Entities of type Home and Vehicle:

(Home;Vehicle)

EXAMPLE 3: Entities of type (Home and Vehicle) or Motorhome:

(Home;Vehicle)|Motorhome

Alternative Representation:

(Home;Vehicle),Motorhome

NOTE: The special characters " , " ; " (" and ") " used in the Entity Type Selection Language are allowed characters in URIs. The use of short names is recommended.

4.18 NGSI-LD Scopes

An NGSI-LD Scope enables putting Entities into a hierarchical structure and restrict results of queries and notifications accordingly. The hierarchical structure is user-defined, e.g. according to (logical) location or organization. The use of Scopes is optional and an Entity can be assigned to one or more Scopes at the same time. The Scope is represented as a special *scope* Property that is non-reified in the normalized Entity representation and reified in the temporal representation of an Entity. In the latter case, it is restricted to having the non-reified Temporal Properties *createdAt*, *modifiedAt* and *deletedAt* as sub-Properties. There shall at most be one instance of the *scope* property per Entity. In case multiple representations of the same Entity have to be merged, e.g. when combining the results of distributed queries, the values of *scope* are merged. The value of *scope* is represented as a JSON array in case there is more than one Scope. For the Temporal Evolution a given Scope is considered valid from the time it has been set until the time it has been explicitly removed by an update or delete operation (for an example see annex C, clause C.5.16).

The grammar that encodes the syntax of the Scope is expressed in ABNF format [12]. It is described below (it has been validated using <https://github.com/ietf-tools/bap>), and it shall be supported by implementations. The special string "urn:ngsi-ld:null" (i.e. the NGSI-LD Null) shall be only used and only appear in case of deleted scopes.

```
Scope = [%x2F] ScopeLevel *(%x2F ScopeLevel)           ; [/] ScopeLevel *(/ScopeLevel)
Scope =/ "urn:ngsi-ld:null"                               ; the literal string "urn:ngsi-ld:null"
ScopeLevel = unicodeLetter *ScopeLevelChar
ScopeLevelChar = unicodeNumber / unicodeLetter
ScopeLevelChar =/ %x5F                                   ; _
```

EXAMPLE 1: /Madrid

EXAMPLE 2: Madrid

EXAMPLE 3: /Madrid/Gardens/ParqueNorte

EXAMPLE 4: /CompanyA/OrganizationB/UnitC

4.19 NGSI-LD Scope Query Language

The NGSI-LD Scope Query Language shall be supported by implementations. It is intended to select only those Entities that are within the specified Scope(s). Scopes are specified as a disjunction of elements, where each element can either directly be a Scope or a conjunction of multiple Scopes. The "+" can be used as a wildcard to match a single scope level within a Scope. The "#" that can be added at the end of a Scope specification serves as a wildcard, which matches the given scope and the whole hierarchy of scopes below. The "/"# matches any non-empty scope, i.e. only explicitly specified scopes. The logical operators are the same as in the NGSI-LD Query Language specified in clause 4.9. As a disjunction of Scopes can also be seen as a list, a comma can be used as an alternative representation of the *or* operator. For logical *and* grouping parenthesis are needed.

```
ScopesQ = OrScopeQ *(orOp OrScopeQ)                   ; OrScopeQ|OrScopeQ
ScopesQ =/ %x2F %23                                     ; / #
OrScopeQ = %x28 ScopeQ *(andOp ScopeQ) %x29            ; (ScopeQ;ScopeQ)
OrScopeQ =/ ScopeQ *1(%x2F %23)                        ; ScopeQ / #
andOp = %x3B                                             ; ;
orOp = %x7C / %x2C                                       ; | ,
ScopeQ = %x2F ScopeQLevel *(%x2F ScopeQLevel)         ; /ScopeQLevel *(/ScopeQLevel)
ScopeQLevel = unicodeLetter *ScopeQLevelChar
```

ScopeQLLevel =/ %x2B ; +
 ScopeQLLevelChar = unicodeNumber / unicodeLetter
 ScopeQLLevelChar =/ %x5F ; _

EXAMPLE 1: Scope /Madrid:

/Madrid

EXAMPLE 2: Scope /Madrid/Gardens and the whole scope tree below:

/Madrid/Gardens/#, e.g. matches the following Scopes:

/Madrid/Gardens, /Madrid/Gardens/ParqueNorte,
 /Madrid/Gardens/ParqueNorte/Parterre1, /Madrid/Gardens/ParqueSur

EXAMPLE 3: Scopes /Madrid/Gardens/ParqueNorte and /Madrid/Sights/ParqueNorte, or any other Scope with Madrid as first scope level and ParqueNorte as third scope level:

/Madrid/+/ParqueNorte

EXAMPLE 4: Scope /Madrid/Districts and /CompanyA:

(/Madrid/Districts;/CompanyA)

EXAMPLE 5: Scope (/Madrid/Districts and /CompanyA) or /CompanyB:

(/Madrid/Districts;/CompanyA)|CompanyB

Alternative Representation:

(/Madrid/Districts;/CompanyA),CompanyB

4.20 NGSI-LD Distributed Operation names

When registering Context Sources (see clause 5.2.9), the registrant NGSI-LD interface endpoint may optionally offer a subset of NGSI-LD operations which it accepts. Table 4.20-1 defines a list of names for each of these operations.

Table 4.20-1: Names of implemented Operations

	Operation name	Implements
Context Information Provision	createEntity	5.6.1 Create Entity
	updateEntity	5.6.2 Update Attributes
	appendAttrs	5.6.3 Append Attributes
	updateAttrs	5.6.4 Partial Attribute update
	deleteAttrs	5.6.5 Delete Attribute
	deleteEntity	5.6.6 Delete Entity
	createBatch	5.6.7 Batch Entity Creation
	upsertBatch	5.6.8 Batch Entity Creation or Update (Upsert)
	updateBatch	5.6.9 Batch Entity Update
	deleteBatch	5.6.10 Batch Entity Delete
	upsertTemporal	5.6.11 Create or Update Temporal Evolution of an Entity
	appendAttrsTemporal	5.6.12 Add Attributes to Temporal Evolution of an Entity
	deleteAttrsTemporal	5.6.13 Delete Attributes from Temporal Evolution of an Entity
	updateAttrInstanceTemporal	5.6.14 Partial Update Attribute Instance in Temporal Evolution of an Entity
	deleteAttrInstanceTemporal	5.6.15 Delete Attribute Instance from Temporal Evolution of an Entity
	deleteTemporal	5.6.16 Delete Temporal Evolution of an Entity
	mergeEntity	5.6.17 Merge Entity
	replaceEntity	5.6.18 Replace Entity
	replaceAttrs	5.6.19 Replace Attribute
	mergeBatch	5.6.20 Batch Entity Merge
	purgeEntity	5.6.21 Purge Entities
Context Information	retrieveEntity	5.7.1 Retrieve Entity
	queryEntity	5.7.2 Query Entities (excluding batch entity queries)

	Operation name	Implements
Consumption	queryBatch	5.7.2 Query Entities (batch entity queries only)
	retrieveTemporal	5.7.3 Retrieve Temporal Evolution of an Entity
	queryTemporal	5.7.4 Query Temporal Evolution of Entities
	retrieveEntityTypes	5.7.5 Retrieve Available Entity Types
	retrieveEntityTypeDetails	5.7.6 Retrieve Details of Available Entity Types
	retrieveEntityTypeInfo	5.7.7 Retrieve Available Entity Type Information
	retrieveAttrTypes	5.7.8 Retrieve Available Attributes
	retrieveAttrTypeDetails	5.7.9 Retrieve Details of Available Attributes
	retrieveAttrTypeInfo	5.7.10 Retrieve Available Attribute Information
Context Information Subscription	createSubscription	5.8.1 Create Subscription
	updateSubscription	5.8.2 Update Subscription
	retrieveSubscription	5.8.3 Retrieve Subscription
	querySubscription	5.8.4 Query Subscription
	deleteSubscription	5.8.5 Delete Subscription
Support operations for distributed operations	retrieveEntityMap	5.14.1 Retrieve EntityMap
	updateEntityMap	5.14.2 Update EntityMap
	deleteEntityMap	5.14.3 Delete EntityMap
	createEntityMapQueryEntity	5.14.4 Create EntityMap for Query Entities
	createEntityMapQueryTemporal	5.14.5 Create EntityMap for Query Temporal Evolution of Entities
	retrieveContextSourceIdentity	5.15.1 Retrieve Context Source Identity Information

In addition to these individual operations, a series of names for common groups of operations have also been defined. Table 4.20-2 defines a list of names for each of these operation groups.

Table 4.20-2: Named Operation Groups

Operation Group name	Implements
federationOps	• retrieveEntity • queryEntity • queryBatch • retrieveEntityTypes • retrieveEntityTypeDetails • retrieveEntityTypeInfo • retrieveAttrTypes • retrieveAttrTypeDetails • retrieveAttrTypeInfo • createSubscription • updateSubscription • retrieveSubscription • querySubscription • deleteSubscription • retrieveEntityMap • updateEntityMap • deleteEntityMap • createEntityMapQueryEntity • retrieveContextSourceIdentity
associationOps	• retrieveEntity • queryEntity • queryBatch • retrieveEntityTypes • retrieveEntityTypeDetails • retrieveEntityTypeInfo • retrieveAttrTypes • retrieveAttrTypeDetails • retrieveAttrTypeInfo • createSubscription • updateSubscription • retrieveSubscription • querySubscription • deleteSubscription • retrieveContextSourceIdentity

Operation Group name	Implements
updateOps	• updateEntity • updateAttrs • replaceEntity • replaceAttrs
retrieveOps	• retrieveEntity • queryEntity
redirectionOps	• createEntity • updateEntity • appendAttrs • updateAttrs • deleteAttrs • deleteEntity • mergeEntity • replaceEntity • replaceAttrs • retrieveEntity • queryEntity • purgeEntity • retrieveEntityTypes • retrieveEntityTypeDetails • retrieveEntityTypeInfo • retrieveAttrTypes • retrieveAttrTypeDetails • retrieveAttrTypeInfo • retrieveEntityMap • updateEntityMap • deleteEntityMap • createEntityMapQueryEntity • retrieveContextSourceIdentity

If no specific subset of operations is defined for a Context Source Registration, the default set of operations matches the group defined as "federationOps".

4.21 NGSI-LD Attribute Projection Language

The NGSI-LD Attribute Projection Language shall be supported by implementations for projection parameters (e.g. *pick* and *omit*). Its aim is to specify the Attributes to be retrieved within an Entity (or associated Linked Entities when Linked Entity Retrieval is used). Projected Attributes are specified as a disjunction of elements, where each element can either directly be an Attribute name, or in the case of Linked Entity Retrieval, a Relationship name followed by an Attribute name within the Linked Entity. Since a disjunction of Attributes is a list, either a comma or a pipe character can be used as alternative representations of the *or* operator. In the following, ABNF grammar for NGSI-LD Attribute Projection Language is given.

```

orOp = %x7C / %x2C ; | ,
ProjectionTerm = AttrName *1(LinkedEntityTerm) *(orOp ProjectionTerm)
LinkedEntityTerm = %x7B ProjectionTerm %x7D ; {ProjectionTerm}

```

See clause 4.9 for the definition of AttrName.

EXAMPLE 1: ?pick=temperature

EXAMPLE 2: ?pick=temperature,humidity

EXAMPLE 3: ?pick=observation{temperature,humidity}

4.22 Transient Storage of Entities and Attributes

In some cases, it is desirable to create an Entity (or Attribute) which is only expected to be stored for a defined period of time. Thereafter such an Entity (or Attribute) should be removed, and can be safely deleted from the context via an automatic garbage collection process.

In this regard NGSI-LD defines the following system Property of type *TemporalProperty* that shall be supported by implementations:

- **expiresAt** is defined as the system temporal Property at which a certain Entity, Property or Relationship shall become invalid and may be automatically removed from the Context Broker. For example, an Alert Entity was created to last for 24 hours and should be removed after this period of time.

It should be noted that clean-up processes will only run periodically, and will be dependent upon the Context Broker implementation, therefore final deletion will always lag the *expiresAt* timestamp to a certain extent. Furthermore, *expiresAt* only applies to the local storage, i.e. the Entity or Attribute is to be deleted locally, but not on other Context Sources or Context Broker hosting such Entity information, where no *expiresAt* timestamp is present. Thus *expiresAt* is not considered to be intrinsic to the Entity or Attribute, but only applies to the storage of the Entity or Attribute respectively. As it pertains to a system function (the deletion from storage after the expiration time), it is considered to be a system attribute.

4.23 Entity Ordering

4.23.1 Introduction

When a Context Consumer queries for Entities retrieved from a single context broker, it shall be possible to request that the array of Entities returned are ordered according to the values held within an Entity Attribute. Context Broker implementations shall interpret such requests and return the array of Entities sorted accordingly.

For each Attribute name listed, one of four types of sort order can be specified:

- Sort in ascending order by **target value** or **target object**
- Sort in descending order by **target value** or **target object**
- Sort by distance in ascending order from a specified GeoJSON geometry
- Sort by distance in descending order from a specified GeoJSON geometry

When sorting by ascending or descending order, the preferred sort order for numbers and datetimes is trivial, however for strings the default collation order shall be defined as using ICU "root" collation order ("und-x-icu") which returns a reasonable language-agnostic sort order (see IETF RFC 6067 [36]).

Sort by distance shall be limited to GeoProperties, and shall be defined as calculating spherical distance using the Haversine formula.

Sort ordering is **never applied** to distributed operations, and the defined sort ordering strategy may depend on implementation specific configurations.

4.23.2 Datatype Comparison Order

When sorting by value, and the values of Attributes are of mixed datatypes, the following comparison order (*null* values last), shall be implicitly applied:

- Numbers
- Strings
- Object
- Array

- Boolean
- *Time*
- *Date*
- *DateTime*
- Null
- Attribute does not exist

Additionally, when sorting by distance, the following comparison order, shall be applied:

- Attribute is a GeoProperty
- Attribute is not a GeoProperty - sorted by value as shown above.

4.23.3 NGSI-LD Entity Ordering Language

The NGSI-LD Entity Ordering Language shall be supported by implementations for the order parameter (i.e. *orderBy*). Its aim is to specify the Attributes to be used when ordering an Array of Entities retrieved. Ordering Attributes are specified as a disjunction of elements, where each element can either directly be an Attribute name, or an Attribute name followed by a direction (either ascending or descending, where ascending shall be the default if not supplied). The comma character shall be used to separate the elements to be used for ordering which shall be applied sequentially.

In the following, ABNF grammar for NGSI-LD Entity Ordering Language is given.

```

thenOp = %x2C                                ; ,
directionOp ::= asc | desc | dist-asc | dist-desc
OrderingTerm = AttrName *1(DirectionTerm) *(thenOp OrderingTerm)
DirectionTerm = %x3B directionOp              ; ;

```

See clause 4.9 for the definition of *AttrName*.

EXAMPLE 1: `?orderBy=name` - applies sort ordering to rank Entities in ascending order using the value of the *name* Attribute.

EXAMPLE 2: `?orderBy=name, age` - applies sort ordering to rank Entities in ascending order using the value of the *name* Attribute and thereafter where multiple Entities have the same *name*, by ascending order using the *age* Attribute.

EXAMPLE 3: `?orderBy=name, age; desc` - applies sort ordering to rank Entities in ascending order using the value of the *name* Attribute and thereafter where multiple Entities have the same *name*, by descending order using the value of *age* Attribute.

EXAMPLE 4: `?orderBy=address[city]` - applies sort ordering using a sub-item within a Property Value defined as a complex JSON object. The trailing path is `[city]`. It is used to refer to a particular subitem within the value of the *address* Property.

EXAMPLE 5: `?orderBy=name.observedAt` - applies sort ordering based on an attribute path that is defined by the production rule *Attribute*, as a dot-separated list of names. Such a list is intended to address a *Property* or *Relationship* included by the matching entities subjacent graph.

When sorting data using a specific ICU collation order as defined by IETF RFC 6067 [36], the collation element (parameter name **collation**) shall represent the preferred ordering.

EXAMPLE 6: `?orderBy=name&collation=und-u-ks-identic` - applies sort ordering to rank Entities in ascending order using the case insensitive value of the *name* Attribute.

EXAMPLE 7: `?orderBy=name&collation=de-u-co-phonebk` - applies sort ordering to rank Entities in ascending order using the value of the *name* Attribute and sorted according to German phonebook collation ordering.

For sort by distance, the coordinates element (parameter name **orderFrom**) shall represent the coordinates of the reference geometry as mandated by IETF RFC 7946 [8], section 3.1.2.

EXAMPLE 8: `?orderBy=location;dist-asc&orderFrom=[8,40]` - applies sort ordering to rank Entities in ascending distance from a Point with respect to the *location GeoProperty*.

EXAMPLE 9: `?orderBy=location;dist-desc&orderFrom=[8,40]` - applies sort ordering to rank Entities in descending distance from a Point with respect to the *location GeoProperty*.

EXAMPLE 10: `?orderBy=location;dist-asc`
`&orderFrom=[[8,40],[9,42],[9,45],[8,40]]`
`&orderGeometry=LineString`
 applies sort ordering to rank Entities in ascending distance from a LineString with respect to the *location GeoProperty*.

5 API Operation Definition

5.1 Introduction

This clause defines data structures and operations of the NGSI-LD API. No specific binding is assumed. Clause 6 maps these operations and data types to the HTTP REST binding.

NOTE: In UML diagrams dotted arrows denote a response to a request.

5.2 Data Types

5.2.1 Introduction

Implementations shall support the data types defined by the clauses below. For each member defined by each data type (including nested ones) a term shall be added to the Core *@context*, as mandated by clause 4.5.

None of the members described admit a *null* value directly, as when a JSON-LD processor encounters *null*, the associated entry or value is always removed when expanding the JSON-LD document.

However, in the context of a partial update or merge operation (see clauses 5.5.8 and 5.5.12), an NGSI-LD Null shall be used to indicate the removal of a target member, as explained in clause 4.5.0. In all other cases, implementations shall raise an error of type *BadRequestData* if an NGSI-LD Null value is encountered.

As *null* cannot be used as a value in JSON-LD, there is still the possibility of using a JSON null literal `{"@type": "@json", "@value": null}` instead. JSON literals are not to be expanded in JSON-LD and thus the respective element is not removed during JSON-LD expansion.

Non-normative JSON Schema [i.11] definitions of the referred data types are also available at [i.13].

The use of URI in the context of the present document also includes the use of International Resource Identifiers (IRIs) as defined in IETF RFC 3987 [23], which extends the use of characters to Unicode characters [22] beyond the ASCII character set, enabling the support of languages other than English.

5.2.2 Common members

The JSON-LD representation of NGSI-LD Entity, Property, Relationship, Context Source Registration and Subscription can include the common members described by Table 5.2.2-1.

Those members are read-only, and shall be automatically generated by NGSI-LD implementations. They shall not be provided by Context Producers. In the event that they are provided (in update or create operations) NGSI-LD implementations shall ignore them.

In query or retrieve operations implementations shall only generate common members (Table 5.2.2-1) when the **Context Consumer** explicitly asks for their inclusion. Clause 6.3.11 defines the mechanism offered by the HTTP binding for such purpose.

Table 5.2.2-1: Common members of NGSI-LD elements

Name	Data Type	Restriction	Cardinality	Description
createdAt	string	<i>DateTime</i> (clause 4.6.3)	0..1	Entity creation timestamp. See clause 4.8.
deletedAt	string	<i>DateTime</i> (clause 4.6.3)	0..1	Entity deletion timestamp. See clause 4.8. It is only used in notifications reporting deletions and in the temporal representation of Entities (clause 4.5.6), Properties (clause 4.5.7), Relationships (clause 4.5.8) and LanguageProperties (clause 5.2.32) and VocabProperties (clause 5.2.35) and JsonProperties (clause 5.2.38).
modifiedAt	string	<i>DateTime</i> (clause 4.6.3)	0..1	Entity last modification timestamp. See clause 4.8.

5.2.3 @context

When encoding NGSI-LD Entities, Context Source Registrations, Subscriptions and Notifications, as pure JSON-LD (MIME type "application/ld+json"), an array (flattened to a single string if necessary), containing a user *@context* where present, and the core *@context* (as described in clause 4.4) shall be included as a special member of the corresponding JSON-LD Object. Table 5.2.3-1 gives a precise definition of this special member.

Table 5.2.3-1: JSON-LD @context tagged member

Name	Data Type	Restriction	Cardinality	Description
@context	URI, JSON Object, or JSON Array	See [2], section 5.1.	0..1	JSON-LD <i>@context</i> .

5.2.4 Entity

This type represents the data needed to define an NGSI-LD Entity as mandated by clause 4.5.1.

The supported JSON members shall follow the requirements provided in Table 5.2.4-1.

Table 5.2.4-1: NGSI-LD Entity data type definition

Name	Data Type	Restriction	Cardinality	Description
id	String	Valid URI	1	Entity ID.
type	String or String[]		1	Entity Type(s). Both short hand string(s) (type name) or URI(s) are allowed.
expiresAt	String	<i>DateTime</i> (see clause 4.6.3)	0..1	System temporal Property representing the expiration date for the storage of the Entity. See clause 4.22.
location	GeoProperty	See datatype definition in clause 5.2.7	0..1	Default geospatial Property of an entity. See clause 4.7.
observationSpace	GeoProperty	See datatype definition in clause 5.2.7	0..1	See clause 4.7.
operationSpace	GeoProperty	See datatype definition in clause 5.2.7	0..1	See clause 4.7.
scope	String or String[]	See clause 4.18	0..1	Scope.

Name	Data Type	Restriction	Cardinality	Description
<Property name>	Property or Property[] (see note 1)	See datatype definitions in clause 5.2.5	0..N	Property as mandated by clause 4.5.2.
	GeoProperty or GeoProperty[] (see note 1)	See datatype definition in clause 5.2.7	0..N	GeoProperty as mandated by clause 4.5.2.
	LanguageProperty or LanguageProperty[] (see note 1)	See datatype definition in clause 5.2.32	0..N	LanguageProperty as mandated by clause 4.5.18.
	JsonProperty or JsonProperty[] (see note 1)	See datatype definition in clause 5.2.38	0..N	JsonProperty as mandated by clause 4.5.24.
	VocabProperty or VocabProperty[] (see note 1)	See datatype definition in clause 5.2.35	0..N	VocabProperty as mandated by clause 4.5.20.
	ListProperty or ListProperty[] (see note 1)	See datatype definition in clause 5.2.36	0..N	ListProperty as mandated by clause 4.5.21.
<Relationship name>	Relationship or Relationship[] (see note 2)	See datatype definition in clause 5.2.6	0..N	Relationship as mandated by clause 4.5.3.
	ListRelationship or ListRelationship[] (see note 2)	See datatype definition in clause 5.2.37	0..N	ListRelationship as mandated by clause 4.5.22.
NOTE 1: For each <i>Property</i> (or subclass of <i>Property</i>) identified by the same Property name, there can be one or more instances separated by <i>datasetId</i> .				
NOTE 2: For each <i>Relationship</i> (or subclass of <i>Relationship</i>) identified by the same Relationship name, there can be one or more instances separated by <i>datasetId</i> .				

5.2.5 Property

This type represents the data needed to define a *Property* as mandated by clause 4.5.2.

The supported JSON members shall follow the requirements provided in Table 5.2.5-1 below. The datatype definition defines all the required attributes for the normalized representation. In the concise representation, the Attribute *type* member can be omitted as *type="Property"* can be inferred from the presence of the *value* member. Furthermore, in the concise representation of a *Property*, the *value* member cannot be a GeoJSON Object (as defined in clause 4.7) as it would be interpreted as a *GeoProperty* (see clause 5.2.7).

Table 5.2.5-1: NGSI-LD Property data type definition

Name	Data Type	Restriction	Cardinality	Description
type	String	It shall be equal to "Property"	1	Node type.
value	Any JSON value as defined by IETF RFC 8259 [6]	See NGSI-LD Value definition in clause 3.1.	1	Property Value.
datasetId	String	Valid URI	0..1	It allows identifying a set or group of property values.
expiresAt	String	<i>DateTime</i> (see clause 4.6.3)	0..1	System temporal Property representing the expiration date for the storage of the Property. See clause 4.22.

Name	Data Type	Restriction	Cardinality	Description
ngsildproof	Property	Property with the non-reified subproperties "entityIdSealed" and "entityTypeSealed" as specified in [35]. The value of its "value" element shall be an object containing the W3C® Data integrity "proof" structure [35].	0..1	Cryptographic signature of the Property guaranteeing its integrity. See clause C.11 for an example.
observedAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	Timestamp. See clause 4.8.
unitCode	String	As mandated by [15]	0..1	Property Value's unit code.
valueType	String		0..1	The native JSON-LD <i>@type</i> for the Property Value. A String Value which shall be type coerced to a URI based on the supplied <i>@context</i> .
<Property name>	Property or Property[] (see note 1)	See datatype definition in clause 5.2.5	0..N	Properties of the Property.
	GeoProperty or GeoProperty[] (see note 1)	See datatype definition in clause 5.2.7	0..N	GeoProperties of the Property.
	LanguageProperty or LanguageProperty[] (see note 1)	See datatype definition in clause 5.2.32	0..N	LanguageProperties of the Property.
	JsonProperty or JsonProperty[] (see note 1)	See datatype definition in clause 5.2.38	0..N	JsonProperties of the Property.
	VocabProperty or VocabProperty[] (see note 1)	See datatype definition in clause 5.2.35	0..N	VocabProperties of the Property.
	ListProperty or ListProperty[] (see note 1)	See datatype definition in clause 5.2.36	0..N	ListProperties of the Property.
<Relationship name>	Relationship or Relationship[] (see note 2)	See datatype definition in clause 5.2.6	0..N	Relationships of the Property.
	ListRelationship or ListRelationship[] (see note 2)	See datatype definition in clause 5.2.37	0..N	ListRelationships of the Property.
NOTE 1: For each <i>Property</i> (or subclass of <i>Property</i>) identified by the same Property name, there can be one or more instances separated by <i>datasetId</i> .				
NOTE 2: For each <i>Relationship</i> (or subclass of <i>Relationship</i>) identified by the same Relationship name, there can be one or more instances separated by <i>datasetId</i> .				

The following **output only** members (defined by Table 5.2.5-2) of the *Property* data structure are also defined. They are read-only and shall be generated by NGSI-LD implementations. They shall not be provided by **Context Producers**. In the event that they are provided (in update or create operations) NGSI-LD implementations shall ignore them.

Table 5.2.5-2: Output only members of the NGSI-LD Property data type

Name	Data Type	Restrictions	Cardinality	Description
createdAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	System generated creation timestamp. See clause 4.8.
deletedAt	String	<i>DateTime</i> (clause 4.6.3) It is only used in notifications reporting deletions	0..1	System generated deletion timestamp. See clause 4.8.
instanceId	String	Valid URI. Only used in temporal representation of Properties	0..1	URI uniquely identifying a Property instance as mandated by clause 4.5.7.
modifiedAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	System generated last modification timestamp. See clause 4.8.
previousValue	Any JSON value as defined by IETF RFC 8259 [6]	Only used in Notifications, if the <i>showChanges</i> option is explicitly requested	0..1	Previous Property Value.

5.2.6 Relationship

This type represents the data needed to define a *Relationship* as mandated by clause 4.5.3.

The supported JSON members shall follow the requirements provided in Table 5.2.6-1 below. The datatype definition defines all the required attributes for the normalized representation. In the concise representation, the Attribute *type* member can be omitted as *type="Relationship"* can be inferred from the presence of the *object* member.

Table 5.2.6-1: NGSI-LD Relationship data type definition

Name	Data Type	Restriction	Cardinality	Description
type	String	It shall be equal to "Relationship"	1	Node type.
object	String or String[]	Valid URI or an Array of Valid URIs	1	Relationship's target object.
datasetId	String	Valid URI	0..1	It allows identifying a set or group of target relationship objects.
expiresAt	String	<i>DateTime</i> (see clause 4.6.3)	0..1	System temporal Property representing the expiration date for the storage of the Relationship. See clause 4.22.
ngsildproof	Property	Property with the non-reified subproperties "entityIdSealed" and "entityTypeSealed" as specified in [35]. The value of its "value" element shall be an object containing the W3C® Data integrity "proof" structure [35].	0..1	Cryptographic signature of the Relationship guaranteeing its integrity.
objectType	String or String[]		0..1	Node Type of the Relationship's target object. Both short hand string(s) (type name) or URI(s) are allowed.
observedAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	Timestamp. See clause 4.8.

Name	Data Type	Restriction	Cardinality	Description
<Property name>	Property or Property[] (see note 1)	See datatype definition in clause 5.2.5	0..N	Properties of the Relationship.
	GeoProperty or GeoProperty[] (see note 1)	See datatype definition in clause 5.2.7	0..N	GeoProperties of the Relationship.
	LanguageProperty or LanguageProperty[] (see note 1)	See datatype definition in clause 5.2.32	0..N	LanguageProperties of the Relationship.
	JsonProperty or JsonProperty[] (see note 1)	See datatype definition in clause 5.2.38	0..N	JsonProperties of the Relationship.
	VocabProperty or VocabProperty[] (see note 1)	See datatype definition in clause 5.2.35	0..N	VocabProperties of the Relationship.
	ListProperty or ListProperty[] (see note 1)	See datatype definition in clause 5.2.36	0..N	ListProperties of the Relationship.
<Relationship name>	Relationship or Relationship[] (see note 2)	See datatype definition in clause 5.2.6	0..N	Relationships of the Relationship.
	ListRelationship or ListRelationship[] (see note 2)	See datatype definition in clause 5.2.37	0..N	ListRelationships of the Relationship.
NOTE 1: For each <i>Property</i> (or subclass of <i>Property</i>) identified by the same Property name, there can be one or more instances separated by <i>datasetId</i> .				
NOTE 2: For each <i>Relationship</i> (or subclass of <i>Relationship</i>) identified by the same Relationship name, there can be one or more instances separated by <i>datasetId</i> .				

The following **output only** members (defined by Table 5.2.6-2) of the *Relationship* data structure are also defined. They are read-only and shall be generated by NGSI-LD implementations. They shall not be provided by Context Producers. In the event that they are provided (in update or create operations) NGSI-LD implementations shall ignore them.

Table 5.2.6-2: Output only members of the NGSI-LD Relationship data type

Name	Data Type	Restrictions	Cardinality	Description
createdAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	System generated creation timestamp. See clause 4.8.
deletedAt	String	<i>DateTime</i> (clause 4.6.3) It is only used in notifications reporting deletions	0..1	System generated deletion timestamp. See clause 4.8.
entity	Entity or Entity[] (see note)	See datatype definition in clause 5.2.4. Only used in Linked Entity Retrieval, if the <i>join=inline</i> option is explicitly requested	0..1	An inline Entity obtained by Linked Entity Retrieval, corresponding to the Relationship's target object. See clause 4.5.23.2.
instanceId	String	Valid URI. Only used in temporal representation of Relationships	0..1	URI uniquely identifying a Relationship instance as mandated by clause 4.5.8.
modifiedAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	System generated last modification timestamp. See clause 4.8.
previousObject	String	Valid URI. Only used in Notifications, if the <i>showChanges</i> option is explicitly requested	0..1	Previous Relationship's target object.
NOTE: one-to-N <i>Relationships</i> expand to an array of Entity elements, since there can be more than one target object URI specified.				

5.2.7 GeoProperty

This type represents the data needed to define a *GeoProperty*.

The supported JSON members shall follow the requirements provided in Table 5.2.7-1 below. The datatype definition defines all the required attributes for the normalized representation. In the concise representation, the Attribute *type* member can be omitted as "GeoProperty" can be inferred from the presence of the *value* member holding a GeoJSON *Geometry* as mandated by clause 4.7.

Table 5.2.7-1: NGS-LD GeoProperty data type definition

Name	Data Type	Restriction	Cardinality	Description
type	String	It shall be equal to "GeoProperty"	1	Node type.
value	JSON Object	As mandated by clause 4.7	1	Geolocation encoded as GeoJSON [8].
datasetId	String	Valid URI	0..1	It allows identifying a set or group of property values.
expiresAt	String	<i>DateTime</i> (see clause 4.6.3)	0..1	System temporal Property representing the expiration date for the storage of the GeoProperty. See clause 4.22.
ngsldproof	Property	Property with the non-reified subproperties "entityIdSealed" and "entityTypeSealed" as specified in [35]. The value of its "value" element shall be an object containing the W3C® Data integrity "proof" structure [35].	0..1	Cryptographic signature of the GeoProperty guaranteeing its integrity.
observedAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	Timestamp. See clause 4.8.
valueType	String		0..1	The native JSON-LD <i>@type</i> for the GeoProperty Value. A String Value which shall be type coerced to a URI based on the supplied <i>@context</i> .
<Property name>	Property or Property[] (see note 1)	See datatype definition in clause 5.2.5	0..N	Properties of the GeoProperty.
	GeoProperty or GeoProperty[] (see note 1)	See datatype definition in clause 5.2.7	0..N	GeoProperties of the GeoProperty.
	LanguageProperty or LanguageProperty[] (see note 1)	See datatype definition in clause 5.2.32	0..N	LanguageProperties of the GeoProperty.
	JsonProperty or JsonProperty[] (see note 1)	See datatype definition in clause 5.2.38	0..N	JsonProperties of the GeoProperty.
	VocabProperty or VocabProperty[] (see note 1)	See datatype definition in clause 5.2.35	0..N	VocabProperties of the GeoProperty.
	ListProperty or ListProperty[] (see note 1)	See datatype definition in clause 5.2.36	0..N	ListProperties of the GeoProperty.
<Relationship name>	Relationship or Relationship[] (see note 2)	See datatype definition in clause 5.2.6	0..N	Relationships of the GeoProperty.
	ListRelationship or ListRelationship[] (see note 2)	See datatype definition in clause 5.2.37	0..N	ListRelationships of the GeoProperty.
NOTE 1: For each <i>Property</i> (or subclass of <i>Property</i>) identified by the same Property name, there can be one or more instances separated by <i>datasetId</i> .				
NOTE 2: For each <i>Relationship</i> (or subclass of <i>Relationship</i>) identified by the same Relationship name, there can be one or more instances separated by <i>datasetId</i> .				

The following **output only** members (defined by Table 5.2.7-2) of the *GeoProperty* data structure are also defined. They are read-only and shall be generated by NGSI-LD implementations. They shall not be provided by **Context Producers**. In the event that they are provided (in update or create operations) NGSI-LD implementations shall ignore them.

Table 5.2.7-2: Output only members of the NGSI-LD GeoProperty data type

Name	Data Type	Restrictions	Cardinality	Description
createdAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	System generated creation timestamp. See clause 4.8.
deletedAt	String	<i>DateTime</i> (clause 4.6.3) It is only used in notifications reporting deletions	0..1	System generated deletion timestamp. See clause 4.8.
instanceId	String	Valid URI. Only used in temporal representation of GeoProperties	0..1	URI uniquely identifying a GeoProperty instance as mandated by clause 4.5.7.
modifiedAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	System generated last modification timestamp. See clause 4.8.
previousValue	Any JSON value as defined by IETF RFC 8259 [6]	Only used in Notifications , if the <i>showChanges</i> option is explicitly requested	0..1	Previous GeoProperty Value.

5.2.8 EntityInfo

This type represents what Entities, Entity Types or group of Entity IDs (as a regular expression pattern mandated by IEEE 1003.2™ [11]) can be provided (by **Context Sources**).

The JSON members shall follow the indications provided in Table 5.2.8-1. *id* takes precedence over *idPattern*.

Notice that Cardinality of *type* being 1 implies that it is not possible to register what Entities can be provided by a **Context Source** just by their *id* or *idPattern* (i.e. without specifying their *type*).

Table 5.2.8-1: EntityInfo data type definition

Name	Data Type	Restrictions	Cardinality	Description
type	String or String[]	Fully Qualified Name of an Entity Type or the Entity Type name as a short-hand string. See clause 4.6.2	1	Entity Type (or JSON array, in case of Entities with multiple Entity Types).
id	String	Valid URI	0..1	Entity identifier.
idPattern	String	Regular expression as per IEEE 1003.2™ [11]	0..1	A regular expression which denotes a pattern that shall be matched by the provided or subscribed Entities.

5.2.9 CSourceRegistration

This type represents the data needed to register a new **Context Source**.

The supported JSON members shall follow the indications provided in Table 5.2.9-1.

Table 5.2.9-1: CSourceRegistration data type definition

Name	Data Type	Restriction	Cardinality	Description
id	String	Valid URI. Unique registration identifier. (JSON-LD <i>@id</i>).	0..1	Generated at creation time, if it is not provided, it will be assigned during registration process and returned to client. It cannot be later modified in update operations.
type	String	It shall be equal to "ContextSourceRegistration"	1	JSON-LD <i>@type</i> Use reserved type for identifying Context Source Registration.
endpoint	String	It shall be a dereferenceable URI	1	Endpoint expressed as dereferenceable URI through which the Context Source exposes its NGSI-LD interface.
contextSourceInfo	KeyValuePair[]		0..1	Generic {key, value} array to convey optional information to provide when contacting the registered Context Source.
information	RegistrationInfo[]	See data type definition in clause 5.2.10. Empty array (0 length) is not allowed	1	Describes the Entities, Properties and Relationships for which the Context Source may be able to provide information.
contextSourceAlias	String	Non-empty string. Pseudonym field as defined in IETF RFC 7230 [27]	0..1	A previously retrieved unique id for a registered Context Source which is used to identify loops. In the multi-tenancy use case (see clause 4.14), this id shall be used to identify a specific Tenant within a registered Context Source.
description	String	Non-empty string	0..1	A description of this Context Source Registration.
datasetId	String[]	Valid URIs, " <i>@none</i> " for including the default Attribute instances.	0..1	Specifies the datasetIds of Attributes that the Context Source can provide, defined as per clause 4.5.5.
expiresAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	Provides an expiration date. When passed the Context Source Registration will become invalid and the Context Source might no longer be available.
location	GeoJSON Geometry as mandated by clause 4.7		0..1	Location for which the Context Source may be able to provide information.

Name	Data Type	Restriction	Cardinality	Description
management	Registration Management Info	See data type definition in clause 5.2.34	0..1	Holds additional optional registration management information that can be used to limit unnecessary distributed operation requests.
managementInterval	TimeInterval	See data type definition in clause 5.2.11	0..1	If present, the Context Source can be queried for Temporal Entity Representations. (If latest Entity information is also provided, a separate Context Registration is needed for this purpose). The <i>managementInterval</i> specifies the time interval for which the Context Source can provide Entity information as specified by the <i>createdAt</i> , <i>modifiedAt</i> and <i>deletedAt</i> Temporal Properties. A temporal query based on the <i>createdAt</i> , <i>modifiedAt</i> or <i>deletedAt</i> Temporal Property is matched against the <i>managementInterval</i> for overlap.
mode	String	It shall be one of: "inclusive", "exclusive", "redirect" or "auxiliary" The mode is assumed to be "inclusive" if not explicitly defined	0..1	The definition of the mode of distributed operation (see clause 4.3.6) supported by the registered Context Source.
observationInterval	TimeInterval	See data type definition in clause 5.2.11	0..1	If present, the Context Source can be queried for Temporal Entity Representations. (If latest Entity information is also provided, a separate Context Registration is needed for this purpose). The <i>observationInterval</i> specifies the time interval for which the Context Source can provide Entity information as specified by the <i>observedAt</i> Temporal Property. A temporal query based on the <i>observedAt</i> Temporal Property, which is the default, is matched against the <i>observationInterval</i> for overlap.

Name	Data Type	Restriction	Cardinality	Description
observationSpace	GeoJSON Geometry as mandated by clause 4.7		0..1	Geographic location that includes the observation spaces of all entities as specified by their respective <i>observationSpace GeoProperty</i> for which the Context Source may be able to provide information.
operations	String[]	Entries are limited to the named API operations and named operation groups (see clause 4.20)	0..1	The definition limited subset of API operations supported by the registered Context Source . If undefined, the default set of operations is "federationOps" (see clause 4.20).
operationSpace	GeoJSON Geometry as mandated by clause 4.7		0..1	Geographic location that includes the operation spaces of all entities as specified by their respective <i>operationSpace GeoProperty</i> for which the Context Source may be able to provide information.
refreshRate	String	String representing a duration in ISO 8601 [17] format	0..1	An indication of the likely period of time to elapse between updates at this registered endpoint. Brokers may optionally use this information to help implement caching.
registrationName	String	Non-empty string	0..1	A name given to this Context Source Registration.
scope	String or String[]	Scope(s)	0..1	Scopes (see clause 4.18) for which the Context Source has Entities.
tenant	String		0..1	Identifies the Tenant that has to be specified in all requests to the Context Source that are related to the information registered in this Context Source Registration. If not present, the default Tenant is assumed. Should only be present in systems supporting multi-tenancy.

The members (defined by Table 5.2.9-2) of the *CSourceRegistration* data structure are also defined. They are read-only and shall be automatically generated by NGSI-LD implementations. In the event that they are provided (in update or create operations) NGSI-LD implementations shall ignore them.

Table 5.2.9-2: Additional members of the CSourceRegistration data type

Name	Data Type	Restrictions	Cardinality	Description
lastFailure	String	<i>DateTime</i> (clause 4.6.3)	0..1	Timestamp corresponding to the instant when the last distributed operation resulting in a failure (for instance, in the HTTP binding, an HTTP response code other than 2xx) was returned.
status	String	Allowed values: "ok" "failed"	0..1	Read-only., Status of the Registration. It shall be "ok" if the last attempt to perform a distributed operation succeeded. It shall be "failed" if the last attempt to perform a distributed operation failed.
timesFailed	Number	0 or greater value	0..1	Number of times that the registration triggered a distributed operation request that failed.
timesSent	Number	0 or greater value	0..1	Number of times that the registration triggered a distributed operation, including failed attempts.
lastSuccess	String	<i>DateTime</i> (clause 4.6.3)	0..1	Timestamp corresponding to the instant when the last successfully distributed operation was sent. Created on first successful operation.

5.2.10 RegistrationInfo

The supported JSON members shall follow the requirements provided in Table 5.2.10-1.

Table 5.2.10-1: RegistrationInfo data type definition

Name	Data Type	Restrictions	Cardinality	Description
entities	EntityInfo[]	See data type definition in clause 5.2.8. Empty array (0 length) is not allowed. Restrictions in clause 4.3.6 apply as well	0..1	Describes the entities for which the CSource may be able to provide information.
propertyNames	String[]	Property names as short hand strings or URIs. Empty array is not allowed. Restrictions in clause 4.3.6 apply as well	0..1	Describes the Properties that the CSource may be able to provide.
relationshipNames	String[]	Relationship names as short hand strings or URIs. Empty array is not allowed. Restrictions in clause 4.3.6 apply as well	0..1	Describes the Relationships that the CSource may be able to provide.

At least one element of *RegistrationInfo* shall be present.

5.2.11 TimeInterval

The supported JSON members shall follow the requirements provided in Table 5.2.11-1.

Table 5.2.11-1: TimeInterval data type definition

Name	Data Type	Restrictions	Cardinality	Description
startAt	String	<i>DateTime</i> (clause 4.6.3)	1	Describes the start of the time interval
endAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	Describes the end of the time interval. If not present the interval is open

5.2.12 Subscription

This datatype represents a Context Subscription.

The supported JSON members shall follow the requirements provided in Table 5.2.12-1.

Table 5.2.12-1: Subscription data type definition

Name	Data Type	Restrictions	Cardinality	Description
id	String	Valid URI	0..1	Subscription identifier (JSON-LD <i>@id</i>). Generated at creation time, if it is not provided, it will be assigned during subscription process and returned to client. It cannot be later modified in update operations.
type	String	It shall be equal to "Subscription"	1	JSON-LD <i>@type</i> .
entities	EntitySelector[]	See data type definition in clause 5.2.33. Empty array (0 length) is not allowed. Mandatory if <i>timeInterval</i> is present, unless the execution of the request is limited to local scope (see clause 5.5.13)	0..1	Entities subscribed.
notification	NotificationParameters	See data type definition in clause 5.2.14	1	Notification details.
notificationTrigger	String[]	Valid notification triggers are "entityCreated", "entityUpdated", "entityDeleted", "attributeCreated", "attributeUpdated", "attributeDeleted"	0..1	The notification triggers listed indicate what kind of changes shall trigger a notification. If not present, the default is the combination "attributeCreated" and "attributeUpdated". "entityUpdated" is equivalent to the combination "attributeCreated", "attributeUpdated" and "attributeDeleted".
watchedAttributes	String[]	Attribute name as short hand strings or URIs. Empty array (0 length) is not allowed. if <i>timeInterval</i> is present it shall not appear (0 cardinality)	0..1	Watched Attributes (Properties or Relationships). If not defined it means any Attribute.
q	String	A valid query string as per clause 4.9	0..1	Query that shall be met by subscribed entities in order to trigger the notification.
geoQ	GeoQuery	See data type definition in clause 5.2.13	0..1	Geoquery that shall be met by subscribed entities in order to trigger the notification.
scopeQ	String	See clause 4.19	0..1	Scope query.

Name	Data Type	Restrictions	Cardinality	Description
temporalQ	TemporalQuery	See data type definition in clause 5.2.21	0..1	Temporal Query to be used only in <i>Context Registration Subscriptions</i> for matching Context Source Registrations of Context Sources providing temporal information.
csf	String	A valid query string as per clause 4.9	0..1	Context source filter that shall be met by Context Source Registrations describing Context Sources to be used for Entity Subscriptions.
datasetId	String[]	Valid URIs, " <i>@none</i> " for including the default Attribute instances.	0..1	Specifies the datasetIds of the Attribute instances to be selected for each matched Attribute as per clause 4.5.5.
description	String		0..1	Subscription description.
expandValues	String	Comma separated list of attribute names	0..1	Values of the identified attributes should be expanded against the supplied <i>@context</i> using JSON-LD type coercion prior to executing the query.
expiresAt	String	<i>DateTime</i> (see clause 4.6.3)	0..1	Expiration date for the subscription.
isActive	Boolean	<i>true</i> by default	0..1	Allows clients to temporarily pause the subscription by making it inactive. <i>true</i> indicates that the Subscription is under operation. <i>false</i> indicates that the subscription is paused, and notifications shall not be delivered.
jsonKeys	String	Comma separate list of attribute names	0..1	Values of the identified attributes are to be considered uninterpretable as JSON-LD and should not be expanded against the supplied <i>@context</i> using JSON-LD type coercion prior to executing the query.
jsonldContext	String	Dereferenceable URI	0..1	The dereferenceable URI of the JSON-LD <i>@context</i> to be used when sending a notification resulting from the subscription. If not provided, the <i>@context</i> used for the subscription shall be used as a default.
lang	String	A natural language filter in the form of a IETF RFC 5646 [28] language code	0..1	Language filter to be applied to the query (clause 4.15).
localOnly	Boolean		0..1	If <i>localOnly=true</i> then the subscription only pertains to the Entities stored locally. In case the Subscription pertains to a Snapshot it is always local, regardless of whether <i>localOnly</i> is set to <i>true</i> or not.
ngsildConformance	String	A semantically versioned string in the form <i>major.minor</i> , which conforms to a version of the NGSI-LD specification	0..1	If provided the notification shall undergo a backwards compatibility operation as defined by clause 4.3.6.8 and be amended to conform to the supplied version of the NGSI-LD specification.

Name	Data Type	Restrictions	Cardinality	Description
splitEntities	Boolean	default decided by implementation; it should be configurable. The parameter does not apply in case <i>localOnly</i> is <i>true</i> .	0..1	If <i>true</i> it is assumed that single Entities are distributed between different Context Brokers and/or Context Sources and this has to be taken into account when applying any kind of filters (q, geoQ, scopeQ, Attributes etc.). If <i>false</i> it is expected that Context Broker and/or Context Source always have complete Entities, which allows applying filters locally.
subscriptionName	String		0..1	A (short) name given to this Subscription.
throttling	Number	Greater than 0. Fractional values are allowed. If <i>timeInterval</i> is present it shall not appear (0 cardinality)	0..1	Minimal period of time in seconds which shall elapse between two consecutive notifications.
timeInterval	Number	Greater than 0 if <i>watchedAttributes</i> is present it shall not appear (0 cardinality)	0..1	Indicates that a notification shall be delivered periodically regardless of attribute changes. Actually, when the time interval (in seconds) specified in this value field is reached.

At least one of (a) *entities* or (b) *watchedAttributes* shall be present, unless the member *localOnly* is set to *true*, in which case the execution of the request is limited to local scope (see clause 5.5.13).

The members (defined by Table 5.2.12-2) of the *Subscription* data structure are also defined. They are read-only and shall be automatically generated by NGSI-LD implementations. They shall not be provided by Context Subscribers. In the event that they are provided (in update or create operations) NGSI-LD implementations shall ignore them.

Table 5.2.12-2: Additional members of the Subscription data type

Name	Data Type	Restrictions	Cardinality	Description
status	String	Allowed values: "active" "paused" "expired"	0..1	Read-only. Provided by the system when querying the details of a subscription.

5.2.13 GeoQuery

This datatype represents a geoquery used for Subscriptions.

The supported JSON members shall follow the requirements provided in Table 5.2.13-1.

Table 5.2.13-1: GeoQuery data type definition

Name	Data Type	Restrictions	Cardinality	Description
coordinates	JSON Array or String	A JSON Array coherent with the geometry type as per IETF RFC 7946 [8]	1	Coordinates of the reference geometry. For the sake of JSON-LD compatibility It can be encoded as a string as described in clause 4.7.1.
geometry	String	A valid GeoJSON [8] geometry type excepting <i>GeometryCollection</i>	1	Type of the reference geometry.
georel	String	A valid geo-relationship as defined by clause 4.10	1	Geo-relationship ("near", "within", etc.).
geoproperty	String	Attribute name as short hand string or URI	0..1	Specifies the GeoProperty to which the GeoQuery is to be applied. If not present, the default GeoProperty is <i>location</i> .

5.2.14 NotificationParams

5.2.14.1 NotificationParams data type definition

This datatype represents the parameters that allow to convey the details of a notification.

The supported JSON members shall follow the requirements provided in Table 5.2.14.1-1.

Table 5.2.14.1-1: NotificationParams data type definition

Name	Data Type	Restrictions	Cardinality	Description
endpoint	Endpoint	See data type definition in clause 5.2.15	1	Notification endpoint details.
attributes	String[]	Attribute name as short hand strings or URIs. Empty array (0 length) is not allowed	0..1	A synonym for <i>pick</i> , except that <i>id</i> , <i>type</i> , <i>scope</i> are not allowed. Deprecated Attribute names to be included in the notification payload body. If undefined it will mean all Attributes.
format	String	It shall be one of: "normalized", "concise", "simplified" (or its synonym "keyValues")	0..1	Conveys the representation format of the entities delivered at notification time. By default, it will be in the normalized format.
join	String	It shall be one of: "flat", "inline", "@none"	0..1	String representing the type of Linked Entity retrieval to apply. By default, it will be "@none".
joinLevel	Number	Positive Integer	0..1	Depth of Linked Entity retrieval to apply. Default is 1. Only applicable if <i>join</i> parameter is "flat", or "inline".
omit	String[]	Entity member ("id", "type", "scope" or a projected Attribute name) as a valid attribute projection language string as per clause 4.21. Empty array (0 length) is not allowed	0..1	When defined, the specified Entity members are removed from each Entity within the payload.
pick	String[]	Entity member ("id", "type", "scope" or a projected Attribute name as a valid attribute projection language string as per clause 4.21). Empty array (0 length) is not allowed	0..1	When defined, every Entity within the payload body is reduced down to only contain the specified Entity members.

Name	Data Type	Restrictions	Cardinality	Description
showChanges	Boolean	false by default	0..1	If <i>true</i> the previous <i>value</i> (<i>previousValue</i>) of Properties or <i>languageMap</i> (<i>previousLanguageMap</i>) of Language Properties or <i>object</i> (<i>previousObject</i>) of Relationships is provided in addition to the current one. This requires that it exists, i.e. in case of modifications and deletions, but not in the case of creations. <i>showChanges</i> cannot be <i>true</i> in case format is "keyValues".
sysAttrs	Boolean	false by default	0..1	If <i>true</i> , the system generated attributes <i>createdAt</i> and <i>modifiedAt</i> and the system attribute <i>expiresAt</i> are included in the response payload body, in the case of a deletion also <i>deletedAt</i> .

5.2.14.2 Output only members

The following **output only** members (defined by Table 5.2.14.2-1) of the *NotificationParams* data structure are also defined. They are read-only and shall be automatically generated by NGSI-LD implementations. They shall not be provided by *Context Subscribers*. In the event that they are provided (in update or create operations) NGSI-LD implementations shall ignore them.

In query or retrieve operations involving Subscriptions, implementations shall generate them as part of their representation.

Table 5.2.14.2-1: Output only members of the NotificationParams data structure

Name	Data Type	Restrictions	Cardinality	Description
lastFailure	String	<i>DateTime</i> (clause 4.6.3)	0..1	Timestamp corresponding to the instant when the last notification resulting in failure (for instance, in the HTTP binding, an HTTP response code different than 200) has been sent. Provided by the system when querying the details of a subscription.
lastNotification	String	<i>DateTime</i> (clause 4.6.3)	0..1	Timestamp corresponding to the instant when the last notification has been sent. Provided by the system when querying the details of a subscription.
lastSuccess	String	<i>DateTime</i> (clause 4.6.3)	0..1	Timestamp corresponding to the instant when the last successful (200 OK response) notification has been sent. Provided by the system when querying the details of a subscription.
status	String	Allowed values: "ok", "failed"	0..1	Status of the Notification. It shall be "ok" if the last attempt to notify the subscriber succeeded. It shall be "failed" if the last attempt to notify the subscriber failed.
timesFailed	Number	Greater than 0	0..1	Number of times an unsuccessful response (or timeout) has been received when delivering the notification. Provided by the system when querying the details of a subscription.
timesSent	Number	Greater than 0	0..1	Number of times that the notification has been sent. Provided by the system when querying the details of a subscription.

5.2.15 Endpoint

This datatype represents the parameters that are required in order to define an endpoint for notifications. This can include, in addition the endpoint's URI, a generic{key, value} array, named *receiverInfo*, which contains, in a generalized form, whatever extra information the **Context Broker** shall convey to the receiver in order for the **Context Broker** to successfully communicate with receiver (e.g. Authorization material), or for the receiver to correctly interpret the received content (e.g. the Link URL to fetch an *@context*). Additionally, it can include another generic{key, value} array, named *notifierInfo*, which contains the configuration that the **Context Broker** needs to know in order to correctly set up the communication channel towards the receiver (e.g. MQTT-Version, MQTT-QoS, in case of MQTT binding, as defined in clause 7.2).

The supported JSON members shall follow the indications provided in Table 5.2.15-1.

Table 5.2.15-1: Endpoint data type definition

Name	Data Type	Restrictions	Cardinality	Description
uri	String	Dereferenceable URI	1	URI which conveys the endpoint which will receive the notification.
accept	String	MIME type. It shall be one of: "application/json" "application/ld+json" "application/geo+json"	0..1	Intended to convey the MIME type of the notification payload body (JSON, or JSON-LD, or GeoJSON). If not present, default is "application/json".
cooldown	Number	Greater than 0	0..1	Once a failure has occurred, minimum period of time in milliseconds which shall elapse before attempting to make a subsequent notification to the same endpoint after failure. If requests are received before the cooldown period has expired, no notification is sent.
notifierInfo	KeyValuePair[]		0..1	Generic {key, value} array to set up the communication channel.
receiverInfo	KeyValuePair[]		0..1	Generic {key, value} array to convey optional information to the receiver.
timeout	Number	Greater than 0	0..1	Maximum period of time in milliseconds which may elapse before a notification is assumed to have failed. The NGSI-LD system can override this value. This only applies if the binding protocol always returns a response.

5.2.16 BatchOperationResult

This datatype represents the result of a batch operation.

The supported JSON members shall follow the indications provided in Table 5.2.16-1.

Table 5.2.16-1: BatchOperationResult data type definition

Name	Data Type	Restrictions	Cardinality	Description
errors	BatchEntityError[]		1	One array item per Element in error. Empty Array if no errors happened.
success	String[]	Array of valid URIs	1	Array of Entity IDs corresponding to the Elements that were successfully treated by the concerned operation. Empty Array if no Element was successfully treated.

5.2.17 BatchEntityError

This datatype represents an error raised (associated to a particular Entity) during the execution of a batch or distributed operation.

The supported JSON members shall follow the indications provided in Table 5.2.17-1.

Table 5.2.17-1: BatchEntityError data type definition

Name	Data Type	Restrictions	Cardinality	Description
entityId	String	Valid URI	1	Entity ID corresponding to the Entity in error.
error	ProblemDetails (see IETF RFC 7807 [10])		1	One instance per Entity in error.
registrationId	String	Valid URI	0..1	Registration Id corresponding to a failed distributed operation (optional).

5.2.18 UpdateResult

This datatype represents the result of Attribute update (append or update) operations in the NGSI-LD API regardless of whether local or distributed.

The supported JSON members shall follow the indications provided in Table 5.2.18-1.

Table 5.2.18-1: UpdateResult data type definition

Name	Data Type	Restrictions	Cardinality	Description
notUpdated	NotUpdatedDetails[]	See clause 5.2.19	1	List which contains the Attributes (represented by their name) that were not updated, together with the reason for not being updated.
updated	String[]		1	List of Attributes (represented by their name) that were appended or updated.

5.2.19 NotUpdatedDetails

This datatype represents additional information provided by an implementation when an Attribute update did not happen. See also clause 5.2.18.

The supported JSON members shall follow the indications provided in Table 5.2.19-1.

Table 5.2.19-1: NotUpdatedDetails data type definition

Name	Data Type	Restrictions	Cardinality	Description
attributeName	String		1	Attribute name.
reason	String		1	Reason for not having changed such Attribute.
registrationId	String	Valid URI	0..1	Registration Id corresponding to a failed distributed operation (optional).

5.2.20 EntityTemporal

This datatype represents the Temporal Evolution of an Entity.

This is the same datatype as mandated by clause 5.2.4 with the only deviation that the representation of Properties and Relationships shall be the temporal one, as defined in clauses 4.5.7 and 4.5.8. Alternatively it is possible to specify the EntityTemporal by using the "Simplified temporal representation of an Entity", as defined in clause 4.5.9.

5.2.21 TemporalQuery

This datatype represents a temporal query.

The supported JSON members shall follow the requirements provided in Table 5.2.21-1.

Table 5.2.21-1: TemporalQuery data type definition

Name	Data Type	Restrictions	Cardinality	Description
timeAt	String representing the <i>timeAt</i> parameter as defined by clause 4.11	It shall be a <i>DateTime</i>	1	
timerel	String	Allowed values: "before", "after" and "between"	1	Represents the temporal relationship as defined by clause 4.11.
aggrMethods	Comma separated list of string	It shall be 1 if <i>aggregatedValues</i> is present in the <i>options</i> parameter	0..1	Each String represents an aggregation method, as defined by clause 4.5.19. Only applicable if "aggregatedValues" is present in the <i>format</i> or <i>options</i> parameter.
aggrPeriodDuration	String		0..1	It represents the duration of each period used for the aggregation as defined by clause 4.5.19. If not specified, it defaults to a duration of 0 seconds and is interpreted as a duration spanning the whole time-range specified by the temporal query. Only applicable if "aggregatedValues" is present in the <i>format</i> or <i>options</i> parameter.
endTimeAt	String representing the <i>endTimeAt</i> parameter as defined by clause 4.11	It shall be a <i>DateTime</i> . Cardinality shall be 1 if <i>timerel</i> is equal to "between"	0..1	
lastN	Positive integer		0..1	Only the last n instances, per Attribute, per Entity (under the specified time interval) shall be retrieved.
timeproperty	String representing a Temporal Property name	Allowed values: "observedAt", "createdAt", "modifiedAt" and "deletedAt". If not specified, the default is "observedAt". (See clause 4.8)	0..1	

5.2.22 KeyValuePair

This datatype represents the optional information that is required when contacting an endpoint for notifications.

The supported members shall follow the indications provided in Table 5.2.22-1. They are intended to represent a key/value pair.

Example optional information includes additional HTTP Headers such as:

- The HTTP Authentication Header.
- The HTTP Prefer Header (IETF RFC 7240 [26]) used when notifying the GeoJSON Endpoint.

Table 5.2.22-1: KeyValuePair data type definition

Name	Data Type	Restrictions	Cardinality	Description
key	String	Binding-dependent	1	The key of the key/value pair.
value	String	Binding-dependent	1	The value of the key/value pair.

5.2.23 Query

This datatype represents the information that is required in order to convey a query when a "Query Entities" operation or a "Query Temporal Evolution of Entities" operation is to be performed (as per clauses 5.7.2 and 5.7.4, respectively).

The supported JSON members shall follow the requirements provided in Table 5.2.23-1.

Table 5.2.23-1: Query data type definition

Name	Data Type	Restrictions	Cardinality	Description
type	String	It shall be equal to "Query"	1	JSON-LD @type
entities	EntitySelector[]	See data type definition in clause 5.2.33. Empty array (0 length) is not allowed	0..1	Entity IDs, id pattern and Entity types that shall be matched by Entities in order to be retrieved.
q	String	A valid query string as per clause 4.9	0..1	Query that shall be matched by Entities in order to be retrieved.
geoQ	GeoQuery	See data type definition in clause 5.2.13	0..1	Geoquery that shall be matched by Entities in order to be retrieved.
scopeQ	String	See clause 4.19	0..1	Scope query.
temporalQ	TemporalQuery	See data type definition in clause 5.2.21	0..1	Temporal Query to be present only for "Query Temporal Evolution of Entities" operation (clause 5.7.4).
attrs	String[]	Attribute name as short hand strings or URIs. Empty array (0 length) is not allowed	0..1	A synonym for a combination of the <i>pick</i> and <i>q</i> members. Deprecated List of Attributes that shall be matched by Entities in order to be retrieved. If not present all Attributes will be retrieved.
omit	String[]	Entity member ("id", "type", "scope" or a projected Attribute name) as a valid attribute projection language string as per clause 4.21. Empty array (0 length) is not allowed	0..1	When defined, the specified Entity members are removed from each Entity within the payload.

Name	Data Type	Restrictions	Cardinality	Description
pick	String[]	Entity member ("id", "type", "scope" or a projected Attribute name) as a valid attribute projection language string as per clause 4.21. Empty array (0 length) is not allowed	0..1	When defined, every Entity within the payload body is reduced down to only contain the specified Entity members.
aggrParams	AggregationParams	See data type definition in clause 5.2.44.	0..1	Aggregation parameters required for supporting aggregation methods in to be present only for "Query Temporal Evolution of Entities" operation (clause 5.7.4). Only applicable if "aggregatedValues" is present in the <i>format</i> or <i>options</i> parameter.
csf	String	A valid query string as per clause 4.9	0..1	Context source filter that shall be matched by Context Source Registrations describing Context Sources to be used for retrieving Entities.
containedBy	String[]	Comma separated list of URIs. Empty array (0 length) is not allowed	0..1	List of entity ids which have previously been encountered whilst retrieving the Entity Graph. Only applicable if <i>joinLevel</i> is present. Only applicable for the "Retrieve Entity" (clause 5.7.1) and "Query Entities" operations (clause 5.7.2).
datasetId	String[]	Valid URIs, "@none" for including the default Attribute instances.	0..1	Specifies the datasetIds of the Attribute instances to be selected for each matched Attribute as per clause 4.5.5.
expandValues	String	Comma separated list of attribute names	0..1	Values of the identified attributes should be expanded against the supplied <i>@context</i> using JSON-LD type coercion prior to executing the query.
entityMap	Boolean		0..1	If <i>true</i> , the location of the EntityMap used in the operation is returned in the response.
entityMapLifetime	String	String representing a duration in ISO 8601 [17] format	0..1	Suggested duration for which the requester wants the EntityMap to exist. The actual expiresAt time of the EntityMap shall be set by the Context Broker or Context Source, possibly overriding the requested duration. Only applicable if entityMap is set to <i>true</i> .
jsonKeys	String	Comma separate list of attribute names	0..1	Values of the identified attributes are to be considered uninterpretable as JSON-LD and should not be expanded against the supplied <i>@context</i> using JSON-LD type coercion prior to executing the query.

Name	Data Type	Restrictions	Cardinality	Description
join	String	It shall be one of: "flat", "inline", "@none"	0..1	String representing the type of Linked Entity retrieval to apply. By default, it will be "@none".
joinLevel	Number	Positive Integer	0..1	Depth of Linked Entity retrieval to apply. Default is 1. Only applicable if <i>join</i> parameter is "flat", or "inline".
lang	String	A natural language filter in the form of a IETF RFC 5646 [28] language code	0..1	Language filter to be applied to the query (clause 4.15).
ordering	OrderingParams	See data type definition in clause 5.2.43	0..1	When defined, the Entities returned in the payload shall be ordered according to the members defined. This only applies if the operation is limited to the local scope.
splitEntities	Boolean	default decided by implementation; it should be configurable. The parameter does not apply in case the parameter <i>local</i> is <i>true</i> or the query applies to a Snapshot	0..1	If <i>true</i> it is assumed that single Entities are distributed between different Context Brokers and/or Context Sources and this has to be taken into account when applying any kind of filters (q, geoQ, scopeQ, Attributes etc.). If <i>false</i> it is expected that Context Broker and/or Context Source always have complete Entities, which allows applying filters locally.

5.2.24 EntityTypeList

This type represents the data needed to define the entity type list representation as mandated by clause 4.5.10.

The supported JSON members shall follow the requirements provided in Table 5.2.24-1.

Table 5.2.24-1: NGS-LD EntityTypeList data type definition

Name	Data Type	Restriction	Cardinality	Description
id	String	Valid URI	1	URI that is unique within the system scope. Identifier for the entity type list.
type	String	It shall be equal to "EntityTypeList"	1	JSON-LD <i>@type</i> .
typeList	String[]		1	List containing the entity type names.

5.2.25 EntityType

This type represents the data needed to define the elements of the detailed entity type list representation as mandated by clause 4.5.11.

The supported JSON members shall follow the requirements provided in Table 5.2.25-1.

Table 5.2.25-1: NGSI-LD EntityType data type definition

Name	Data Type	Restriction	Cardinality	Description
id	String	Valid URI	1	Fully Qualified Name (FQN) of the entity type being described.
type	String	It shall be equal to "EntityType"	1	JSON-LD <i>@type</i> .
attributeNames	String[]		1	List containing the names of attributes that instances of the entity type can have.
typeName	String		1	Name of the entity type, short name if contained in <i>@context</i> .

5.2.26 EntityTypeInfo

This type represents the data needed to define the detailed entity type information representation as mandated by clause 4.5.12.

The supported JSON members shall follow the requirements provided in Table 5.2.26-1.

Table 5.2.26-1: NGSI-LD EntityTypeInfo data type definition

Name	Data Type	Restriction	Cardinality	Description
id	String	Valid URI	1	Fully Qualified Name (FQN) of the entity type being described.
type	String	It shall be equal to "EntityTypeInfo"	1	JSON-LD <i>@type</i> .
attributeDetails	Attribute[]	See data type definition in clause 5.2.28. Attribute with only the elements "id", "type", "attributeName" and "attributeTypes"	1	List of attributes that entity instances with the specified entity type can have.
entityCount	Number	Unsigned integer	1	Number of entity instances of this entity type.
typeName	String		1	Name of the entity type, short name if contained in <i>@context</i> .

5.2.27 AttributeList

This type represents the data needed to define the attribute list representation as mandated by clause 4.5.13.

The supported JSON members shall follow the requirements provided in Table 5.2.27-1.

Table 5.2.27-1: NGSI-LD AttributeList data type definition

Name	Data Type	Restriction	Cardinality	Description
id	String	Valid URI	1	URI that is unique within the system scope. Identifier for the attribute list.
type	String	It shall be equal to "AttributeList"	1	JSON-LD <i>@type</i> .
attributeList	String[]		1	List containing the attribute names.

5.2.28 Attribute

This type represents the data needed to define the attribute information needed as:

- part of the entity type information representation as mandated by clause 4.5.12;
- the detailed attribute list representation as mandated by clause 4.5.14;

- the attribute information representation as mandated by clause 4.5.15.

The supported JSON members shall follow the requirements provided in Table 5.2.28-1.

Table 5.2.28-1: NGSI-LD Attribute data type definition

Name	Data Type	Restriction	Cardinality	Description
id	String	Valid URI	1	Full URI of attribute name.
type	String	It shall be equal to "Attribute"	1	JSON-LD <i>@type</i> .
attributeName	String		1	Name of the attribute, short name if contained in <i>@context</i> .
attributeCount	Number	Unsigned integer	0..1	Number of attribute instances with this attribute name.
attributeTypes	String[]		0..1	List of attribute types (e.g. <i>Property</i> , <i>Relationship</i> , <i>GeoProperty</i>) for which entity instances exist, which contain an attribute with this name.
typeNameNames	String[]		0..1	List of entity type names for which entity instances exist containing attributes that have the respective name.

5.2.29 Feature

This data type represents a spatially bounded Entity in GeoJSON format, as mandated by IETF RFC 7946 [8]. The supported JSON members shall follow the requirements provided in Table 5.2.29-1.

Table 5.2.29-1: Feature data type definition

Name	Data Type	Restriction	Cardinality	Description
id	String	Valid URI	1	Entity ID.
type	String	It shall be equal to "Feature"	1	GeoJSON Type.
geometry	GeoJSON Object	The value field from the matching <i>GeoProperty</i> (as specified in clause 4.5.16) or <i>null</i>	1	<i>null</i> if no matching <i>GeoProperty</i> .
properties	FeatureProperties	See data type definition	1	List of attributes as mandated by clause 5.2.31.
@context	URI, JSON Object, or JSON Array	See [2], section 5.1.	0..1	JSON-LD <i>@context</i> . This field is only present if requested in the payload by the HTTP Prefer Header (IETF RFC 7240 [26]).

5.2.30 FeatureCollection

This data type represents a list of spatially bounded Entities in GeoJSON format, as mandated by IETF RFC 7946 [8]. The supported JSON members shall follow the requirements provided in Table 5.2.30-1.

Table 5.2.30-1: FeatureCollection data type definition

Name	Data Type	Restriction	Cardinality	Description
type	String	It shall be equal to "FeatureCollection"	1	GeoJSON Type.
features	Feature[]	See data type definition	1..N	In the case that no matches are found, <i>features</i> will be an empty array.
@context	URI, JSON Object, or JSON Array	See [2], section 5.1.	0..1	JSON-LD <i>@context</i> . This field is only present if requested in the payload by the HTTP Prefer Header (IETF RFC 7240 [26]).

5.2.31 FeatureProperties

This data type represents the type and the associated attributes (*Properties* and *Relationships*) of an Entity in GeoJSON format.

Table 5.2.31-1: FeatureProperties data type definition

Name	Data Type	Restriction	Cardinality	Description
type	String or String[]	Entity Type	1	Entity Type (or JSON array, in case of Entities with multiple Entity Types). Both short hand string (type name) or URI are allowed.
<Property name>	Property or Property[], see note 1	See data type definition in clause 5.2.5	0..N	Property as mandated by clause 4.5.2.
	GeoProperty or GeoProperty[], see note 1	See datatype definition in clause 5.2.7	0..N	GeoProperty as mandated by clause 4.5.2.
	LanguageProperty or LanguageProperty[], see note 1	See datatype definition in clause 5.2.32	0..N	LanguageProperty as mandated by clause 4.5.18.
	JsonProperty or JsonProperty[] see note 1	See datatype definition in clause 5.2.38	0..N	JsonProperties as mandated by clause 4.5.24.
	VocabProperty or VocabProperty[] see note 1	See datatype definition in clause 5.2.35	0..N	VocabProperty as mandated by clause 4.5.20.
	ListProperty or ListProperty[], see note 1	See datatype definition in clause 5.2.36	0..N	ListProperty as mandated by clause 4.5.21.
<Relationship name>	Relationship or Relationship[], see note 2	See data type definition in clause 5.2.6	0..N	Relationship as mandated by clause 4.5.3.
	ListRelationship or ListRelationship[], see note 2	See datatype definition in clause 5.2.37	0..N	ListRelationship as mandated by clause 4.5.22.
NOTE 1: For each <i>Property</i> (or subclass of <i>Property</i>) identified by the same Property name, there can be one or more instances separated by <i>datasetId</i> .				
NOTE 2: For each <i>Relationship</i> (or subclass of <i>Relationship</i>) identified by the same Relationship name, there can be one or more instances separated by <i>datasetId</i> .				

5.2.32 LanguageProperty

This type represents the data needed to define a *LanguageProperty* as mandated by clause 4.5.18. Note that in case of concise representation, the *type* can be omitted (see clause 4.5.18.3).

The supported JSON members shall follow the requirements provided in Table 5.2.32-1.

Table 5.2.32-1: NGSI-LD LanguageProperty data type definition

Name	Data Type	Restriction	Cardinality	Description
type	string	It shall be equal to "LanguageProperty"	1	Node type.
languageMap	JSON object	A set of key-value pairs whose keys shall be strings representing IETF RFC 5646 [28] language codes and whose values shall be JSON strings or arrays of JSON strings	1	String Property Values defined in multiple natural languages.
datasetId	String	Valid URI	0..1	It allows identifying a set or group of property values.
expiresAt	String	<i>DateTime</i> (see clause 4.6.3)	0..1	System temporal Property representing the expiration date for the storage of the LanguageProperty. See clause 4.22.
ngsildproof	Property	Property with the non-reified subproperties "entityIdSealed" and "entityTypeSealed" as specified in [35]. The value of its "value" element shall be an object containing the W3C® Data integrity "proof" structure [35].	0..1	Cryptographic signature of the LanguageProperty guaranteeing its integrity.
observedAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	Timestamp. See clause 4.8.
valueType	String	It shall be equal to "langString" (see note 1)	0..1	The native JSON-LD @type for the languageMap attribute. It shall align with the RDF datatype [34].
<Property name>	Property or Property[] (see note 2)	See datatype definition in clauses 5.2.5	0..N	Properties of the LanguageProperty.
	GeoProperty or GeoProperty[] (see note 2)	See datatype definition in clause 5.2.7	0..N	GeoProperties of the LanguageProperty.
	LanguageProperty or LanguageProperty[] (see note 2)	See datatype definition in clause 5.2.32	0..N	LanguageProperties of the LanguageProperty.
	JsonProperty or JsonProperty[] (see note 2)	See datatype definition in clause 5.2.38	0..N	JsonProperties of the LanguageProperty.
	VocabProperty or VocabProperty[] (see note 2)	See datatype definition in clause 5.2.35	0..N	VocabProperties of the LanguageProperty.
	ListProperty or ListProperty[] (see note 2)	See datatype definition in clause 5.2.36	0..N	ListProperties of the LanguageProperty.
<Relationship name>	Relationship or Relationship[] (see note 3)	See datatype definition in clause 5.2.6	0..N	Relationships of the LanguageProperty.
	ListRelationship or ListRelationship[] (see note 3)	See datatype definition in clause 5.2.37	0..N	ListRelationships of the LanguageProperty.
NOTE 1: The assigned @type for language tagged strings is datatype URI: http://www.w3.org/1999/02/22-rdf-syntax-ns#langString [34].				
NOTE 2: For each <i>Property</i> (or subclass of <i>Property</i>) identified by the same Property name, there can be one or more instances separated by <i>datasetId</i> .				
NOTE 3: For each <i>Relationship</i> (or subclass of <i>Relationship</i>) identified by the same Relationship name, there can be one or more instances separated by <i>datasetId</i> .				

The following **output only** members (defined by Table 5.2.32-2) of the *LanguageProperty* data structure are also defined. They are read-only and shall be generated by NGSI-LD implementations. They shall not be provided by **Context Producers**. In the event that they are provided (in update or create operations) NGSI-LD implementations shall ignore them.

Table 5.2.32-2: Output only members of the NGSI-LD LanguageProperty data type

Name	Data Type	Restrictions	Cardinality	Description
createdAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	System generated creation timestamp. See clause 4.8.
deletedAt	String	<i>DateTime</i> (clause 4.6.3) It is only used in notifications reporting deletions	0..1	System generated deletion timestamp. See clause 4.8.
instanceId	String	Valid URI. Only used in temporal representation of LanguageProperties	0..1	URI uniquely identifying a LanguageProperty instance as mandated by clause 4.5.7.
modifiedAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	System generated last modification timestamp. See clause 4.8.
previousLanguageMap	JSON object	A set of key-value pairs whose keys shall be strings representing IETF RFC 5646 [28] language codes and whose values shall be JSON strings. Only used in Notifications, if the <i>showChanges</i> option is explicitly requested	0..1	Previous Language Property's <i>languageMap</i> .

5.2.33 EntitySelector

This type selects which entity or group of entities are queried or subscribed to by **Context Consumers**. Entities can be specified by their id, Entity Types or group of Entity IDs (as a regular expression pattern mandated by IEEE 1003.2™ [11]).

The JSON members shall follow the indications provided in Table 5.2.33-1. *id* takes precedence over *idPattern*.

Table 5.2.33-1: EntitySelector data type definition

Name	Data Type	Restrictions	Cardinality	Description
type	String	A valid type selection string as per clause 4.17. To indicate a request for all Entities (with implied local scope), "*" is also allowed as a value.	1	Selector of Entity Type(s); If type is specified as "*", implying local scope, local scope shall not be explicitly set to be <i>false</i> (clause 5.5.13) for the execution of the corresponding operation.
id	String or String[]	Valid URI(s)	0..1	Entity identifier(s).
idPattern	String	Regular expression as per IEEE 1003.2™ [11]	0..1	A regular expression which denotes a pattern that shall be matched by the provided or subscribed Entities.

5.2.34 RegistrationManagementInfo

This type represents the data to alter the default behaviour of a **Context Broker** when making a distributed operation request to a registered **Context Source**. The supported JSON members shall follow the indications provided in Table 5.2.34-1. Brokers may override these recommendations.

Table 5.2.34-1: RegistrationManagementInfo data type definition

Name	Data Type	Restrictions	Cardinality	Description
cacheDuration	String	String representing a duration in ISO 8601 [17] format	0..1	Minimal period of time which shall elapse between two consecutive context information consumption operations (as defined in clause 5.7) related to the same context data will occur. If the <i>cacheDuration</i> latency period has not been reached, a cached value for the entity or its attributes shall be returned where available.
cooldown	Number	Greater than 0	0..1	Minimum period of time in milliseconds which shall elapse before attempting to make a subsequent forwarded request to the same endpoint after failure. If requests are received before the cooldown period has expired, a timeout error response for the registration is automatically returned.
localOnly	Boolean		0..1	If <i>localOnly=true</i> then distributed operations associated to this Context Source Registration will act only on data held directly by the registered Context Source itself (see clause 4.3.6.4).
timeout	Number	Greater than 0	0..1	Maximum period of time in milliseconds which may elapse before a forwarded request is assumed to have failed.

5.2.35 VocabProperty

This type represents the data needed to define a **VocabProperty** as mandated by clause 4.5.20. In case of concise representation, the *type* can be omitted (see clause 4.5.20.3).

The supported JSON members shall follow the requirements provided in Table 5.2.35-1.

Table 5.2.35-1: NGSI-LD VocabProperty data type definition

Name	Data Type	Restriction	Cardinality	Description
type	String	It shall be equal to "VocabProperty"	1	Node type.
vocab	String or string[]		1	String Values which shall be type coerced to URIs based on the supplied <i>@context</i> .
datasetId	String	Valid URI	0..1	It allows identifying a set or group of property values.
expiresAt	String	<i>DateTime</i> (see clause 4.6.3)	0..1	System temporal Property representing the expiration date for the storage of the VocabProperty . See clause 4.22.

Name	Data Type	Restriction	Cardinality	Description
ngsildproof	Property	Property with the non-reified subproperties "entityIdSealed" and "entityTypeSealed" as specified in [35]. The value of its "value" element shall be an object containing the W3C® Data integrity "proof" structure [35].	0..1	Cryptographic signature of the VocabProperty guaranteeing its integrity.
observedAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	Timestamp. See clause 4.8.
valueType	String		0..1	The native JSON-LD <i>@type</i> for the <i>vocab</i> attribute. A String Value which shall be type coerced to a URI based on the supplied <i>@context</i> .
<Property name>	Property or Property[] (see note 1)	See datatype definitions in clauses 5.2.5	0..N	Properties of the VocabProperty.
	GeoProperty or GeoProperty[] (see note 1)	See datatype definition in clause 5.2.7	0..N	GeoProperties of the VocabProperty.
	LanguageProperty or LanguageProperty[] (see note 1)	See datatype definition in clause 5.2.32	0..N	LanguageProperties of the VocabProperty.
	JsonProperty or JsonProperty[] (see note 1)	See datatype definition in clause 5.2.38	0..N	JsonProperties of the VocabProperty.
	VocabProperty or VocabProperty[] (see note 1)	See datatype definition in clause 5.2.35	0..N	VocabProperties of the VocabProperty.
	ListProperty or ListProperty[] (see note 1)	See datatype definition in clause 5.2.36	0..N	ListProperties of the VocabProperty.
<Relationship name>	Relationship or Relationship[] (see note 2)	See datatype definition in clause 5.2.6	0..N	Relationships of the VocabProperty.
	ListRelationship or ListRelationship[] (see note 2)	See datatype definition in clause 5.2.37	0..N	ListRelationships of the VocabProperty.
NOTE 1: For each <i>Property</i> (or subclass of <i>Property</i>) identified by the same Property name, there can be one or more instances separated by <i>datasetId</i> .				
NOTE 2: For each <i>Relationship</i> (or subclass of <i>Relationship</i>) identified by the same Relationship name, there can be one or more instances separated by <i>datasetId</i> .				

The following **output only** members (defined by Table 5.2.35-2) of the VocabProperty data structure are also defined. They are read-only and shall be generated by NGSI-LD implementations. They shall not be provided by *Context Producers*. In the event that they are provided (in update or create operations) NGSI-LD implementations shall ignore them.

Table 5.2.35-2: Output only members of the NGSI-LD VocabProperty data type

name	Data Type	Restrictions	Cardinality	Description
createdAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	System generated creation timestamp. See clause 4.8.
deletedAt	String	<i>DateTime</i> (clause 4.6.3) It is only used in notifications reporting deletions	0..1	System generated deletion timestamp. See clause 4.8.
instanceId	String	Valid URI. Only used in temporal representation of VocabProperties	0..1	URI uniquely identifying a VocabProperty instance as mandated by clause 4.5.7.
modifiedAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	System generated last modification timestamp. See clause 4.8.
previousVocab	String or String[]	Only used in Notifications	0..1	Previous <i>VocabProperty's</i> vocab.

5.2.36 ListProperty

This type represents the data needed to define a *ListProperty* as mandated by clause 4.5.21. Note that in case of concise representation, the *type* can be omitted (see clause 4.5.21.3).

The supported JSON members shall follow the requirements provided in Table 5.2.36-1.

Table 5.2.36-1: NGSI-LD ListProperty data type definition

Name	Data Type	Restriction	Cardinality	Description
type	String	It shall be equal to "ListProperty"	1	Node type.
valueList	An array of JSON values as defined by IETF RFC 8259 [6]	See NGSI-LD Value definition in clause 3.1	1	Ordered array of Property Values.
datasetId	String	Valid URI	0..1	It allows identifying a set or group of property list values.
expiresAt	String	<i>DateTime</i> (see clause 4.6.3)	0..1	System temporal Property representing the expiration date for the storage of the ListProperty. See clause 4.22.
ngsildproof	Property	Property with the non-reified subproperties "entityIdSealed" and "entityTypeSealed" as specified in [35]. The value of its "value" element shall be an object containing the W3C® Data integrity "proof" structure [35].	0..1	Cryptographic signature of the ListProperty guaranteeing its integrity.
observedAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	Timestamp. See clause 4.8.

Name	Data Type	Restriction	Cardinality	Description
valueType	String		0..1	The native JSON-LD <i>@type</i> for the <i>valueList</i> attribute. A String Value which shall be type coerced to a URI based on the supplied <i>@context</i> .
<Property name>	Property or Property[] (see note 1)	See datatype definition in clauses 5.2.5	0..N	Properties of the ListProperty.
	GeoProperty or GeoProperty[] (see note 1)	See datatype definition in clause 5.2.7	0..N	GeoProperties of the ListProperty.
	LanguageProperty or LanguageProperty[] (see note 1)	See datatype definition in clause 5.2.32	0..N	LanguageProperties of the ListProperty.
	JsonProperty or JsonProperty[] (see note 1)	See datatype definition in clause 5.2.38	0..N	JsonProperties of the ListProperty.
	VocabProperty or VocabProperty[] (see note 1)	See datatype definition in clause 5.2.35	0..N	VocabProperties of the ListProperty.
	ListProperty or ListProperty[] (see note 1)	See datatype definition in clause 5.2.36	0..N	ListProperties of the ListProperty.
<Relationship name>	Relationship or Relationship[] (see note 2)	See datatype definition in clause 5.2.6	0..N	Relationships of the ListProperty.
	ListRelationship or ListRelationship[] (see note 2)	See datatype definition in clause 5.2.37	0..N	ListRelationships of the ListProperty.
NOTE 1: For each <i>Property</i> (or subclass of <i>Property</i>) identified by the same Property name, there can be one or more instances separated by <i>datasetId</i> .				
NOTE 2: For each <i>Relationship</i> (or subclass of <i>Relationship</i>) identified by the same Relationship name, there can be one or more instances separated by <i>datasetId</i> .				

The following **output only** members (defined by Table 5.2.36-2) of the *ListProperty* data structure are also defined. They are read-only and shall be generated by NGSI-LD implementations. They shall not be provided by Context Producers. In the event that they are provided (in update or create operations) NGSI-LD implementations shall ignore them.

Table 5.2.36-2: Additional members of the NGSI-LD ListProperty data type

Name	Data Type	Restrictions	Cardinality	Description
createdAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	System generated creation timestamp. See clause 4.8.
deletedAt	String	<i>DateTime</i> (clause 4.6.3) It is only used in notifications reporting deletions	0..1	System generated deletion timestamp. See clause 4.8.
instanceId	String	Valid URI. Only used in temporal representation of ListProperties	0..1	URI uniquely identifying a ListProperty instance as mandated by clause 4.5.7.
modifiedAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	System generated last modification timestamp. See clause 4.8.
previousValueList	An array of JSON values as defined by IETF RFC 8259 [6]	See NGSI-LD Value definition in clause 3.1	0..1	Ordered array of previous Property Values.

5.2.37 ListRelationship

This type represents the data needed to define a *ListRelationship* as mandated by clause 4.5.22. Note that in case of concise representation, the *type* can be omitted (see clause 4.5.22.3) and the *objectList* may also be represented as an ordered array of Strings.

The supported JSON members shall follow the requirements provided in Table 5.2.37-1.

Table 5.2.37-1: NGSI-LD ListRelationship data type definition

Name	Data Type	Restriction	Cardinality	Description
type	String	It shall be equal to "ListRelationship"	1	Node type.
objectList	JSON Object[] or String[]	In the normalized form, each array element holds a JSON object containing a single Attribute with a key called "object" and where the value is a valid URI. In the concise form, each string in the array holds a valid URI	1	Ordered array of Relationship target objects.
datasetId	String	Valid URI	0..1	It allows identifying a set or group of relationship list objects.
expiresAt	String	<i>DateTime</i> (see clause 4.6.3)	0..1	System temporal Property representing the expiration date for the storage of the ListRelationship. See clause 4.22.

Name	Data Type	Restriction	Cardinality	Description
ngsildproof	Property	Property with the non-reified subproperties "entityIdSealed" and "entityTypeSealed" as specified in [35]. The value of its "value" element shall be an object containing the W3C® Data integrity "proof" structure [35].	0..1	Cryptographic signature of the ListRelationship guaranteeing its integrity.
objectType	String or String[]		0..1	Node Type of the Relationship's target object. Both short hand string(s) (type name) or URI(s) are allowed.
observedAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	Timestamp. See clause 4.8.
<Property name>	Property or Property[] (see note 1)	See datatype definition in clause 5.2.5	0..N	Properties of the ListRelationship.
	GeoProperty or GeoProperty[] (see note 1)	See datatype definition in clause 5.2.7	0..N	GeoProperties of the ListRelationship.
	LanguageProperty or LanguageProperty[] (see note 1)	See datatype definition in clause 5.2.32	0..N	LanguageProperties of the ListRelationship.
	JsonProperty or JsonProperty[] (see note 1)	See datatype definition in clause 5.2.38	0..N	JsonProperties of the ListProperty.
	VocabProperty or VocabProperty[] (see note 1)	See datatype definition in clause 5.2.35	0..N	VocabProperties of the ListRelationship.
	ListProperty or ListProperty[] (see note 1)	See datatype definition in clause 5.2.36	0..N	ListProperties of the ListRelationship.
<Relationship name>	Relationship or Relationship[] (see note 2)	See datatype definition in clause 5.2.6	0..N	Relationships of the ListRelationship.
	ListRelationship or ListRelationship[] (see note 2)	See datatype definition in clause 5.2.37	0..N	ListRelationships of the ListRelationship.
NOTE 1: For each <i>Property</i> (or subclass of <i>Property</i>) identified by the same Property name, there can be one or more instances separated by <i>datasetId</i> .				
NOTE 2: For each <i>Relationship</i> (or subclass of <i>Relationship</i>) identified by the same Relationship name, there can be one or more instances separated by <i>datasetId</i> .				

The following **output only** members (defined by Table 5.2.37-2) of the *ListRelationship* data structure are also defined. They are read-only and shall be generated by NGSI-LD implementations. They shall not be provided by Context Producers. In the event that they are provided (in update or create operations) NGSI-LD implementations shall ignore them.

Table 5.2.37-2: Additional members of the NGSI-LD ListRelationship data type

Name	Data Type	Restrictions	Cardinality	Description
createdAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	System generated creation timestamp. See clause 4.8.
deletedAt	String	<i>DateTime</i> (clause 4.6.3) It is only used in notifications reporting deletions	0..1	System generated deletion timestamp. See clause 4.8.
entityList	Entity[]	See datatype definition in clause 5.2.4 Only used in Linked Entity Retrieval, if the <i>join=inline</i> option is explicitly requested	0..1	An array of inline Entities obtained by Linked Entity retrieval, corresponding to the <i>ListRelationship</i> 's target objects See clause 4.5.23.2.
instanceId	String	Valid URI. Only used in temporal representation of <i>ListRelationships</i>	0..1	URI uniquely identifying a <i>ListRelationship</i> instance as mandated by clause 4.5.8.
modifiedAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	System generated last modification timestamp. See clause 4.8.
previousObjectList	JSON Object[] or String[]	In the normalized form, each array element holds a JSON object containing a containing a single Attribute with a key called "object" and where the value is a valid URI. In the concise form, each string in the array holds a valid URI	0..1	Ordered array of previous <i>Relationship</i> target objects.

5.2.38 JsonProperty

This type represents the data needed to define a *JsonProperty* as mandated by clause 4.5.24. In case of concise representation, the *type* can be omitted (see clause 4.5.24.3).

The supported JSON members shall follow the requirements provided in Table 5.2.38-1.

Table 5.2.38-1: NGSI-LD JsonProperty data type definition

Name	Data Type	Restriction	Cardinality	Description
type	String	It shall be equal to "JsonProperty"	1	Node type.
json	JSON		1	Raw unexpandable JSON which shall not be interpreted as JSON-LD using the supplied <i>@context</i> .
datasetId	String	Valid URI	0..1	It allows identifying a set or group of property values.
expiresAt	String	<i>DateTime</i> (see clause 4.6.3)	0..1	System temporal Property representing the expiration date for the storage of the <i>JsonProperty</i> . See clause 4.22

Name	Data Type	Restriction	Cardinality	Description
ngsildproof	Property	Property with the non-reified subproperties "entityIdSealed" and "entityTypeSealed" as specified in [35]. The value of its "value" element shall be an object containing the W3C® Data integrity "proof" structure [35].	0..1	Cryptographic signature of the JsonProperty guaranteeing its integrity.
observedAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	Timestamp. See clause 4.8.
valueType	String		0..1	The native JSON-LD <i>@type</i> for the <i>json</i> attribute. A String Value which shall be type coerced to a URI based on the supplied <i>@context</i> .
<Property name>	Property or Property[] (see note 1)	See datatype definitions in clauses 5.2.5	0..N	Properties of the JsonProperty.
	GeoProperty or GeoProperty[] (see note 1)	See datatype definition in clause 5.2.7	0..N	GeoProperties of the JsonProperty.
	LanguageProperty or LanguageProperty[] (see note 1)	See datatype definition in clause 5.2.32	0..N	LanguageProperties of the JsonProperty.
	JsonProperty or JsonProperty[] (see note 1)	See datatype definition in clause 5.2.38	0..N	JsonProperties of the JsonProperty.
	VocabProperty or VocabProperty[] (see note 1)	See datatype definition in clause 5.2.35	0..N	VocabProperties of the JSONProperty.
	ListProperty or ListProperty[] (see note 1)	See datatype definition in clause 5.2.36	0..N	ListProperties of the JsonProperty.
<Relationship name>	Relationship or Relationship[] (see note 2)	See datatype definition in clause 5.2.6	0..N	Relationships of the JsonProperty.
	ListRelationship or ListRelationship[] (see note 2)	See datatype definition in clause 5.2.37	0..N	ListRelationships of the JsonProperty.
NOTE 1: For each <i>Property</i> (or subclass of <i>Property</i>) identified by the same Property name, there can be one or more instances separated by <i>datasetId</i> .				
NOTE 2: For each <i>Relationship</i> (or subclass of <i>Relationship</i>) identified by the same Relationship name, there can be one or more instances separated by <i>datasetId</i> .				

The following **output only** members (defined by Table 5.2.38-2) of the JsonProperty data structure are also defined. They are read-only and shall be generated by NGSI-LD implementations. They shall not be provided by *Context Producers*. In the event that they are provided (in update or create operations) NGSI-LD implementations shall ignore them.

Table 5.2.38-2: Output only members of the NGSI-LD JsonProperty data type

name	Data Type	Restrictions	Cardinality	Description
createdAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	System generated creation timestamp. See clause 4.8.
deletedAt	String	<i>DateTime</i> (clause 4.6.3) It is only used in notifications reporting deletions	0..1	System generated deletion timestamp. See clause 4.8.
instanceId	String	Valid URI. Only used in temporal representation of <i>JsonProperties</i>	0..1	URI uniquely identifying a <i>JsonProperty</i> instance as mandated by clause 4.5.7.
modifiedAt	String	<i>DateTime</i> (clause 4.6.3)	0..1	System generated last modification timestamp. See clause 4.8.
previousJson	JSON	Only used in Notifications	0..1	Previous <i>JsonProperty's json</i> .

5.2.39 EntityMap

This type represents the data needed to define an *EntityMap* as mandated by clause 4.5.25.

The supported JSON members shall follow the requirements provided in Table 5.2.39-1.

Table 5.2.39-1: NGSI-LD EntityMap data type definition

Name	Data Type	Restriction	Cardinality	Description
id	String	Valid URI	1	EntityMap ID.
type	String	It shall be equal to "EntityMap"	1	Node type.
expiresAt	String	<i>DateTime</i> (see clause 4.6.3)	1	Expiration date for the EntityMap.

The following **output only** members (defined by Table 5.2.39-2) of the *EntityMap* data structure are also defined. They are read-only and shall be generated by NGSI-LD implementations. They shall not be provided by *Context Consumers*. In the event that they are provided in update operations NGSI-LD implementations shall ignore them.

Table 5.2.39-2: Additional members of the NGSI-LD EntityMap data type

Name	Data Type	Restrictions	Cardinality	Description
entityMap	JSON	A set of key-value pairs whose keys shall be strings representing Entity ids and whose values shall be an array holding every <i>CSourceRegistration</i> id which is relevant to the ongoing <i>Context Information Consumption</i> request (see clause 4.21). The key "@none" shall be used to refer to an Entity that is held locally.	1	System generated mapping of Entities to <i>CSourceRegistrations</i> .

Name	Data Type	Restrictions	Cardinality	Description
linkedMaps	JSON	A set of key-value pairs whose keys shall be strings representing CSourceRegistration ids which are relevant to the ongoing Context Information request and whose values shall represent the associated EntityMap id used by the ContextSource.	1	System generated mapping of Context CSourceRegistrations to a URI indicating which EntityMaps was used by the Context Source.

5.2.40 Context Source Identity

This type represents the data uniquely identifying a Context Source, and if the Context Source supports multi-tenancy (see clause 4.14) uniquely identifying a Tenant within that Context Source.

The supported JSON members shall follow the requirements provided in Table 5.2.40-1.

Table 5.2.40-1: Context Source Identity data type definition

Name	Data Type	Restriction	Cardinality	Description
id	String	Valid URI	1	Context Source ID.
type	String	It shall be equal to "ContextSourceIdentity"	1	JSON-LD @type.
contextSourceAlias	String	Non-empty string. Pseudonym field as defined in IETF RFC 7230 [27]	1	A unique id for a Context Source which can be used to identify loops. In the multi-tenancy use case (see clause 4.14), this id shall be identifying a specific Tenant within a registered Context Source.
contextSourceUptime	String	String representing a duration in ISO 8601 [17] format	1	Total Duration that the Context Source has been available.
contextSourceTimeAt	String	<i>DateTime</i> (clause 4.6.3)	1	Current time observed at the Context Source. Timestamp See clause 4.8.
contextSourceExtras	JSON		0..1	Instance specific information relevant to the configuration of the Context Source itself in raw un-expandable JSON which shall not be interpreted as JSON-LD using the supplied @context.

5.2.41 Snapshot

This type represents the data needed to create and manage a Snapshot.

The supported JSON members shall follow the requirements provided in Table 5.2.41-1.

Table 5.2.41-1: Snapshot data type definition

Name	Data Type	Restriction	Cardinality	Description
id	String	Valid URI	0..1	System-generated at creation time, if it is not provided. It cannot be later modified in update operations.
type	String	It shall be equal to "Snapshot"	1	JSON-LD @type.
snapshotQueries	Query[]	List of Query (clause 5.2.23), where no Query no "temporalQ" element.	0..1	Allows clients to specify a list of queries whose results are to be part of the Snapshot. This member can only be set when creating a snapshot, after that it becomes read-only. At least one of "snapshotQueries" or "snapshotTemporalQueries" shall be present when creating the snapshot and both shall be omitted when updating the Snapshot status or cloning the Snapshot.
snapshotTemporalQueries	Query[]	List of Query (clause 5.2.23), where each Query has "temporalQ" element.	0..1	Allows clients to specify a list of temporal queries whose results are to be part of the Snapshot. This member can only be set when creating a snapshot, after that it becomes read-only. At least one of "snapshotQueries" or "snapshotTemporalQueries" shall be present when creating the snapshot and both shall be omitted when updating the Snapshot status.
snapshotLifetime	String	String representing a duration in ISO 8601 [17] format	0..1	Suggested duration for which the requester wants the Snapshot to exist with respect to the current point in time. The actual expiresAt time of the Snapshot shall be set by the NGSI-LD system, possibly overriding the requested duration.
snapshotPriority	Number	Positive Integer between 1 and 10	0..1	The priority of a snapshot ranges from 1 (minimum) to 10 (maximum) with 5 being the default priority. It can be used when deciding which snapshot is to be deleted if an implementation determines that it is low on resources.
endpoint	String	Dereferenceable URI	0..1	URI to which notifications regarding changes in the status of the Snapshot are to be sent. The information contained in the receiverInfo member shall be considered.
receiverInfo	KeyValuePair[]		0..1	Generic {key, value} array to convey optional information to the receiver.

The following **output only** members (defined by Table 5.2.41-2) of the *Snapshot* data structure are also defined. They are read-only and shall be generated by NGSI-LD implementations. They shall not be provided by *Context Consumers*. In the event that they are provided (in update or create operations) NGSI-LD implementations shall ignore them.

Table 5.2.41-2: Output only members of the Snapshot data type

Name	Data Type	Restrictions	Cardinality	Description
snapshotStatus	String	<i>It shall be one of: "preparing", "success", "partial", "failure" or "empty"</i> <i>The initial status shall be "preparing".</i>	1	Allows clients to check the status of the Snapshot.
snapshotQueriesDetails	ExecutionResultDetails[]	<i>List of ExecutionResultDetails (clause 5.2.42)</i>	0..1	List with one result per Snapshot query execution, in the same order as Snapshot queries.
snapshotTemporalQueriesDetails	ExecutionResultDetails[]	<i>List of ExecutionResultDetails (clause 5.2.42)</i>	0..1	List with one result per temporal Snapshot query execution, in the same order as temporal Snapshot queries.
createdAt	String	<i>DateTime (clause 4.6.3)</i>	0..1	System generated creation timestamp. See clause 4.8.
modifiedAt	String	<i>DateTime (clause 4.6.3)</i>	0..1	System generated last modification timestamp. See clause 4.8.
expiresAt	String	<i>DateTime (clause 4.6.3)</i>	1	Expiration time for the snapshot - it is possible to indirectly update the member by updating snapshotLifetime. See clause 4.8.
lastUsedAt	String	<i>DateTime (clause 4.6.3)</i>	0..1	System generated timestamp identifying the point in time when the snapshot was most recently used. It is initialized at creation time. See clause 4.8.

5.2.42 ExecutionResultDetails

This type represents the details of the result of execution a request, e.g. a Query.

The supported JSON members shall follow the requirements provided in Table 5.2.42-1.

Table 5.2.42-1: ExecutionResultDetails data type definition

Name	Data Type	Restriction	Cardinality	Description
problemDetails	@JSON	ProblemDetails (see IETF RFC 7807 [10])	0..1	Provides more details regarding the result status, especially when reporting an error.
resultStatus	String	<i>It shall be one of: "success", "failure" or "empty"</i>	1	Describes the status of the result. "failure", if an error is reported, "success" in case of a non-empty result and "empty" otherwise.

5.2.43 OrderingParams

This datatype represents the parameters that convey the definition used whilst ordering Entities.

The supported JSON members shall follow the requirements provided in Table 5.2.43-1.

Table 5.2.43-1: OrderingParams data type definition

Name	Data Type	Restriction	Cardinality	Description
collation	String	An ICU collation (see IETF RFC 6067 [36]).	0..1	The Entities returned in the payload shall be ordered according to the collation given. By default, it shall be ICU "root" collation order ("und-x-icu").
coordinates	JSON Array	A JSON Array coherent with the geometry type as per IETF RFC 7946 [8].	0..1 It shall be one if <i>orderBy</i> uses order by distance	Coordinates of a Geometry used when sorting by distance.
geometry	String	A valid GeoJSON [8] geometry type excepting <i>GeometryCollection</i> .	0..1	The type of geometry whose coordinates are used when sorting by distance. By default, it shall be a "Point" geometry.
orderBy	String[]	Each String is an Entity member ("id", "type", "scope" or an Attribute name) appended with an optional sorting style ("asc", "desc", "dist-asc", "dist-desc") as per clause 4.23.	1	When defined, the Entities returned in the payload shall be ordered according to members defined.

5.2.44 AggregationParams

This datatype represents the parameters required for supporting aggregation methods in temporal requests.

The supported JSON members shall follow the requirements provided in Table 5.2.44-1.

Table 5.2.44-1: AggregationParams data type definition

Name	Data Type	Restriction	Cardinality	Description
aggrMethods	Comma separated list of strings	Each String represents an aggregation method, as defined by clause 4.5.19.	1	-
aggrPeriodDuration	String	It represents the duration of each period used for the aggregation as defined by clause 4.5.19. If not specified, it defaults to a duration of 0 seconds and is interpreted as a duration spanning the whole time-range specified by the temporal query.	0..1	-

5.3 Notification data types

5.3.1 Notification

This datatype represents the parameters that allow building a notification to be sent to a subscriber. How to build this notification is detailed in clause 5.8.6.

The supported JSON members shall follow the indications provided in Table 5.3.1-1.

Table 5.3.1-1: Notification data type definition

Name	Data Type	Restrictions	Cardinality	Description
id	String	Valid URI	1	Notification identifier (JSON-LD <i>@id</i>). It shall be automatically generated by the implementation.
type	String	It shall be equal to "Notification"	1	JSON-LD <i>@type</i> .
data	NGSI-LD Entity[] or FeatureCollection		1	The content of the notification as NGSI-LD Entities. See clause 5.2.4. If the notification has been triggered from a Subscription that has the <i>notification</i> . <i>endpoint.accept</i> field set to <i>application/geo+json</i> then data is returned as a <i>FeatureCollection</i> . In this case, if the <i>notification.endpoint.receiverInfo</i> contains the key "Prefer" and it is set to the value "body=json", then the <i>FeatureCollection</i> will not contain an <i>@context</i> field. If <i>endpoint.accept</i> is not set or holds another value then Entity[] is returned.
notifiedAt	String	<i>DateTime</i> (clause 4.6.3)	1	Timestamp corresponding to the instant when the notification was generated by the system.
subscriptionId	String	Valid URI	1	Identifier of the subscription that originated the notification.

5.3.2 CSourceNotification

This datatype represents the parameters that allow building a Context Source Notification to be sent to a subscriber. How to build this notification is detailed in clause 5.11.7.

The supported JSON members shall follow the indications provided in Table 5.3.2-1.

Table 5.3.2-1: CSourceNotification data type definition

Name	Data Type	Restrictions	Cardinality	Description
id	String	Valid URI	1	CSource notification identifier (JSON-LD <i>@id</i>).
type	String	It shall be equal to "ContextSourceNotification"	1	JSON-LD <i>@type</i> .
data	CSourceRegistration[]		1	The content of the notification as NGSI-LD CSourceRegistrations. See clause 5.2.9.
notifiedAt	String	<i>DateTime</i> (see clause 4.6.3)	1	Timestamp corresponding to the instant when the notification was generated by the system.
subscriptionId	String	Valid URI	1	Identifier of the subscription that originated the notification.
triggerReason	String	<i>TriggerReasonEnumeration</i> (see clause 5.3.3)	1	Indicates whether the CSources in the CSourceRegistration(s) in data are newly matching (initial notification or creation), have been updated (but still match) or do not match any longer.

5.3.3 TriggerReasonEnumeration

The enumeration can take one of the following values:

- **"newlyMatching"** - describes the case that the notified Context Source Registration(s) newly match(es) the identified subscription. This value is used in the first notification and whenever a new Context Source Registration matching the Subscription has been registered, or an existing Context Source Registration that did not match before has been updated in such a way that it matches now.
- **"updated"** - describes the case that the notified Context Source Registration that was part of a previous notification has been updated, but still matches the Subscription.
- **"noLongerMatching"** - describes the case that the notified Context Source Registration that was part of a previous notification no longer matches the Subscription, i.e. as a result of an update or because it was deleted.

5.3.4 SnapshotNotification

This datatype represents the parameters that allow building a snapshot notification to be sent to a subscriber. How to build this notification is detailed in clause 5.16.6.

The supported JSON members shall follow the indications provided in Table 5.3.4-1.

Table 5.3.4-1: SnapshotNotification data type definition

Name	Data Type	Restrictions	Cardinality	Description
Id	String	Valid URI	1	Notification identifier (JSON-LD <i>@id</i>). It shall be automatically generated by the implementation.
type	String	It shall be equal to "SnapshotNotification"	1	JSON-LD <i>@type</i> .
expiresAt	String	<i>DateTime</i> (clause 4.6.3)	1	Expiration time for the snapshot - if expiresAt is before notifiedAt, the notification indicates that the snapshot has been deleted. In this case, snapshotReady shall be set to <i>false</i> .
notifiedAt	String	<i>DateTime</i> (clause 4.6.3)	1	Timestamp corresponding to the instant when the notification was generated by the system.
snapshotId	String	Valid URI	1	Identifier of the snapshot whose status change triggered the notification.
snapshotPriority	Number	Positive Integer between 1 and 10	1	The priority of a snapshot ranges from 1 (minimum) to 10 (maximum) with 5 being the default priority. It is used when deciding which snapshot is to be deleted if an implementation determines that it is low on resources.
snapshotQueriesDetails	ExecutionResultDetails[]	List of <i>ExecutionResultDetails</i> (clause 5.2.42)	0..1	List with one result per Snapshot query execution, in the same order as Snapshot queries.
snapshotStatus	String	It shall be one of: "success" "partial" "failure" or "empty"	1	Indicates the status of the Snapshot, enabling the user to decide whether the snapshot is ready for querying.
temporalSnapshotQueriesDetails	ExecutionResultDetails[]	List of <i>ExecutionResultDetails</i> (clause 5.2.42)	0..1	List with one result per temporal Snapshot query execution, in the same order as temporal Snapshot queries.

5.4 NGSI-LD Fragments

When updating NGSI-LD elements (Entities, Attributes, Context Source Registrations or Subscriptions) it is necessary to have a means of describing a set of modifications to their content.

An NGSI-LD Fragment is a JSON Merge Patch document [16] and [i.10] which describes changes to be made to a target JSON-LD document using a syntax that closely mimics the document being modified.

An NGSI-LD Fragment is a JSON-LD Object which shall include the following members:

- *id* (optional for certain bindings where it can be determined from the operation signature). It shall be equal to the id of the target (mutated) NGSI-LD element. Attribute Fragments do not contain explicit ids.
- *type* (optional for certain bindings where it can be determined from the operation signature). It shall contain the Type name(s) of the target NGSI-LD element.

- A member (following the same data representation and nesting structure) for each new member to be added to the target NGSI-LD element.
- A member (following the same data representation and nesting structure) for each new member to be modified in the target NGSI-LD element, which value shall correspond to the new member value to be given.

EXAMPLE 1: The following Subscription Fragment allows the modification of a Subscription by changing its endpoint's URI:

```
{
  "id": "urn:ngsi-ld:Subscription:MySubscription",
  "type": "Subscription",
  "endpoint": {
    "uri": "http://example.org/newNotificationEndPoint"
  }
}
```

- A member (following the same data representation and nesting structure) with value equal to an NGSI-LD Null shall cause for the member to be removed from the target NGSI-LD element.

EXAMPLE 2: The following NGSI-LD Fragment allows the modification of an Entity by changing its *batteryLevel* Attribute, updating the *observedAt* sub-Attribute, removing the *providedBy* sub-Attribute and removing the *uncharged* Attribute from the Entity:

```
{
  "id": "urn:ngsi-ld:TemperatureSensor:001",
  "type": "TemperatureSensor",
  "batteryLevel": {
    "type": "Property",
    "value": 7,
    "observedAt": "2022-03-14T12:51:02.000Z",
    "providedBy": "urn:ngsi-ld:null"
  },
  "uncharged" : {
    "type": "Property",
    "value": "urn:ngsi-ld:null"
  }
}
```

5.5 Common Behaviours

5.5.1 Introduction

This clause defines common behaviours for the API operations.

When comparing URIs, implementations shall follow the recommendations of IETF RFC 3986 [5], section 6.

5.5.2 Error types

Table 5.5.2-1 details a list of error types defined by NGSI-LD. The particular conditions under which error type shall be raised are defined when describing each operation supported by the API.

Table 5.5.2-1: Error types in NGSI-LD

Error Type	Description
https://uri.etsi.org/ngsi-ld/errors/AlreadyExists	The referred element already exists.
https://uri.etsi.org/ngsi-ld/errors/BadRequestData	The request includes input data which does not meet the requirements of the operation.
https://uri.etsi.org/ngsi-ld/errors/Conflict	The operation conflicts with the current state of the system.
https://uri.etsi.org/ngsi-ld/errors/InternalServerError	There has been an error during the operation execution.
https://uri.etsi.org/ngsi-ld/errors/InvalidRequest	The request associated to the operation is syntactically invalid or includes wrong content.
https://uri.etsi.org/ngsi-ld/errors/LdContextNotAvailable	A remote JSON-LD @context referenced in a request cannot be retrieved by the NGSI-LD Broker and expansion or compaction cannot be performed.
https://uri.etsi.org/ngsi-ld/errors/NoMultiTenantSupport	The NGSI-LD API implementation does not support multiple tenants.
https://uri.etsi.org/ngsi-ld/errors/NonexistentTenant	The addressed tenant does not exist.
https://uri.etsi.org/ngsi-ld/errors/OperationNotSupported	The operation is not supported.
https://uri.etsi.org/ngsi-ld/errors/ResourceNotFound	The referred resource has not been found.
https://uri.etsi.org/ngsi-ld/errors/TooComplexQuery	The query associated to the operation is too complex and cannot be resolved.
https://uri.etsi.org/ngsi-ld/errors/TooManyResults	The query associated to the operation is producing so many results that can exhaust client or server resources. It should be made more restrictive.

5.5.3 Error response payload body

When reporting errors back to clients, NGSI-LD implementations shall generate a JSON object in accordance with IETF RFC 7807 [10], section 3.1, including, at least the following terms:

- **type:** Error type as per clause 5.5.2.
- **title:** Error title which shall be a short string summarizing the error.
- **detail:** A detailed message that should convey enough information about the error.

Even though IETF RFC 7807 [10] defines a specific MIME type for error payloads, NGSI-LD implementations shall use the standard JSON MIME type ("application/json") when reporting errors, so that old clients or existing tools are not broken.

EXAMPLE:

```
{
  "type": "https://uri.etsi.org/ngsi-ld/errors/ResourceNotFound",
  "title": "Resource not found.",
  "detail": "urn:ngsi-ld:Device:widget001 was not found",
  "status": 404,
  "instance": "urn:ngsi-ld:Device:widget001"
}
```

5.5.4 General NGSI-LD validation

All the operations that take a JSON-LD document as input shall process such JSON-LD document as follows:

- If the request payload body is not a valid JSON document then an error of type *InvalidRequest* shall be raised.
- If the data included by the JSON-LD document is not syntactically correct, according to the @context or the API data type definitions, then an error of type *BadRequestData* shall be raised.
- A Context Producer may supply an additional hint regarding the overall validity of the payload body and the version of the NGSI-LD specification that the API datatype definitions conform to. When receiving such an annotated context production request, a Context Broker that it is only partially capable of interpreting the datatypes held within the payload body may use this information to amend the data held within the payload through applying fallbacks (see clause 4.3.6.8) prior to validation.

- Any attempt to use "urn:ngsi-ld:null" as a first level member value ("**<key>**": "urn:ngsi-ld:null"), with the exception of NGSI-LD Fragments (see clause 5.4) used in partial update and merge operations (as mandated by clause 5.5.8 and clause 5.5.12) or to represent deleted Properties in concise representation as part of notifications, shall result in an error of type *BadRequestData*.
- Any attempt to use "urn:ngsi-ld:null" as the right-hand side of *value* in a *Property*, as the right-hand side of *object* in a *Relationship* or to use {"@none": "urn:ngsi-ld:null"} as the right-hand side of *languageMap*, with the exception of NGSI-LD Fragments (see clause 5.4) used in update and merge operations (as mandated by clauses 5.5.8 and 5.5.12) and the representation of deleted Properties, Relationships or Language Properties in notifications and the temporal evolution, shall result in an error of type *BadRequestData*.
- Any attempt to use "urn:ngsi-ld:null" as the value of a key value pair within a JSON object, which is the right-hand side of the *value* of a *Property*, with the exception of NGSI-LD Fragments used in merge operations (see clause 5.5.12), shall result in an error of type *BadRequestData*.

5.5.5 Default @context assignment

If the input provided by an API client does not include any *@context*, then the implementation shall at minimum assign the Core *@context* to such an input. In addition, the Context Broker implementation may allow configuring a default user *@context* (with default terms), to be used when no user *@context* is provided. The Core *@context* shall always take precedence.

5.5.6 Operation execution and generic error handling

When executing an operation if an unexpected error happens and the operation cannot be completed, implementations shall raise an error of type *InternalError*. This includes, as well, situations such as database timeouts, etc.

If the NGSI-LD endpoint is not capable of executing the requested operation, an error of type *OperationNotSupported* shall be raised. This may happen in a distributed architecture where a Context Broker might not be able to store Entities (only to forward queries to Context Sources), and as a result, certain operations such as "Create Entity" might not be supported.

When a query operation is so complex that cannot be resolved by an NGSI-LD system, implementations shall raise an error of type *TooComplexQuery*.

When a query operation is producing so many results that can potentially exhaust client or server resources, or it can be just impractical to be managed, implementations shall raise an error of type *TooManyResults*. The threshold conditions used as criteria to raise such error is up to each implementation.

When a remote JSON-LD *@context* referenced by an incoming request is not available, implementations shall raise an error of type *LdContextNotAvailable*. If the remote JSON-LD *@context* is invalid, implementations shall raise an error of type *BadRequestData*.

5.5.7 Term to URI expansion or compaction

NGSI-LD API operations allow clients to use short-hand strings as non-qualified names, particularly for *Property*, *Relationship* or Type names and *VocabProperty* values. For instance, an API client can refer to the term "Vehicle" as a non-qualified type name. When executing API update-related operations, NGSI-LD systems shall expand terms to URIs, in order to obtain and store Fully Qualified Names. Likewise, when executing query-related operations, NGSI-LD systems shall compact URIs (Fully Qualified Names) to short terms in order to provide short-hand strings to context consumers.

The term to URI expansion or compaction shall be performed using a *@context* as described by the JSON-LD specification [2], section 5.1, and in clause 4.4. **In the absence of a user *@context*, the term expansion or compaction shall be performed according to clause 5.5.5.**

For the avoidance of doubt, the *@context* used to perform compaction or expansion of terms **shall be the one provided by each API call** (or the default *@context* in its absence), and not any other *@context* which might have been supplied previously. For instance, when performing "Query Entity" operations (clause 5.7.2), the *@context* used to perform URI expansion and compaction shall be the one provided by the request.

In case of HTTP binding via GET (clause 6.4.3.2) of the "Query Entity" operation, this means using the JSON-LD Link Header as described by the JSON-LD specification [2], section 6.2. In case of HTTP binding via POST (clause 6.23.3.1), of the "Query Entity" operation, this means giving the *@context* either via Link Header or within the payload body, depending on the Content-Type Header being "application/json" or "application/ld+json", respectively.

It is important to warn users that updating a *@context* could lead to behaviour that might be perceived as inconsistent. If, for instance, a fully qualified name that qualified a given short-hand name is changed, from that moment onwards, the short-hand name is referencing a different Attribute. This will effectively change the results of queries that use the given short name, possibly not giving back anymore the expected set of results.

Moreover, this user *@context* shall not:

- Contain JSON-LD Scoped Contexts (see [2], section 4.1.8).
- Be embedded into NGSI-LD Attributes, i.e. there cannot be parts of the user *@context* other than at the top level of the NGSI-LD document.

Parts of user *@context* that are not following the two points above should result in an error of type *BadRequestData*, because JSON-LD Scoped Contexts and nested embedded *@context* could be used to modify terms defined in the Core *@context* or to reshape NGSI-LD Elements during the expansion of terms.

As the Core *@context* is protected and cannot be overridden, when performing term to URI expansion or compaction, implementations **shall always consider the Core *@context* as having absolute precedence**, regardless of the position of the Core *@context* in the *@context* array of elements. Nonetheless, NGSI-LD data providers may use appropriate term prefixing to ensure that a proper term to URI expansion or compaction is performed.

At compaction time, in the event that no matching term is found in the current *@context*, implementations shall render Fully Qualified Names.

EXAMPLE: An entity of type "Vehicle" bound to a certain *@context*, C, will match a query by "Vehicle" type if and only if the supplied query *@context*, Q, maps the term "Vehicle" to the same URI as C.

Note that the *JsonProperty* is designed to hold native JSON values which are by definition not available for expansion and compaction via an *@context*. Therefore, the given short-hand name is always used for accessing JSON keys within a *JsonProperty json* element.

5.5.8 Partial Update Patch Behaviour

The Partial Update Patch procedure modifies an existing NGSI-LD element by overwriting the data at the Attribute level, replacing it with the data provided in the NGSI-LD Fragment.

When updating NGSI-LD elements (Entities, Context Source Registrations or Context Subscriptions) using NGSI-LD Fragments, implementations shall determine the exact set of changes being requested by comparing the content of the provided Fragment (patch) against the current content (a JSON-LD object) of the target element.

With respect to update operations, implementations shall perform an algorithm equivalent to the one described below (adapted from IETF RFC 7396 [16]), in order to observe the name to URI expansion rules and the JSON-LD *null* processing):

- For each member of the Fragment perform the term to URI expansion.
- If the provided Fragment (a JSON Merge Patch document) contains members that do not appear within the target (their URIs do not match), those members are added to the target.
- For each member of the Fragment contained by the target, the target member value is replaced by the value given in the Fragment. In the case of a member representing a reified Property or Relationship including a *datasetId*, such member is only replaced if the *datasetId* is the same, otherwise the member of the Fragment is added as a new instance to the target. If no *datasetId* is present, the default Attribute instance is targeted and replaced if present and otherwise added. In case of a member *type* (of an entity) in Entity Fragments, all included Entity Types are added, if they are not already contained in the *type* member of the target.

- For each member of the Fragment, whose value is an NGSI-LD Null, contained by the target, the target member is deleted. In the case of deleting a specific Attribute instance with a *datasetId*, the handling shall be in accordance with the description found in clause 5.6.5. A *datasetId* cannot be deleted by setting it to the value "urn:ngsi-ld:null".

EXAMPLE 1: Given an Entity containing the following Property:

```
{
  "temperature": {
    "type": "Property",
    "value": 25,
    "unitCode": "CEL"
    "observedAt": "2022-03-14T01:59:26.535Z"
  }
}
```

Applying **partial attribute update** operation (as defined in clause 5.6.4) at the Attribute level onto the *temperature* Attribute, with the following Attribute Fragment payload:

```
{
  "type": "Property",
  "value": 100,
  "observedAt": "2022-03-14T13:00:00.000Z"
}
```

Results in an overwrite of the *value* and *observedAt* sub-Attributes, leaving the *unitCode* sub-Attribute untouched as shown:

```
{
  "temperature": {
    "type": "Property",
    "value": 100,
    "unitCode": "CEL"
    "observedAt": "2022-03-14T13:00:00.000Z"
  }
}
```

EXAMPLE 2: Given an Entity containing the following *Property*:

```
{
  "temperature": {
    "type": "Property",
    "value": 25,
    "unitCode": "CEL"
    "observedAt": "2022-03-14T01:59:26.535Z"
  }
}
```

Applying an **update attributes** operation (as defined in clause 5.6.2) onto the Entity as a whole with the following Entity Fragment payload:

```
{
  "temperature": {
    "type": "Property",
    "value": 100,
    "observedAt": "2022-03-14T13:00:00.000Z"
  }
}
```

Results in an overwrite of the whole *temperature* Attribute - other Attributes would remain untouched. The result is that the *value* and *observedAt* sub-Attributes are updated and the *unitCode* sub-Attribute is removed as shown:

```
{
  "temperature": {
    "type": "Property",
    "value": 100,
    "observedAt": "2022-03-14T13:00:00.000Z"
  }
}
```


EXAMPLE 3: Given an Entity containing the following *Property*:

```
{
  "temperature": {
    "type": "Property",
    "value": 25,
    "unitCode": "CEL"
    "observedAt": "2022-03-14T01:59:26.535Z"
  }
}
```

Applying an **update attributes** operation (as defined in clause 5.6.2) onto the Entity as a whole, with the following Entity Fragment payload:

```
{
  "temperature": {
    "type": "Property",
    "value": "urn:ngsi-ld:null"
  }
}
```

Results in the deletion of the whole *temperature* Attribute - all other Attributes remain untouched.

5.5.9 Pagination Behaviour

5.5.9.1 General Pagination Behaviour

When resolving NGSI-LD Query operations, NGSI-LD Systems shall exhibit the behaviour described by the present clause:

- Let **Md** be equal to the default maximum number of NGSI-LD Elements to be retrieved by the API during each query pagination iteration, as defined by the NGSI-LD implementation.
- Let **Mc** be equal to the maximum number of NGSI-LD Elements to be retrieved as requested by the NGSI-LD Client. If **Mc** is undefined then it shall be equal to **Md**.
- Let **L** be the maximum number of NGSI-LD Elements to be retrieved by the API during each query pagination iteration. **L** shall be equal to **Mc**.
- During query execution and for each pagination iteration, the query resolution mechanisms of the NGSI-LD System shall ensure that only up to a maximum of **L** NGSI-LD Elements are retrieved and returned to the NGSI-LD client, i.e. the maximum page size per iteration shall not overpass **L**. Nonetheless, implementations shall take care of not overpassing a maximum size of response payload body, which, in practice, implies that, under certain circumstances, the number of Elements retrieved per page can be lower than **L**.
- After the retrieval of each page (containing at most **L NGSI-LD Elements**) implementations shall check whether there are pending NGSI-LD Elements to be retrieved in the context of the current query. If that is the case, implementations shall flag NGSI-LD Clients of the existence of such NGSI-LD Elements. Ultimately, the flagging mechanisms used shall depend on each API binding but shall be present as mandated by the present clause.
- When flagging the existence of additional NGSI-LD Elements (pages) pending to be retrieved, generally, implementations shall provide NGSI-LD Clients pointers to get access to both the following page of NGSI-LD Elements and the previous one, according to the current pagination iteration.
- The pointer to the previous page of NGSI-LD Elements shall be included for all pagination iterations excepting the first one, as, obviously, there will be no previous NGSI-LD Elements.
- When the last page of NGSI-LD Elements is reached, only the pointer to the previous page shall be provided to NGSI-LD Clients, so that they can detect that no more NGSI-LD Elements are available.
- The pointers to NGSI-LD Elements shall contain all the parameters needed to allow NGSI-LD Clients to retrieve the next and previous page, without further interactions with the API.

While iterating over a set of pages, there might be changes in the target result set, due to additions or removals of NGSI-LD Elements occurring in between. Implementations may detect those situations and may warn NGSI-LD Clients appropriately. During pagination, the same `@context` shall be used. An attempt to use a different `@context` should result in an *BadRequestData* error.

5.5.9.2 Pagination option using limit and offset

The general pagination behaviour described in clause 5.5.9.1 only requires pointers to the following and previous pages, which can be implemented in a completely opaque way. Pagination may also be implemented in a transparent way, giving the Context Consumer more control over the process. In this case, the parameters *Limit* and *offset* should be used, allowing the Context Consumer to adapt the size of the page using *Limit* and jump to a desired set of elements in the results using *offset*.

5.5.9.3 Pagination with Entity maps

In the case of queries based on Entity maps, the set of Entities considered for the result is fixed with the initial query creating the Entity map. However, the Entity information itself can be dynamic, so filters shall be rechecked before returning results. In the case of split Entities, the Entities in the Entity map can only be considered as candidate Entities, since, at the time of Entity map creation, the Entities have not been aggregated and thus the filters could not be applied. This can only be done when preparing a page for pagination. Thus, Entities not or no longer fitting the query shall be removed from the Entity map during pagination. Pages shall always be filled to the maximum, as long as Entities are available. When using the previous link, the set of Entities on the previous page may not be the same as before, due to dynamic changes resulting in Entities no longer fulfilling the filter criteria of the query. As a result, the first page when going backward, and the last page, when going forward, may have less than the maximum number of Entities.

5.5.10 Multi-Tenant Behaviour

If a `Tenant` is specified for an NGSI-LD operation, the operation shall only be applied to information related to the specified `Tenant`. If no `Tenant` is specified, the operation shall apply to the implicitly existing default `Tenant`. If a `Tenant` is explicitly specified, but the system implementing the NGSI-LD API does not support multi-tenancy, an error of type *NoMultiTenantSupport* should be raised.

In case an operation applies to a `Tenant`, this information shall also be provided in the response to the operation. This also applies to notifications sent as a result of subscriptions (clauses 5.8 and 5.11).

A `Tenant` is represented in form of a String. How the `Tenant` is specified for an API operation is protocol binding specific. How `Tenants` are created, is implementation-specific.

One implementation option is to support the implicit creation of `Tenants`. This means that a `Tenant` is implicitly created when an NGSI-LD operation for creating information targets a new `Tenant`; this is the case for:

- Create Entity (clause 5.6.1).
- Batch Entity Creation (clause 5.6.7).
- Create or Update Temporal Evolution of an Entity (clause 5.6.11).
- Create Subscription (clause 5.8.1).
- Register Context Source (clause 5.9.2).
- Create Context Source Registration (clause 5.11.2).

All other NGSI-LD operations, e.g. for retrieving, updating, appending or deleting information that target a non-existing `Tenant` should raise an error of type *NonexistentTenant*.

If the system implementing the NGSI-LD API does not support multiple `Tenants`, the attempt to register a `Context Source` with `Tenant` information in the `Context Source Registration` should also result in an error of type *NoMultiTenantSupport*.

5.5.11 More than one instance of the same Entity in an Entity array

5.5.11.0 Foreword

The following operations operate on an array of entities (as input payload):

- Batch Entity Creation (clause 5.6.7).
- Batch Entity Creation or Update (Upsert) (clause 5.6.8).
- Batch Entity Update (clause 5.6.9).
- Batch Entity Delete (clause 5.6.10).
- Batch Entity Merge (clause 5.6.20).

It is allowed for such an input Entity array to contain more than one instance of the same entity (those instances have identical ids).

In order for such a request to be correctly handled, those instances that have the same *id* are processed by the Broker in the order they have in the array: the higher the index in the array, the later it will be processed. If the order is altered, the outcome may be altered.

All Entities and Attributes in the batch will get the same *modifiedAt* timestamp, so it makes sense to distinguish them via the *observedAt* temporal property.

Implementations shall treat the entity instances as if they had all arrived in separate requests.

The following clauses specify the behaviour in each case.

5.5.11.1 Batch Entity Creation case

The first occurrence of an entity in the input array (the oldest one) is used for the creation of the entity. Any subsequent instance of the same entity is reported as an error (entity already exists) in the response.

5.5.11.2 Batch Entity Creation or Update (Upsert) case

This operation has two modes of operation, with an optional flag to select between the two. The default behaviour is to replace any already existing entities, while the optional behaviour is to update already existing entities. Non existing entities are created in both modes.

If the entity does not yet exist, the first occurrence of an entity is used to create the entity, and subsequent instances of that same entity are used to either replace (default behaviour) or to update (optional behaviour) the entity. These replace or update operations shall be done **in chronological order**.

Only the entity resulting from merging all of the entity instances, in the correct order, is maintained in the current state (as defined in clause 4.3.1). For Temporal Evolution of Entities (as defined in clause 4.3.1), all entity instances shall be taken into account, in the correct order.

5.5.11.3 Batch Entity Update case

This operation has two modes of operation, with an optional flag to select between the two. The default behaviour is to replace any already existing attributes of the entities, while the optional behaviour is to preserve already existing attributes of the entities.

Brokers shall send separate notifications for each individual update, taking throttling into account.

5.5.11.4 Batch Entity Delete case

The Batch Entity Delete operation has as input an array of Entity IDs, for the entities to be deleted. If an Entity ID is replicated in the array, the first occurrence will delete the entity, while subsequent occurrences of the same Entity ID will provoke an error in the response (entity does not exist).

5.5.11.5 Batch Entity Merge case

The Batch Entity Merge operation has as input an array of Entity IDs, for the entities to be merged. If an Entity ID is replicated in the array, these merge operations shall be done **in chronological order**. Only the entity resulting from merging all of the entity instances, in the correct order, is maintained in the current state (as defined in clause 4.3.1). For **Temporal Evolution of Entities** (as defined in clause 4.3.1), all entity instances shall be taken into account, in the correct order.

5.5.12 Merge Patch Behaviour

The merge patch procedure modifies an existing NGSI-LD element by applying the set of changes found in an NGSI-LD Fragment data to the target resource. Unlike the partial update patch behaviour (described in clause 5.5.8), which replaces the complete element on the first level, e.g. a whole Attribute, the procedure described in this clause merges the provided information with the existing information up to an arbitrary depth, e.g. including going into JSON objects representing a Property value.

When merging NGSI-LD Entities using NGSI-LD Fragments, implementations shall determine the exact set of changes being requested by comparing the content of the provided Fragment (patch) against the current content (a JSON-LD object) of the target element.

With respect to merge operations, implementations shall perform an algorithm equivalent to the one described below (adapted from IETF RFC 7396 [16], in order to observe the name to URI expansion rules and JSON-LD *null* processing):

- For each member of the Fragment perform the term to URI expansion.
- If the provided Fragment (a JSON Merge Patch document) contains members that do not appear within the target (their URIs do not match), those members are added to the target.
- For each member of the Fragment contained by the target, the target member value is **merged** with the value given in the Fragment. NGSI-LD Nulls within the Fragment are given special meaning to indicate the removal of existing values within the target member value. In the case of a member representing a reified Property or Relationship including a *datasetId*, such member is only updated if the *datasetId* is the same, otherwise the member of the Fragment is added as a new instance to the target. If no *datasetId* is present, the default Attribute instance is targeted and merged if present and otherwise added. In case of a member *type* (of an Entity) in Entity Fragments, all included Entity Types are added, if they are not already contained in the *type* member of the target.
- For each member of the Fragment, whose value is an NGSI-LD Null, contained by the target, the target member is removed. In the case of deleting a specific Attribute instance with a *datasetId*, the handling shall be according to the description in clause 5.6.5. A *datasetId* cannot be deleted by setting it to the value "urn:ngsi-ld:null".

EXAMPLE 1: Given an Entity containing the following *Property*:

```
{
  "temperature": {
    "type" : "Property",
    "value" : 25,
    "unitCode": "CEL"
    "observedAt": "2022-03-14T01:59:26.535Z"
  }
}
```

Applying a **merge entity** operation (as defined in clause 5.6.17) onto the Entity as a whole, with the following Entity Fragment payload:

```
{
  "temperature": {
    "type" : "Property",
    "value" : 100,
    "observedAt": "2022-03-14T13:00:00.000Z"
  }
}
```

```
}
```

Results in the update of the *value* and *observedAt* sub-Attributes and leaves the *unitCode* sub-Attribute untouched, as shown:

```
{
  "temperature": {
    "type" : "Property",
    "value" : 100,
    "unitCode": "CEL",
    "observedAt": "2022-03-14T13:00:00.000Z"
  }
}
```

EXAMPLE 2: Given an Entity containing the following *Property*:

```
{
  "address": {
    "type" : "Property",
    "value" : {
      "street": "Straße des 17. Juni",
      "city": "Berlin",
      "country": "Germany"
    }
  }
}
```

Applying a **merge entity** operation (as defined in clause 5.6.17) onto the Entity as a whole with the following Entity Fragment payload:

```
{
  "address": {
    "type" : "Property",
    "value" : {
      "street": "Pariser Platz",
      "country": "urn:ngsi-ld:null"
    }
  }
}
```

Results in the updating of the *address* Attribute value applying the JSON Merge Patch processing rules as defined in IETF RFC 7396 [16], updating *street* and removing *country* resulting as shown:

```
{
  "address": {
    "type" : "Property",
    "value": {
      "street": "Pariser Platz",
      "city": "Berlin"
    }
  }
}
```

5.5.13 Limiting operations to local scope

The API provides a binding-specific mechanism to limit the execution of operations to a local scope, i.e. to only execute it based on information available in the *Context Source* or *Context Broker* directly targeted by the request. This enables limiting cascading distributed operations (clause 4.3.6.4) and, in some cases, also enables broader local operations, e.g. pertaining to **all** entities, which would be too expensive in the distributed case.

The *localOnly* member in the *Subscription* (clause 5.5.12) requests that the subscription only matches Entities stored locally.

The *localOnly* member in the *RegistrationManagementInfo* (clause 5.2.34) of a *Context Source* *Registration* request that, when contacting the registered *Context Source*, a local scope shall be specified.

Thus, the registered `Context Source` only provides its local information, not information from further `Context Sources` in its own `Context Registry`.

If the request limits the execution of the operation to the local scope, `Context Source Registrations` from the `Context Registry` will not be used.

Operations on a `Snapshot` (see clause 5.5.15) are always implicitly local scope, overriding setting a local scope to *false*, i.e. such a setting is to be ignored by implementations.

5.5.14 Distributed Transactional Behaviour

The following operations may occur as part of a distributed transactional request:

- Retrieve Entity (clause 5.7.1).
- Query Entities (clause 5.7.2).
- Retrieve Temporal Evolution of an Entity (clause 5.7.3).
- Query Temporal Evolution of Entities (clause 5.7.4).

In the case that a client wishes to indicate to a `Context Broker` that a request is part of a unitary sequence of requests, the `Context Broker` shall create and cache an `EntityMap` for future use and should return the location of said `EntityMap` in a specific field. The caching strategy and expiry time shall take into account a suggested expiry time, if present, and depend on implementation specific configurations. `Context Sources` should indicate that they do not support `EntityMaps`, through declining to return the location of an `EntityMap` when requested to do so.

If a subsequent request references an existent `EntityMap`, it shall be used for the purposes of Entity registration matching, and queries shall be filtered to only consider the Entities listed. Subsequent requests referencing an `EntityMap` shall use the same parameters as in the original request that created the `EntityMap`, except for the specification of Entity identifiers or parameters related to pagination, or, in the case of temporal requests, the temporal query. If an `EntityMap` has expired, or cannot be accessed, no inference can be made as to which entities are held within the `Context Sources` and a new one shall be created. An `EntityMap` fixes the Entities to be considered for subsequent requests based on it. The creating `Context Source` shall remove Entities from the `EntityMap` that do not match the filters of the query at the time of processing. Other components shall only be allowed to update the expiry timestamp of the `EntityMap`, which can optionally be extended if the `Context Sources` implementation allows for it.

5.5.15 Snapshot Behaviour

If a `Snapshot` (see clause 4.3.7) is specified for an NGSI-LD operation, the operation shall only be applied to information related to the specific `Snapshot`. Operations permitted on `Snapshots` are the same as those specified for the Core API and, if supported, the Temporal API (see Table 4.3.5-1). This means all operations are implicitly limited to a local scope with all relaxations allowing broader operations (see clause 5.5.13) and there is no need to explicitly set the local scope as a request parameter. The implicit limitation to local scope also means that `Context Source Registrations` from the `Context Registry` will not be used.

`Snapshots` are explicitly created, and the `Snapshot` identifier is provided on creation. How the `Snapshot` identifier is specified for an API operation is protocol binding specific.

The `Snapshot` concept is orthogonal to the `Tenant` concept (see clause 4.14). `Snapshots` can also be created on `Tenants`. To execute an operation on a `Snapshot` created on a `Tenant`, both `Snapshot` and `Tenant` (see clause 5.5.10) have to be specified for an API operation.

If an implementation determines that it is low on resources, it may delete one or more `snapshots`. For determining the order in which `snapshots` are to be deleted, the `snapshotPriority`, ranging from 1 (minimum) to 10 (maximum) with 5 being the default priority, should be considered. An implementation may also consider other aspects like the `expiresAt`, or `lastUsedAt` timestamp for such a decision.

5.6 Context Information Provision

5.6.1 Create Entity

5.6.1.1 Description

This operation allows creating a new NGSI-LD Entity.

5.6.1.2 Use case diagram

A Context Producer can create an Entity within an NGSI-LD system as shown in Figure 5.6.1.2-1.




Figure 5.6.1.2-1: Create entity use case

5.6.1.3 Input data

A JSON-LD document representing an NGSI-LD Entity as mandated by clause 5.2.4.

5.6.1.4 Behaviour

Implementations shall exhibit the following behaviour:

- Execute the behaviour defined in clause 5.5.4 on JSON-LD validation.
- If an **exclusive** Context Source Registration already exists for this Entity id (URI), Attributes from matching input data are forwarded for remote processing:
 - For matching Registrations where the Create Entity operation is supported, the operation is forwarded to the registration endpoint. If the endpoint then raises an error, this shall result in an error in case the complete create failed or in a partial success if some parts of the create succeeded.
 - For matching Registrations where the Create Entity operation is not supported, this shall result in an error of type *Conflict* if the complete Create Entity operation failed or in a partial success if some parts of it succeeded.

The matching Attributes are then removed from the Fragment and not processed further.

- If any **redirect** Context Source Registrations exist that match against the input data, that input data is forwarded for remote processing by one or more matching endpoints:
 - For matching Registrations where the Create Entity operation is supported, matching input data is forwarded. If any such endpoint then raises an error, this shall result in an error in case the complete create has failed or in a partial success if some parts of the create has succeeded.
 - For matching **redirect** Registrations where the Create Entity operation is not supported, this shall result in an error of type **Conflict** if the complete Create Entity operation failed or in a partial success if some parts of it succeeded.

The matching Attributes are then removed from the Fragment and not processed further.

- For any **inclusive** Context Source Registrations that match against the remaining input data, that input data is also forwarded for remote processing by matching endpoints in case the Create Entity operation is supported.
- If the Entity already exists locally this shall result in an error of type **AlreadyExists**, if the complete Create Entity operation has failed or in a partial success if some parts of it has succeeded.
- Any remaining input data shall be used to create the Entity locally.

5.6.1.5 Output data

A URI identifying the newly created Entity.

5.6.2 Update Attributes

5.6.2.1 Description

This operation allows modifying an existing NGSI-LD Entity by updating **already existing** Attributes (Properties or Relationships) and by appending non-existing ones.

5.6.2.2 Use case diagram

A Context Producer can update Attributes within an NGSI-LD system as shown in Figure 5.6.2.2-1.




Figure 5.6.2.2-1: Update Attributes use case

5.6.2.3 Input data

- A URI representing the id of the Entity to be updated (target Entity).
- A selector of Entity types as specified by clause 4.17 (optional).
- A JSON-LD document representing an NGSI-LD Entity Fragment.

5.6.2.4 Behaviour

- If the Entity ID is not present or it is not a valid URI then an error of type *BadRequestData* shall be raised.
- If the NGSI-LD endpoint does not know about the target Entity, because there is no existing Entity whose id (URI), and where specified type, is equivalent to the target entity held locally, and no matching registrations apply (see clause 5.12), an error of type *ResourceNotFound* shall be raised.
- Execute the behaviour defined in clause 5.5.4 on JSON-LD validation. NGSI-LD Nulls should be supported by this operation. If NGSI-LD Nulls are found in the payload, but are not supported, an error of type *OperationNotSupported* shall be raised.
- If an **exclusive** or **redirect** Context Source Registration matches against the input data, Attributes from matching input data are forwarded for remote processing. For each matching registration:
 - If the Update Attributes operation is supported by the registration (see clause 4.3.6), matching input data is forwarded to the Registration endpoint.
 - If the Update Attributes operation is not supported by the registration, this shall result in an error of type *Conflict* if the complete update failed or in a partial success if some parts of the update succeeded.
- The matching Attributes are then removed from the Fragment and not processed further.
- If there are remaining Attributes, for any **inclusive** Context Source Registrations that match against the remaining input data, that input data is also forwarded for remote processing to matching endpoints in case the Update Attributes operation is supported by the matched registration.
- Then, implementations shall perform a partial update patch operation over the remains of the target Entity as mandated by clause 5.5.8, using the following procedure.
- For each Attribute (Property or Relationship) included by the Entity Fragment at root level:
 - If the target Entity does not include a matching Attribute (considering term expansion rules as mandated by clause 5.5.7) then such Attribute shall be appended to the target Entity.
 - If the target Entity already includes a matching Attribute (considering term expansion rules as mandated by clause 5.5.7):
 - If a *datasetId* is present in the Attribute included by the Entity Fragment:
 - If an Attribute instance in the target Entity has the same *datasetId* and the Attribute value is not NGSI-LD Null, then the existing Attribute instance with the specified *datasetId* in the target Entity shall be replaced by the new one supplied. The system generated *createdAt* Temporal Property as defined in clause 4.8 shall remain unchanged.
 - If an Attribute instance in the target Entity has the same *datasetId* and the Attribute value is NGSI-LD Null then the existing Attribute instance with the specified *datasetId* in the target Entity shall be deleted.
 - Otherwise the Attribute instance with the specified *datasetId* shall be appended to the target Entity.

- If no *datasetId* is present in the Attribute included by the Entity Fragment, the default Attribute instance is targeted:
 - If the default Attribute instance is present and the Attribute value is not NGSI-LD Null, then the existing Attribute in the target Entity shall be replaced by the new one supplied. The system generated *createdAt* Temporal Property as defined in clause 4.8 shall remain unchanged.
 - If the default Attribute instance is present and the Attribute value is NGSI-LD Null, then the existing Attribute in the target Entity shall be deleted.
 - Otherwise the default Attribute instance shall be appended to the target Entity.
- If *type* is included in the Fragment and it includes Entity Type names that are not yet in the target Entity, add them to the list of Entity Type names of the target Entity.
- If *scope* is included in the Fragment and the target entity includes *scope*, replace the scope by the one included in the Fragment, otherwise ignore it.

5.6.2.5 Output data

- A status code indicating whether all the new Attributes were updated or only some of them.
- List of Attributes (Properties or Relationships) actually updated.

5.6.3 Append Attributes

5.6.3.1 Description

This operation allows modifying an NGSI-LD Entity by adding new attributes (Properties or Relationships).

5.6.3.2 Use case diagram

A Context Producer can append new Attributes to an existing Entity within an NGSI-LD system as shown in Figure 5.6.3.2-1.




Figure 5.6.3.2-1: Append Attributes use case

5.6.3.3 Input data

- A URI representing the id of the E to be modified (target Entity).
- A selector of Entity types as specified by clause 4.17 (optional).
- A JSON-LD document representing an NGSI-LD Entity Fragment.
- An optional flag indicating whether overwriting existing Attributes within the append operation should be permitted or denied. By default, Attribute overwrites are permitted.

5.6.3.4 Behaviour

The following behaviour shall be exhibited by compliant implementations:

- If the Entity ID is not present or it is not a valid URI then an error of type *BadRequestData* shall be raised.
- If the NGSI-LD endpoint does not know about this Entity, because there is no existing Entity which id (URI), and where specified type, is equivalent held locally to the one passed as a parameter, and no matching registrations apply (see clause 5.12), an error of type *ResourceNotFound* shall be raised.
- The behaviour defined in clause 5.5.4 on JSON-LD validation.
- If an **exclusive** or **redirect** Context Source Registration matches against the input data, the Attributes from matching input data are forwarded for remote processing. For each matching registration:
 - If the Append Attributes operation is supported by the registration (see clause 4.3.6), matching input data is forwarded to the Registration endpoint.
 - If the Append Attributes operation is not supported by the registration, this shall result in an error of type *Conflict* if the complete append failed or in a partial success if some parts of the append succeeded.
- The matching Attributes are then removed from the Fragment and not processed further.
- For any **inclusive** Context Source Registrations that match against the remaining input data, that input data is also forwarded for remote processing to matching endpoints in case the Append Attributes operation is supported.
- Then, implementations shall perform an Append Attributes operation over the remains of the target Entity as using the following procedure:
 - For each Attribute (Property or Relationship) included by the Entity Fragment at root level:
 - If a *datasetId* is present in the Attribute included by the Entity Fragment:
 - If no Attribute instance of the same target Entity exists that has the same *datasetId*, then such an Attribute shall be appended to the target Entity.
 - If an Attribute instance of the same target Entity exists that has the same *datasetId*:
 - If overwrite is allowed, then the existing Attribute with the specified *datasetId* in the target Entity shall be replaced by the new one supplied. The system generated *createdAt* Temporal Property as defined in clause 4.8 shall remain unchanged.
 - If overwrite is not allowed, the existing Attribute with the specified *datasetId* in the target Entity shall be left untouched.
 - If no *datasetId* is present in the Attribute included by the Entity Fragment:
 - If no default Attribute instance of the same target Entity exists, then such Attribute shall be appended to the target Entity.
 - If a default Attribute instance of the same target Entity exists:

- If *overwrite* is allowed, then the existing default Attribute in the target Entity shall be replaced by the new one supplied. The system generated *createdAt* Temporal Property as defined in clause 4.8 shall remain unchanged.
- If *overwrite* is not allowed the existing default Attribute in the target Entity shall be left untouched.
- If *type* is included in the Fragment and it includes Entity Type names that are not yet in the target Entity, add them to the list of Entity Type names of the target Entity.
- If *scope* is included in the Fragment and *overwrite* is allowed, the scope of the target Entity will become the one included in the Fragment. Otherwise, the Scopes in the Fragment that are not part of the value of *scope* of the target Entity will be appended to the value of the *scope* of the target Entity. If there is more than one Scope, the value of *scope* is represented as a JSON array containing all Scopes.

5.6.3.5 Output data

- A status code indicating whether all the new Attributes were appended or only some of them.
- List of Attributes (Properties and/or Relationships) actually appended.

5.6.4 Partial Attribute update

5.6.4.1 Description

This operation allows performing a **partial update on an NGSI-LD Entity's Attribute** (Property or Relationship). A partial update only changes the elements provided in an Entity Fragment, leaving the rest as they are. This operation supports the deletion of sub-Attributes but not the deletion of the whole Attribute itself.

5.6.4.2 Use case diagram

A Context Producer can carry out a partial Attribute update of an Entity within an NGSI-LD System as shown in Figure 5.6.4.2-1.




Figure 5.6.4.2-1: Partial Attribute update use case

5.6.4.3 Input data

- Entity ID (URI) of the concerned Entity, the target Entity.
- A selector of Entity types as specified by clause 4.17 (optional).

- Target Attribute (Property or Relationship) to be modified, identified by a name.
- A JSON-LD document representing an NGSI-LD Attribute Fragment.

5.6.4.4 Behaviour

- If the target Entity ID is not a valid URI or it is not present, then an error of type *BadRequestData* shall be raised.
- If the target Attribute name is not valid or it is not present, then an error of type *BadRequestData* shall be raised.
- If the target Attribute is *scope*, then an error of type *BadRequestData* shall be raised.
- The behaviour defined in clause 5.5.4 on JSON-LD validation. NGSI-LD Nulls should be supported by this operation. If NGSI-LD Nulls are found in the payload, but are not supported, an error of type *OperationNotSupported* shall be raised.
- If the NGSI-LD endpoint does not know about the target Entity, because there is no existing Entity whose id (URI), and where specified type, is equivalent held locally, and no matching registrations apply (see clause 5.12), then an error of type *ResourceNotFound* shall be raised.
- If an **exclusive** or **redirect** Context Source Registration matches against the input data, the Attributes from matching input data are forwarded for remote processing. For each matching registration:
 - If the Partial Attribute update operation is supported by the registration (see clause 4.3.6), matching input data is forwarded to the Registration endpoint.
 - If the Partial Attribute update operation is not supported by the registration, this shall result in an error of type *Conflict* in case the complete partial Attribute update failed, or in a partial success if some parts of the partial Attribute update succeeded.

No further processing is required.

- For any **inclusive** Context Source Registrations that match against the input data, that input data is also forwarded for remote processing to matching endpoints in case the Partial Attribute update operation is supported.
- Apply term expansion as mandated by clause 5.5.7, so that the fully qualified name (URI) associated to the target Attribute is properly obtained.
- If the target Entity does not contain the target Attribute:
 - as a default instance in case no *datasetId* is present;
 - as an instance with the specified *datasetId* if present;
 then an error of type *ResourceNotFound* shall be raised.
- Perform a partial update patch operation on the target Attribute following the algorithm mandated by clause 5.5.8. If present in the provided NGSI-LD Entity Fragment, the type of the Attribute has to be the same as the type of the targeted Attribute fragment, i.e. it is not allowed to change the type of an Attribute.

5.6.4.5 Output data

None.

5.6.5 Delete Attribute

5.6.5.1 Description

This operation allows deleting an NGSI-LD Attribute (Property or Relationship). The Attribute itself and all its children shall be deleted.

5.6.5.2 Use case diagram

A Context Producer can delete a specific Attribute within an NGSI-LD system as shown in Figure 5.6.5.2-1.




Figure 5.6.5.2-1: Delete Attribute use case

5.6.5.3 Input data

- Entity ID (URI) of the concerned Entity, the target Entity.
- A selector of Entity types as specified by clause 4.17 (optional).
- Target Attribute (Property or Relationship) to be deleted, identified by a name.
- An optional parameter identifying the *datasetId* of the target Attribute instance to be deleted. Otherwise the default Attribute instance is targeted.
- An optional flag *deleteAll* indicating whether also all target Attribute instances with a *datasetId* are to be deleted.
- An optional JSON-LD *@context*.

5.6.5.4 Behaviour

- If the target Entity ID is not a valid URI or it is not present, then an error of type **BadRequestData** shall be raised.
- If the target Attribute name is not a valid name or it is not present, then an error of type **BadRequestData** shall be raised.
- If an **exclusive** or **redirect** Context Source Registration matches against the input data, the input data (see clause 5.12), the input data is forwarded. For each matching registration:
 - If the Delete Attribute operation is supported by the registration (see clause 4.3.6), matching input data is forwarded to the Registration endpoint.
 - If the Delete Attribute update operation is not supported by the matched registration, this shall result in an error of type **Conflict** in case the complete delete Attribute failed, or in a partial success if some parts of the delete Attribute succeeded.

No further processing is required.

- If the NGSI-LD endpoint does not know about the target Entity, because there is no existing Entity whose id (URI), and where specified type, is equivalent held locally, and no matching registrations apply, then an error of type *ResourceNotFound* shall be raised.
- For any **inclusive** Context Source Registrations that match against the input data, that input data is also forwarded for remote processing to matching endpoints in case the Delete Attribute operation is supported by the matched registration.
- Apply term expansion as mandated by clause 5.5.7 so that the fully qualified name (URI) associated to the target Attribute is properly obtained.
- If the target Entity does not contain the target Attribute then an error of type *ResourceNotFound* shall be raised.
- If the target Attribute is *scope*, remove the *scope* Attribute from the target Entity.
- If the *deleteAll* flag is set, remove all target Attribute instances from the target Entity.
- Otherwise:
 - if a *datasetId* parameter is provided, remove only the target Attribute instance from the given dataset whose *datasetId* matches the parameter;
 - if no *datasetId* parameter is provided, remove the default target Attribute instance from the target Entity.

5.6.5.5 Output data

None.

5.6.6 Delete Entity

5.6.6.1 Description

This operation allows deleting an NGSI-LD Entity.

5.6.6.2 Use case diagram

A Context Producer can completely delete an Entity within an NGSI-LD system as shown in Figure 5.6.6.2-1.




Figure 5.6.6.2-1: Delete Entity use case

5.6.6.3 Input data

- Entity ID (URI) of the Entity to be deleted, the target Entity.
- A selector of Entity types as specified by clause 4.17 (optional).

5.6.6.4 Behaviour

- If the target Entity ID is not present or it is not a valid URI, then an error of type **BadRequestData** shall be raised.
- If the NGSI-LD endpoint does not know about the target Entity, because there is no existing Entity whose id (URI), and where specified type, is equivalent held locally, and no matching registrations apply (see clause 5.12), then an error of type **ResourceNotFound** shall be raised.
- If an **exclusive** or **redirect** Context Source Registration matches against the id, the request is forwarded for remote processing. For each matching registration:
 - If the Delete Entity operation is supported by the registration (see clause 4.3.6), the request is forwarded to the Registration endpoint.
 - If the Delete Entity update operation is not supported by the matched registration, this shall result in an error of type **Conflict** in case the complete delete Entity failed, or in a partial success if some parts of the delete Entity succeeded.
- For any **inclusive** Context Source Registrations that match against the remaining input data, that input data is also forwarded in the case the Delete Entity operation is supported for remote processing by matching endpoints.
- The input data shall be used to remove the entity locally if it exists.

5.6.6.5 Output data

None.

5.6.7 Batch Entity Creation

5.6.7.1 Description

This operation allows creating a batch of NGSI-LD Entities.

5.6.7.2 Use case diagram

A Context Producer can create a batch of NGSI-LD Entities within an NGSI-LD system as shown in Figure 5.6.7.2-1.




Figure 5.6.7.2-1: Create a batch of Entities use case

5.6.7.3 Input data

- A JSON-LD Array containing one or more JSON-LD documents each one representing an NGSI-LD Entity as mandated by clause 5.2.4. See clause 5.5.11.1 for information on behaviour when there is more than one instance of the same entity in the input Array.

5.6.7.4 Behaviour

Implementations shall exhibit the following behaviour:

- If the input Array is empty or contains a *null* value in any of its items an error of type **BadRequestData** shall be raised.
- Execute the behaviour defined in clause 5.5.4 on JSON-LD validation.
- Let S be an array which shall contain a list of Entity IDs, one for each NGSI-LD Entity successfully created. S shall be initialized to the empty array.
- Let E be an array which shall contain a list of **BatchEntityError** as defined by clause 5.2.17, one for each NGSI-LD Entity which resulted in error. E shall be initialized to the empty array.
- For each Context Source Registration CSR in the Context Registry:
 - Let IN be a copy of the original input array.
 - Remove from IN all Entities not matched by CSR and remove non-matching Attributes from the remaining Entities.
 - Remove all Attributes from the remaining Entities in IN for which there is a matching **exclusive** Context Source Registration, which is not CSR itself.
 - Remove all Attributes from the remaining Entities in IN for which there is a matching **redirect** Context Source Registration, unless CSR is a redirect Context Source Registration itself.
 - If IN is empty, continue with the next Context Source Registration if there is any.
 - If the Batch Entity Creation operation is supported by CSR:
 - Forward the Batch Entity Creation request with IN as input Array.

- Merge the returned list of Entities successfully created with S.
- Merge the returned list of Entities in Error with E.
- Otherwise, if the Create Entity operation (clause 5.6.1) is supported by CSR:
 - For each Entity EN in the input array:
 - Forward a Create Entity request for Entity EN.
 - Merge any successful result(s) for Entity EN created with S.
 - Merge any error result(s) for Entity EN created with E.
- Otherwise:
 - In case CSR is an **exclusive** or **redirect** Context Source Registration, add an Error of type *Conflict* for each Entity in IN to E.
- For each of the NGSI-LD Entities included in the input Array execute the behaviour defined by clause 5.6.1, but limited to a local operation, as follows:
 - If the Entity was successfully created, then add the corresponding Entity ID to the S array.
 - If the Entity creation failed, then a new *BatchEntityError* shall be added to E containing the failed Entity ID and the related *ProblemDetails*.

5.6.7.5 Output data

- The list of Entities successfully created (S Array), if all Entities were created correctly; or
- the list of Entities successfully created (S Array) and the list of Entities in error (E Array), if only some or none of the Entities were created.

5.6.8 Batch Entity Creation or Update (Upsert)

5.6.8.1 Description

This operation allows creating a batch of NGSI-LD Entities, updating each of them if they already existed. In some database jargon this kind of operation is known as "upsert".

5.6.8.2 Use case diagram

A Context Producer can create or update a batch of Entities within an NGSI-LD system as shown in Figure 5.6.8.2-1.




Figure 5.6.8.2-1: Upsert a batch of Entities use case

5.6.8.3 Input data

- A JSON-LD Array containing one or more JSON-LD documents each one representing an Entity as mandated by clause 5.2.4. See clause 5.5.11.2 for information on behaviour when there is more than one instance of the same entity in the input Array.
- An optional flag indicating the update mode (only applies in case the Entity already exists):
 - Replace. All the existing Entity content shall be replaced (default mode).
 - Update. Existing Entity content shall be updated.

5.6.8.4 Behaviour

Implementations shall exhibit the following behaviour:

- If the input Array is empty or contains a *null* value in any of its items, an error of type **BadRequestData** shall be raised.
- Execute the behaviour defined in clause 5.5.4 on JSON-LD validation.
- Let S be an array which shall contain a list of Entity IDs, one for each NGSI-LD Entity which was successfully processed. S shall be initialized to the empty array.
- Let E be an array which shall contain a list of **BatchEntityError** as defined by clause 5.2.17, one for each NGSI-LD Entity which resulted in error. E shall be initialized to the empty array.
- For each **Context Source Registration** CSR in the **Context Registry**:
 - Let IN be a copy of the original input array.
 - Remove from IN all Entities not matched by CSR and remove non-matching Attributes from the remaining Entities.
 - Remove all Attributes from the remaining Entities in IN for which there is a matching **exclusive** **Context Source Registration**, which is not CSR itself.
 - Remove all Attributes from the remaining Entities in IN for which there is a matching **redirect** **Context Source Registration**, unless CSR is a redirect **Context Source Registration** itself.

- If IN is empty, continue with the next Context Source Registration if there is any.
- If the Batch Entity Creation or Update (Upsert) operation is supported by CSR:
 - Forward the Batch Entity Creation or Update (Upsert) request with IN as input Array.
 - Merge the returned list of Entities successfully created with S.
 - Merge the returned list of Entities in Error with E.
- Otherwise, if the Create Entity operation (clause 5.6.1) is supported by CSR:
 - For each Entity EN in the input array:
 - Forward a Create Entity request for Entity EN.
 - If an error of type *AlreadyExists* is returned:
 - If the Replace Entity operation (clause 5.6.18) is supported by CSR and the value of the update mode flag is *Replace* or the flag is not set, forward a Replace Entity request for Entity EN.
 - Otherwise, if the Update Attributes operation (clause 5.6.2) is supported by CSR and the value of the update mode flag is *Update*, forward an Update Attributes request for Entity EN.
 - Otherwise add an *OperationNotSupported* Error for Entity EN related to CSR to E.
 - Merge any successful result(s) for Entity EN created or updated with S.
 - Merge any error result(s) for Entity EN created or updated with E.
- Otherwise, if the Replace Entity operation (clause 5.6.18) is supported by CSR and the value of the update mode flag is *Replace* or the flag is not set:
 - Forward a Replace Entity request for Entity EN.
 - Merge any successful result(s) for Entity EN updated with S.
 - Merge any error result(s) for Entity EN updated with E.
- Otherwise, if the Update Attributes operation (clause 5.6.2) is supported by CSR and the value of the update mode flag is *Update*:
 - Forward an Update Attributes request for Entity EN.
 - Merge any successful result(s) for Entity EN updated with S.
 - Merge any error result(s) for Entity EN updated with E.
- Otherwise:
 - In case CSR is an **exclusive** or **redirect** Context Source Registration, add an Error of type *Conflict* for each Entity in IN to E.
- For each of the NGSI-LD Entities included in the input Array implementations shall:
 - Create the Entity locally if it does not exist (i.e. no Entity with the same Entity ID is present) executing the behaviour defined by clause 5.6.1, but limited to a local operation.
 - If there were an existing Entity with the same Entity ID, it shall be completely replaced by the new Entity content provided, if the requested update mode is 'replace'.
 - If there were an existing Entity with the same Entity ID, the behaviour defined by clause 5.6.2 shall be executed, but limited to a local operation, if the requested update mode is "update".

- If successful, the local creation or update shall be added to S. If while processing an Entity there is any kind of error or abnormal situation, a **BatchEntityError** shall be added to E containing the failed Entity ID and the related *ProblemDetails*.

5.6.8.5 Output data

- none (if all Entities already existed and are successfully updated); or
- the list of Entities successfully created (S Array), if all Entities not existing prior to this request have been successfully created and the others have been successfully updated; or
- the list of Entities successfully created or updated (S Array), and the list of Entities in error (E Array), if only some or none of the Entities have been successfully created or updated.

5.6.9 Batch Entity Update

5.6.9.1 Description

This operation allows updating a batch of NGSI-LD Entities.

5.6.9.2 Use case diagram

A Context Producer can update a batch of Entities within an NGSI-LD system as shown in Figure 5.6.9.2-1.




Figure 5.6.9.2-1: Update a batch of Entities use case

5.6.9.3 Input data

- A JSON-LD Array containing one or more JSON-LD documents each one representing an Entity as mandated by clause 5.2.4. See clause 5.5.11.3 for information on behaviour when there is more than one instance of the same entity in the input Array.
- An optional flag indicating whether Attributes shall be overwritten or not. By default, Attributes will be overwritten.

5.6.9.4 Behaviour

Implementations shall exhibit the following behaviour:

- If the input Array is empty or contains a *null* value in any of its items, an error of type ***BadRequestData*** shall be raised.
- Execute the behaviour defined in clause 5.5.4 on JSON-LD validation.
- Let S be an array which shall contain a list of Entity IDs, one for each NGSI-LD Entity which was successfully processed. S shall be initialized as the empty array.
- Let E be an array which shall contain a list of ***BatchEntityError*** as defined by clause 5.2.17, one for each NGSI-LD Entity which resulted in error. E shall be initialized as the empty array.
- For each **Context Source Registration** CSR in the **Context Registry**:
 - Let IN be a copy of the original input array.
 - Remove from IN all Entities not matched by CSR and remove non-matching Attributes from the remaining Entities.
 - Remove all Attributes from the remaining Entities in IN for which there is a matching **exclusive Context Source Registration**, which is not CSR itself.
 - Remove all Attributes from the remaining Entities in IN for which there is a matching **redirect Context Source Registration**, unless CSR is a redirect **Context Source Registration** itself.
 - Remove all Entities without Attributes from IN.
 - If IN is empty, continue with the next **Context Source Registration** if there is any.
 - If the Batch Entity Update operation is supported by CSR:
 - Forward the Batch Entity Update request with IN as input Array.
 - Merge the returned list of Entities successfully created with S.
 - Merge the returned list of Entities in Error with E.
 - Otherwise, if the Update Attributes operation (clause 5.6.2) is supported by CSR and Attribute overwrite is permitted:
 - For each Entity EN in the input array:
 - Forward an Update Attributes request for Entity EN.
 - Merge any successful result(s) for Entity EN updated with S.
 - Merge any error result(s) for Entity EN updated with E.
 - Otherwise, if the Append Attributes operation (clause 5.6.3) is supported by CSR and Attribute overwrite is not permitted:
 - For each Entity EN in the input array:
 - Forward an Append Attributes request for Entity EN with Attribute overwrite disabled:
 - Merge any successful result(s) for Entity EN updated with S.
 - Merge any error result(s) for Entity EN updated with E.
 - Otherwise:
 - In case CSR is an **exclusive** or **redirect Context Source Registration**, add an Error of type ***Conflict*** for each Entity in IN to E.

- For each of the NGSI-LD Entities included in the input Array execute the behaviour defined by clause 5.6.3, but limited to a local operation, as follows:
 - If the Entity was successfully updated (Attributes appended), then add the corresponding Entity ID to the S array.
 - If the Entity update failed, then a new **BatchEntityError** shall be added to E containing the failed Entity ID and the *ProblemDetails* associated.

5.6.9.5 Output data

- none (if all Entities are successfully updated); or
- the list of Entities successfully updated (S Array), and the list of Entities in error (E Array), if only some or none of the Entities have been successfully updated.

5.6.10 Batch Entity Delete

5.6.10.1 Description

This operation allows deleting a batch of NGSI-LD Entities.

5.6.10.2 Use case diagram

A Context Producer can delete a batch of Entities within an NGSI-LD system as shown in Figure 5.6.10.2-1.




Figure 5.6.10.2-1: Delete a batch of Entities use case

5.6.10.3 Input data

- A JSON-LD Array containing a list of Entity IDs (URIs) that are requested to be deleted. See clause 5.5.11.4 for information on behaviour when there is more than one instance of the same Entity ID in the input Array.

5.6.10.4 Behaviour

Implementations shall exhibit the following behaviour:

- If the input Array is empty or contains a *null* value in any of its items, an error of type ***BadRequestData*** shall be raised.
- Let S be an array which shall contain a list of Entity IDs, one for each NGSI-LD Entity which was successfully processed. S shall be initialized to the empty array.
- Let E be an array which shall contain a list of ***BatchEntityError*** as defined by clause 5.2.17, one for each NGSI-LD Entity which resulted in error. E shall be initialized to the empty array.
- For each Context Source Registration CSR in the Context Registry:
 - Let IN be a copy of the original input array.
 - Remove from IN all Entities not matched by CSR and remove non-matching Attributes from the remaining Entities.
 - Remove all Attributes from the remaining Entities in IN for which there is a matching **exclusive** Context Source Registration, which is not CSR itself.
 - Remove all Attributes from the remaining Entities in IN for which there is a matching **redirect** Context Source Registration, unless CSR is a redirect Context Source Registration itself.
 - Remove all Entities without Attributes from IN.
 - If IN is empty, continue with the next Context Source Registration if there is any.
 - If the Batch Entity Delete operation is supported by CSR:
 - Forward the Batch Entity Delete request with IN as input Array.
 - Merge the returned list of Entities successfully created with S.
 - Merge the returned list of Entities in Error with E.
 - Otherwise, if the Delete Entity operation (clause 5.6.6) is supported by CSR:
 - For each Entity EN in the input array:
 - Forward a Delete Entity request for Entity EN.
 - Merge any successful result(s) for Entity EN deleted with S.
 - Merge any error result(s) for Entity EN deleted with E.
 - Otherwise:
 - In case CSR is an **exclusive** or **redirect** Context Source Registration, add an Error of type ***Conflict*** for each Entity in IN to E.
 - For each of the NGSI-LD Entity IDs included in the input Array execute the behaviour defined by clause 5.6.6, but limited to a local operation, as follows:
 - If the Entity corresponding to an Entity ID was successfully deleted, then add such Entity ID to the S array.
 - If the Entity deletion failed, then a new ***BatchEntityError*** shall be added to E containing the failed Entity ID and the related *ProblemDetails*.

5.6.10.5 Output data

- none (if all Entities that already existed are successfully deleted); or
- the list of Entities successfully deleted (S Array), and the list of Entities in error (E Array), if some or all of the Entities have not been successfully deleted.

5.6.11 Create or Update (Upsert) Temporal Evolution of an Entity

5.6.11.1 Description

This operation allows creating or updating (by adding new Attribute instances) the Temporal Evolution of an Entity.

5.6.11.2 Use case diagram

A Context Producer can create the Temporal Evolution of an Entity within an NGSI-LD system as shown in Figure 5.6.11.2-1.

Similarly, if the Entity already exists then an Update scenario will be in place.




Figure 5.6.11.2-1: Create or Update (Upsert) Temporal Evolution of an Entity use case

5.6.11.3 Input data

A JSON-LD document representing the Temporal Evolution of an Entity as mandated by clause 5.2.20.

5.6.11.4 Behaviour

Implementations shall exhibit the following behaviour:

- Execute the behaviour defined in clause 5.5.4 on JSON-LD validation.
- If an **exclusive** Context Source Registration already exists for this Entity id (URI), Attributes from matching input data are forwarded for remote processing:
 - For matching Registrations where the "Create or Update (Upsert) Temporal Evolution of an Entity" operation is supported, the operation is forwarded to the registration endpoint. If the endpoint then raises an error, this shall result in an error in case the complete "Create or Update (Upsert) Temporal Evolution of an Entity" operation failed or in a partial success if some parts of it succeeded.

- For matching Registrations where the "Create Entity" operation is not supported, this shall result in an error of type **Conflict** in case the complete "Create or Update (Upsert) Temporal Evolution of an Entity" operation failed or in a partial success if some parts of it succeeded.

The matching Attributes are then removed from the Fragment and not processed further.

- If any **redirect** Context Source Registrations exist that match against the input data, that input data is forwarded for remote processing by one or more matching endpoints:
 - For matching Registrations where the "Create or Update (Upsert) Temporal Evolution of an Entity" operation is supported, matching input data is forwarded. If any such endpoint then raises an error, this shall result in an error in case the complete "Create or Update (Upsert) Temporal Evolution of an Entity" operation failed or in a partial success if some parts of it succeeded.
 - For matching **redirect** Registrations where the "Create or Update (Upsert) Temporal Evolution of an Entity" operation is not supported, this shall result in an error of type **Conflict** in case the complete "Create or Update (Upsert) Temporal Evolution of an Entity" operation failed or in a partial success if some parts of it succeeded.

The matching Attributes are then removed from the Fragment and not processed further.

- For any **inclusive** Context Source Registrations that match against the remaining input data, that input data is also forwarded for remote processing by matching endpoints.
- If the NGSI-LD endpoint already knows about this Temporal Evolution of an Entity, because there is an existing Temporal Evolution of an Entity whose id (URI) is the same, then all the Attribute instances included by the Temporal Evolution shall be added to the existing Entity as mandated by clause 5.6.12. If *type* is included in the EntityTemporal Fragment and it includes Entity Type names that are not yet in the target Temporal Evolution of an Entity, add them to the list of Entity Type names of the target Temporal Evolution of an Entity.
- Otherwise, implementations shall create the provided Temporal Evolution of an Entity.

5.6.11.5 Output data

On creation, a URI identifying the newly created Temporal Evolution of an Entity. None otherwise.

5.6.12 Add Attributes to Temporal Evolution of an Entity

5.6.12.1 Description

This operation allows modifying the Temporal Evolution of an Entity by adding new Attribute instances.

5.6.12.2 Use case diagram

A Context Producer can add new Attributes or Attribute instances to an existing Temporal Evolution of an Entity within an NGSI-LD system as shown in Figure 5.6.12.2-1.




Figure 5.6.12.2-1: Add Attributes to Temporal Evolution of an Entity use case

5.6.12.3 Input data

- Entity ID (URI) of the target Temporal Evolution of an Entity to be modified with additional Attributes.
- A JSON-LD document representing an NGSI-LD Fragment of *EntityTemporal*, including only the new Attribute instance(s), and contained by an Array.

5.6.12.4 Behaviour

The following behaviour shall be exhibited by compliant implementations:

- If the Entity ID is not present or it is not a valid URI then an error of type *BadRequestData* shall be raised.
- If the NGSI-LD endpoint does not know about the target Temporal Evolution of an Entity, because there is no existing Temporal Evolution of an Entity whose id (URI) is equivalent to the one passed as a parameter held locally and no matching registrations apply, an error of type *ResourceNotFound* shall be raised.
- The behaviour defined in clause 5.5.4 on JSON-LD validation.
- If an **exclusive** or **redirect** Context Source Registration matches against the input data, the Attributes from matching input data are forwarded for remote processing. For each matching registration:
 - If the "Add Attributes to Temporal Evolution of an Entity" operation is supported by the matched registration, matching input data is forwarded to the Registration endpoint.
 - If the "Add Attributes to Temporal Evolution of an Entity" operation is not supported by the matched registration, this shall result in an error of type *Conflict* if the complete "Add Attributes to Temporal Evolution of an Entity" operation failed or in a partial success if some parts of it succeeded.

The matching Attributes are then removed from the Fragment and not processed further.

- For any **inclusive** Context Source Registrations that match against the remaining input data, that input data is also forwarded for remote processing to matching endpoints.

- If the target Temporal Evolution of an Entity exists locally and matches against the remaining input data, implementations shall do the following:
 - For each Attribute (Property or Relationship) instance included by the *EntityTemporal* Fragment at root level:
 - The Attribute (considering term expansion rules as mandated by clause 5.5.7) instance(s) shall be added to the target Temporal Evolution of an Entity.
- If *type* is included in the EntityTemporal Fragment and it includes Entity Type names that are not yet in the target Temporal Evolution of an Entity, add them to the list of Entity Type names of the target Temporal Evolution of the Entity.

5.6.12.5 Output data

None.

5.6.13 Delete Attribute from Temporal Evolution of an Entity

5.6.13.1 Description

This operation allows deleting an Attribute (Property or Relationship) of the Temporal Evolution of an Entity. The Attribute itself and all its child NGSI-LD elements shall be deleted.

5.6.13.2 Use case diagram

A Context Producer can delete a specific Attribute of the Temporal Evolution of an Entity within an NGSI-LD system as shown in Figure 5.6.13.2-1.




Figure 5.6.13.2-1: Delete Attribute from Temporal Evolution of an Entity use case

5.6.13.3 Input data

- Entity ID (URI) of the target Temporal Evolution of an Entity to be modified by removing the target Attribute.
- Target Attribute (Property or Relationship) to be deleted, identified by a name.
- An optional parameter identifying the dataset (*datasetId*) of the target Attribute instance to be deleted.

- An optional parameter, a flag, (*deleteAll*) indicating whether all target Attribute instances are to be deleted, regardless of *datasetId*.
- An optional JSON-LD *@context*.

5.6.13.4 Behaviour

- If the target Entity ID is not a valid URI or it is not present, then an error of type **BadRequestData** shall be raised.
- If the target Attribute name is not a valid name, then an error of type **BadRequestData** shall be raised.
- If the NGSI-LD endpoint does not know about the target Entity, because there is no existing **Temporal Evolution of an Entity** whose id (URI) is equivalent held locally and no matching registrations apply, then an error of type **ResourceNotFound** shall be raised.
- If an **exclusive** or **redirect** Context Source Registration matches against the input data, the input data is forwarded. For each matching registration:
 - If the "Delete Attribute from Temporal Evolution of an Entity" operation is supported by the matched registration, matching input data is forwarded to the Registration endpoint.
 - If the "Delete Attribute from Temporal Evolution of an Entity" operation is not supported by the matched registration, this shall result in an error of type **Conflict** in case the complete "Delete Attribute from Temporal Evolution of an Entity" failed, or in a partial success if some parts of it succeeded.

No further processing is required.

- For any **inclusive** Context Source Registrations that match against the input data, that input data is also forwarded for remote processing to matching endpoints.
- If the target **Temporal Evolution of an Entity** exists locally, implementations shall do the following:
 - Apply term expansion as mandated by clause 5.5.7 so that the fully qualified name (URI) associated to the target Attribute is properly obtained.
 - If the target **Temporal Evolution of an Entity** does not contain the target Attribute then an error of type **ResourceNotFound** shall be raised.
 - If the *deleteAll* flag is set, remove all target Attribute instances from the target Entity.
- Otherwise:
 - if a *datasetId* parameter is provided, remove only any target Attribute instance from the given dataset;
 - if no *datasetId* parameter is provided, remove only the default target Attribute instance *datasetId* from the target Entity.

5.6.13.5 Output data

None.

5.6.14 Modify Attribute instance in Temporal Evolution of an Entity

5.6.14.1 Description

This operation allows modifying a specific Attribute (Property or Relationship) instance, identified by its *instanceId*, of the **Temporal Evolution of an Entity**.

This operation enables the correction of wrong information that could have been previously added to the **Temporal Evolution of an Entity**.

5.6.14.2 Use case diagram

A Context Producer can modify a specific Attribute instance, identified by a given *instanceId*, of the Temporal Evolution of an Entity within an NGSI-LD system as shown in Figure 5.6.14.2-1.




Figure 5.6.14.2-1: Modify Attribute Instance in Temporal Evolution of an Entity use case

5.6.14.3 Input data

- Entity ID (URI) of the target Temporal Evolution of an Entity to be modified by changing an instance of the target Attribute.
- Target Attribute (Property or Relationship) to be modified, identified by a name.
- Attribute instance to be modified, identified by its *instanceId*.
- A JSON-LD document representing an NGSI-LD Fragment of *EntityTemporal*, including only the new Attribute instance, contained by an Array of exactly one item.
- An optional JSON-LD @context.

5.6.14.4 Behaviour

- If the target Entity ID is not a valid URI or it is not present, then an error of type **BadRequestData** shall be raised.
- If the target Attribute name is not a valid name or it is not present, then an error of type **BadRequestData** shall be raised.
- If the target *instanceId* is not a valid URI or it is not present, then an error of type **BadRequestData** shall be raised.
- If the NGSI-LD endpoint does not know about the target Entity, because there is no existing Entity whose id (URI) is equivalent held locally and no matching registrations apply, then an error of type **ResourceNotFound** shall be raised.
- If an **exclusive** or **redirect** Context Source Registration matches against the input data, the input data is forwarded. For each matching registration:
 - If the "Modify Attribute instance in Temporal Evolution of an Entity" operation is supported by the matched registration, matching input data is forwarded to the Registration endpoint.

- If the "Modify Attribute instance in Temporal Evolution of an Entity" operation is not supported by the matched registration, this shall result in an error of type **Conflict** in case the complete Modify Attribute instance in Temporal Evolution of an Entity operation failed, or in a partial success if some parts of it succeeded.

No further processing is required.

- For any **inclusive Context Source Registrations** that match against the input data, that input data is also forwarded for remote processing to matching endpoints.
- If the target Temporal Evolution of an Entity exists locally, implementations shall do the following:
 - Apply term expansion as mandated by clause 5.5.7 so that the fully qualified name (URI) associated to the target Attribute is properly obtained.
 - If the target Temporal Evolution of an Entity does not contain the target Attribute then an error of type **ResourceNotFound** shall be raised.
 - If for the target Attribute no instance with the specified *instanceId* exists, an error of type **ResourceNotFound** shall be raised.
- Replace the target Attribute instance identified by the *instanceId* with the Attribute instance in the *EntityTemporal* Fragment. The *createdAt* property of the concerned instance shall remain unchanged, but the *modifiedAt* property shall be set to the timestamp corresponding to this modification.

5.6.14.5 Output data

None.

5.6.15 Delete Attribute instance from Temporal Evolution of an Entity

5.6.15.1 Description

This operation allows deleting one Attribute instance (Property or Relationship), identified by its *instanceId*, of the Temporal Evolution of an Entity. The Attribute itself and all its child elements shall be deleted. This operation enables the removal of individual Attribute instances that could have been previously added to the Temporal Evolution of an Entity.

5.6.15.2 Use case diagram

A Context Producer can delete an Attribute instance, identified by a given *instanceId*, of the Temporal Evolution of an Entity within an NGSI-LD system as shown in Figure 5.6.15.2-1.




Figure 5.6.15.2-1: Delete Attribute Instance from Temporal Evolution of an Entity use case

5.6.15.3 Input data

- Entity ID (URI) of the target Temporal Evolution of an Entity to be modified by deleting an instance of the target Attribute.
- Target Attribute (Property or Relationship) to be deleted, identified by a name.
- Attribute instance to be deleted, identified by its *instanceId*.
- An optional JSON-LD *@context*.

5.6.15.4 Behaviour

- If the target Entity ID is not a valid URI or it is not present, then an error of type **BadRequestData** shall be raised.
- If the target Attribute name is not a valid name or it is not present, then an error of type **BadRequestData** shall be raised.
- If the target *instanceId* is not a valid URI or it is not present, then an error of type **BadRequestData** shall be raised.
- If the NGSI-LD endpoint does not know about the target Entity, because there is no existing Entity whose id (URI) is equivalent held locally and no matching registrations apply, then an error of type **ResourceNotFound** shall be raised.
- If an **exclusive** or **redirect** Context Source Registration matches against the input data, the input data is forwarded. For each matching registration:
 - If the "Delete Attribute instance from Temporal Evolution of an Entity" operation is supported by the matched registration, matching input data is forwarded to the Registration endpoint.
 - If the "Delete Attribute instance from Temporal Evolution of an Entity" operation is not supported by the matched registration, this shall result in an error of type **Conflict** in case the complete "Delete Attribute instance from Temporal Evolution of an Entity" failed, or in a partial success if some parts of it succeeded.

No further processing is required.

- For any **inclusive** Context Source Registrations that match against the input data, that input data is also forwarded for remote processing to matching endpoints.

- If the target Temporal Evolution of an Entity exists locally, implementations shall do the following:
 - Apply term expansion as mandated by clause 5.5.7 so that the fully qualified name (URI) associated to the target Attribute is properly obtained.
- If the target Temporal Evolution of the Entity does not contain the target Attribute then an error of type *ResourceNotFound* shall be raised.
- If for the target Attribute no instance with the specified *instanceId* exists, an error of type *ResourceNotFound* shall be raised.
- Remove the instance, with the specified *instanceId*, of the target Attribute from the target Entity.

5.6.15.5 Output data

None.

5.6.16 Delete Temporal Evolution of an Entity

5.6.16.1 Description

This operation allows deleting the Temporal Evolution of an Entity.

5.6.16.2 Use case diagram

A Context Producer can completely delete the Temporal Evolution of an Entity within an NGSI-LD system as shown in Figure 5.6.16.2-1.




Figure 5.6.16.2-1: Delete Temporal Evolution of an Entity use case

5.6.16.3 Input data

- Entity ID (URI) of the target Temporal Evolution of an Entity to be deleted.

5.6.16.4 Behaviour

- If the target Entity ID is not present or it is not a valid URI, then an error of type *BadRequestData* shall be raised.

- If the NGSI-LD endpoint does not know about the target Entity because there is no existing Entity whose id (URI) is equivalent held locally and no matching registrations apply, then an error of type *ResourceNotFound* shall be raised.
- If an **exclusive** or **redirect** Context Source Registration matches against the input data, the input data is forwarded. For each matching registration:
 - If the "Delete Temporal Evolution of Entity" operation is supported by the matched registration, matching input data is forwarded to the Registration endpoint.
 - If the "Delete Temporal Evolution of Entity" operation is not supported by the matched registration, this shall result in an error of type *Conflict* in case the complete "Delete Temporal Evolution of Entity" failed, or in a partial success if some parts of it succeeded.

No further processing is required.

- For any **inclusive** Context Source Registrations that match against the input data, that input data is also forwarded for remote processing to matching endpoints.
- If the target Temporal Evolution of an Entity exists locally, the entire Temporal Evolution of the Entity shall be removed.

5.6.16.5 Output data

None.

5.6.17 Merge Entity

5.6.17.1 Description

This operation allows modification of an existing NGSI-LD Entity aligning to the JSON Merge Patch processing rules defined in IETF RFC 7396 [16] by adding new Attributes (Properties or Relationships) or modifying or deleting existing Attributes associated with an existing Entity.

5.6.17.2 Use case diagram

A Context Producer can perform a merge on an Entity within an NGSI-LD system as shown in Figure 5.6.17.2-1.




Figure 5.6.17.2-1: Merge Entity use case

5.6.17.3 Input data

- A URI representing the id of the Entity to be merged (target Entity).
- A selector of Entity types as specified by clause 4.17 (optional).
- A JSON-LD document representing an NGSI-LD Entity Fragment.
- An optional flag indicating whether the JSON-LD document contains a simplified representation of the entity.
- An optional parameter indicating a common *observedAt* timestamp to use across merged Attributes.
- An optional parameter representing a common IETF RFC 5646 [28] language tag to use across merged LanguageMap Attributes.

5.6.17.4 Behaviour

The following behaviour shall be exhibited by compliant implementations:

- If the Entity ID is not present or it is not a valid URI then an error of type *BadRequestData* shall be raised.
- If the NGSI-LD endpoint does not know about the target Entity, because there is no existing Entity whose id (URI), and where specified type, is equivalent held locally, and no matching registrations apply (see clause 5.12), then an error of type *ResourceNotFound* shall be raised.
- If an **exclusive** or **redirect** Context Source Registration matches against the input data, Attributes from matching input data are forwarded. For each matching registration:
 - If the Merge Entity operation is supported by the matched registration (see clause 4.3.6), matching input data is forwarded to the Registration endpoint.
 - If the Merge Entity operation is not supported by the matched registration, this shall result in an error of type *Conflict* in case the complete Merge Entity operation failed, or in a partial success if some parts of it succeeded.

The matching Attributes are then removed from the Fragment and not processed further.

- For any **inclusive** Context Source Registrations that match against the remaining input data, that input data is also forwarded in the case the Merge Entity operation is supported for remote processing to matching endpoints.
- The behaviour defined in clause 5.5.4 on JSON-LD validation. NGSI-LD Nulls should be supported by this operation. If NGSI-LD Nulls are found in the payload, but are not supported, an error of type *OperationNotSupported* shall be raised.

Then, implementations shall perform a merge operation over the target Entity as mandated by clause 5.5.12, using the following procedure:

For each Attribute (Property or Relationship) included by the Entity Fragment:

- If the target Entity does not include a matching Attribute (considering term expansion rules as mandated by clause 5.5.7), then such Attribute shall be appended to the target Entity.
- If the target Entity already includes a matching Attribute (considering term expansion rules as mandated by clause 5.5.7):
 - If the Attribute (Property or Relationship) to be merged is represented in a simplified representation, the *type* of any pre-existing Attribute in the target entity shall be preserved.
 - If a common language tag is defined and a *LanguageProperty* Attribute to be merged is represented as a string, the pre-existing *languageMap* JSON object shall be preserved. The string value shall only replace the value associated to the language tag key found within the *languageMap*.

- If a common *observedAt* timestamp is defined and an existing Attribute to be merged previously contained an *observedAt* sub-Attribute, the *observedAt* sub-Attribute is also updated using the common timestamp, unless the Entity Fragment itself contains an explicit updated value for the *observedAt* sub-Attribute.
- If a *datasetId* is present in the Attribute included by the Entity Fragment:
 - If an Attribute instance in the target Entity has the same *datasetId*:
 - If overwrite is allowed and the Attribute value is not NGSI-LD Null, then the existing Attribute with the specified *datasetId* in the target Entity shall be merged with the new one supplied.
 - If overwrite is allowed and the Attribute value is NGSI-LD Null, then the existing Attribute with the specified *datasetId* in the target Entity shall be deleted.
 - If overwrite is not allowed, the existing Attribute with the specified *datasetId* in the target Entity shall be left untouched.
 - Otherwise the Attribute instance with the specified *datasetId* shall be appended to the target Entity.
- If no *datasetId* is present in the Attribute included by the Entity Fragment, the default Attribute instance is targeted:
 - If the default Attribute instance is present:
 - If overwrite is allowed and the Attribute value is not NGSI-LD Null, then the existing Attribute in the target Entity shall be merged with the new one supplied.
 - If overwrite is allowed and the Attribute value is NGSI-LD Null, then the existing Attribute with the specified *datasetId* in the target Entity shall be deleted.
 - If overwrite is not allowed, the existing Attribute in the target Entity shall be left untouched.
 - Otherwise if value is not NGSI-LD Null, the default Attribute instance shall be appended to the target Entity.
- If *type* is included in the Fragment and it includes Entity Type names that are not yet in the target Entity, add them to the list of Entity Type names of the target Entity.
- If *scope* is included in the Fragment and overwrite is allowed, the scope of the target Entity will become the one included in the Fragment. Otherwise, the Scopes in the Fragment that are not part of the value of *scope* of the target Entity will be appended to the value of the *scope* of the target Entity. If there is more than one Scope, the value of *scope* is represented as a JSON array containing all Scopes.

5.6.17.5 Output data

- A status code indicating whether all the Attributes were merged successfully.
- List of Attributes (Properties and/or Relationships) actually merged.

5.6.18 Replace Entity

5.6.18.1 Description

This operation allows the modification of an existing NGSI-LD Entity by replacing all of the Attributes (Properties or Relationships).

5.6.18.2 Use case diagram

A Context Producer can replace an entire Entity within an NGSI-LD system as shown in Figure 5.6.18.2-1.




Figure 5.6.18.2-1: Replace Entity use case

5.6.18.3 Input data

- A URI representing the id of the Entity to be replaced (target Entity).
- A selector of Entity types as specified by clause 4.17 (optional).
- A JSON-LD document representing an NGSI-LD Entity.

5.6.18.4 Behaviour

- If the target Entity ID is not a valid URI or it is not present, then an error of type **BadRequestData** shall be raised.
- If the NGSI-LD endpoint does not know about the target Entity, because there is no existing Entity whose id (URI), and where specified type, is equivalent held locally, and no matching registrations apply (see clause 5.12), then an error of type **ResourceNotFound** shall be raised.
- The behaviour defined in clause 5.5.4 on JSON-LD validation. NGSI-LD Nulls are not supported by this operation.
- If an **exclusive** or **redirect** Context Source Registration matches against the input data, Attributes from matching input data are forwarded. For each matching registration:
 - If the Replace Entity operation is supported by the registration (see clause 4.3.6), matching input data is forwarded to the Registration endpoint.
 - If the Replace Entity operation is not supported by the registration, this shall result in an error of type **Conflict** in case the complete Replace Entity operation failed, or in a partial success if some parts of it succeeded.
- The matching Attributes are then removed from the Fragment and not processed further.
- For any **inclusive** Context Source Registrations that match against the remaining input data, that input data is also forwarded in the case the Replace Entity operation is supported for remote processing to matching endpoints.
- If the target Entity exists locally, completely replace the existing Entity with the same Entity ID with the new Entity content provided. The system generated *createdAt* Temporal Property of the Entity as defined in clause 4.8 remain unchanged.

5.6.18.5 Output data

None.

5.6.19 Replace Attribute

5.6.19.1 Description

This operation allows the replacement of a single Attribute (Property or Relationship) within an NGSI-LD Entity.

5.6.19.2 Use case diagram

A Context Producer can carry out a replacement of an Attribute within an Entity within an NGSI-LD System as shown in Figure 5.6.19.2-1.




Figure 5.6.19.2-1: Replace Attribute use case

5.6.19.3 Input data

- Entity ID (URI) of the concerned Entity, the target Entity.
- A selector of Entity types as specified by clause 4.17 (optional).
- Target Attribute (Property or Relationship) to be replaced, identified by a name.
- A JSON-LD document representing an NGSI-LD Attribute Fragment.

5.6.19.4 Behaviour

- If the target Entity ID is not a valid URI, then an error of type **BadRequestData** shall be raised.
- If the target Attribute name is not valid, then an error of type **BadRequestData** shall be raised.
- If the target Attribute is *scope*, then an error of type **BadRequestData** shall be raised.
- If the NGSI-LD endpoint does not know about the target Entity, because there is no existing Entity whose id (URI), and where specified type, is equivalent held locally, and no matching registrations apply (see clause 5.12), then an error of type **ResourceNotFound** shall be raised.
- The behaviour defined in clause 5.5.4 on JSON-LD validation.

- If an **exclusive** or **redirect** Context Source Registration matches against the input data, the input data is forwarded. For each matching registration:
 - If the Replace Attribute operation is supported by the registration (see clause 4.3.6), matching input data is forwarded to the Registration endpoint.
 - If the Replace Attribute operation is not supported by the registration, this shall result in an error of type **Conflict** in case the complete Replace Attribute operation failed, or in a partial success if some parts of it succeeded.

No further processing is required.

- For any **inclusive** Context Source Registrations that match against the remaining input data, that input data is also forwarded in the case the Replace Attribute operation is supported for remote processing to matching endpoints.
- Apply term expansion as mandated by clause 5.5.7, so that the fully qualified name (URI) associated to the target Attribute is properly obtained.
- If the target Entity does not contain the target Attribute:
 - as a default instance in case no *datasetId* is present;
 - as an instance with the specified *datasetId* if present;
 - then this shall result in an error of type **ResourceNotFound** in case the complete Replace Attribute operation failed, or in a partial success if some parts of it succeeded.
- Completely replace the existing Attribute instance with the new Attribute content provided. The system generated *createdAt* Temporal Property as defined in clause 4.8 remains unchanged.

5.6.19.5 Output data

None.

5.6.20 Batch Entity Merge

5.6.20.1 Description

This operation allows modification of a batch of NGSI-LD Entities according to the JSON Merge Patch processing rules defined in IETF RFC 7396 [16] by adding new attributes (Properties or Relationships) or modifying or deleting existing attributes associated with an existing Entity.

5.6.20.2 Use case diagram

A Context Producer can merge a batch of Entities within an NGSI-LD system as shown in Figure 5.6.20.2-1.




Figure 5.6.20.2-1: Merge a batch of Entities use case

5.6.20.3 Input data

A JSON-LD Array containing one or more JSON-LD documents each one representing an Entity as mandated by clause 5.2.4. See clause 5.5.11.5 for information on behaviour when there is more than one instance of the same entity in the input Array.

5.6.20.4 Behaviour

Implementations shall exhibit the following behaviour:

- Execute the behaviour defined in clause 5.5.4 on JSON-LD validation.
- Let S be an array which shall contain a list of Entity IDs, one for each NGSI-LD Entity which was successfully processed. S shall be initialized as the empty array.
- Let E be an array which shall contain a list of **BatchEntityError** as defined by clause 5.2.17, one for each NGSI-LD Entity which resulted in error. E shall be initialized as the empty array.
- For each Context Source Registration CSR in the Context Registry:
 - Let IN be a copy of the original input array.
 - Remove from IN all Entities not matched by CSR and remove non-matching Attributes from the remaining Entities.
 - Remove all Attributes from the remaining Entities in IN for which there is a matching **exclusive** Context Source Registration, which is not CSR itself.
 - Remove all Attributes from the remaining Entities in IN for which there is a matching **redirect** Context Source Registration, unless CSR is a redirect Context Source Registration itself.
 - Remove all Entities without Attributes from IN.
 - If IN is empty, continue with the next Context Source Registration if there is any.
 - If the Batch Entity Merge operation is supported by CSR:
 - Forward the Batch Entity Merge request with IN as input Array.
 - Merge the returned list of Entities successfully created with S.

- Merge the returned list of Entities in Error with E.
- Otherwise, if the Merge Entity operation (clause 5.6.17) is supported by CSR:
 - For each Entity EN in the input array:
 - Forward a Merge Entity request for Entity EN.
 - Merge any successful result(s) for Entity EN merged with S.
 - Merge any error result(s) for Entity EN merged with E.
- Otherwise:
 - In case CSR is an **exclusive** or **redirect** Context Source Registration, add an Error of type *Conflict* for each Entity in IN to E.
- For each of the NGSI-LD Entities included in the input Array execute the behaviour defined by clause 5.6.17, but limited to a local operation, as follows:
 - If the Entity was successfully merged (Attributes updated, appended or deleted), then add the corresponding Entity ID to the S array.
 - If the Entity merge failed, then a new **BatchEntityError** shall be added to E containing the failed Entity ID and the *ProblemDetails* associated.

5.6.20.5 Output data

- none (if all Entities already existed and are successfully merged); or
- the list of Entities successfully merged (S Array), and the list of Entities in error (E Array), if only some or none of the Entities have been successfully merged.

5.6.21 Purge Entities

5.6.21.1 Description

This operation allows the deletion of entities within an NGSI-LD system based upon a query.

5.6.21.2 Use case diagram

A Context Producer can delete a set of entities which matches a specific query from an NGSI-LD system as shown in Figure 5.6.21.2-1.




Figure 5.6.21.2-1: Purge Entities use case

5.6.21.3 Input data

- A reference to a JSON-LD *@context* (optional).
- A selector of Entity types as specified by clause 4.17 (optional). Both type names (short hand string) and fully qualified type names (URI) are allowed in the selector.
- A list (one or more) of Entity identifiers (optional).
- A restrictive list of Attribute names to be deleted as attributes (optional).
- An exclusionary list Attribute names to be kept as attributes (optional).
- An id pattern as a regular expression (optional).
- An NGSI-LD query (to filter out Entities by Attribute values) as per clause 4.9 (optional).
- An NGSI-LD geoquery (to filter out Entities by spatial relationships) as mandated by clause 4.10 (optional).
- A NGSI-LD Scope query (to filter out Entities based on their Scope) as mandated by clause 4.19 (optional).
- An NGSI-LD query (called context source filter, to filter out Context Sources by the values of properties that describe them) as per clause 4.9 (optional).

In the general case, it is not possible to purge a set of entities by only specifying desired Entity identifiers, without further specifying restrictions on the entities' types or attributes, either explicitly, via selector of Entity types or of Attribute names, or implicitly, within an NGSI-LD Query or GeoQuery. If the execution of the operation is limited to the local scope (see clause 5.5.13), no further restrictions have to be provided.

5.6.21.4 Behaviour

- At least one of the following input data shall be provided:
 - a) selector of Entity Types;
 - b) list of Attribute names, including at least one non-system Attribute;
 - c) NGSI-LD Query, including at least one non-system Attribute;
 - d) NGSI-LD GeoQuery;
 - e) local scope (see clause 5.5.13).

If none of the above is provided, then an error of type *BadRequestData* shall be raised (too wide query).

- If the list of Entity identifiers includes a URI which it is not valid, or the query, geoquery or context source filter are not syntactically valid (as per the referred clauses 4.9 and 4.10) an error of type *BadRequestData* shall be raised.
- If projection attributes are present and indicate the use of **Linked Entity** retrieval, then an error of type *BadRequestData* shall be raised.
- If the filter conditions specified by the query includes **Linked Entity** attributes then an error of type *BadRequestData* shall be raised.
- Term to URI expansion of type and Attribute names shall be performed, as mandated by clause 5.5.7.
- Otherwise, implementations shall run a Query Entities operation (clause 5.6.2) to retrieve a list of ids of Entities for deletion, that meet all of the following conditions (given the respective parameter is provided):
 - id is equal to any of the id(s) passed as a parameter;
 - the Entity Type names match the selector of Entity Types (expanded) that is passed as a parameter;
 - Attribute matches any of the expanded attribute(s) in the list that is passed as a parameter;
 - id matches the id pattern passed as a parameter;
 - the filter conditions specified by the query are met (as mandated by clause 4.9);
 - the geospatial restrictions imposed by the geoquery are met (as mandated by clause 4.10); if there are multiple instances of the *GeoProperty* on which the geoquery is based, it is sufficient if any of these instances meets the geospatial restrictions;
 - if the Scope query is present, it shall match a present Entity Scope (as mandated by clause 4.19, for an example see annex C, clause C.5.15).

And thereafter:

- when no restrictive list of Entity member names is present, the implementation shall delete all Entities that can be found locally using retrieved list of Entity ids;
- when a restrictive list of Entity member names is present, the implementation shall delete the given set of Attributes with those member names from all Entities that can be found locally using retrieved list of Entity ids;
- when an exclusionary list of Entity member names is present, the implementation shall delete all but the given set of Attributes with those member names from all Entities that can be found locally using retrieved list of Entity ids.
- Unless local scope is specified (see clause 5.5.13), for **Context Source Registrations** that match the query and support the "purgeEntity" operation (see operations and operation groups in clause 4.20), implementations shall do the following:
- If an **inclusive**, **exclusive** or **redirect** **Context Source Registration** matches against the filter conditions specified, the request is forwarded for remote processing. For each matching registration:
 - If the Purge Entity operation is supported by the registration (see clause 4.3.6), matching input data is forwarded to the Registration endpoint.
 - If the Purge Entity operation is not supported by the matched registration, this shall result in an error of type *Conflict* in case the complete purge Entities failed, or in a partial success if some parts of purge Entities succeeded.

5.6.21.5 Output data

None.

5.7 Context Information Consumption

5.7.1 Retrieve Entity

5.7.1.1 Description

This operation allows retrieving an NGSI-LD Entity.

5.7.1.2 Use case diagram

A *Context Consumer* can retrieve a specific Entity from an NGSI-LD system as shown in Figure 5.7.1.2-1.




Figure 5.7.1.2-1: Retrieve Entity use case

5.7.1.3 Input data

- Entity ID (URI) of the Entity to be retrieved (target Entity).
- A selector of Entity types as specified by clause 4.17 (optional).
- List of Attribute (*Properties* or *Relationships*) names to be retrieved (projection attributes) (optional).
- A restrictive list of Entity member names ("id", "type", "scope" or an Attribute name) to be retrieved (projection attributes as defined by clause 4.21) (optional).
- An exclusionary list of Entity member names ("id", "type", "scope" or an Attribute name) to be removed (projection attributes as defined by clause 4.21) (optional).
- A language filter as defined by clause 4.15 (optional).
- An optional JSON-LD context.
- In the case of a GeoJSON representation, the name of the *GeoProperty* attribute to use as the geometry for the GeoJSON representation as mandated by clause 4.5.16 (optional).

- An optional flag indicating whether to include additional **Linked Entities** corresponding to the **Relationships** retrieved and how to format those **Linked Entities**. See clause 4.5.23 (optional).
- A limit to the depth of **Linked Entities** to search whilst traversing an Entity graph. See clause 4.5.23 (optional).
- A list (one or more) of **Linked Entity identifiers** previously encountered whilst traversing an Entity graph. See clause 4.5.23 (optional).
- A flag indicating whether to return the location of the EntityMap used within the operation (optional).
- A suggested lifetime for the EntityMap, if EntityMap is to be created (optional).
- The location of a resource holding an EntityMap of matching Entity registrations (optional).
- A *datasetId* parameter that specifies which Attribute instances are to be selected as defined by clause 4.5.5 (optional).

5.7.1.4 Behaviour

- If the Entity ID is not present or it is not a valid URI, then an error of type **BadRequestData** shall be raised.
- If geometryProperty parameter is present and the Accept Header is not set to "application/geo+json", then an error of type **BadRequestData** shall be raised.
- If projection attributes are present and indicate the use of **Linked Entity** retrieval and the use of **Linked Entity** retrieval is not specified, or the projected attribute depth exceeds the **Linked Entity** retrieval depth, an error of type **BadRequestData** shall be raised.
- If the NGSI-LD endpoint does not know about the target Entity, because there is no existing Entity held locally whose id (URI), and where specified type, is equivalent, and no matching registrations apply, then an error of type **ResourceNotFound** shall be raised.
- The implementation shall retrieve any Attribute data held locally which is associated with the Entity defined by the Entity ID.
- If the ContextBroker implementation supports the use of EntityMaps then:
 - If the location of a resource holding an EntityMap of matching Entity registrations is present it shall be retrieved:
 - If the resource cannot be found, or the data has expired, a new EntityMap shall be created.
 - If the data has not expired, **only** the retrieved Entity Map shall be used to determine which **Context Source Registrations** match the Entity ID.
 - If a flag to return an EntityMap was present in the request, and no EntityMap currently exists, then a new EntityMap shall be created.
- For **Context Source Registrations** that match and support the retrieveEntity operation (see operations and operation groups in clause 4.20), implementations shall do the following:
 - If an EntityMap is in use for this operation, and an EntityMap entry linked to a **Context Source Registration** is found, the location of the linked EntityMap shall be passed as part of any forwarded request.
 - For any **exclusive**, **redirect** and **inclusive** **Context Source Registrations**, the request is forwarded for remote retrieval by matching endpoints, and remote Attribute data for the Entity is received. If an *expiresAt DateTime* is present on the Entity and the date lies in the past, the Attribute data shall be discarded, otherwise the Attribute data is then merged together according to the algorithm defined in clause 4.5.5.
 - For any **auxiliary** **Context Source Registrations** the remote Attribute data received is added to the payload only when an Attribute is not present in any of the Attribute data received elsewhere.

- If an EntityMap is in use for this operation, the EntityMap's linked maps are updated to hold the location of every EntityMap used by the Context Source Registrations.
- Term to URI expansion of Attribute names shall be observed as mandated by clause 5.5.7.
- For each Attribute found in the target Entity, when the *datasetId* parameter is provided in the request:
 - Filter the Attribute instances based on the *datasetId* parameter, i.e. keep only the Attribute instances whose *datasetId* is specified. The default Attribute instance is matched, if "@none" is specified.
 - If there is no Attribute instance whose *datasetId* matches the value of the parameter, the Attribute shall not be returned with the Entity.
- If the optional Attribute list is present and the NGSI-LD endpoint does know about a matching Entity for the Entity ID, but this Entity does not have any of the Attributes in the Attribute list, then an error of type *ResourceNotFound* shall be raised.
- If the Accept Header is set to "application/json" or "application/ld+json", return a JSON-LD object representing the Entity as mandated by clause 5.2.4 and containing only the Attributes requested (if present).
 - If the Prefer Header [26] is set to "ngsi-ld=<version>" then the ContextBroker shall endeavour to amend the JSON-LD object to conform to the specified version of the NGSI-LD specification as mandated in clause 4.3.6.8, and return a Preference-Applied Header set to "ngsi-ld=<conformant-version>" in the response.
- If the Accept Header is set to "application/geo+json", a GeoJSON *Feature* object representing the entity as mandated by clause 5.2.29 and containing only the Attributes requested (if present):
 - If the Prefer Header is omitted or set to "body=ld+json" then the *Feature* object will also contain an *@context* field.
 - If the Prefer Header is set to "body=json" the *@context* is set as a Link Header and removed from the Feature object.

5.7.1.5 Output data

A JSON-LD object representing the target Entity as mandated by clause 5.2.4 or a GeoJSON *Feature* as mandated by clause 5.2.29.

If a restrictive list of Entity member names is present, the Entity is reduced down to only contain the defined Entity members (applies to all Entities in the payload in the case of **Linked Entity** retrieval).

If an exclusionary list of Entity member names is present, the defined Entity members listed are removed from the Entity (applies to all Entities in the payload in the case of **Linked Entity** retrieval).

If any of the returned Attributes corresponds to a **VocabProperty**, the returned value shall be compacted according to the supplied *@context*.

If a language filter is specified and any of the returned Attributes corresponds to a *LanguageProperty*, the *LanguageProperty* in question shall be converted into a *Property*. The value of this *Property* shall correspond to the value of the string or strings from the matching key-value pair of the *languageMap* where the key matches the language filter. A non-reified subproperty *lang* shall be included in the response indicating the chosen language.

If no match can be made for a *LanguageProperty* then a single language shall be chosen, up to the implementation.

If **inline Linked Entity retrieval** (see clause 4.5.23.2) is specified, and any of the returned Attributes corresponds to an annotated *Relationship*, then unless a URI has been previously encountered, an *entity* sub-Property shall be included in the response holding the **Linked Entity** data for each URI corresponding to that *Relationship*'s target *object* URI. If any of the returned Attributes corresponds to an annotated *ListRelationship*, then an *entityList* subproperty shall be included in the response holding the **ordered** array of **Linked Entities** corresponding to that *ListRelationship*'s target *objectList* URIs unless a URI has been previously encountered.

If **flattened Linked Entity retrieval** (see clause 4.5.23.3) is specified, the response shall be an array of JSON-LD objects representing the target entity itself and the targets of its *Relationships*. If any of the returned Attributes corresponds to an annotated *Relationship*, unless a target URI has been previously encountered, the data corresponding to each URI of the Relationship's target *object* URIs is appended to the array. If any of the returned Attributes corresponds to an annotated *ListRelationship*, then an **ordered** array of additional **Linked Entities** is appended to the array which hold the data corresponding to each of the URIs found within *ListRelationship*'s target *objectList* unless a URI has been previously encountered.

Flattened **Linked Entity** retrieval output changes to a GeoJSON *FeatureCollection* as mandated by clause 5.2.30 if the Accept Header is set to "application/geo+json".

If the location of a previously generated EntityMap was passed into the request, or a flag to return an EntityMap was present in the request, the location of the EntityMap used in the operation shall be returned in a specific field in the response.

5.7.2 Query Entities

5.7.2.1 Description

This operation allows querying an NGSI-LD system.

5.7.2.2 Use case diagram

A *Context Consumer* can retrieve a set of entities which matches a specific query from an NGSI-LD system as shown in Figure 5.7.2.2-1.




Figure 5.7.2.2-1: Query Entities use case

5.7.2.3 Input data

- A reference to a JSON-LD *@context* (optional).
- A selector of Entity types as specified by clause 4.17 (optional). Both type names (short hand string) and fully qualified type names (URI) are allowed in the selector.

- A list (one or more) of Entity identifiers (optional).
- An id pattern as a regular expression (optional).
- An NGSI-LD query (to filter out Entities by Attribute values) as per clause 4.9 (optional).
- An NGSI-LD geoquery (to filter out Entities by spatial relationships) as mandated by clause 4.10 (optional).
- A NGSI-LD Scope query (to filter out Entities based on their Scope) as mandated by clause 4.19 (optional).
- A restrictive list of Entity member names ("id", "type", "scope" or an Attribute name) to be retrieved (projection attributes as defined by clause 4.21) (optional).
- An exclusionary list member names ("id", "type", "scope" or an Attribute name) to be removed (projection attributes as defined by clause 4.21) (optional).
- A list (one or more) of Attribute names of which at least one shall exist in order for an Entity to be selected, and also used as query projection attributes (optional, deprecated).
- An NGSI-LD query (called context source filter, to filter out Context Sources by the values of properties that describe them) as per clause 4.9 (optional).
- A limit to the number of Entities to be retrieved. See clause 5.5.9.
- A specified language filter as per clause 4.15 (optional).
- A list (one or more) of Attribute names whose values shall be expanded to URIs prior to executing a query (optional).
- An optional flag indicating whether to include additional Linked Entities corresponding to the Relationships retrieved and how to format those Linked Entities. See clause 4.5.23 (optional).
- A limit to the depth of Linked Entities to search whilst traversing an Entity Graph. See clause 4.5.23 (optional).
- A list (one or more) of Linked Entity identifiers previously encountered whilst traversing an Entity Graph. See clause 4.5.23 (optional).
- A flag indicating whether to return the location of the EntityMap used within the operation (optional).
- A suggested lifetime for the EntityMap, if EntityMap is to be created (optional).
- The location of a resource holding an EntityMap of matching Entity registrations (optional).
- A *datasetId* parameter that specifies which Attribute instances are to be selected as defined by clause 4.5.5 (optional).
- A list of Entity member names ("id", "type", "scope" or an Attribute name) to be used to define the preferred ordering of Entities retrieved (Entity ordering attributes as defined by clause 4.23) (optional).
- A preferred collation setting to be used when applying ordering of Entities (optional).
- A location to be used when applying ordering of Entities (optional).
- A defined geometry type to be used when applying ordering of Entities (optional).
- A flag indicating whether split Entities are to be expected, i.e. Entities whose information is distributed across different Context Sources (optional).
- In the case of GeoJSON representation, the name of the *GeoProperty* attribute to use as the geometry for the GeoJSON representation as mandated by clause 4.5.16 (optional).

In the general case, it is not possible to retrieve a set of entities by only specifying desired Entity identifiers, without further specifying restrictions on the entities' types or attributes, either explicitly, via selector of Entity types or of Attribute names, or implicitly, within an NGSI-LD Query or GeoQuery. If the execution of the operation is limited to the local scope (see clause 5.5.13), no further restrictions have to be provided.

5.7.2.4 Behaviour

- At least one of the following input data shall be provided:
 - a) selector of Entity Types;
 - b) list of Attribute names, including at least one non-system Attribute;
 - c) NGSI-LD Query, including at least one non-system Attribute;
 - d) NGSI-LD GeoQuery;
 - e) local scope (see clause 5.5.13).

If none of the above is provided, then an error of type *BadRequestData* shall be raised (too wide query).

- If the list of Entity identifiers includes a URI which it is not valid, or the query, geoquery or context source filter are not syntactically valid (as per the referred clauses 4.9 and 4.10) an error of type *BadRequestData* shall be raised.
- If projection attributes are present and indicate the use of *Linked Entity* retrieval, and the use of *Linked Entity* retrieval is not specified, or the projected attribute depth exceeds the *Linked Entity* retrieval depth, then an error of type *BadRequestData* shall be raised.
- If the filter conditions specified by the query includes *Linked Entity* attributes, and the use of *Linked Entity* retrieval is not specified, or the *Linked Entity* attribute query depth exceeds the *Linked Entity* retrieval depth, an error of type *BadRequestData* shall be raised (too deep query).
- If geometryProperty parameter is present and the Accept Header is not set to "application/geo+json", then an error of type *BadRequestData* shall be raised.
- If the ordering parameter is present and the execution of the operation is not limited to the local scope (see clause 5.5.13) then an error of type *BadRequestData* shall be raised.
- If the ordering parameter is present and refers to ordering by distance and no location is present, then an error of type *BadRequestData* shall be raised.
- If a preferred collation setting is present and it does not conform to a valid ICU collation (see IETF RFC 6067 [36]) then an error of type *BadRequestData* shall be raised.
- If a location to be used when applying ordering of Entities setting is present and it is not syntactically valid (as per the referred clauses 4.9 and 4.10) or an error of type *BadRequestData* shall be raised.
- Otherwise,
 - Term to URI expansion of type and Attribute names shall be performed, as mandated by clause 5.5.7.
 - If a list of Attribute names whose values shall be expanded to URIs has been supplied, JSON-LD type coercion shall be performed as mandated by clause 5.5.7.
- If split entities flag is explicitly set to true or, if not explicitly set, the default setting of the deployment allows split entities, and local scope is not specified, implementations shall run a query that shall return an Entity Array containing all the Entities found locally, that meet all of the following conditions (given the respective parameter is provided):
 - id is equal to any of the id(s) passed as a parameter;
 - the Entity Type names match the selector of Entity Types (expanded) that is passed as a parameter;
 - id matches the id pattern passed as a parameter.

- Otherwise, implementations shall run a query that shall return an Entity Array containing all the Entities found locally, that meet **all** of the following conditions (given the respective parameter is provided):
 - id is equal to any of the id(s) passed as a parameter;
 - the Entity Type names match the selector of Entity Types (expanded) that is passed as a parameter;
 - Attribute matches any of the expanded attribute(s) in the list that is passed as a parameter;
 - id matches the id pattern passed as a parameter;
 - the filter conditions specified by the query are met (as mandated by clause 4.9);
 - the geospatial restrictions imposed by the geoquery are met (as mandated by clause 4.10); if there are multiple instances of the *GeoProperty* on which the geoquery is based, it is sufficient if any of these instances meets the geospatial restrictions;
 - if the Scope query is present, it shall match a present Entity Scope (as mandated by clause 4.19, for an example see annex C, clause C.5.15);
 - if the Attribute list is present, in order for an Entity to match, it shall contain at least one of the Attributes in the projection Attribute list.
- If the ContextBroker implementation supports the use of EntityMaps then:
 - If the location of a resource holding an EntityMap of matching Entity registrations is present it shall be retrieved:
 - If the resource cannot be found, or the data has expired, a new EntityMap shall be created.
 - If the data has not expired, **only** the retrieved EntityMap shall be used to determine which **Context Source Registrations** match the Entity ID.
 - If a flag to return an EntityMap was present in the request, and no EntityMap currently exists, then a new EntityMap shall be created.
- Unless local scope is specified (see clause 5.5.13), for **Context Source Registrations** that match the query and support the "queryEntity" operation (see operations and operation groups in clause 4.20), implementations shall do the following:
 - If an EntityMap is in use for this operation, and an EntityMap entry linked to a **Context Source Registration** is found, the location of the EntityMap shall be passed as part of the forwarded request.
 - If split entities flag is explicitly set to true or, if not explicitly set, the default setting of the deployment allows split entities, the filters (filter conditions specified by the query, geospatial restrictions imposed by the geoquery, Scope query, Attributes) shall be removed before forwarding the request. These filters then have to be applied after the Entity information from different Context Sources and local information, if there is any, has been aggregated.
 - For any **exclusive**, **redirect** and **inclusive** **Context Source Registrations**, the request is forwarded for remote querying by matching endpoints. The result of each remote query is an Entity Array. Within each Entity Array, any Entities where an *expiresAt DateTime* is present and the date lies in the past shall be discarded. The Entity Arrays are then merged together with the locally queried result according to the algorithm defined in clause 4.5.5.
 - For any **auxiliary** **Context Source Registrations**, the request is forwarded for remote querying by matching endpoints. Data from the Entity Array received is added to the payload only when an Attribute is not already present in the merged Entity Arrays are received elsewhere.
- If split Entities flag is explicitly set to true or, if not explicitly set, the default setting of the deployment allows split Entities, and local scope is not specified, the following filters shall be applied on the aggregated Entities:
 - the filter conditions specified by the query are met (as mandated by clause 4.9);

- the geospatial restrictions imposed by the geoquery are met (as mandated by clause 4.10); if there are multiple instances of the *GeoProperty* on which the geoquery is based, it is sufficient if any of these instances meets the geospatial restrictions;
- if the Scope query is present, it shall match a present Entity Scope (as mandated by clause 4.19, for an example see annex C, clause C.5.15);
- if the Attribute list is present, in order for an Entity to match, it shall contain at least one of the Attributes in the projection Attribute list.
- Pagination logic shall be in place as mandated by clause 5.5.9.
- If in the process of obtaining the query result it is necessary to issue a **Context Source** discovery operation, the same **Context Source** filter input parameter (if present) shall be propagated.
- For each Attribute found in the target Entity, when *datasetId* parameter is provided in the request:
 - Filter the Attribute instances based on the *datasetId* parameter, i.e. keep only the Attribute instances whose *datasetId* is specified. The default Attribute instance is matched, if "@none" is specified.
 - If there is no Attribute instance whose *datasetId* matches the value of the parameter, the Attribute shall not be returned with the Entity.
- If the Accept Header is set to "application/json" or "application/ld+json", a JSON-LD array is returned, representing the Entities as mandated by clause 5.2.4 and containing only the Attributes requested (if present).
 - If the Prefer Header [26] is set to "ngsi-ld=<version>" then the ContextBroker shall endeavour to amend the elements of the JSON-LD array to conform to the specified version of the NGSI-LD specification as mandated in clause 4.3.6.8, and return a Preference-Applied Header set to "ngsi-ld=<conformant-version>" in the response.
- If the Accept Header is set to "application/geo+json", the response shall be a GeoJSON *FeatureCollection* as mandated by clause 5.2.30, with each *Feature* within the *FeatureCollection* containing only the Attributes requested (if present):
 - If the Prefer Header is omitted or set to "body=ld+json" then the *FeatureCollection* will also contain an @context field.
 - If the Prefer Header is set to "body=json" the @context is sent as a Link Header and removed from the *FeatureCollection* object.

5.7.2.5 Output data

A JSON-LD array representing the matching entities as defined by clause 5.2.4 or in the case of GeoJSON requests a *FeatureCollection* as mandated by clause 5.2.30. If Entity ordering is specified (see clause 4.23), then the JSON-LD array or *FeatureCollection* returned shall be ordered according to the list of member names specified. For each matching Entity, only the Attributes specified by the Attribute list parameter shall be included. If such parameter is not present, then all Attributes shall be included.

If a restrictive list of Entity member names is present, every Entity within the payload body is reduced down to only contain the defined Entity members.

If an exclusionary list of Entity member names is present, the defined Entity members listed are removed from each Entity within the payload.

If any of the returned Attributes corresponds to a **VocabProperty**, the returned value shall be compacted according to the supplied @context.

If a language filter is specified and any of the returned Attributes corresponds to a *LanguageProperty*, the *LanguageProperty* in question shall be converted into a *Property*. The value of this *Property* shall correspond to the value of the string or strings from matching key-value pair of the *languageMap* where the key matches the language filter. A non-reified subproperty *lang* shall be included in the response indicating the chosen language.

If no match can be made for a *LanguageProperty*, then the default identified by the JSON-LD "@none" shall be chosen if present, otherwise the choice of a single language is up to the implementation.

If **inline Linked Entity retrieval** (see clause 4.5.23.2) is specified, and any of the returned Attributes corresponds to an annotated *Relationship*, then unless a URI has been previously encountered, an *entity* subproperty shall also be included in the response holding the data for each URI corresponding to that *Relationship*'s target *object* URIs. If any of the returned Attributes corresponds to an annotated *ListRelationship*, then an *entityList* subproperty shall also be included in the response holding the **ordered** data corresponding to that *ListRelationship*'s target *objectList* URIs, unless a URI has been previously encountered.

If **flattened Linked Entity retrieval** (see clause 4.5.23.3) is specified, the response shall be an array of JSON-LD objects representing the matching entities and the targets of their *Relationships*. If any of the returned Attributes corresponds to an annotated *Relationship*, unless a URI has been previously encountered, then data for each URI corresponding to the *Relationship*'s target *object* URIs is appended to the array. If any of the returned Attributes corresponds to an annotated *ListRelationship*, then an array of additional **Linked Entities** is appended to the array which hold the data corresponding to each of the URIs found within *ListRelationship*'s target *objectList* unless a URI has been previously encountered.

If the location of a previously generated EntityMap was passed into the request, or a flag to return an EntityMap was present in the request, the location of the EntityMap used in the operation shall be returned in a specific field in the response.

5.7.3 Retrieve Temporal Evolution of an Entity

5.7.3.1 Description

This operation allows retrieving the Temporal Evolution of an Entity.

5.7.3.2 Use case diagram

A Context Consumer can retrieve the Temporal Evolution of an Entity (in the form of a temporal representation) from an NGSI-LD system as shown in Figure 5.7.3.2-1.




Figure 5.7.3.2-1: Retrieve Temporal Evolution of an Entity use case

5.7.3.3 Input data

- Entity ID (URI) of the target Temporal Evolution of an Entity to be retrieved.
- List of Attribute (Properties or Relationships) names to be retrieved (projection attributes) (optional).

- A restrictive list of Entity member names ("id", "type", "scope" or an Attribute name) to be retrieved (projection attributes as defined by clause 4.21) (optional).
- An exclusionary list of Entity member names ("id", "type", "scope" or an Attribute name) to be removed (projection attributes as defined by clause 4.21) (optional).
- An NGSI-LD temporal query as mandated by clause 4.11 (optional).
- A parameter (*lastN*) conveying that only the last N instances (per Attribute) within the concerned temporal interval shall be retrieved (optional).
- An optional JSON-LD context.
- A flag indicating whether to return the location of the EntityMap used within the operation (optional).
- A suggested lifetime for the EntityMap, if EntityMap is to be created (optional).
- The location of a resource holding an EntityMap of matching Entity registrations (optional).
- A *datasetId* parameter that specifies which Attribute instances are to be selected as defined by clause 4.5.5 (optional).

5.7.3.4 Behaviour

- If the Entity ID is not present or it is not a valid URI, then an error of type **BadRequestData** shall be raised.
- If projection attributes are present and indicate the use of **Linked Entity** retrieval, an error of type **BadRequestData** shall be raised.
- If the NGSI-LD endpoint does not know about the target Entity, because there is no existing Entity whose id (URI) is equivalent held locally, and no matching registrations apply, then an error of type **ResourceNotFound** shall be raised.
- Term to URI expansion of Attribute names shall be observed as mandated by clause 5.5.7.
- The *lastN* parameter refers to a number, *n*, of Attribute instances which shall correspond to the last *n* timestamps (in descending ordering) of the temporal property (by default *observedAt*) within the concerned temporal interval.
- Let *S* be the **Temporal Evolution of the Entity** as mandated by clause 5.2.20 with the specified Entity ID as it is available locally. *S* is empty in case no **Temporal Evolution of the Entity** is available locally.
- From *S*, select only those Attribute instances (corresponding to the Attributes specified by the query or all if none are specified) match the temporal restrictions imposed by the temporal query (as mandated by clause 4.11); i.e. if the time series, for all the concerned Attributes of an Entity, does not include data corresponding to the temporal query interval, then such Entity shall be removed from *S*, thus it shall not appear in the final result set. Let *S1* be this new subset.
- If the ContextBroker implementation supports the use of EntityMaps then:
 - If the location of a resource holding an EntityMap of matching Entity registrations is present it shall be retrieved:
 - If the resource cannot be found, or the data has expired, a new EntityMap shall be created.
 - If the data has not expired, **only** the retrieved Entity Map shall be used to determine which Context Source Registrations match the Entity ID.
 - If a flag to return an EntityMap was present in the request, and no EntityMap currently exists, then a new EntityMap shall be created.

- For **Context Source Registrations** that match the Entity ID and support the "retrieveTemporal" operation (see operations and operation groups in clause 4.20), implementations shall do the following:
 - If an **EntityMap** is in use for this operation, and an **EntityMap** entry linked to a **Context Source Registration** is found, the location of the linked **EntityMap** shall be passed as part of any forwarded request.
 - For any **exclusive**, **redirect** and **inclusive Context Source**, the request is forwarded for remote retrieval by matching endpoints. and remote Attribute data for the Entity is received. The result is then merged together with S1 according to the algorithm defined in clause 4.5.5.
 - For any **auxiliary Context Source Registrations** the remote Attribute data received is added to S1 only when the Attribute instance, whose value of the *timeproperty*, which is used for the temporal query (*observedAt* as default), is not present in any of the Attribute instances received from elsewhere.
 - If an **EntityMap** is in use for this operation, the **EntityMap's** linked maps are updated to hold the location of every **EntityMap** used by the **Context Source Registrations**.
- For the set of **Attribute Instances** that are in S1, when *datasetId* parameter is provided in the request:
 - Filter the Attribute instances based on the *datasetId* parameter, i.e. keep only the Attribute instances whose *datasetId* is specified. The default Attribute instance is matched, if "@none" is specified.
 - If there is no Attribute instance whose *datasetId* matches the value of the parameter, remove the complete Attribute from S1.
- From the set of **Attribute Instances** that are in S1, include in their temporal representation only the Attribute instances (up to *lastN*) corresponding to the query's projection Attributes, or aggregated values of Attribute instances (if aggregated temporal representation is requested).

In the general case, it is not possible to retrieve a set of entities by only specifying desired Entity identifiers, without further specifying restrictions on the entities' types or attributes, either explicitly, via selector of Entity types or of Attribute names, or implicitly, within an NGSI-LD Query or GeoQuery. If the execution of the operation is limited to the local scope (see clause 5.5.13), no further restrictions have to be provided.

5.7.3.5 Output data

A JSON-LD object representing the **Temporal Evolution of the target Entity** as mandated by clause 5.2.20.

If a restrictive list of Entity member names is present, every Entity within the payload body is reduced down to only contain the defined Entity members.

If an exclusionary list of Entity member names is present, the defined Entity members listed are removed from each Entity within the payload.

5.7.4 Query Temporal Evolution of Entities

5.7.4.1 Description

This operation allows querying the **Temporal Evolution of Entities** present in an NGSI-LD system. It is similar to the operation defined by clause 5.7.2 (Query Entities) with the addition of a temporal query.

5.7.4.2 Use case diagram

A **Context Consumer** can retrieve the Temporal Evolution of a set of Entities which matches a specific query from an NGSI-LD system as shown in Figure 5.7.4.2-1.




Figure 5.7.4.2-1: Query Temporal Evolution of Entities use case

5.7.4.3 Input data

- An NGSI-LD temporal query as mandated by clause 4.11.
- A reference to a JSON-LD *@context* (optional).
- A selector of Entity types as specified by clause 4.17 (optional).
Both type name (short hand string) and fully qualified type name (URI) are allowed.
- A restrictive list of Entity member names ("id", "type", "scope" or an Attribute name) to be retrieved (projection attributes as defined by clause 4.21) (optional).
- An exclusionary list of Entity member names ("id", "type", "scope" or an Attribute name) to be removed (projection attributes as defined by clause 4.21) (optional).
- A list (one or more) of Entity identifiers (optional).
- A list (one or more) of Attribute names of which at least one shall exist in order for an Entity to be selected, and also used as query projection attributes (optional).
- An id pattern as a regular expression (optional).
- An NGSI-LD query (to filter out Entities by Attribute values) as per clause 4.9 (optional).
- An NGSI-LD geoquery (to filter out Entities by spatial relationships) as mandated by clause 4.10 (optional).
- A NGSI-LD Scope query (to filter out Entities based on their Scope) as mandated by clause 4.19 (optional).
- An NGSI-LD query (called context source filter, to filter out Context Sources by the values of properties that describe them) as per clause 4.9 (optional).
- A limit to the number of Entities to be retrieved. See clause 5.5.9.
- A parameter (lastN) conveying that only the last N instances (per Attribute) within the concerned temporal interval shall be retrieved (optional).
- A specified language filter as per clause 4.15 (optional).
- A list (one or more) of Attribute names whose values shall be expanded to URIs prior to executing a query (optional).

- A flag indicating whether to return the location of the EntityMap used within the operation (optional).
- A suggested lifetime for the EntityMap, if EntityMap is to be created (optional).
- The location of a resource holding an EntityMap of matching Entity registrations (optional).
- A *datasetId* parameter that specifies which Attribute instances are to be selected as defined by clause 4.5.5 (optional).
- A preferred ordering of Entities retrieved - only the Entity member name "id" can be used (Entity ordering attributes as defined by clause 4.23) (optional).
- A preferred collation setting to be used when applying ordering of Entities (optional).
- A flag indicating whether split Entities are to be expected, i.e. Entities whose information is distributed across different Context Sources (optional).

In the general case, it is not possible to retrieve a set of entities by only specifying desired Entity identifiers, without further specifying restrictions on the entities' types or attributes, either explicitly, via selector of Entity types or of Attribute names, or implicitly, within an NGSI-LD query or geoquery. If the execution of the operation is limited to the local scope (see clause 5.5.13), no further restrictions have to be provided.

5.7.4.4 Behaviour

- If a temporal query is not provided then an error of type *BadRequestData* shall be raised.
- At least one of the following input data shall be provided:
 - a) selector of Entity Types;
 - b) list of Attribute names, including at least one non-system Attribute;
 - c) NGSI-LD Query, including at least one non-system Attribute;
 - d) NGSI-LD GeoQuery;
 - e) local scope (see clause 5.5.13).

If none of the above is provided, then an error of type *BadRequestData* shall be raised (too wide query).

- If projection attributes or filter conditions indicate the use of **Linked Entity** retrieval, an error of type *BadRequestData* shall be raised.
- If the list of Entity identifiers includes a URI which it is not valid, or the query, geoquery or context source filter are not syntactically valid (as per the referred clauses 4.9 and 4.10) an error of type *BadRequestData* shall be raised.
- If the ordering parameter is present and the execution of the operation is not limited to the local scope (see clause 5.5.13) then an error of type *BadRequestData* shall be raised.
- If the ordering parameter is present and refers an entity name other than "id", then an error of type *BadRequestData* shall be raised.
- If a preferred collation setting is present and it does not conform to a valid ICU collation (see IETF RFC 6067 [36]) then an error of type *BadRequestData* shall be raised.
- Term to URI expansion of type and Attribute names shall be observed mandated by clause 5.5.7.
- If a list of Attribute names whose values shall be expanded to URIs has been supplied, JSON-LD type coercion shall be performed as mandated by clause 5.5.7.
- The *lastN* parameter refers to a number, *n*, of Attribute instances which shall correspond to the last *n* timestamps (in descending ordering) of the temporal property (by default *observedAt*) within the concerned temporal interval.

- Otherwise,
 - Let S be the set of selected Entities i.e. the query result set.
 - If split entities flag is explicitly set to true or, if not explicitly set, the default setting of the deployment allows split entities, and local scope is not specified, implementations shall run a query that shall return an Entity Array containing all the Entities found locally, that meet all of the following conditions (given the respective parameter is provided):
 - If id(s) is provided, keep in S only those Entities whose id is equivalent to any of the id(s) passed as a parameter.
 - If an id pattern is provided, keep in S only those Entities whose id matches the id pattern.
 - If a selector of Entity Types is provided, keep in S only those Entities whose Entity Type names match the selector of Entity Types.
 - From S, select only those Entities any of whose Attribute instances (corresponding to the Attributes specified by the query or all if none are specified) match the temporal restrictions imposed by the temporal query (as mandated by clause 4.11); i.e. if the time series, for all the concerned Attributes of an Entity, does not include data corresponding to the temporal query interval, then such Entity shall be removed from S, thus it shall not appear in the final result set. Let S4 be this new subset.
 - Implementations shall run a query that shall return the Temporal Evolution of the matching Entities (all Entities stored locally, in case only a local scope is specified); the logical steps to select the final result set of Entities, and the Attribute instances included as part of their temporal representation, are enumerated as follows:
 - If id(s) is provided, keep in S only those Entities whose id is equivalent to any of the id(s) passed as a parameter.
 - If an id pattern is provided, keep in S only those Entities whose id matches the id pattern.
 - If a selector of Entity Types is provided, keep in S only those Entities whose Entity Type names match the selector of Entity Types.
 - From S, select only those Entities any of whose Attribute instances (corresponding to the Attributes specified by the query or all if none are specified) match the temporal restrictions imposed by the temporal query (as mandated by clause 4.11); i.e. if the time series, for all the concerned Attributes of an Entity, does not include data corresponding to the temporal query interval, then such Entity shall be removed from S, thus it shall not appear in the final result set. Let S1 be this new subset.
 - If a values filter query is provided, from S1, select those Entities whose Attribute instances (during the interval defined by the temporal query) meet the matching conditions specified by the query (as mandated by clause 4.9); i.e. the values filter query shall be checked against all the Attribute instances resulting from the initial filtering performed by the temporal query. Let S2 be this new subset.
 - If no values filter query is provided, then S2 is equal to S1.
 - If geoquery is present, from S2, select those Entities whose *GeoProperty* instances meet the geospatial restrictions imposed by the geoquery (as mandated by clause 4.10); those geospatial restrictions shall be checked against the *GeoProperty* instances that are within the interval defined by the temporal query. Let S3 be this new subset.
 - If no geoquery is provided, then S3 is equal to S2.
 - If the Scope query is present, from S3, select those Entities whose Entity Scope instances match the Scope query (as mandated by clause 4.19, for an example see annex C, clause C.5.16). Let S4 be the new subset.
 - If no Scope query is provided, then S4 is equal to S3.

- If the ContextBroker implementation supports the use of EntityMaps then:
 - If the location of a resource holding an EntityMap of matching Entity registrations is present it shall be retrieved:
 - If the resource cannot be found, or the data has expired, a new EntityMap shall be created.
 - If the data has not expired, only the retrieved EntityMap shall be used to determine which Context Source Registrations match the Entity ID.
 - If a flag to return an EntityMap was present in the request, and no EntityMap currently exists, then a new EntityMap shall be created.
- Unless local scope is specified (see clause 5.5.13), for **Context Source Registrations** that match the query and support the "queryTemporal" operation (see operations and operation groups in clause 4.20), implementations shall do the following:
 - If an EntityMap is in use for this operation, and an EntityMap entry linked to a Context Source Registration is found, the location of the EntityMap shall be passed as part of the forwarded request.
 - If split entities flag is explicitly set to true or, if not explicitly set, the default setting of the deployment allows split entities, the filters (filter conditions specified by the query, geospatial restrictions imposed by the geoquery, Scope query, Attributes) shall be removed before forwarding the request. These filters then have to be applied after the Entity information from different Context Sources and local information, if there is any, has been aggregated.
 - For any **exclusive**, **redirect** and **inclusive** **Context Source Registrations** that match against the query, the request is forwarded for remote querying by matching endpoints. The result of each remote query is an Entity Array. The returned result is then merged into S4 according to the algorithm defined in clause 4.5.5.
 - For any **auxiliary** **Context Source Registrations** that match against the query, the request is forwarded for remote querying by matching endpoints. Data from the Entity Array received is merged only into S4 when an Attribute instance, whose value of the timeproperty used for the temporal query, is not already present in S4.
- If split entities flag is explicitly set to true or, if not explicitly set, the default setting of the deployment allows split entities, and local scope is not specified, the following filters have to be applied on the aggregated Entities:
 - If a values filter query is provided, from S4, select those Entities whose Attribute instances (during the interval defined by the temporal query) meet the matching conditions specified by the query (as mandated by clause 4.9); i.e. the values filter query shall be checked against all the Attribute instances resulting from the initial filtering performed by the temporal query. Let S5 be this new subset.
 - If no values filter query is provided, then S5 is equal to S4.
 - If geoquery is present, from S5, select those Entities whose GeoProperty instances meet the geospatial restrictions imposed by the geoquery (as mandated by clause 4.10); those geospatial restrictions shall be checked against the GeoProperty instances that are within the interval defined by the temporal query. Let S6 be this new subset.
 - If no geoquery is provided, then S6 is equal to S5.
 - If the Scope query is present, from S6, select those Entities whose Entity Scope instances match the Scope query (as mandated by clause 4.19, for an example see annex C, clause C.5.16). Let S7 be the new subset.
 - If no Scope query is provided, then S7 is equal to S6.

- Otherwise S7 is equal to S4.
- If a *datasetId* parameter is provided, from S7, for all Entities, filter the Attribute instances based on the *datasetIds* specified in the parameter, i.e. keep only the Attribute instances whose *datasetId* is specified. The default Attribute instance is matched, if "@none" is specified. Remove all Attributes without remaining Attribute instances. Let S8 be this new subset.
- If no *datasetId* is provided then S8 equals to S7.
- From the set of Entities that are in S8, include in their temporal representation only the Attribute instances (up to *lastN*) corresponding to the query's projection Attributes, or aggregated values of Attribute instances (if aggregated temporal representation is requested), and which meet the temporal, query and geoquery restrictions.

If an aggregated temporal representation is requested and any of the requested Attributes is not eligible for at least one of the aggregation methods specified in the request parameters, then an error of type *InvalidRequest* shall be raised.

- Pagination logic shall be in place as mandated by clause 5.5.9.
- If in the process of obtaining the query result it is necessary to issue a **Context Source** discovery operation, the same **Context Source** filter input parameter (if present) shall be propagated.

5.7.4.5 Output Data

A JSON-LD array representing the matching entities as defined by clause 5.2.21 and selected according to the behaviour described by clause 5.7.4.4.

If Entity ordering is specified (see clause 4.23), then the JSON-LD array returned shall be ordered according to the member names specified.

If a restrictive list of Entity member names is present, every Entity within the payload body is reduced down to only contain the defined Entity members.

If an exclusionary list of Entity member names is present, the defined Entity members listed are removed from each Entity within the payload.

5.7.5 Retrieve Available Entity Types

5.7.5.1 Description

This operation allows retrieving a list of NGSI-LD entity types for which entity instances exist within the NGSI-LD system.

5.7.5.2 Use case diagram

A *Context Consumer* can retrieve a list of NGSI-LD entity types from the system as shown in Figure 5.7.5.2-1.




Figure 5.7.5.2-1: Retrieve Available Entity Types use case

5.7.5.3 Input data

- An optional JSON-LD context.

5.7.5.4 Behaviour

- Return a JSON-LD object representing the list of entity types, as mandated by clause 5.2.24, for which entity instances exist within the NGSI-LD system. See clause 5.7.11 for architecture-related implementation aspects, in particular the distributed behaviour.

5.7.5.5 Output data

A JSON-LD object representing the list of available entity types, as mandated by clause 5.2.24.

5.7.6 Retrieve Details of Available Entity Types

5.7.6.1 Description

This operation allows retrieving a list with a detailed representation of NGSI-LD entity types for which entity instances exist within the NGSI-LD system. The detailed representation includes the type name (as short name if available in the provided *@context*) and the attribute names that existing instances of this entity type have.

5.7.6.2 Use case diagram

A *Context Consumer* can retrieve a list with a detailed representation of NGSI-LD entity types from the system as shown in Figure 5.7.6.2-1.




Figure 5.7.6.2-1: Retrieve Details of Available Entity Types use case

5.7.6.3 Input data

- An optional JSON-LD context.

5.7.6.4 Behaviour

- Return a list of JSON-LD objects representing the details of available entity types as mandated by clause 5.2.25 for which entity instances exist within the NGSI-LD system. See clause 5.7.11 for architecture-related implementation aspects, in particular the distributed behaviour.

5.7.6.5 Output data

A list of JSON-LD objects representing the details of available entity types as mandated by clause 5.2.25.

5.7.7 Retrieve Available Entity Type Information

5.7.7.1 Description

This operation allows retrieving detailed entity type information about a specified NGSI-LD entity type for which entity instances exist within the NGSI-LD system. The detailed representation includes the type name (as short name if available in the provided *@context*), the count of available entity instances and details about attributes that existing instances of this entity type have, including their name (as short name if available in the provided *@context*) and a list of types the attribute can have (e.g. *Property*, *Relationship* or *GeoProperty*).

5.7.7.2 Use case diagram

A *Context Consumer* can retrieve a detailed representation of a specified NGSI-LD entity type from the system as shown in Figure 5.7.2-1.




Figure 5.7.7.2-1: Retrieve Available Entity Type Information use case

5.7.7.3 Input data

- Entity type name for which detailed information is to be retrieved.
- An optional JSON-LD context.

5.7.7.4 Behaviour

- Return a JSON-LD object representing the details of the specified entity type as mandated by clause 5.2.26, for which instances exist within the NGSI-LD system. See clause 5.7.11 for architecture-related implementation aspects.

5.7.7.5 Output data

A JSON-LD object representing the details of the specified entity type as mandated by clause 5.2.26.

5.7.8 Retrieve Available Attributes

5.7.8.1 Description

This operation allows retrieving a list of NGSI-LD attributes that belong to entity instances existing within the NGSI-LD system.

5.7.8.2 Use case diagram

A *Context Consumer* can retrieve a list of NGSI-LD attributes from the system as shown in Figure 5.7.8.2-1.




Figure 5.7.8.2-1: Retrieve Available Attributes use case

5.7.8.3 Input data

- An optional JSON-LD context.

5.7.8.4 Behaviour

- Return a JSON-LD object representing the list of attributes as mandated by clause 5.2.27 that belong to entity instances existing within the NGSI-LD system. See clause 5.7.11 for architecture-related implementation aspects, in particular the distributed behaviour.

5.7.8.5 Output data

A JSON-LD object representing the list of available attributes as mandated by clause 5.2.27.

5.7.9 Retrieve Details of Available Attributes

5.7.9.1 Description

This operation allows retrieving a list with a detailed representation of NGSI-LD attributes that belong to entity instances existing within the NGSI-LD system. The detailed representation includes the attribute name (as short name if available in the provided *@context*) and the type names for which entity instances exist that have the respective attribute.

5.7.9.2 Use case diagram

A context consumer can retrieve a list with a detailed representation of NGSI-LD attributes from the system as shown in Figure 5.7.9.2-1.




Figure 5.7.9.2-1: Retrieve Details of Available Attributes use case

5.7.9.3 Input data

An optional JSON-LD context.

5.7.9.4 Behaviour

Return a list of JSON-LD objects representing the details of available attributes as mandated by clause 5.2.28 (restricted to the elements *id*, *type*, *attributeName* and *typeNames*) that belong to entity instances existing within the NGSI-LD system. See clause 5.7.11 for architecture-related implementation aspects, in particular the distributed behaviour.

5.7.9.5 Output data

A list of JSON-LD objects representing the details of available attributes as mandated by clause 5.2.28 (restricted to the elements *id*, *type*, *attributeName* and *typeNames*).

5.7.10 Retrieve Available Attribute Information

5.7.10.1 Description

This operation allows retrieving detailed attribute information about a specified NGSI-LD attribute that belongs to entity instances existing within the NGSI-LD system. The detailed representation includes the attribute name (as short name if available in the provided *@context*) and the type names for which entity instances exist that have the respective attribute, a count of available attribute instances and a list of types the attribute can have (e.g. *Property*, *Relationship* or *GeoProperty*).

5.7.10.2 Use case diagram

A context consumer can retrieve a list with a detailed representation of NGSI-LD attributes from the system as shown in Figure 5.7.10.2-1.




Figure 5.7.10.2-1: Retrieve Available Attribute Information use case

5.7.10.3 Input data

- Name of the attribute for which detailed information is to be retrieved.
- An optional JSON-LD context.

5.7.10.4 Behaviour

Return a JSON-LD object representing the details of available attributes as mandated by clause 5.2.28 that belong to entity instances existing within the NGSI-LD system. See clause 5.7.11 for architecture-related implementation aspects, in particular the distributed behaviour.

5.7.10.5 Output data

A JSON-LD object representing the details of available attributes as mandated by clause 5.2.28.

5.7.11 Architecture-related aspects of retrieval of Entity Types and Attributes

Retrieving information about available types or attributes can be an expensive operation depending on the scale and architectural design decisions of the NGSI-LD system. This is in particular the case for retrieving the information about all available entity types and attributes related to all entity information available in an NGSI-LD system. Especially in the case of distributed architecture (clause 4.3.3) and federated architecture (clause 4.3.4) checking all entities can be so expensive that it can become practically infeasible.

Therefore, implementations may only take into account only information that is available or can be derived from a local datastore and the `Context Registry`, when implementing the retrieval of available entity types and attributes, as described in clauses 5.7.5, 5.7.6, 5.7.7, 5.7.8, 5.7.9 and 5.7.10. Context registrations do not always reflect which entity instances are actually available from a `Context Source` at a particular point in time, but only which entity instances are possibly available from a `Context Source`, thus in this case the information about available entity types and attributes is to be interpreted as "possibly available". Also, context registrations can have different granularities, i.e. they possibly only contain entity type or attribute information, and thus the provided information about available entity types and attributes is possibly incomplete as a result. In particular the *attributeNames* in the *EntityType* data structure (clause 5.2.25), the *attributeDetails* in the *EntityTypeInfo* data structure (clause 5.2.26), and the *attributeTypes* and *typeName*s in the *Attribute* data structure (clause 5.2.27) may be provided as empty arrays if the information is not included in the respective context registration. Implementations may also provide estimates for the entity count or attribute count instead of the accurate count.

As an alternative to relying on local information only, the request can be forwarded to all **Context Sources** which support the respective operation according to the **Context Source Registration** describing them. In this case the returned lists are merged with the local list of entity types before returning them. This approach is more expensive but leads to a more accurate result.

5.8 Context Information Subscription

5.8.1 Create Subscription

5.8.1.1 Description

This operation allows creating a new subscription.

5.8.1.2 Use case diagram

A *Context Subscriber* can create a subscription to receive context updates within an NGSI-LD system as shown in Figure 5.8.1.2-1.




Figure 5.8.1.2-1: Create subscription use case

5.8.1.3 Input data

- A data structure (represented in JSON-LD) conforming to the *Subscription* data type as mandated by clause 5.2.12.

5.8.1.4 Behaviour

- If the data types, cardinalities and restrictions expressed by clause 5.2.12 are not met, then an error of type *BadRequestData* shall be raised.
- If the NGSI-LD endpoint already knows about this Subscription, as there is an existing Subscription whose id (URI) is equivalent, an error of type *AlreadyExists* shall be raised.
- If the subscription document does not include a Subscription identifier, a new locally unique identifier (URI) shall be automatically generated by the implementation.

- Then, implementations shall add a new Subscription. The behaviour for corresponding notifications shall be as per clause 5.8.6. The parameters of the created Subscription shall be configured as follows:
 - The Subscription expiration date shall be equal to the value of the *expiresAt* member. If the expiration timestamp provided represents a moment before the current date and time, then an error of type *BadRequestData* shall be raised. If there is no *expiresAt* member the Subscription shall be considered as perpetual.
 - If the value of the *isActive* field is not included or is *true* then the initial status of the Subscription shall be set to "active".
 - If the value of the *isActive* field is *false*, then the initial status of the Subscription shall be set to "paused".
 - If present, the subscribed entities shall be those matching the conditions expressed under the *EntitySelector*, as defined in clause 5.2.33.
 - Watched Attributes shall be those Attributes (subject to clause 5.5.7 Term to URI expansion) pertaining to the subscribed entities (if present) and conveyed through the *watchedAttributes* member. Watched Attributes are those that trigger a new notification when they are changed. A non-present *watchedAttributes* member means that all Attributes shall be watched. If no subscribed entities have been specified, all entities with attributes matching at least one member of *watchedAttributes* are subscribed to.
 - The *@context* to be used for sending Notifications related to this Subscription shall be the one specified in the *jsonldContext* field. If not present, the *jsonldContext* field shall be initialized with the *@context* applicable for the Subscription (if *@context* is also not present in the Subscription, see clause 5.5.5). When the remote JSON-LD *@context* referenced by the *jsonldContext* field is not available implementations shall raise an error of type *LdContextNotAvailable*. If the remote JSON-LD *@context* referenced by the *jsonldContext* field is invalid, implementations shall raise an error of type *BadRequestData*.
 - Based on the content of the Subscription, a Context Source Registration Subscription shall be created (clause 5.11.2). The mapping of the *id* of the Subscription to the returned *subscriptionId* of the Context Source Registration Subscription shall be stored to enable updating or deleting the Context Source Registration Subscription in case of changes to the Subscription.
- If the subscription defines a *timeInterval* member, a Notification shall be sent periodically, when the time interval (in seconds) specified in such value field is reached, regardless of Attribute changes.
- If *timeInterval* is not defined, whenever there is a change in the watched Attribute nodes (Properties or Relationships) of the concerned Entities, implementations shall post a new Notification as per the rules defined by clause 5.8.6.
- If *localOnly=false*, each time a Context Source Notification with the *subscriptionId* of the previously created Context Source Registration Subscription is received, implementations shall do the following:
 - For any **exclusive**, **redirect** and **inclusive** Context Source Registration received as part of the notification, implementation shall do the following depending on the *triggerReason*:
 - If the *triggerReason* is "newlyMatching" and the Context Source Registration indicates support for the Create Subscription operation, a copy of the original Subscription shall be reduced to what is matched by the registration information. If the *splitEntities* member is explicitly set to true or, if not explicitly set, the default setting of the deployment allows split entities, the members *q*, *geoQ* and *scopeQ* shall be removed from the created copy of the original Subscription. Also from the *notification* member, the *attributes*, *pick* and *omit* members are to be removed. The copied Subscription is then forwarded to the Context Source as a new Subscription where the notification endpoint is set to that of the local Broker. The mapping of the received *subscriptionId* with the own Subscription identifier is stored to enable forwarding received notifications to the original subscriber. In addition, a mapping of the *id* of the Context Source Registration to the received *subscriptionId* is stored, to enable updating or deleting the subscription in case of changes.

- If the *triggerReason* is "updated" and the Context Source Registration indicates support for the Update Subscription operation, an update of the original Subscription, reduced to what is matched by the registration information, shall be created. If the *splitEntities* member is explicitly set to true or, if not explicitly set, the default setting of the deployment allows split entities, the members *q*, *geoQ* and *scopeQ* shall be removed from the updated Subscription. Also from the *notification* member the *attributes*, *pick* and *omit* members are to be removed. The updated Subscription is then forwarded to the Context Source with the *subscriptionId* related to the Context Source Registration. As an optimization, an implementation may keep the originally forwarded Context Source Registration and compare with the new one to only forward the update, if there was a relevant change.
 - If the *triggerReason* is "noLongerMatching" and the Context Source Registration indicates support for the delete Subscription operation, a delete Subscription shall be forwarded to the Context Source with the *subscriptionId* related to the Context Source Registration.
- Implementations shall ensure that, when the Subscription expiration date is due, the status of the Subscription changes automatically to "expired", so that notifications will no longer be sent.

5.8.1.5 Output data

A URI identifying the newly created Subscription.

5.8.2 Update Subscription

5.8.2.1 Description

This operation allows updating an existing subscription.

5.8.2.2 Use case diagram

A Context Subscriber can update an existing subscription within an NGSI-LD system as shown in Figure 5.8.2.2-1.




Figure 5.8.2.2-1: Update subscription use case

5.8.2.3 Input data

- Subscription identifier (URI), the target subscription.
- A JSON-LD document representing a Subscription Fragment.

5.8.2.4 Behaviour

- If the Subscription id is not present or it is not a valid URI, then an error of type **BadRequestData** shall be raised.
- If the NGSI-LD System does not know about the target Subscription, because there is no existing Subscription whose id (URI) is equivalent, an error of type **ResourceNotFound** shall be raised.
- Execute the behaviour defined in clause 5.5.4 on JSON-LD validation.
- If the data types and restrictions expressed by clause 5.2.12 are not met by the *Subscription Fragment*, then an error of type **BadRequestData** shall be raised.
- Term to URI expansion of Attribute names shall be observed as mandated by clause 5.5.7.
- If the *jsonldContext* field is present and the referenced JSON-LD *@context* is not available, implementations shall raise an error of type **LdContextNotAvailable**. If the referenced JSON-LD *@context* is invalid, implementations shall raise an error of type **BadRequestData**.
- Then, implementations shall modify the target Subscription as mandated by clause 5.5.8.
- Finally, the following extra behaviour shall be observed when updating Subscriptions:
 - If *isActive* is equal to *true* and *expiresAt* is not present, then *status* shall be updated to "active", if and only if, the previous value of *status* was different than "expired".
 - If *isActive* is equal to *true* and *expiresAt* corresponds to a *DateTime* in the future, then *status* shall be updated to "active".
 - If *isActive* is equal to *false* and *expiresAt* is not present, then *status* shall be updated to "paused", if and only if, the previous value of *status* was different than "expired".
 - If only *expiresAt* is included and refers to a *DateTime* in the future, then *status* shall be updated to "active", if and only if the previous value of *status* was "expired".
 - If *expiresAt* is included but referring to a *DateTime* in the past, then a **BadRequestData** error shall be raised, regardless the value of *isActive*.
 - Based on the mapping of the Subscription to its respective Context Source Registration Subscription (see clause 5.8.1.4), that Context Source Registration Subscription shall be updated (clause 5.11.3).

5.8.2.5 Output data

None.

5.8.3 Retrieve Subscription

5.8.3.1 Description

This operation allows retrieving an existing subscription.

5.8.3.2 Use case diagram

A *Context Subscriber* can retrieve a specific subscription from an NGSI-LD system as shown in Figure 5.8.3.2-1.




Figure 5.8.3.2-1: Retrieve subscription use case

5.8.3.3 Input data

Id (URI) of the subscription to be retrieved (target subscription).

5.8.3.4 Behaviour

- If the subscription Id is not present or it is not a valid URI, then an error of type *BadRequestData* shall be raised.
- If the identifier provided does not correspond to any existing subscription in the system then an error of type *ResourceNotFound* shall be raised.
- Otherwise, implementations shall query the subscriptions and obtain the subscription data to be returned to the caller.

5.8.3.5 Output data

A JSON-LD object representing the subscription details as mandated by clause 5.2.12.

5.8.4 Query Subscriptions

5.8.4.1 Description

This operation allows querying existing Subscriptions.

5.8.4.2 Use case diagram

A Context Consumer can query the existent Subscriptions from an NGSI-LD system as shown in Figure 5.8.4.2-1.




Figure 5.8.4.2-1: Query subscriptions use case

5.8.4.3 Input data

A limit to the number of subscriptions to be retrieved. See clause 5.5.9.

5.8.4.4 Behaviour

- The NGSI-LD system shall list all the existing subscriptions up to the limit specified as input data. If no limit is specified the number of subscriptions retrieved may depend on the implementation.
- Pagination logic shall be in place as mandated by clause 5.5.9.

5.8.4.5 Output data

A list (represented as a JSON array) of JSON-LD objects each one representing subscription details as mandated by clause 5.2.12.

5.8.5 Delete Subscription

5.8.5.1 Description

This operation allows deleting an existing subscription.

5.8.5.2 Use case diagram

A *Context Subscriber* can delete a subscription within an NGSI-LD system as shown in Figure 5.8.5.2-1.




Figure 5.8.5.2-1: Delete subscription use case

5.8.5.3 Input data

A subscription identifier (URI).

5.8.5.4 Behaviour

- If the subscription Id is not present or it is not a valid URI, then an error of type *BadRequestData* shall be raised.
- If the subscription id provided does not correspond to any existing subscription in the system, then an error of type *ResourceNotFound* shall be raised.
- Otherwise, implementations shall delete the Subscription and no longer perform notifications concerning such Subscription.
- Using the mapping of the own Subscription identifier to each of the *subscriptionId* of a subscription to a Context Source, a delete Subscription shall be forwarded to each such Context Source, if the delete Subscription operation is supported as indicated in the corresponding Context Source Registration:
 - Based on the mapping of the Subscription to its respective Context Source Registration Subscription (see clause 5.8.1.4), that Context Source Registration Subscription shall be deleted (clause 5.11.6).

5.8.5.5 Output data

None.

5.8.6 Notification behaviour

A notification is a message that allows a subscriber to be aware of the changes in subscribed Entities. Implementations shall exhibit the following behaviour:

- Notifications shall only be sent if and only if the status of the corresponding subscription (*subscription.status*) is "active", i.e. not "paused" nor "expired".

- If a Subscription defines a *timeInterval* member, a Notification shall be sent periodically, when the time interval (in seconds) specified in such value field is reached, regardless of Attribute changes. The notification message shall include all the subscribed Entities that match the query, geoquery and Scope query conditions. If none of query, geoquery and Scope query are defined, then all subscribed Entities shall be included. For each entity in the notification, only the Attribute instances are to be included that match the *datasetId* member in Subscription. If there are no matching Entities, no Notification is sent.
- If a Subscription does not define a *timeInterval* term, the notification shall be sent whenever there is a change in the watched Attributes. An Attribute is considered to change when any of the members (including children) in its corresponding JSON-LD node is updated with a value different than the existing one. The notification message shall include all the subscribed Entities that changed and that match (as mandated by clauses 4.9, 4.10 and 4.19) the query, geoquery and Scope query conditions. If query, geoquery and scopeQuery are all not defined then all subscribed Entities that changed shall be included. If, for an Entity, there are multiple instances of the *GeoProperty* on which the geoquery is based, it is sufficient if any of these instances meets the geospatial restrictions. For each entity in the notification, only the Attribute instances are to be included that match the *datasetId* member in Subscription. Finally, if a *Context Source* filter is defined, then only the subscribed Entities whose origin *Context Source* matches the referred filter shall be included.
- If a Notification with a *subscriptionId* is received that has a mapping to a local Subscription identifier, the Notification shall be copied.
If the *splitEntities* member of the local Subscription is explicitly set to true or, if not explicitly set, the default setting of the deployment allows split entities:
 - The Entities contained in the data member of the Notification shall be retrieved locally and from all Context Sources that have information about these Entities, except for the one from which the Notification has been received.
 - The retrieved Entities then shall be merged with the Entities in the data member.
 - All Entities that do not match the query, geoquery and Scope query conditions of the local Subscription shall be removed from the data member.
- If there are Entities in the *data* member of the Notification copy, the Notification copy shall be forwarded to the Notification endpoint specified in the local Subscription, using this local Subscription identifier instead of the *subscriptionId* **received**.
- A Notification shall be sent as follows:
 - The structure of the notification message shall be as mandated by clause 5.3.1.
 - The *@context* to be used is the one specified in the *jsonldContext* field of the corresponding Subscription.
 - The Attributes included (*Properties* or *Relationships*) shall be those specified by the *notification.attributes* member in the Subscription data type (clause 5.2.12). Term to URI expansion shall be observed (clause 5.5.7). The absence of the *notification.attributes* member of a Subscription means that all Attributes shall be included. If the notification was triggered by the deletion of an Entity and the *notification.showChanges* member is not set to *true*, only the *deletedAt* system property shall be provided. In case *notification.sysAttrs* is set to *true*, also *createdAt* and *modifiedAt* shall be provided.
 - If an Attribute has been deleted, only the name of the attribute as key and the URI "*urn:ngsi-ld:null*" as value shall be provided, unless more information is required. The latter is the case, if:
 - a *datasetId* needs to be provided;
 - the *notification.sysAttrs* is set to *true* and thus the system generated sub-attributes (see clause 4.8) have to be provided;
 - *notification.showChanges* is set to *true* and thus a previous value or object has to be provided.

In all such cases, a JSON object with all the required information is provided, where the *value* or the *object* is set to the URI "*urn:ngsi-ld:null*" respectively or, in case of a *LanguageProperty*, the *languageMap* is set to {"@none": "*urn:ngsi-ld:null*"}

- If the *notification.format* member value is set, the representation of the entities changes:
 - If the *notification.format* is set to "concise" then a concise representation of the entities shall be provided as defined by clause 4.5.1.
 - If the *notification.format* is set to "simplified" (or "keyValues" as a synonym), then a simplified representation of the entities (as mandated by clause 4.5.4) shall be provided.
 - Otherwise the normalized format as defined by clause 4.5.1 shall be used.
- If the *notification.pick* member value is set, every Entity within the payload body is reduced down to only contain the defined entity members listed.
- If the *notification.omit* member value is set, the defined entity members listed are removed from each Entity within the payload.
- If the *notification.join* member value is set then **Linked Entity retrieval** (as mandated by clause 4.5.23) shall be provided.
- If the *ngsildConformance* field of the corresponding Subscription is set, the notification shall undergo a backwards compatibility operation as defined by clause 4.3.6.8 and be amended to conform to the supplied version of the NGSI-LD specification.
- A Notification shall be sent (as mandated by each concrete binding and including any optional *endpoint.receiverInfo* defined by clause 5.2.22) to the endpoint specified by the *endpoint.uri* member of the notification structure defined by clause 5.2.14. The Notification content shall be JSON by default. However, this can be changed to JSON-LD or GeoJSON by means of the *endpoint.accept* member.
- The *notification.timesSent* member shall be incremented by one.
- The *notification.lastNotification* member shall be updated with a timestamp representing the current date and time.
- If the response to the notification request is 200 OK then implementations shall:
 - Update *notification.lastSuccess* with a timestamp representing the current date and time.
 - Update *notification.status* to "ok".
- If the response to the notification request is different than 200 OK then implementations shall:
 - Update *notification.lastFailure* with a timestamp representing the current date and time.
 - Update *notification.Status* to "failed".

5.9 Context Source Registration

5.9.1 Introduction

As described in clause 5.2.9, **Context Source Registrations** have a similar structure as Entities and are generally handled in the same way. However, there are some aspects that are specific to Registrations, in particular with respect to the handling of required properties. Thus, the operation descriptions for Registrations reference the respective operations for Entities and in addition specify any deviations and additions that are necessary for handling **Context Source Registrations**.

Context Source Registrations either contain information about Context Sources providing the latest information or about Context Sources providing temporal information, but not both. Context Sources that can provide both thus have to use two separate Context Source Registrations. If no temporal query is present, only Context Source Registrations for Context Sources providing latest information are returned, i.e. those which do not specify time intervals used for temporal queries. If a temporal query is present in a request for Context Source Registrations, only those Context Source Registrations that have a matching time interval are returned.

5.9.2 Register Context Source

5.9.2.1 Description

This operation allows registering a context source within an NGSI-LD system.

5.9.2.2 Use case diagram

A *Context Provider* can register one or more context sources within an NGSI-LD system as shown in Figure 5.9.2.2-1.




Figure 5.9.2.2-1: Register context source use case

5.9.2.3 Input data

A data structure conforming to the *CSourceRegistration* data type as mandated by clause 5.2.9.

5.9.2.4 Behaviour

Implementations shall generally exhibit the behaviour described in clause 5.6.1.4, but instead of any type of entities only Context Source Registrations can be provided and the following exceptions shall apply:

- If the data types and restrictions expressed by clause 5.2.9 are not met by the Context Source Registration, then an error of type **BadRequestData** shall be raised.
- If a *contextSourceInfo* array is defined and the restrictions expressed by clause 4.3.6.6 are not met by the Context Source Registration, then an error of type **BadRequestData** shall be raised.
- If the Context Source to be registered has its *mode* property defined as **exclusive**, the following additional restrictions apply:
 - If an **exclusive** or **redirect** Context Source Registration already matches against the Entity ID (URI) and any of the Attributes defined in the registration, an error of type **Conflict** shall be raised.

- If an Entity already exists for the supplied Entity ID (URI) and the existing Entity contains any of the Attributes defined in the registration, an error of type **Conflict** shall be raised.
- If the Context Source to be registered has its *mode* property defined as **redirect**, the following additional restriction applies:
 - If an existing Entity already matches the Context Source Registration, an error of type **Conflict** shall be raised.
- If the Context Source to be registered has its *mode* property defined as **auxiliary**, the following additional restriction applies:
 - If the *operations* property is not defined as one of: "retrieveOps", "retrieveEntity" or "queryEntity", or a combination thereof, an error of type **BadRequestData** shall be raised.
- If the property *expiresAt* is not defined then the Context Source Registration shall last forever (or until it is deleted from the system).
- If *expiresAt* is a date and time in the past, an error of type **BadRequestData** shall be raised.
- If *expiresAt* is a date and time in the future, implementations shall delete the Registration when this point in time is reached.
- If the NGSI-LD endpoint knows about this Context Source Registration, as there is an existing Context Source Registration whose id (URI) is equivalent, an error of type **AlreadyExists** shall be raised.
- If the subscription document does not include a Context Source Registration identifier, a new identifier (URI) shall be automatically generated by the implementation.
- Implementations shall add the concerned Context Source Registration and return an 'ok' response together with a registration identifier (*id*).
- This *id* shall be used if NGSI-LD clients need to manage the registration later.

5.9.2.5 Output data

A URI identifying the newly created ContextSourceRegistration.

5.9.3 Update Context Source Registration

5.9.3.1 Description

This operation allows updating a Context Source Registration in an NGSI-LD system.

5.9.3.2 Use case diagram

A Context Provider can update a Context Source Registration in an NGSI-LD system as shown in Figure 5.9.3.2-1.




Figure 5.9.3.2-1: Update context source registration use case

5.9.3.3 Input data

- Context Source Registration identifier (URI), the target Context Source Registration.
- A JSON-LD document representing a Context Source Registration Fragment (clause 5.4).

5.9.3.4 Behaviour

- If the target Context Source Registration id (*id*) is not present or it is not a valid URI, then an error of type *BadRequestData* shall be raised.
- If a "contextSourceInfo" array is defined and the restrictions expressed by clause 4.3.6.6 are not met by the Context Source Registration, then an error of type *BadRequestData* shall be raised.
- If the NGSI-LD System does not know about the target Context Source Registration, because there is no existing Context Source Registration whose id (URI) is equivalent, an error of type *ResourceNotFound* shall be raised.
- Execute the behaviour defined in clause 5.5.4 on JSON-LD validation.
- If the data types and restrictions expressed by clause 5.2.9 are not met by the Context Source Registration *Fragment*, then an error of type *BadRequestData* shall be raised.
- Term to URI expansion of Attribute names shall be observed as mandated by clause 5.5.7.
- If the Context Source Registration to be updated has its *mode* property defined as **exclusive**, the following additional restrictions apply:
 - If an **exclusive** or **redirect** Context Source Registration already matches against the Entity ID (URI) and any of the Attributes defined in the registration, an error of type *Conflict* shall be raised.
 - If an Entity already exists for the supplied Entity ID (URI) and the existing Entity contains any of the Attributes defined in the registration, an error of type *Conflict* shall be raised.
- If the Context Source Registration to be updated has its *mode* property defined as **redirect**, the following additional restriction applies:
 - If an existing Entity already matches the Context Source Registration, an error of type *Conflict* shall be raised.

- If the Context Source to be updated has its *mode* property defined as **auxiliary**, the following additional restriction applies:
 - If the *operations* property is not defined as one of: "retrieveOps", "retrieveEntity" or "queryEntity", an error of type **BadRequestData** shall be raised.
- Then, implementations shall modify the target Context Source Registration as mandated by clause 5.5.8.

5.9.3.5 Output data

None.

5.9.4 Delete Context Source Registration

5.9.4.1 Description

This operation allows deleting a Context Source Registration from an NGSI-LD system.

5.9.4.2 Use case diagram

A context provider can delete a context source registration from an NGSI-LD system as shown in Figure 5.9.4.2-1.




Figure 5.9.4.2-1: Delete context source registration use case

5.9.4.3 Input data

Registration identifier (URI) of the context source registration to be deleted (target registration).

5.9.4.4 Behaviour

- If the target context source registration id is not present or it is not a valid URI, then an error of type **BadRequestData** shall be raised.
- If the NGSI-LD endpoint does not know about the target context source registration, because there is no existing context source registration whose id (URI) is equivalent, then an error of type **ResourceNotFound** shall be raised.
- Otherwise the context source registration shall be removed.

5.9.4.5 Output data

None.

5.10 Context Source Discovery

5.10.1 Retrieve Context Source Registration

5.10.1.1 Description

This operation allows retrieving a specific context source registration from an NGSI-LD system.

5.10.1.2 Use case diagram

A context consumer or a context provider can retrieve a specific context source registration from an NGSI-LD system as shown in Figure 5.10.1.2-1.




Figure 5.10.1.2-1: Retrieve context source registration use case

5.10.1.3 Input data

Context source registration identifier (*id*) of the context source registration to be retrieved (target registration).

5.10.1.4 Behaviour

- If the context source registration id (*id*) is not present or it is not a valid URI, then an error of type *BadRequestData* shall be raised.
- If the NGSI-LD endpoint does not know about the target context source registration, because there is no existing context source registration whose id (URI) is equivalent, then an error of type *ResourceNotFound* shall be raised.
- Term to URI expansion of Attribute names shall be observed as mandated by clause 5.5.7.
- Otherwise return a JSON-LD object representing the Context Source Registration as mandated by clause 5.2.9.

5.10.1.5 Output data

A JSON-LD object representing the target context source registration as mandated by clause 5.2.9.

5.10.2 Query Context Source Registrations

5.10.2.1 Description

This operation allows discovering context source registrations from an NGSI-LD system. The behaviour of the discovery of context source registrations differs significantly from the querying of entities as described in clause 5.7.2. The approach is that the client submits a query for entities as described in clause 5.7.2, but instead of receiving the Entity information, it receives a list of **Context Source Registrations** describing **Context Sources** that possibly have some of the requested Entity information. This means that the requested Entities and Attributes are matched against the 'information' property as described in clause 5.12.

If no temporal query is present, only **Context Source Registrations** for **Context Sources** providing latest information, i.e. without specified time intervals, are considered. If a temporal query is present only **Context Source Registrations** with matching time intervals, i.e. *observationInterval* or *managementInterval*, are considered.

5.10.2.2 Use case diagram

A *Context Consumer* can discover context source registrations that may be able to provide (part of) the context information specified in the query from an NGSI-LD system as shown in Figure 5.10.2.2-1.




Figure 5.10.2.2-1: Discover context source registrations use case

5.10.2.3 Input data

- A reference to a JSON-LD *@context* (optional).
- A selector of Entity types as specified by clause 4.17. Both type name (short hand string) and fully qualified type name (URI) are allowed (optional).
- A list (one or more) of Entity identifiers (optional).
- A list (one or more) of Attribute names (called query projection attributes) (optional).
- An id pattern as a regular expression (optional).

- An NGS-LD query (to filter out Entities by Attribute values, used here to identify relevant attributes) as per clause 4.9 (optional).
- An NGS-LD geoquery (to filter out Entities by spatial relationships, used here to identify relevant GeoProperties and for geographical scoping) as per clause 4.10 (optional).
- In the case of GeoJSON representation:
 - The name of the *GeoProperty* attribute to use as the geometry for the GeoJSON representation as mandated by clause 4.5.16 (optional).
 - A *datasetId* specifying which instance of the value is to be selected if the *GeoProperty* value has multiple instances as defined by clause 4.5.5 (optional).
- An NGS-LD temporal query as per clause 4.11 (optional).
- An NGS-LD context source query as per clause 4.9 (optional).
- A NGS-LD Scope query as mandated by clause 4.19 (optional).
- A limit to the number of **Context Source Registrations** to be retrieved. See clause 5.5.9.
- A specified language filter as per clause 4.15 (optional).

It is not possible to retrieve a set of context source registrations related to entities by only specifying desired Entity identifiers, without further specifying restrictions on the entities' types or attributes, either explicitly, via lists of Entity types or of Attribute names, or implicitly, within an NGS-LD query or geoquery.

5.10.2.4 Behaviour

- Execute the behaviour defined in clause 5.5.4 on JSON-LD validation.
- At least one of the following input data shall be provided:
 - a) selector of Entity Types;
 - b) list of Attribute names;
 - c) NGS-LD Query;
 - d) NGS-LD GeoQuery.

If none of them is provided, then an error of type ***BadRequestData*** shall be raised (too wide query). Attributes specified in NGS-LD Query or NGS-LD GeoQuery shall be used for matching *RegistrationInfo* elements in the same way as the attributes in the *list of Attribute names*.

- If the list of Entity identifiers includes a URI which it is not valid, or the query, geoquery or temporal query are not syntactically valid (as per clauses 4.9, 4.10 and 4.11) an error of type ***BadRequestData*** shall be raised.
- Term to URI expansion of type and Attribute names shall be performed, as mandated by clause 5.5.7.
- Otherwise, implementations shall run a query that shall return context source registrations that meet **all** the applicable conditions:
 - If present, the entity specification in the query consisting of a combination of entity type selector and entity id/entity id pattern (optional) matches an *EntityInfo* specified in a *RegistrationInfo* of the information property in a context source registration. If there is no *EntityInfo* specified in the *RegistrationInfo*, the entity specification is considered matching. This matching is further described in clause 5.12.
 - If present, at least one Attribute name specified in the query matches one Property or Relationship in the *RegistrationInfo* element of the information property in a context source registration. If no Properties or Relationships are specified in the *RegistrationInfo*, the Attribute names are considered matching. This matching is further described in clause 5.12.

- If present, the geoquery is matched against the *GeoProperty* identified in the geoquery. If no *GeoProperty* is specified in the geoquery, the default property is *location*. The geoquery matches the *GeoProperty* specified in the **Context Source Registration**, if the location directly matches or if the location possibly contains locations that would match the geoquery.
 - If no temporal query is present, only **Context Source Registrations** for **Context Sources** providing latest information, i.e. without specified time intervals, are considered.
 - If a temporal query is present, only **Context Source Registrations** with specified time intervals, i.e. *observationInterval* or *managementInterval* are considered. If the *timeproperty* is *observedAt* or no *timeproperty* is specified in the temporal query (default: *observedAt*), the temporal query is matched against the *observationInterval* (if present). If the *timeproperty* is *createdAt*, *modifiedAt* or *deletedAt*, the temporal query is matched against the *managementInterval* (if present). If the relevant interval is not present, there is no match:
 - The semantics of the match is that the *timeAt* in the case of the "before" and "after" relation is contained in or is an endpoint of a time period included in the specified time interval. In the case of the "between" relation there is a match if there is an overlap between the interval specified by the *timeAt* and *endTimeAt* and the specified time interval.
 - If present, the conditions specified by the context source query filter match the respective **Context Source Properties** (as mandated by clause 4.9).
 - If present, the Scope query (as mandated by clause 4.19) is matched against the *scope* property.
- Pagination logic shall be in place as mandated by clause 5.5.9.

5.10.2.5 Output data

A JSON-LD array of matching **Context Source Registrations** as defined by clause 5.2.9. Instead of the original **Context Source Registration** which may contain a lot of irrelevant information, implementations should return filtered **Context Source Registrations**, which only contain context source registration information relevant for the query, in particular only matching *RegistrationInfo* elements.

5.11 Context Source Registration Subscription

5.11.1 Introduction

Context Source Registration Subscriptions in general work like context information subscriptions; however, instead of resulting in notifications with context information, the notifications contain **Context Source Registrations** describing **Context Sources** that can potentially provide the requested context information. If no temporal query is present, only **Context Source Registrations** for **Context Sources** providing latest information, i.e. without such time intervals, are considered. If a temporal query is present only **Context Source Registrations** with matching time intervals, i.e. *observationInterval* or *managementInterval*, are considered.

5.11.2 Create Context Source Registration Subscription

5.11.2.1 Description

This operation allows creating a new **Context Source Registration Subscription**.

5.11.2.2 Use case diagram

A **Context Source Subscriber** can subscribe to a new **Context Source Registration** as shown in Figure 5.11.2.2-1.




Figure 5.11.2.2-1: Subscribe Context Source Registration use case

5.11.2.3 Input data

- A data structure (represented in JSON-LD) conforming to the Subscription data type as mandated by clause 5.2.12.

5.11.2.4 Behaviour

- The behaviour shall be as described in clause 5.8.1.4, restricted to the local case (where Subscriptions cannot be received from remote Brokers), with the following exceptions:
 - If all checks described in clause 5.8.1.4 pass, implementations shall add a new **Context Source Registration Subscription**. The parameters of the created subscription shall be configured as described in clause 5.2.12.
 - Instead of directly matching the entities and watched Attributes from the Subscription with the **Context Source** registrations, the entities specified in the subscription, the watched Attributes and the Attributes specified in the notification parameter are matched against the respective *information* property of the **Context Source** registrations. If either the watched Attributes or the Attributes in the notification are not present or of length 0, all possible Attributes (if present in the **Context Source Registrations**) for matching entities match. This matching is further described in clause 5.12.
 - If present, the geoquery in the geoQ element is matched against the *GeoProperty* of the subscription identified in the geoQ element. If no *GeoProperty* is specified in the geoquery, the default property is 'location'. The geoquery matches the *GeoProperty* specified in the **Context Source Registration**, if the location directly matches or if the location possibly contains locations that would match the geoquery.
 - If no temporal query is present in the *temporalQ* element, only **Context Source Registrations** for **Context Sources** providing latest information, i.e. without specified time intervals for *observationInterval* or *managementInterval*, are considered.

- If a temporal query in the *temporalQ* element is present, only Context Source Registrations with specified time intervals are considered. If the *timeproperty* is "observedAt" or no *timeproperty* is specified in the temporal query (default: *observedAt*), the temporal query is matched against the *observationInterval* (if present). If the *timeproperty* is "createdAt", "modifiedAt" or "deletedAt", the temporal query is matched against the *managementInterval* (if present). If the relevant interval is not present, there is no match:
 - The semantics of the match is that the *timeAt* in the case of the "before" and "after" relation is contained in or is an endpoint of a time period included in the specified time interval. In the case of the "between" relation there is a match if there is an overlap between the interval specified by the *timeAt* and *endTimeAt* and the specified time interval.
- If present, the conditions specified by the context source filter match the respective Context Source Properties (as mandated by clause 4.9).
- If the subscription defines a *timeInterval* term, a *CsourceNotification* (clause 5.3.2) with all matching Context Source Registrations shall be sent periodically, initially on subscription and when the time interval (in seconds) specified in such value field is reached, independent of any changes to the set of Context Source registrations.
- If *timeInterval* is not defined, initially on subscription and whenever there is a change of a matching Context Source Registration (creation, update, deletion), implementations shall post a new *CsourceNotification* to the endpoint specified in the notification parameters informing about this change by providing the Context Source Registration(s) together with the appropriate trigger reason in the *triggerReason* member.
- If present, the conditions specified by the context source query match the respective Context Source Properties (as mandated by clause 4.9).
- If present, the Scope query (as mandated by clause 4.19) is matched against the *scope* property.

5.11.2.5 Output data

A URI identifying the newly created Subscription.

5.11.3 Update Context Source Registration Subscription

5.11.3.1 Description

This operation allows updating an existing Context Source Registration Subscription.

5.11.3.2 Use case diagram

A context source subscriber can update a Context Source Registration Subscription as shown in Figure 5.11.3.2-1.




Figure 5.11.3.2-1: Update Context Source Registration Subscription use case

5.11.3.3 Input data

- Subscription identifier (URI), the target Context Source Registration Subscription.
- A JSON-LD document representing a Subscription Fragment.

5.11.3.4 Behaviour

- If the Subscription Id is not present or it is not a valid URI, then an error of type *BadRequestData* shall be raised.
- If the data types and restrictions expressed by clause 5.2.12 are not met by the Subscription Fragment, then an error of type *BadRequestData* shall be raised.
- Then, implementations shall modify the target subscription as mandated by clause 5.5.8.
- Finally, send a notification with all currently matching Context Source Registrations.

5.11.3.5 Output data

None.

5.11.4 Retrieve Context Source Registration Subscription

5.11.4.1 Description

This operation allows retrieving an existing Context Source Registration Subscription.

5.11.4.2 Use case diagram

A Context Source subscriber can retrieve a specific Context Source Registration Subscription as shown in Figure 5.11.4.2-1.




Figure 5.11.4.2-1: Retrieve Context Source Registration Subscription use case

5.11.4.3 Input data

Id (URI) of the subscription to be retrieved (target subscription).

5.11.4.4 Behaviour

- If the subscription Id is not present or it is not a valid URI, then an error of type *BadRequestData* shall be raised.
- If the identifier provided does not correspond to any existing subscription in the system then an error of type *ResourceNotFound* shall be raised.
- Otherwise implementations shall query the Context Source Registration Subscriptions and obtain the subscription data to be returned to the caller.

5.11.4.5 Output data

A JSON-LD object representing the subscription details as mandated by clause 5.2.12.

5.11.5 Query Context Source Registration Subscriptions

5.11.5.1 Description

This operation allows querying existing Context Source Registration Subscriptions.

5.11.5.2 Use case diagram

A context source subscriber can query all existing Context Source Registration Subscriptions as shown in Figure 5.11.5.2-1.




Figure 5.11.5.2-1: Query Context Source Registration Subscriptions use case

5.11.5.3 Input data

A limit to the number of Context Source Registration Subscriptions to be retrieved. See clause 5.5.9.

5.11.5.4 Behaviour

- The NGSI-LD System shall list all the existing Context Source Registration Subscriptions.
- Pagination logic shall be in place as mandated by clause 5.5.9.

5.11.5.5 Output data

A list (represented as a JSON array) of JSON-LD objects each one representing subscription details as mandated by clause 5.2.12.

5.11.6 Delete Context Source Registration Subscription

5.11.6.1 Description

This operation allows deleting an existing Context Source Registration Subscription.

5.11.6.2 Use case diagram

A context source subscriber can delete a Context Source Registration Subscription as shown in Figure 5.11.6.2-1.




Figure 5.11.6.2-1: Delete Context Source Registration Subscriptions use case

5.11.6.3 Input data

- A subscription identifier (URI).

5.11.6.4 Behaviour

- If the subscription Id is not present or it is not a valid URI, then an error of type *BadRequestData* shall be raised.
- If the subscription id provided does not correspond to any existing subscription in the system then an error of type *ResourceNotFound* shall be raised.
- Otherwise implementations shall delete the Context Source Registration Subscription and no longer perform notifications concerning that Subscription.

5.11.6.5 Output data

None.

5.11.7 Notification behaviour

A Context Source Notification is a message that allows a subscriber to be aware of the changes in the set of Context Source Registrations describing Context Sources that can potentially provide the requested context information. Implementations shall exhibit the behaviour described in clause 5.8.6 with the following exceptions:

- If a subscription defines a *timeInterval* member, a *CsourceNotification* (clause 5.3.2) shall be sent on initial subscription and periodically, when the time specified time interval (in seconds) has elapsed, regardless of any changes to the set of context source registrations. The *CsourceNotification* message shall include all the Context Source Registrations whose *information* property matches the entities and watched Attributes or Attributes specified in the notification parameter and, if present, have a matching geoquery. If either the watched Attributes or the Attributes in the notification are not present or of length 0, all possible Attributes (if present in the Context Source Registrations) for fitting entities match.
- If a subscription does not define a *timeInterval* term, the csource notification shall be sent on initial subscription and whenever there is a change in a matching csource registration. Such a change may be triggered by the creation of a new matching csource registration, the update of a csource registration (whether matching before the update, after the update or in both cases) or the deletion of a matching csource registration. The notification message shall include the matching csource registration(s) together with the appropriate trigger reason in the *triggerReason* member.

- Instead of providing the original **Context Source Registration** which may contain a lot of irrelevant information, implementations should return filtered **Context Source Registrations**, which only contain context source registration information relevant for the subscription, in particular only matching *RegistrationInfo* elements.
- A csource notification shall be sent as follows:
 - The structure of the csource notification message shall be as mandated by clause 5.3.2.
 - A csource notification shall be sent to the *endpoint*.
 - The *notification.timesSent* member shall be incremented by one.
 - The *notification.lastNotification* member shall be updated with the current timestamp.
 - If the notification is sent successfully:
 - Update *notification.lastSuccess* with the current timestamp.
 - If the notification is not sent successfully:
 - Update *notification.lastFailure* with the current timestamp.
 - Update the subscription *status* to "failed".

5.12 Matching Context Source Registrations

When querying **Context Source Registrations** as described in clause 5.10.2 and subscribing to **Context Source Registrations** as described in clause 5.11.2, the Entities and/or Attributes specified in the request have to be matched against the set of **Context Source Registrations**, extracting the matching ones. This clause describes this matching.

The relevant specification information in the query for **Context Source Registrations** are the selector of Entity Types (if present), the list of Entity identifiers (if present), the id pattern (if present) and the list of Attribute names (if present). In the case of subscriptions to **Context Source Registrations**, it is the Entities as specified in the array of type *EntitySelector* in the Subscription, the *watchedAttributes* element of the *Subscription* and the attributes specified as part of the *NotificationParams* element of the Subscription. If the attributes in the *NotificationParams* element are empty or not present, the matching is done as if no attribute identifiers have been specified, otherwise the combination of the *watchedAttributes* and the attributes in the *NotificationParams* element are used as the specified attribute identifiers for the matching.

Even though the way relevant Entities are specified differs in queries and subscriptions, they consist of the same information, so for the purpose of this clause, the specification of Entity Types or Attributes refers to the relevant elements for matching, i.e. Entity Types, Entity identifiers, id pattern and Attribute names. A specification of Entity Types or Attributes shall contain at least one of:

- a) selector of Entity Types; or
- b) list of Attribute names.

A specification of Entity Types or Attributes matches a **Context Source Registration** if at least one of the *RegistrationInfo* elements in the *information* element matches. An Entity specification matches a *RegistrationInfo* if the following conditions hold:

- If present, the selector of Entity Types, Entity identifiers and id pattern match at least one of the *EntityInfo* elements of the *RegistrationInfo* (see below).
- If present, the Attribute identifiers match the combination of Properties and Relationships specified in the *RegistrationInfo* (see below).

An Entity specification consisting of selector of Entity Types, Entity identifiers and id pattern matches an *EntityInfo* element of the *RegistrationInfo* if the type selector matches the entity types in the *EntityInfo* element and one of the following conditions holds:

- The *EntityInfo* contains neither an *id* nor an *idPattern*.
- One of the specified Entity identifiers matches the *id* in the *EntityInfo*.
- At least one of the specified Entity identifiers matches the *idPattern* in the *EntityInfo*.
- The specified id pattern matches the *id* in the *EntityInfo*.
- Both a specified id pattern and an *idPattern* in the *Entity Info* are present (since in the general case it is not easily feasible to determine if there can be identifiers matching both patterns).

Attribute names match the combination of Properties and Relationships if one of the following conditions hold:

- No Attribute names have been specified (as this means all Attributes are requested).
- The combination of Properties and Relationships is empty (as this means only Entities have been registered and the **Context Sources** may have matching Property or Relationship instances).
- If at least one of the specified attribute names matches a Property or Relationship specified in the *RegistrationInfo*.

If the request that triggered the matching includes a *datasetId* parameter and the CSourceRegistration to be matched contains a *datasetId* element, the CSourceRegistration should only be considered matching, if both have at least one value in common. If only one of them specifies a *datasetId*, it is considered a match.

In the case of distributed operations (see clause 4.3.6.4), where a listing of all previously encountered **Context Sources** is supplied with the request, no registration shall match if the CSourceRegistration *contextSourceAlias* can be found within the listing of previously encountered **Context Sources**.

Note that distributed queries (see clause 4.3.6.7), can be supplied with an EntityMap (see clause 4.5.25) which lists all Entity ids successfully matched during a previous request. If the location of an EntityMap is passed into a subsequent request, the retrieved EntityMap shall be used in preference to the matching algorithm described above, provided that the EntityMap is valid and has not expired.

5.13 Storing, Managing and Serving @contexts

5.13.1 Introduction

Context Brokers optionally (see clause 4.3.5) offer the capability to store and serve *@contexts* to clients. The stored *@contexts* may be managed by clients directly, via the APIs specified in clause 5.13. Clients can store custom user *@contexts* at the Context Broker, effectively using the Context Broker as a *@context* server.

Moreover, in order to optimize performance, brokers may automatically store and use the stored copies of common *@contexts* as a local cache, downloading them just once, thus avoiding fetching them over and over again at each NGSI-LD request. In order for the broker to understand if a needed *@context* is already in the local storage or not, the broker uses the URL, where the *@context* is originally hosted, as an identifier for it in the local storage. Consequently, the broker has no ability to cache *@contexts* that arrive to it as **embedded** parts within the NGSI-LD documents, since they are not uniquely (and implicitly) identified by any URL; **Context Brokers** only cache *@contexts* that are referred to by means of explicit URLs (either in the HTTP Link header or as URLs in the payload body). Thus, the **recommended best-practice, in order to exploit caching, is that clients do not embed their user @contexts into their NGSI-LD documents**; instead clients should explicitly host their user *@contexts* at their premises, or use the broker's capability to host user *@contexts* on their behalf.

When an external *@context* is stored, either explicitly upon a client's request or implicitly downloaded for caching purposes, the Context Broker generates a unique local *@context* identifier. The original *@context*'s URL, if any, is stored alongside the generated local id. The local id is then used for subsequent managing operations on the specific *@context*, that are specified in clauses 5.13.2 to 5.13.5. Moreover, the broker tags the entry with the current timestamp, (createdAt) so that, subsequently, clients can check the timestamp before deciding whether to force a refresh of the stored copy of the *@context*. This is primarily intended as a means for clients to well-behave, thus avoiding triggering continuous refresh of a stored *@context* on the broker, for fear that it is not at the latest version.

Stored *@contexts* are flagged as one of three kinds: "Cached", "Hosted", "ImplicitlyCreated":

- **Cached:**
@contexts implicitly and automatically fetched by the broker from external URLs during normal NGSI-LD operations are flagged as "Cached". A locally unique identifier is generated for each *@context* not already in the internal storage. The downloaded content, its URL and the current time in UTC are stored alongside the locally unique identifier. Implementations shall periodically invalidate the "Cached" *@contexts*. Depending on the binding of the NGSI-LD API to a specific protocol, that specific protocol may provide explicit indications about expiration times of cached content. In such cases, implementations shall comply with the indications provided by the protocol. Implementations should assign a heuristic expiration time when an explicit time is not specified. **Entries flagged as "Cached" shall not be served by brokers on-demand, but only be used as a local cache to improve performance.**
- **Hosted:**
@contexts that are explicitly added by users are flagged as "Hosted". These entries shall be served by brokers on-demand.
- **ImplicitlyCreated:**
@contexts that are implicitly, but *ex-novo*, created by the broker as a result of a user request are flagged as "ImplicitlyCreated". For instance, when a client creates a subscription using an *@context* that is an array, and the broker has to notify with Content-Type "application/json", then the broker needs this *@context* array to be hosted at a URL. Hence the broker has to create a new *@context* that is an array, and it is going to be served from an own URL. These entries shall be served by brokers on-demand.

5.13.2 Add @context

5.13.2.1 Description

With this operation, a client can ask the broker to store the full content of a specific *@context*, by giving it to the broker.

5.13.2.2 Use case diagram

A client can add an *@context* to be stored within an NGSI-LD system as shown in Figure 5.13.2.2-1.




Figure 5.13.2.2-1: Add @context use case

5.13.2.3 Input data

A JSON object that has a top-level field named *@context*, i.e. a JSON object representing a JSON-LD "local context". As specified in the JSON-LD specification [2], all extra information located outside of the *@context* subtree in the referenced object shall be discarded.

5.13.2.4 Behaviour

A new entry is created in the local storage and a locally unique identifier (URI) is generated for it. The JSON object representing the client-supplied *@context* and the current UTC time are stored alongside the locally unique identifier. That identifier shall be given back as a result in the output data. The entry is flagged as being of kind "Hosted".

The behaviour described in clause 5.5.4 about JSON and JSON-LD validation shall be applied in case of invalid *@context*.

5.13.2.5 Output data

A locally unique URI identifying the *@context* in the broker's internal storage.

5.13.3 List @contexts

5.13.3.1 Description

With this operation a client can obtain a list of URLs that represent all of the *@contexts* stored in the local context store of the broker. Each URL can be used to download the corresponding *@context*, and, in case the *@context's* kind is "Cached", it shall be the original URL the broker downloaded the *@context* from.

In case a *details* flag is set to *true*, the client obtains a list of JSON objects, each representing information (metadata) about an *@context* currently stored by the broker. Each JSON object contains information about the *@context's* original URL (if any), its local identifier in the broker's storage, its kind ("Cached", "Hosted" and "ImplicitlyCreated"), its creation timestamp, its expiry date, and additional optional information.

5.13.3.2 Use case diagram

A client can list all *@contexts* stored within an NGSI-LD system as shown in Figure 5.13.3.2-1.




Figure 5.13.3.2-1: List @contexts use case

5.13.3.3 Input data

- A *kind* filter indicating the kind of stored *@contexts* that are to be included in the output list. Currently, possible kinds are "Cached", "Hosted" and "ImplicitlyCreated" (optional).
- A boolean *details* flag indicating that detailed JSON objects representing metadata about the stored *@contexts* instead of simple URLs are requested (optional).

5.13.3.4 Behaviour

The broker shall provide a URL or JSON object for each *@context* currently stored in the internal broker's storage, that match the filter. If no filter is specified, all kinds are included.

5.13.3.5 Output data

A list of URLs, or a list of resulting JSON objects containing the following fields:

- URL;
- localId;
- kind;
- createdAt;
- expiresAt [OPTIONAL];
- lastUsage [OPTIONAL];
- numberOfHits [OPTIONAL, number of times the *@context* was found in the storage];
- extraInfo [OPTIONAL, used by implementations to report any kind of custom information].

5.13.4 Serve @context

5.13.4.1 Description

With this operation a client can obtain the full content of a specific @context (only for @contexts of kind "Hosted" or "ImplicitlyCreated"), which is currently stored in the broker's internal storage, or its metadata (for all kinds of stored @contexts).

5.13.4.2 Use case diagram

A client can request the broker to serve a specific @context stored within the NGSI-LD system as shown in Figure 5.13.4.2-1.




Figure 5.13.4.2-1: Serve @context use case

5.13.4.3 Input data

- The locally unique identifier that identifies the desired @context in the broker's internal storage. Such unique identifiers are obtained by the client as a result of either a "Add @context" (clause 5.13.2) API operation or of a "List @contexts" (clause 5.13.3) API operation. For @contexts of kind "Cached" this can also be the original URL the broker downloaded the @context from.
- A boolean *details* flag indicating that a JSON object representing metadata about the @context, instead of the full content, is requested (optional).

5.13.4.4 Behaviour

- If *details* is set to *false*, or *details* is not present, the broker shall give back the full content of the @context that corresponds to the indicated local identifier, serving it from its internal storage, if the @context that corresponds to the indicated local identifier is of kind "Hosted" or "ImplicitlyGenerated". It shall give back **OperationNotSupported** error if it is of kind "Cached". It shall give back **ResourceNotFound** if the identifier is not found.
- Otherwise, if *details* is set to *true*, the broker shall give back metadata about the @context that corresponds to the indicated local identifier. It shall give back **ResourceNotFound** error if the identifier is not found.

5.13.4.5 Output data

The full content of the indicated *@context* (or its metadata as specified in clause 5.13.3.5), or *ResourceNotFound/OperationNotSupported* errors.

5.13.5 Delete and Reload *@context*

5.13.5.1 Description

With this operation, a client supplies a local identifier to the broker, indicating a stored *@context*, that the broker shall remove from its storage. For *@contexts* of kind "Cached" this can also be the original URL the broker downloaded the *@context* from. If the entry in the local storage that corresponds to the identifier is itself an array of *@contexts*, this operation will **not** delete the children, i.e. the *@contexts* in the array, but just the entry.

5.13.5.2 Use case diagram

A client can request the broker to delete (and optionally reload) a specific *@context* stored within the NGSI-LD system as shown in Figure 5.13.5.2-1.




Figure 5.13.5.2-1: Delete and Reload *@context* use case

5.13.5.3 Input data

- The locally unique identifier that identifies the desired *@context* in the broker's internal storage. For *@contexts* of kind "Cached" this can also be the original URL the broker downloaded the *@context* from.
- A *reload* boolean flag indicating that reloading of the *@context* shall be attempted (optional).

5.13.5.4 Behaviour

- If the *@context* identifier is not supplied, then an error of type *BadRequestData* shall be raised.
- If the *@context* identifier does not correspond to any existing entry in the *@context* storage, then an error of type *ResourceNotFound* shall be raised.

- If *reload* is *true* and the kind of the *@context* is "Cached", implementations shall try to re-download the identified *@context* from its original URL, before removing it from the internal storage. If downloading fails, or the downloaded *@context* is invalid according to JSON and JSON-LD validation of clause 5.5.4, then an error of type *LdContextNotAvailable* shall be raised. More detailed information about the errors shall be specified in the *ProblemDetails* (see IETF RFC 7807 [10]) field of the response. In case of any error, the operation ends without removing the existing *@context*. Otherwise, the existing *@context* is replaced with the newly downloaded one.
- If *reload* is *true* and the kind of the *@context* is **not** "Cached", implementations shall return a *BadRequestData* error.
- If *reload* is *false* (or *reload* is not supplied), implementations shall remove from the internal storage the *@context* that corresponds to the given identifier. The local identifier is used for finding the *@contexts* in the internal broker's storage. If the local identifier is not in the storage, a *ResourceNotFound* error shall be raised.

5.13.5.5 Output data

None.

5.14 Context Source Entity Mapping

5.14.1 Retrieve EntityMap

5.14.1.1 Description

With this operation a client can obtain a cached EntityMap which is currently stored in the broker's internal storage, or memory.

5.14.1.2 Use case diagram

A client can request the broker to retrieve a specific EntityMap within the NGSI-LD system as shown in Figure 5.14.1.2-1.




Figure 5.14.1.2-1: Retrieve EntityMap

5.14.1.3 Input data

EntityMap ID (URI) of the EntityMap to be retrieved (target EntityMap).

5.14.1.4 Behaviour

- If the Entity ID is not present or it is not a valid URI, then an error of type *BadRequestData* shall be raised.
- If the NGSI-LD endpoint does not know about a matching EntityMap for the EntityMap ID, then an error of type *ResourceNotFound* shall be raised.
- Otherwise, return a JSON-LD object representing the EntityMap as mandated by clause 5.2.39.

5.14.1.5 Output data

A JSON-LD object representing the target EntityMap as mandated by clause 5.2.39.

5.14.2 Update EntityMap

5.14.2.1 Description

This operation allows performing a **partial update on an NGSI-LD EntityMap**. A partial update only changes the elements provided in the EntityMap Fragment, leaving the rest as they are.

5.14.2.2 Use case diagram

A client can request the broker to update an EntityMap which is currently stored in the broker's internal storage as shown in Figure 5.14.2.2-1.




Figure 5.14.2.2-1: Update EntityMap

5.14.2.3 Input data

- EntityMap ID (URI) of the EntityMap to be retrieved (target EntityMap).
- A JSON-LD document representing an EntityMap.

5.14.2.4 Behaviour

- If the EntityMap ID is not present or it is not a valid URI, then an error of type *BadRequestData* shall be raised.
- If the NGSI-LD endpoint does not know about a matching EntityMap for the EntityMap ID, then an error of type *ResourceNotFound* shall be raised.

- Perform an update operation on the target EntityMap using the fields specified within then JSON-LD document. Any provided output-only fields shall be ignored.

5.14.2.5 Output data

None.

5.14.3 Delete EntityMap

5.14.3.1 Description

This operation allows deleting an NGSI-LD EntityMap.

5.14.3.2 Use case diagram

A client can request the broker to completely delete an EntityMap held within the NGSI-LD system as shown in Figure 5.14.3.2-1.




Figure 5.14.3.2-1: Delete EntityMap

5.14.3.3 Input data

EntityMap ID (URI) of the EntityMap to be retrieved (target EntityMap).

5.14.3.4 Behaviour

- If the EntityMap ID is not present or it is not a valid URI, then an error of type *BadRequestData* shall be raised.
- If the NGSI-LD endpoint does not know about a matching EntityMap for the EntityMap ID, then an error of type *ResourceNotFound* shall be raised.
- The EntityMap shall be removed from the broker's internal storage, or memory.

5.14.3.5 Output data

None.

5.14.4 Create EntityMap for Query Entities

5.14.4.1 Description

This operation is very similar to the Query Entities operation in clause 5.7.2, except that it does not directly return Entities, but creates and returns an Entity map including the identifiers of the Entities that are candidates to be part of the query results. The Entity map can then be used for paginating through the included Entities.

5.14.4.2 Use case diagram

A *Context Consumer* can retrieve an Entity map with the Entities that match a specific query from an NGSI-LD system as shown in Figure 5.14.4.2-1.




Figure 5.14.4.2-1: Query Entities for creating EntityMap use case

5.14.4.3 Input data

To simplify the operation, the same parameters as for the Query Entities operation (clause 5.7.2) are allowed, but some of these are irrelevant for creating an Entity map and thus shall be ignored.

- A reference to a JSON-LD *@context* (optional).
- A selector of Entity types as specified by clause 4.17 (optional). Both type names (short hand string) and fully qualified type names (URI) are allowed in the selector.
- A list (one or more) of Entity identifiers (optional).
- A list (one or more) of Attribute names of which at least one shall exist in order for an Entity to be selected, and also used as query projection attributes (optional).
- A restrictive list of Entity member names ("id", "type", "scope" or an Attribute name) to be retrieved (projection attributes as defined by clause 4.21) (optional).
- An exclusionary list member names ("id", "type", "scope" or an Attribute name) to be removed (projection attributes as defined by clause 4.21) (optional).
- An id pattern as a regular expression (optional).
- An NGSI-LD query (to filter out Entities by Attribute values) as per clause 4.9 (optional).
- An NGSI-LD geoquery (to filter out Entities by spatial relationships) as mandated by clause 4.10 (optional).

- In the case of GeoJSON representation:
 - The name of the *GeoProperty* attribute to use as the geometry for the GeoJSON representation as mandated by clause 4.5.16 (**ignored**).
 - A *datasetId* specifying which instance of the value is to be selected if the *GeoProperty* value has multiple instances as defined by clause 4.5.5 (**ignored**).
- A NGSI-LD Scope query (to filter out Entities based on their Scope) as mandated by clause 4.19 (optional).
- An NGSI-LD query (called context source filter, to filter out Context Sources by the values of properties that describe them) as per clause 4.9 (optional).
- A limit to the number of Entities to be retrieved. See clause 5.5.9 (**ignored**).
- A specified language filter as per clause 4.15 (optional).
- A list (one or more) of Attribute names whose values shall be expanded to URIs prior to executing a query (optional).
- An optional flag indicating whether to include additional Linked Entities corresponding to the Relationships retrieved and how to format those Linked Entities. See clause 4.5.23 (optional).
- A limit to the depth of Linked Entities to search whilst traversing an Entity Graph. See clause 4.5.23 (optional).
- A list (one or more) of Linked Entity identifiers previously encountered whilst traversing an Entity Graph. See clause 4.5.23 (optional).
- A flag indicating whether to return the location of the EntityMap used within the operation (**ignored**, EntityMap is always returned).
- A suggested expiry time for the EntityMap (optional).
- The location of a resource holding an EntityMap of matching Entity registrations (**ignored**).
- A *datasetId* parameter that specifies which Attribute instances are to be selected as defined by clause 4.5.5 (optional).
- A flag indicating whether split Entities are to be expected, i.e. Entities whose information is distributed across different Context Sources (optional).

In the general case, it is not possible to retrieve a set of entities by only specifying desired Entity identifiers, without further specifying restrictions on the Entities' types or attributes, either explicitly, via selector of Entity types or of Attribute names, or implicitly, within an NGSI-LD Query or GeoQuery. If the execution of the operation is limited to the local scope (see clause 5.5.13), no further restrictions have to be provided.

5.14.4.4 Behaviour

- At least one of the following input data shall be provided:
 - a) selector of Entity Types;
 - b) list of Attribute names, including at least one non-system Attribute;
 - c) NGSI-LD Query, including at least one non-system Attribute;
 - d) NGSI-LD GeoQuery;
 - e) local scope (see clause 5.5.13).

If none of the above is provided, then an error of type *BadRequestData* shall be raised (too wide query).

- If the list of Entity identifiers includes a URI which it is not valid, or the query, geoquery or context source filter are not syntactically valid (as per the referred clauses 4.9 and 4.10) an error of type *BadRequestData* shall be raised.
- If projection attributes are present and indicate the use of *Linked Entity* retrieval, and the use of *Linked Entity* retrieval is not specified, or the projected attribute depth exceeds the *Linked Entity* retrieval depth, then an error of type *BadRequestData* shall be raised.
- If the filter conditions specified by the query includes *Linked Entity* attributes, and the use of *Linked Entity* retrieval is not specified, or the *Linked Entity* attribute query depth exceeds the *Linked Entity* retrieval depth, an error of type *BadRequestData* shall be raised (too deep query).
- If geometryProperty parameter is present and the Accept Header is not set to "application/geo+json", then an error of type *BadRequestData* shall be raised.
- Otherwise,
 - Term to URI expansion of type and Attribute names shall be performed, as mandated by clause 5.5.7.
 - If a list of Attribute names whose values shall be expanded to URIs has been supplied, JSON-LD type coercion shall be performed as mandated by clause 5.5.7.
 - If split entities flag is explicitly set to true or, if not explicitly set, the default setting of the deployment allows split entities, and local scope is not specified, implementations shall run a query that shall return an EntityMap containing the identifiers of all the Entities found locally that meet all of the following conditions:
 - id is equal to any of the id(s) passed as a parameter;
 - the Entity Type names match the selector of Entity Types (expanded) that is passed as a parameter;
 - id matches the id pattern passed as a parameter.
 - Otherwise, implementations shall run a query that shall return an EntityMap containing the identifiers of all Entities found locally that meet all of the following conditions (given the respective parameter is provided):
 - id is equal to any of the id(s) passed as a parameter;
 - the Entity Type names match the selector of Entity Types (expanded) that is passed as a parameter;
 - Attribute matches any of the expanded attribute(s) in the list that is passed as a parameter;
 - id matches the id pattern passed as a parameter;
 - the filter conditions specified by the query are met (as mandated by clause 4.9);
 - the geospatial restrictions imposed by the geoquery are met (as mandated by clause 4.10); if there are multiple instances of the GeoProperty on which the geoquery is based, it is sufficient if any of these instances meets the geospatial restrictions;
 - if the Scope query is present, it shall match a present Entity Scope (as mandated by clause 4.19, for an example see annex C, clause C.5.15);
 - if the Attribute list is present, in order for an Entity to match, it shall contain at least one of the Attributes in the projection Attribute list.
- Unless local scope is specified (see clause 5.5.13), for Context Source Registrations that match the query and support the "createEntityMapQueryEntity" operation (see operations and operation groups in clause 4.20), implementations shall do the following:
 - If split entities flag is explicitly set to true or, if not explicitly set, the default setting of the deployment allows split entities, the filters (filter conditions specified by the query, geospatial restrictions imposed by the geoquery, Scope query, Attributes) shall be removed before forwarding the request.

- For each matching Context Source Registration, the request is forwarded for remote querying by matching endpoints. The result of each remote query is an EntityMap. The mapping between the Context Source Registration and the EntityMap Id is added to the *linkedMaps* element of the local EntityMap and for the Entity ids included in the returned EntityMaps a mapping to the Context Source Registration is added to the *entityMap* element of the local EntityMap.
- The local EntityMap is stored and made accessible based on its identifier.

5.14.4.5 Output data

The location of the local EntityMap shall be returned in a specific field in the response, as well as the EntityMap itself.

5.14.5 Create EntityMap for Query Temporal Evolution of Entities

5.14.5.1 Description

This operation is very similar to the Query Temporal Evolution of Entities operation in clause 5.7.4, except that it does not directly return the Temporal Evolution of Entities but creates and returns an Entity map including the identifiers of the Entities that are candidates to be part of the query results. The Entity map can then be used for paginating through Temporal Evolution of the included Entities.

5.14.5.2 Use case diagram

A *Context Consumer* can retrieve an Entity map with the Temporal Evolution of a set of Entities which matches a specific query from an NGSI-LD system as shown in Figure 5.14.5.2-1.




Figure 5.14.5.2-1: Query Temporal Evolution for creating EntityMap use case

5.14.5.3 Input data

To simplify the operation, the same parameters as for the Query Temporal Evolution of Entities operation (clause 5.7.4) are allowed, but some of these are irrelevant for creating an Entity map and thus will be ignored.

- An NGSI-LD temporal query as mandated by clause 4.11.
- A reference to a JSON-LD *@context* (optional).
- A selector of Entity types as specified by clause 4.17 (optional).
Both type name (short hand string) and fully qualified type name (URI) are allowed.

- A restrictive list of Entity member names ("id", "type", "scope" or an Attribute name) to be retrieved (projection attributes as defined by clause 4.21) (optional).
- An exclusionary list of Entity member names ("id", "type", "scope" or an Attribute name) to be removed (projection attributes as defined by clause 4.21) (optional).
- A list (one or more) of Entity identifiers (optional).
- A list (one or more) of Attribute names of which at least one shall exist in order for an Entity to be selected, and also used as query projection attributes (optional).
- An id pattern as a regular expression (optional).
- An NGSI-LD query (to filter out Entities by Attribute values) as per clause 4.9 (optional).
- An NGSI-LD geoquery (to filter out Entities by spatial relationships) as mandated by clause 4.10 (optional).
- A NGSI-LD Scope query (to filter out Entities based on their Scope) as mandated by clause 4.19 (optional).
- An NGSI-LD query (called context source filter, to filter out Context Sources by the values of properties that describe them) as per clause 4.9 (optional).
- A limit to the number of Entities to be retrieved. See clause 5.5.9 (**ignored**).
- A parameter (lastN) conveying that only the last N instances (per Attribute) within the concerned temporal interval shall be retrieved (optional).
- A specified language filter as per clause 4.15 (optional).
- A list (one or more) of Attribute names whose values shall be expanded to URIs prior to executing a query (optional).
- A flag indicating whether to return the location of the EntityMap used within the operation (**ignored**, EntityMap is always returned).
- A suggested expiry time for the EntityMap (optional).
- The location of a resource holding an EntityMap of matching Entity registrations (**ignored**).
- A datasetId parameter that specifies which Attribute instances are to be selected as defined by clause 4.5.5 (optional).
- A flag indicating whether split Entities are to be expected, i.e. Entities whose information is distributed across different Context Sources (optional).
- In the general case, it is not possible to retrieve a set of entities by only specifying desired Entity identifiers, without further specifying restrictions on the entities' types or attributes, either explicitly, via selector of Entity types or of Attribute names, or implicitly, within an NGSI-LD query or geoquery. If the execution of the operation is limited to the local scope (see clause 5.5.13), no further restrictions have to be provided.

5.14.5.4 Behaviour

- If a temporal query is not provided then an error of type *BadRequestData* shall be raised.
- At least one of the following input data shall be provided:
 - a) selector of Entity Types;
 - b) list of Attribute names, including at least one non-system Attribute;
 - c) NGSI-LD Query, including at least one non-system Attribute;
 - d) NGSI-LD GeoQuery;
 - e) local scope (see clause 5.5.13).

If none of the above is provided, then an error of type ***BadRequestData*** shall be raised (too wide query).

- If projection attributes or filter conditions indicate the use of **Linked Entity** retrieval, an error of type ***BadRequestData*** shall be raised.
- If the list of Entity identifiers includes a URI which it is not valid, or the query, geoquery or context source filter are not syntactically valid (as per the referred clauses 4.9 and 4.10) an error of type ***BadRequestData*** shall be raised.
- Term to URI expansion of type and Attribute names shall be observed mandated by clause 5.5.7.
- If a list of Attribute names whose values shall be expanded to URIs has been supplied, JSON-LD type coercion shall be performed as mandated by clause 5.5.7.
- The *lastN* parameter refers to a number, *n*, of Attribute instances which shall correspond to the last *n* timestamps (in descending ordering) of the temporal property (by default *observedAt*) within the concerned temporal interval.
- Otherwise,
 - Let S be the set of selected Entities i.e. the query result set.
 - If the split entities flag is explicitly set to true or, if not explicitly set, the default setting of the deployment allows split entities, and local scope is not specified, implementations shall run a query so that the result set contains all the Entities found locally, that meet **all** of the following conditions (given the respective parameter is provided):
 - If id(s) is provided, keep in S only those Entities whose id is equivalent to any of the id(s) passed as a parameter.
 - If an id pattern is provided, keep in S only those Entities whose id matches the id pattern.
 - If a selector of Entity Types is provided, keep in S only those Entities whose Entity Type names match the selector of Entity Types.
 - From S, select only those Entities any of whose Attribute instances (corresponding to the Attributes specified by the query or all if none are specified) match the temporal restrictions imposed by the temporal query (as mandated by clause 4.11); i.e. if the time series, for all the concerned Attributes of an Entity, does not include data corresponding to the temporal query interval, then such Entity shall be removed from S, thus it shall not appear in the final result set. Let S4 be this new subset.
- Otherwise implementations shall run a query that shall return the set of matching Entities (all Entities stored locally, in case only a local scope is specified); the logical steps to select the final result set of Entities, and the Attribute instances included as part of their temporal representation, are enumerated as follows:
 - If id(s) is provided, keep in S only those Entities whose id is equivalent to any of the id(s) passed as a parameter.
 - If an id pattern is provided, keep in S only those Entities whose id matches the id pattern.
 - If a selector of Entity Types is provided, keep in S only those Entities whose Entity Type names match the selector of Entity Types.
 - From S, select only those Entities any of whose Attribute instances (corresponding to the Attributes specified by the query or all if none are specified) match the temporal restrictions imposed by the temporal query (as mandated by clause 4.11); i.e. if the time series, for all the concerned Attributes of an Entity, does not include data corresponding to the temporal query interval, then such Entity shall be removed from S, thus it shall not appear in the final result set. Let S1 be this new subset.
 - If a values filter query is provided, from S1, select those Entities whose Attribute instances (during the interval defined by the temporal query) meet the matching conditions specified by the query (as mandated by clause 4.9); i.e. the values filter query shall be checked against all the Attribute instances resulting from the initial filtering performed by the temporal query. Let S2 be this new subset.
 - If no values filter query is provided, then S2 is equal to S1.

- If geoquery is present, from S2, select those Entities whose *GeoProperty* instances meet the geospatial restrictions imposed by the geoquery (as mandated by clause 4.10); those geospatial restrictions shall be checked against the *GeoProperty* instances that are within the interval defined by the temporal query. Let S3 be this new subset.
- If no geoquery is provided, then S3 is equal to S2.
- If the Scope query is present, from S3, select those Entities whose Entity Scope instances match the Scope query (as mandated by clause 4.19, for an example see annex C, clause C.5.16). Let S4 be the new subset.
- If no Scope query is provided, then S4 is equal to S3.
- The local EntityMap is created based on S4.
- Unless local scope is specified (see clause 5.5.13), for **Context Source Registrations** that match the query and support the "createEntityMapQueryTemporal" operation (see operations and operation groups in clause 4.20), implementations shall do the following:
 - If split entities flag is explicitly set to true or, if not explicitly set, the default setting of the deployment allows split entities, the filters (filter conditions specified by the query, geospatial restrictions imposed by the geoquery, Scope query, Attributes) shall be removed before forwarding the request.
 - For each matching **Context Source Registration**, the request is forwarded for remote querying by matching endpoints. The result of each remote query is an EntityMap. The mapping between the Context Source Registration and the EntityMap Id is added to the *linkedMaps* element of the local EntityMap and for the Entity ids included in the returned Entity Maps a mapping to the Context Source Registration is added to the *entityMap* element of the local EntityMap.
- The local EntityMap is stored and made accessible based on its identifier.

5.14.5.5 Output Data

The location of the local EntityMap shall be returned in a specific field in the response, as well as the EntityMap itself.

5.15 Context Source Identity Information

5.15.1 Retrieve Context Source Identity Information

5.15.1.1 Description

With this operation, a client can obtain **Context Source** identity information which uniquely defines the **Context Source** itself. In the multi-tenancy use case (see clause 4.14), a client can obtain identify information about a specific Tenant within a **Context Source**.

5.15.1.2 Use case diagram

A client can request the broker to retrieve identity information about a specific **Context Source** within the NGSI-LD system as shown in Figure 5.15.1.2-1.




Figure 5.15.1.2-1: Retrieve Context Source Identity Information

5.15.1.3 Input data

None.

5.15.1.4 Behaviour

- If the **Context Source** is unable to supply identity information about itself, then error of type *NotImplemented* shall be raised.
- Return a JSON-LD object representing the identity of the **Context Source** itself as mandated by clause 5.2.40. This can also include additional configurational data dependent on the specific **Context Source** implementation.

5.15.1.5 Output data

A JSON-LD object representing the identity of the **Context Source** as mandated by clause 5.2.40.

5.16 Snapshot Functionality

5.16.1 Create Snapshot

5.16.1.1 Description

This operation allows creating a new snapshot.

5.16.1.2 Use case diagram

A *Context Consumer* can create a new snapshot to have a more consistent basis for queries on the persisted Entity information as shown in Figure 5.16.1.2-1.




Figure 5.16.1.2-1: Create Snapshot use case

5.16.1.3 Input data

- A JSON-LD document conforming to the *Snapshot* data type as mandated by clause 5.2.41.

5.16.1.4 Behaviour

- If the data types, cardinalities and restrictions expressed by clause 5.2.41 are not met, then an error of type *BadRequestData* shall be raised.
- If the NGSI-LD endpoint already knows about this Snapshot, as there is an existing Snapshot whose id (URI) is equivalent, an error of type *AlreadyExists* shall be raised.
- If the Snapshot document does not include a Snapshot identifier, a new locally unique identifier (URI) shall be automatically generated by the implementation.
- The *createdAt* and *modifiedAt* timestamps are set to the current time of the NGSI-LD system.
- The *snapshotStatus* member is set to *preparation*.
- The *expiresAt* member shall be set, taking into account the *snapshotLifetime* requested, but applying the configured limit of the NGSI-LD system.
- If the Snapshot document does not include a *snapshotPriority*, the *snapshotPriority* shall be set to 5.
- The request returns, confirming the creation of the snapshot to the requesting *Context Consumer*, providing the *Snapshot identifier*. The status is accessible to the *Context Consumer* as described in clause 5.16.3. The following is executed in the background.
- Implementations shall execute the Queries specified in the *snapshotQueries* member, in each case following the behaviour described in clause 5.7.2.4. If the size of the respective results require pagination, all pages are to be retrieved completely. The response details for each query execution (aggregated across paginated ones) shall be stored in the *snapshotQueryDetails* list, at the same position as the query in *snapshotQueries*.
- Implementations shall execute the Temporal Queries specified in the *snapshotTemporalQueries* member, in each case following the behaviour described in clause 5.7.4.4. If the size of the respective results require pagination, all pages are to be retrieved completely. The response details for each query execution (aggregated across paginated ones) shall be stored in the *snapshotTemporalQueryDetails* list, at the same position as the query in *snapshotTemporalQueries*.

- The retrieved Entity information shall be stored by the NGSI-LD system, so that local NGSI-LD operations can then be executed on it.
- After all information has been stored, the *snapshotStatus* member is set to *success*, if all queries and temporal queries were executed successfully and yielded at least one result, to *partial*, if at least one query or temporal query was executed successfully and yielded a result, to *empty*, if at least one query or temporal query was executed successfully, but without yielding a result, and to *failure* otherwise.
- If the *endpoint* member is present, a Snapshot notification is sent as described in clause 5.16.6.

5.16.1.5 Output data

A URI identifying the newly created Snapshot.

5.16.2 Clone Snapshot

5.16.2.1 Description

This operation allows cloning a snapshot stored in an NGSI-LD system.

5.16.2.2 Use case diagram

A *Context Consumer* can clone an existing snapshot stored in an NGSI-LD system as shown in Figure 5.16.2.2-1.




Figure 5.16.2.2-1: Clone Snapshot use case

5.16.2.3 Input data

- A URI identifying the Snapshot to be cloned.
- A JSON-LD document conforming to the *Snapshot* data type as mandated by clause 5.2.41, but without *snapshotQueriesDetails* and *snapshotQueriesTemporalDetails* elements.

5.16.2.4 Behaviour

- If the data types, cardinalities and restrictions expressed by clause 5.2.41 are not met, then an error of type ***BadRequestData*** shall be raised.
- If the identifier of the Snapshot to be cloned is not found, an error of type ***ResourceNotFound*** shall be raised.

- If the Snapshot document includes a Snapshot identifier and there is an existing Snapshot whose id (URI) is equivalent, an error of type *AlreadyExists* shall be raised.
- If the Snapshot document does not include a Snapshot identifier, a new locally unique identifier (URI) shall be automatically generated by the implementation.
- The *createdAt* and *modifiedAt* timestamps of the clone are set to the current time of the NGSI-LD system.
- The *snapshotStatus* member of the clone is set to *preparation*.
- The *expiresAt* member shall be set, taking into account the *snapshotLifetime* requested, but applying the configured limit of the NGSI-LD system.
- The request returns, confirming the creation of the cloned Snapshot to the requesting *Context Consumer*. The status is accessible to the *Context Consumer* as described in clause 5.16.3. The following is executed in the background.
- All Entity and Temporal Evolution of Entity data that is part of the Snapshot to be cloned is copied to the clone Snapshot.
- After all information has been stored, the *snapshotStatus* member is set to *success*, if cloning was successful, and to *failure* otherwise.
- If the *endpoint* member is present, a Snapshot notification is sent as described in clause 5.16.6.

5.16.2.5 Output data

A URI identifying the newly created Snapshot.

5.16.3 Retrieve Snapshot Status

5.16.3.1 Description

This operation allows retrieving the status of a Snapshot stored in an NGSI-LD system.

5.16.3.2 Use case diagram

A *Context Consumer* can retrieve the status of a Snapshot from an NGSI-LD system as shown in Figure 5.16.3.2-1.




Figure 5.16.3.2-1: Retrieve Snapshot Status use case

5.16.3.3 Input data

Snapshot Id (URI) of the Snapshot whose status is to be retrieved.

5.16.3.4 Behaviour

- If the Snapshot Id is not present or it is not a valid URI, then an error of type ***BadRequestData*** shall be raised.
- If the identifier provided does not correspond to any existing Snapshot in the system, then an error of type ***ResourceNotFound*** shall be raised.
- Otherwise, implementations shall retrieve the Snapshot status and return it to the caller.

5.16.3.5 Output data

A JSON-LD object representing the Snapshot status as mandated by clause 5.2.41.

5.16.4 Update Snapshot Status

5.16.4.1 Description

This operation allows updating an existing Snapshot.

5.16.4.2 Use case diagram

A *Context Consumer* can update an existing Snapshot within an NGSI-LD system as shown in Figure 5.16.4.2-1.




Figure 5.16.4.2-1: Update Snapshot Status use case

5.16.4.3 Input data

- Snapshot identifier (URI), the target Snapshot.
- A JSON-LD document representing a Snapshot Fragment

5.16.4.4 Behaviour

- If the Snapshot id is not present or it is not a valid URI, then an error of type ***BadRequestData*** shall be raised.

- If the NGSI-LD System does not know about the target Snapshot, because there is no existing Snapshot whose id (URI) is equivalent, an error of type **ResourceNotFound** shall be raised.
- Execute the behaviour defined in clause 5.5.4 on JSON-LD validation.
- If the data types and restrictions expressed by clause 5.2.41 are not met by the *Snapshot Fragment* - in particular whether elements can be updated - then an error of type **BadRequestData** shall be raised.
- Term to URI expansion of Attribute names shall be observed as mandated by clause 5.5.7.
- Then, implementations shall modify the target Snapshot as mandated by clause 5.5.8.

5.16.4.5 Output data

A JSON-LD object representing the Snapshot status as mandated by clause 5.2.41.

5.16.5 Delete Snapshot

5.16.5.1 Description

This operation allows deleting an existing Snapshot.

5.16.5.2 Use case diagram

A *Context Consumer* can delete a Snapshot within an NGSI-LD system as shown in Figure 5.16.5.2-1.




Figure 5.16.5.2-1: Delete Snapshot use case

5.16.5.3 Input data

- A Snapshot identifier (URI), the target Snapshot.

5.16.5.4 Behaviour

- If the Snapshot Id is not present or it is not a valid URI, then an error of type **BadRequestData** shall be raised.
- If the Snapshot id provided does not correspond to any existing Snapshot in the system, then an error of type **ResourceNotFound** shall be raised.
- Otherwise implementations shall delete the Snapshot.

5.16.5.5 Output data

None.

5.16.6 Snapshot status notification behaviour

A Snapshot status notification allows the subscriber, typically the creator of the Snapshot, to be notified when the Snapshot is ready, or in case of any problems or updates. Implementations shall exhibit the following behaviour:

- SnapshotNotification (clause 5.3.4) messages can only be sent, if the *endpoint* member is set.
- SnapshotNotification messages are sent to the URI specified in the *endpoint* member of the Snapshot status.
- The information in the *receiverInfo* member of the Snapshot status shall be added to the SnapshotNotification message in the way required for the binding protocol.
- Snapshot Status Notifications shall be sent after all the specified Snapshot queries have been executed, the query results have been integrated and the Snapshot status has been updated accordingly.
- Snapshot status Notifications shall also be sent after any Snapshot status update, e.g. informing about the actual updated expiresAt timestamp.
- The SnapshotNotification shall be as mandated by clause 5.3.4.

5.16.7 Purge Snapshots

5.16.7.1 Description

This operation allows purging selected Snapshots.

5.16.7.2 Use case diagram

A *Context Consumer* can purge Snapshots within an NGSI-LD system as shown in Figure 5.16.7.2-1.




Figure 5.16.7.2-1: Purge Snapshots use case

5.16.7.3 Input data

- An NGSI-LD Query to select fitting Snapshots to be purged based on members of the Snapshot data type (see Table 5.2.41-1), as per clause 4.9.

5.16.7.4 Behaviour

- If the NGSI-LD Query is not present or it is not a valid as per clause 4.9, restricted to members of the Snapshot data type, then an error of type *BadRequestData* shall be raised.
- Implementations shall purge the Snapshots fitting the NGSI-LD Query.

5.16.7.5 Output data

None.

6 API HTTP Binding

6.1 Introduction

This clause defines the resources and operations of the NGSI-LD API. The NGSI-LD API is structured in terms of HTTP [3], [4] verbs, request and response payload bodies.

A non-normative OAS specification [i.12] of the referred HTTP binding can be found at [i.14].

6.2 Global Definitions and Resource Structure

All resource URIs of this API shall have the following root:

- {apiRoot}/{apiName}/{apiVersion}/

NOTE 1: The *apiRoot* discovery process is out of the scope of the present document.

NOTE 2: The *apiRoot* for **Context Source** related aspects and the *apiRoot* for general Entity-related aspects can be different, e.g. the **Context Source** related aspects can be implemented by a **Context Registry** as shown for the distributed and federated architectures (see clause 4.3), whereas the Entity-related aspects would be implemented by a **Context Broker**.

NOTE 3: The *apiRoot* for **Context Source** related aspects and the *apiRoot* for general Entity-related aspects can be different than the *apiRoot* for temporal aspects, e.g. the temporal aspects can be implemented by an NGSI-LD subsystem specialized in historical data.

The *apiRoot* includes the scheme ("http" or "https"), host and optional port, and an optional prefix string. The API shall support HTTP over TLS (also known as HTTPS - see IETF RFC 2818 [18]). TLS version 1.2 as defined by IETF RFC 5246 [19] shall be supported. HTTP without TLS is not recommended.

The *apiName* shall be set to "ngsi-ld" and the *apiVersion* shall be set to "v1" for the present document.

All resource URIs are defined relative to the above root URI. The structure of the resources under the root URI is shown in Figure 6.2-1 and methods defined on them are shown in Table 6.2-1.

An OpenAPI companion specification is available, see [37].




Figure 6.2-1: Resource URI structure of the NGSI-LD API

Table 6.2-1: Resources and HTTP methods defined on them

Resource Name	Resource URI	HTTP Method	Meaning	Clauses
Entity List	/entities/	POST	Entity creation	5.6.1; 6.4.3.1
		GET	Query entities	5.7.2; 6.4.3.2
		DELETE	Purge entities	5.6.21; 6.4.3.3
Entity by id	/entities/{entityId}	GET	Entity retrieval by id	5.7.1; 6.5.3.1
		DELETE	Entity deletion by id	5.6.6; 6.5.3.2
		PATCH	Entity merge by id	5.6.17; 6.5.3.4
		PUT	Entity replacement by id	5.6.18; 6.5.3.3
Attribute List	/entities/{entityId}/attrs/	POST	Append Attributes to Entity	5.6.3; 6.6.3.1
		PATCH	Update Attributes of an Entity	5.6.2; 6.6.3.2
Attribute by id	/entities/{entityId}/attrs/{attrId}	PATCH	Partial Attribute update	5.6.4; 6.7.3.1
		DELETE	Attribute deletion	5.6.5; 6.7.3.2
		PUT	Attribute replacement	5.6.19; 6.7.3.3
Subscriptions List	/subscriptions/	POST	Create Subscription	5.8.1; 6.10.3.1
		GET	Retrieve list of Subscriptions	5.8.4; 6.10.3.2
Subscription by Id	/subscriptions/{subscriptionId}	GET	Subscription retrieval by id	5.8.3; 6.11.3.1
		PATCH	Subscription update by id	5.8.2; 6.11.3.2
		DELETE	Subscription deletion by id	5.8.5; 6.11.3.3
Entity Types	/types/	GET	Retrieve available entity types	5.7.5 and 5.7.6; 6.25.3.1
Entity Type	/types/{type}	GET	Details about available entity type	5.7.7; 6.26.3.1
Attributes	/attributes/	GET	Available attributes	5.7.8 and 5.7.9; 6.27.3.1
Attribute	/attributes/{attrId}	GET	Details about available attribute	5.7.10; 6.28.3.1
Context source registration list	/csourceRegistrations/	POST	Csource registration creation	5.9.2; 6.8.3.1
		GET	Discover Csource registrations	5.10.2; 6.8.3.2
Context source registration by Id	/csourceRegistrations/{registrationId}	GET	Csource registration retrieval by id	5.10.1; 6.9.3.1
		PATCH	Csource registration update by id	5.9.3; 6.9.3.2
		DELETE	Csource registration deletion by id	5.9.4; 6.9.3.3
Context source registration subscription list	/csourceSubscriptions/	POST	Create subscription to Csource registration	5.11.2; 6.12.3.1
		GET	Retrieval of list of subscription to Csource registration	5.11.5; 6.12.3.2
Context source registration subscription by Id	/csourceSubscriptions/{subscriptionId}	GET	Csource registration subscription retrieval by id	5.11.4; 6.13.3.1
		PATCH	Csource registration subscription update by id	5.11.3; 6.13.3.2
		DELETE	Csource registration subscription deletion by id	5.11.6; 6.13.3.3
Entity Operations. Create	/entityOperations/create	POST	Batch Entity creation	5.6.7; 6.14.3.1
Entity Operations. Upsert	/entityOperations/upsert	POST	Batch Entity creation or update (<i>upsert</i>)	5.6.8; 6.15.3.1
Entity Operations. Update	/entityOperations/update	POST	Batch Entity update	5.6.9; 6.16.3.1
Entity Operations. Delete	/entityOperations/delete	POST	Batch Entity deletion	5.6.10; 6.17.3.1
Entity Operations. Query	/entityOperations/query	POST	Query entities based on POST	5.7.2; 6.23.3.1
Entity Operations. Merge	/entityOperations/merge	POST	Batch Entity merge	5.6.20; 6.31.3.1
Temporal Evolution of Entities	/temporal/entities/	POST	Temporal Evolution of an Entity creation	5.6.11; 6.18.3.1
		GET	Query Temporal Evolution of Entities	5.7.4; 6.18.3.2

Resource Name	Resource URI	HTTP Method	Meaning	Clauses
Temporal Evolution of an Entity by id	/temporal/entities/{entityId}	GET	Temporal Evolution of an Entity retrieval by id	5.7.3; 6.19.3.1
		DELETE	Temporal Evolution of an Entity deletion by id	5.6.16; 6.18.3.2
Temporal Representation of Attribute List	/temporal/entities/{entityId}/attrs/	POST	Temporal Evolution of Attribute of an Entity instance addition	5.6.12; 6.20.3.1
Temporal Representation of Attribute by id	/temporal/entities/{entityId}/attrs/{attrId}	DELETE	Attribute from Temporal Evolution of an Entity deletion	5.6.13; 6.21.3.1
Temporal Representation of Attribute Instance by id	/temporal/entities/{entityId}/attrs/{attrId}/{instanceId}	PATCH	Attribute Instance from Temporal Evolution of an Entity update by instance id	5.6.14; 6.22.3.1
		DELETE	Attribute Instance from Temporal Evolution of an Entity deletion by instance id	5.6.15; 6.22.3.2
Create EntityMap for Query Temporal Evolution of Entities	/temporal/entityMaps/	GET	Query temporal evolution of entities for creating entity map	5.15.4; 6.35.3.1
		POST	Query temporal evolution of entities for creating entity map based on POST	5.15.4; 6.35.3.2
Temporal Query Operation	/temporal/entityOperations/query	POST	Query Temporal Evolution of Entities based on POST	5.7.4; 6.24.3.1
Add and List @context	/jsonldContexts	POST	Add a user @context to the internal cache	5.13.2; 6.29.3.1
		GET	List all cached @contexts	5.13.3; 6.29.3.2
Serve, Delete and Reload @context	/jsonldContexts/{contextId}	GET	Serve one specific user @context	5.13.4; 6.30.3.1
		DELETE	Delete one specific @context from internal cache, possibly re-inserting a freshly downloaded copy of it	5.13.5; 6.30.3.2
Create EntityMap for Query Entities	/entityMaps/	GET	Query entities for creating entity map	5.14.4; 6.34.3.1
		POST	Query entities for creating entity map based on POST	5.14.4; 6.34.3.2
Retrieve, Update and Delete Entity Maps	/entityMaps/{entityMapId}	GET	EntityMap Retrieval by id	5.14.1; 6.32.3.1
		PATCH	EntityMap Update by id	5.14.2; 6.32.3.2
		DELETE	EntityMap Deletion by id	5.14.3; 6.32.3.3
Retrieve Context Source Identity Information	/info/sourceIdentity	GET	Context Source Identity Retrieval	5.15.1; 6.33.3.1
Create Snapshot or Purge Snapshots	/snapshots/	POST	Create Snapshot	5.16.1; 6.36.3.1
		DELETE	Purge Snapshots	5.16.7; 6.36.3.2
Retrieve and Update Snapshot Status or Delete Snapshot	/snapshots/{snapshotId}	GET	Retrieve Snapshot Status	5.16.3; 6.37.3.1
		PATCH	Update Snapshot Status	5.16.4; 6.37.3.2
		DELETE	Delete Snapshot	5.16.5; 6.37.3.3
Clone Snapshot	/snapshots/{snapshotId}/clone	POST	Clone Snapshot	5.16.2; 6.38.3.1

6.3 Common Behaviours

6.3.1 Introduction

This clause extends the API common behaviours to the particularities of the HTTP REST binding. For each operation implementations shall exhibit the common behaviours as specified by clause 5.5 and the behaviours defined by the present clause.

6.3.2 Error Types

This clause associates API error types (which shall be contained in the response payload body) defined by clause 5.5.2 with HTTP status codes as shown in Table 6.3.2-1.

Table 6.3.2-1: Mapping of error types to HTTP status codes

Error Type	HTTP status
https://uri.etsi.org/ngsi-ld/errors/AlreadyExists	409
https://uri.etsi.org/ngsi-ld/errors/BadRequestData	400
https://uri.etsi.org/ngsi-ld/errors/InternalServerError	500
https://uri.etsi.org/ngsi-ld/errors/InvalidRequest	400
https://uri.etsi.org/ngsi-ld/errors/LdContextNotAvailable	504
https://uri.etsi.org/ngsi-ld/errors/NoMultiTenantSupport	501
https://uri.etsi.org/ngsi-ld/errors/NonexistentTenant	404
https://uri.etsi.org/ngsi-ld/errors/OperationNotSupported	422
https://uri.etsi.org/ngsi-ld/errors/ResourceNotFound	404
https://uri.etsi.org/ngsi-ld/errors/TooComplexQuery	403
https://uri.etsi.org/ngsi-ld/errors/TooManyResults	403

In addition, implementations shall support the standard specific errors of HTTP bindings, such as the following:

- "Method Not Allowed" (405) which shall be raised when a client invokes a wrong HTTP verb over a resource. Implementations shall provide the allowed HTTP methods as mandated by IETF RFC 7231 [3] in section 6.5.5.
- "Request Entity too large" (413) which shall be raised when the HTTP input data stream provided by a client was too large i.e. too many bytes.
- "Length required" (411) which shall be raised when an HTTP request provided by a client does not define the "Content-Length" HTTP header.
- "Unsupported Media Type" (415) which shall be raised when an HTTP request payload body (as per the "Content-Type" header) it is not "application/json" nor "application/ld+json".
- "Not Acceptable" (406) which shall be raised when the response media types that are acceptable by a client (as per the "Accept" header) do not include or expand to "application/json" nor "application/ld+json".

6.3.3 Reporting errors

When an API operation results in an error, implementations shall return an HTTP response as follows:

- Content-Type: application/json.
- HTTP Status Code: As per clause 6.3.2 depending on error type.
- Payload body: A JSON object including all the terms defined by clause 5.5.3.

6.3.4 HTTP request preconditions

For HTTP POST, PATCH and PUT HTTP requests implementations shall check the following preconditions:

- Content-Type header shall be "application/json" or "application/ld+json".
- Content-Length header shall include the length of the request payload body.

For HTTP PATCH requests "application/merge-patch+json" is allowed as Content-Type, as mandated by IETF RFC 7396 [16]. Implementations shall interpret such MIME type as equivalent to "application/json".

For HTTP GET requests and for HTTP POST operations corresponding to "Query Entities" (see clause 5.7.2) and "Query Temporal Evolutions of Entities" (see clause 5.7.4), implementations shall check the following preconditions:

- Accept header shall include (or define a media range that can be expanded to) at least one of:
 - "application/json"
 - "application/ld+json"
 - "application/geo+json"

The order of the list above is significant. If the Accept header can be expanded to more than one of the options of the list, the first one of the list shall be selected, unless amended by the HTTP Accept header processing rules, e.g. the presence of a "q" parameter indicating a relative weight, (as mandated by IETF RFC 7231 [3], section 5.3.2) require otherwise.

If the Accept header is not present, "application/json" shall be assumed.

If an incoming HTTP request does not meet the preconditions stated above, an HTTP error response of type *InvalidRequest* shall be returned, with the following exceptions:

- "Content-Length" HTTP header absence, shall result in just a **411** HTTP status code (without any payload body).
- Unsupported Media Type, i.e. "Content-Type" header is not "application/json" nor "application/ld+json", shall result in just a **415** HTTP status code (without any payload body).
- Not Acceptable Media Type, i.e. "Accept" header does not imply "application/json" nor "application/ld+json", shall result in a **406** HTTP status code and the body of the message shall contain the list of the available representations of the resources.

Notwithstanding the above, if the Accept Header is set to "application/geo+json":

- For Context Information Consumption operations only, specifically "Retrieve Entity" (see clause 5.7.1) and "Query Entity" (clause 5.7.2) GeoJSON is considered as an acceptable content type and a GeoJSON payload will be returned.
- For all other operations, the request will result in a Not Acceptable Media Type error, returning a **406** HTTP status code and the body of the message shall contain the list of the available representations of the resources.

6.3.5 JSON-LD @context resolution

In the HTTP REST binding, implementations shall resolve the JSON-LD @context associated to an incoming HTTP request as follows:

- If the request verb is GET or DELETE, then the associated JSON-LD @context shall be obtained from a Link header [7] as mandated by JSON-LD [2], section 6.2. In the absence of such Link header, then the associated @context shall be the default JSON-LD @context.

EXAMPLE: The structure of the referred Link header is shown below. The first component (between < >) is a dereferenceable URI pointing to the JSON-LD document which contains the *@context* to be used to expand the terms used by the corresponding operation. The second parameter is a fixed, non-dereferenceable URI used to denote a unique identifier and semantics for this header (marking it as a link to a JSON-LD *@context*). The third and final parameter flags the MIME type of the linked resource (JSON-LD).

Link: <http://json-ld.org/contexts/person.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json".

- If the request verb is POST, PATCH or PUT and the Content-Type header is "application/json", then the *@context* shall be obtained from a Link Header as mandated by JSON-LD [2], section 6.2. In the absence of such Link Header, then the *@context* shall be the default *@context*. In any case, if the request payload body (as JSON) contains a *@context* term, then an HTTP error response of type *BadRequestData* shall be raised.
- If the request verb is POST, PATCH or PUT and the Content-Type header is "application/ld+json", then the associated *@context* shall be obtained from the request payload body itself. If no *@context* can be obtained from the request payload body, then an HTTP error response of type *BadRequestData* shall be raised. In any case, the presence of a JSON-LD Link header in the incoming HTTP request when the Content-Type header is "application/ld+json" shall result in an HTTP error response of type *BadRequestData*.

In summary, from a developer's perspective, for POST, PATCH and PUT operations, if MIME type is "application/ld+json", then the associated *@context* shall be provided only as part of the request payload body. Likewise, if MIME type is "application/json", then the associated *@context* shall be provided only by using the JSON-LD Link header. No mixes are allowed, i.e. mixing options shall result in HTTP response errors. Implementations should provide descriptive error messages when these situations arise.

In contrast, GET and DELETE operations always take their input *@context* from the JSON-LD Link Header.

6.3.6 HTTP response common requirements

Implementations shall honour the Accept header provided by HTTP requests as mandated by clause 6.3.4:

- If the target response's MIME type is "application/json" such response shall include a Link to the associated JSON-LD *@context* as mandated by [2], section 6.2.
 - If the Prefer header [26] is set to "ngsi-ld=<version>" then implementations shall set the Preference-Applied header to "ngsi-ld=<conformant-version>" in the returned response indicating which version of the NGSI-LD specification the payload body conforms to, as mandated in clause 4.3.6.8.
- If the target response's MIME type is "application/ld+json", then the response payload body provided by the HTTP response shall include a JSON-LD *@context*.
 - If the Prefer Header [26] is set to "ngsi-ld=<version>" then implementations shall set Preference-Applied header to "ngsi-ld=<conformant-version>" in the returned response indicating which version of the NGSI-LD specification the payload body conforms to, as mandated in clause 4.3.6.8.
- If the target response's MIME type is "application/geo+json" and the Prefer Header [26] is omitted or set to "body=ld+json", then the response payload body provided by the HTTP response shall include a JSON-LD *@context*, and the representation of the entities shall be in GeoJSON format in the response payload body.
- If the target response's MIME type is "application/geo+json" and the Prefer Header [26] is set to "body=json" such response shall include a Link to the associated JSON-LD *@context* as mandated by [2], section 6.2 and the representation of the entities shall be in GeoJSON format in the response payload body, and *@context* shall be omitted from the payload body.

Operations where the response payload body is not present such as successful HTTP POST, PATCH, PUT or DELETE operations and all error responses, do not include the Link header in the response.

Operations that result in an error that return a payload or that result in a partial success (207 Multi-Status) shall always respond with MIME type "`application/json`", regardless of the Accept header. It is assumed that if a client application understands any of the supported MIME types, the application shall understand "`application/json`" errors. Only Fully Qualified Names shall be used in the payload body of error or partial success responses, as there is no context present.

No Content-Length HTTP header shall be present if the response code is 204.

6.3.7 Representation of Entities

For HTTP GET and POST operations corresponding to "Retrieve Entity" (see clause 5.7.1) and "Query Entities" (see clause 5.7.2), **Context Broker** implementations shall support the parameter specified in Table 6.3.7-1, which specifies all possible supported representations formats.

In contrast, at a minimum, registered **Context Source** implementations shall support the normalized representation of Entities as default. When a registered **Context Source** is unable to support additional representations, a 501 Not Implemented Error shall be raised.

Table 6.3.7-1: Entity representations: format + options parameter

Name	Data Type	Cardinality	Remarks
format	String	0..1	When its value is "normalized", a normalized representation of Entities shall be provided as defined by clause 4.5.1, with Attributes returned in the normalized representation as defined in clauses 4.5.2.2, 4.5.3.2, 4.5.18.2 and 4.5.20.2. When its value is "concise", a concise lossless representation of Entities shall be provided as defined by clause 4.5.1. with Attributes returned in the concise representation as defined in clauses 4.5.2.3, 4.5.3.3, 4.5.18.3 and 4.5.20.3. In this case the Context Broker will return data in the most concise lossless representation possible, for example removing all Attribute <i>type</i> members. When its value is "simplified" (or its synonym "keyValues"). a simplified representation of Entities shall be provided as defined by clause 4.5.4. If the Accept Header is set to "application/geo+json" the response will be in simplified GeoJSON format as defined by clause 4.5.17.
options	Comma separated list of strings	0..1	An alternative mechanism to include the <i>format</i> parameter. Deprecated When its value includes the keyword "normalized", a normalized representation of Entities shall be provided as defined by clause 4.5.1, with Attributes returned in the normalized representation as defined in clauses 4.5.2.2, 4.5.3.2, 4.5.18.2 and 4.5.20.2. When its value includes the keyword "concise", a concise lossless representation of Entities shall be provided as defined by clause 4.5.1. with Attributes returned in the concise representation as defined in clauses 4.5.2.3, 4.5.3.3, 4.5.18.3 and 4.5.20.3. In this case the Context Broker will return data in the most concise lossless representation possible, for example removing all Attribute <i>type</i> members. When its value includes the keyword "simplified" (or its synonym "keyValues"). a simplified representation of Entities shall be provided as defined by clause 4.5.4. If the Accept Header is set to "application/geo+json" the response will be in simplified GeoJSON format as defined by clause 4.5.17.

6.3.8 Notification behaviour

In the HTTP binding a notification that is triggered by a subscription shall be sent by issuing an HTTP POST request targeted to the value of *notification.endpoint.uri* member of the subscription structure (defined by clauses 5.2.12, 5.2.14 and 5.2.15). For the HTTP binding, the protocol part of the endpoint URI is http or https. In case the optional MQTT notification binding (clause 7) is supported, the protocol part of the endpoint URI can also be *mqtt* or *mqtt.s*. The MIME type associated to the POST request shall be "application/json" by default. However, this can be changed to "application/ld+json", or "application/geo+json" by means of the *endpoint.accept* member.

If the target MIME type is "application/json" then the HTTP notification request shall include a Link header with a reference to the corresponding JSON-LD *@context* as mandated by the JSON-LD specification [2], section 6.2 (to the default JSON-LD *@context* if none available).

If the optional array (of *KeyValuePair* type, as defined by clause 5.2.22) *notification.endpoint.receiverInfo* of the subscription is present, then a new custom HTTP header for each member named "key" of the key, value pairs that make up the array shall be generated and included in the HTTP POST's list of headers. The content of each custom header shall be set equal to the content of the corresponding "value" member of the *KeyValuePair*. "Key" and "value" members shall adhere to IETF RFC 7230 [27] Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing definitions concerning HTTP headers.

If the target MIME type is "application/geo+json" and the *notification.endpoint.receiverInfo* member contains a key "Prefer" whose value is set to "body=json" then the HTTP notification request shall include a Link header with a reference to the corresponding JSON-LD @context as mandated by the JSON-LD specification [2], section 6.2 (to the default JSON-LD @context if none available).

If the target MIME type is "application/geo+json" and the *notification.endpoint.receiverInfo* contains a key "Prefer" whose value is set to "body=ld+json" or the "Prefer" key is omitted or *notification.endpoint.receiverInfo* does not exist, then the HTTP notification request includes an @context element in the payload body.

6.3.9 Csource Notification behaviour

In the HTTP binding a csource notification that is triggered by a csource subscription shall be sent by issuing an HTTP POST request targeted to the value of *notification.endpoint.uri* member of the csource subscription structure (defined by clauses 5.2.12 and 5.2.14). The MIME type associated to the POST request shall be "application/json" by default. However, this can be changed to "application/ld+json" by means of the *endpoint.accept* member.

If the target MIME type is "application/json" then the HTTP notification request shall include a Link header with a reference to the corresponding JSON-LD @context as mandated by the JSON-LD specification [2], section 6.2 (to the default JSON-LD @context if none available).

If the optional array (of *KeyValuePair* type, as defined by clause 5.2.22) *notification.endpoint.receiverInfo* of the subscription is present, then a new custom HTTP Header for each member named "key" of the key, value pairs that make up the array shall be generated and included in the HTTP POST's list of headers. The content of each custom header shall be set equal the content of the corresponding "value" member of the *KeyValuePair*. "Key" and "value" members shall adhere to IETF RFC 7230 [27] Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing definitions concerning HTTP headers.

6.3.10 Pagination behaviour

For HTTP operations corresponding to the operations listed in clause 4.12 (see Table 6.2-1 for a list of HTTP operations with their corresponding clauses), implementations shall support the HTTP query parameter specified in Table 6.3.10-1.

Table 6.3.10-1: Pagination: limit parameter

Name	Data Type	Cardinality	Remarks
limit	Integer	0..1	Only values greater or equal to 0. It defines the limit to the number of NGSI-LD Elements that shall be retrieved at a maximum as mandated by clause 5.5.9. The value 0 is only allowed in combination with the <i>count</i> URI parameter.

This clause defines the specific HTTP binding mechanisms that shall be used in conjunction with the behaviours defined by clause 5.5.9. Particularly, to flag the existence of related pages that could be retrieved when dealing with query operations involving pagination, NGSI-LD Systems implementing the HTTP binding shall use the HTTP Link header field as mandated by IETF RFC 8288 [7], clause 3, as follows:

- The pointers to the next and previous pages (when needed as mandated by clause 5.5.9) shall be serialized as link-value elements. The content of such link-value(s) shall be:
 - For the next page, the Link Target shall be a URI-reference that could be dereferenced by an NGSI-LD Client to retrieve the next page of NGSI-LD Elements. In addition, the Link Relation Type shall be equal to "next", registered under the IANA Registry of Link Relation Types [20].

- For the previous page, the Link Target shall be a URI-reference that could be dereferenced by an NGSI-LD Client to retrieve the previous page of NGSI-LD Elements. In addition, the Link Relation Type shall be equal to "prev", registered under the IANA Registry of Link Relation Types [20].
- At least, the "type" Link Target Attribute shall be included by the previously described serialized Link Header, as mandated by IETF RFC 8288 [7], clause 3.4, and its value shall be exactly equal to the media type resulting from the original request made by the NGSI-LD Client (the request that triggered the current pagination iteration).

EXAMPLE: If the media type requested originally was "application/json" then during the entire pagination iteration the value of the Link Target Attribute "type" shall be "application/json".

If the transparent pagination option as described in clause 5.5.9.2 is used, the URI-references for the "next" and "prev" link to the next and previous pages respectively should include the *Limit* parameter as specified in Table 6.3.10-1 and the *offset* parameter as specified in Table 6.3.10-2, giving the Context Consumer more control, i.e. to adapt the size of the page using *Limit* and jump to a desired set of elements in the results using *offset*.

Table 6.3.10-2: Pagination: offset parameter

Name	Data Type	Cardinality	Remarks
offset	Integer	0..1	Only values greater or equal to 0. It defines the offset of the first NGSI-LD Element that shall be retrieved from the beginning of the result set. If offset is set to a value larger than the result set, the offset should be assumed to be equal to the size of the result set, i.e. only the last element of the result set is to be returned if there are any results.

Temporal representation of resources adds an additional dimension to the pagination. Depending on the requested time range, the response will contain multiple instances of the requested Attribute, and therefore an additional pagination mechanism for those temporal representations is required, in order to limit the time range of the response. If no limits are specified, a default limit is enforced, depending on implementation specific configurations. For HTTP operations on **Temporal Evolution of Entities**, implementations shall use the Partial Content Response (206) as specified by IETF RFC 7233 [31], clause 4.1, if the implementation is not able to respond with the full representation at once. In this case, for requests where the parameter *lastN* is present, pagination shall happen "backwards" (from the most recent to the least recent timestamp in the requested time range). For requests without the parameter *lastN*, pagination shall happen "forwards" (from the least recent to the most recent timestamp in the requested time range).

This is achieved by including the "Content - Range" header field with the following contents:

- "unit" shall be equal to "DateTime";
- "range-start" and "range-end" shall be of type *DateTime*. They depend on the requested time relationship *timerel* (as defined by clause 4.11), as follows:
 - If the *lastN* parameter is present, pagination shall happen "backwards":
 - "range-start" shall be equal to "timeAt" for requests with *timerel=before*, "endTimeAt" for requests with *timerel=between*, or the most recent timestamp in the range of the response, for requests with *timerel=after*;
 - "range-end" shall be equal to the least recent timestamp in the range of the response;
 - "size" shall be equal to the requested *lastN*.
 - If the *lastN* parameter is **not** present, pagination shall happen "forwards":
 - "range-start" shall be equal to *timeAt* for requests with *timerel=after* or *timerel=between*, or the least recent timestamp in the range of the response, for requests with *timerel=before*;
 - "range-end" shall be equal to the most recent timestamp in the range of the response;

- "size" shall be equal to "*".

6.3.11 Including system Attributes

For HTTP GET operations performed over the resources /entities/, /subscriptions/, /csourceRegistrations/, /csourceSubscriptions/, /temporal/entities/ and all of its sub-resources, and for HTTP POST operations corresponding to "Query Entities" (see clause 5.7.2) and "Query Temporal Evolutions of Entities" (see clause 5.7.4), implementations shall support the parameter specified in Table 6.3.11-1. Implementations shall not raise an error if they do not hold system generated temporal attributes.

Table 6.3.11-1: Including system generated temporal attributes: options parameter

Name	Data Type	Cardinality	Remarks
options	Comma separated list of strings	0..1	When its value includes the keyword "sysAttr", a representation of NGSI-LD Elements shall be provided so that the system generated temporal attributes <i>createdAt</i> , <i>modifiedAt</i> and the system temporal attribute <i>expiresAt</i> are included in the response payload body where known. In the case of temporal representations, also the system generated temporal attribute <i>deletedAt</i> is included, if the NGSI-LD Element has been deleted.

6.3.12 Simplified or aggregated temporal representation of Entities

For HTTP GET and POST operations corresponding to "Retrieve Temporal Evolution of an Entity" (see clause 5.7.3) and "Query Temporal Evolution of Entities" (see clause 5.7.4), implementations shall support the parameter specified in Table 6.3.12-1.

Table 6.3.12-1: Temporal Entity representations: format + options parameters

Name	Data Type	Cardinality	Remarks
format	String	0..1	It shall be one of: "temporalValues", "aggregatedValues". When its value is "temporalValues", a simplified temporal representation of entities shall be provided as defined by clause 4.5.8. When its value is "aggregatedValues", an aggregated temporal representation of entities shall be provided as defined by clause 4.5.19.
options	Comma separated list of strings	0..1	An alternative mechanism to include the <i>format</i> parameter. Deprecated When its value includes the keyword "temporalValues", a simplified temporal representation of entities shall be provided as defined by clause 4.5.8. When its value includes the keyword "aggregatedValues", an aggregated temporal representation of entities shall be provided as defined by clause 4.5.19. Only one of the two keywords can be present in the values of the parameter. If both <i>format</i> and <i>options</i> are present, the value of the <i>format</i> parameter shall take precedence.

6.3.13 Counting number of results

This clause implements the behaviour described in clause 4.13, in case of HTTP binding.

For HTTP operations corresponding to the operations listed in clause 4.12 (see Table 6.2-1 for a list of HTTP operations with their corresponding clauses), implementations shall support the HTTP query parameter specified in Table 6.3.13-1.

Table 6.3.13-1: Counting number of results: count parameter

Name	Data Type	Cardinality	Remarks
count	Boolean	0..1	If <i>true</i> , then a special HTTP header (NGSILD-Results-Count) is set in the response. Regardless of how many entities are actually returned (maybe due to the <i>limit</i> URI parameter), the total number of matching results (e.g. number of Entities) is returned.

This clause defines the specific HTTP binding mechanisms that can be useful to plan the *limit* and *offset* URI parameters for pagination, thus allowing to convey an overview of the number of entities in a system.

To get only the count itself, and no entities, the URI parameter *limit* may have the value "0", and an empty array shall be returned as payload body.

Setting the URI parameter *limit* to zero without including the *count* URI parameter will result in a 400 *BadRequest* error.

6.3.14 Tenant specification

If the system implementing the NGSI-LD API supports multi-tenancy as described in clause 4.14 and clause 5.5.10, the **Tenant**, to which the NGSI-LD HTTP operation is targeted, is specified as the HTTP header "NGSILD-Tenant", whose value is the String identifying the **Tenant**. In case the target **Tenant** is the default **Tenant**, the HTTP header is omitted. If the HTTP header "NGSILD-Tenant" is present in the HTTP request, it shall also be present in HTTP response. This also applies to HTTP notifications sent as a result of subscriptions with an "NGSILD-Tenant" HTTP header (clause 6.3.8).

6.3.15 GeoJSON representation of spatially bound entities

For HTTP GET and POST operations corresponding to "Retrieve Entity" (see clause 5.7.1) and "Query Entities" (see clause 5.7.2), if the GeoJSON Accept header ("application/geo+json") is present, implementations shall render the entities of the response in the GeoJSON format, as described in clause 5.2.29.

For GeoJSON representations, a *GeoProperty* may be selected as the geolocation to be used as the geometry within the GeoJSON payload. If no *geometryProperty* parameter is specified, then the *location GeoProperty* of the Entity is used.

Table 6.3.15-1: Selecting a geometry

Name	Data Type	Cardinality	Remarks
geometryProperty	String	0..1	If not present, "location" is used.

6.3.16 Expiration time for cached @contexts

Implementations shall comply with the Expires header field (section 5.3 of IETF RFC 7234 [30]) or a max-age or s-maxage response directive of Cache-Control header field (section 5.2.2 of IETF RFC 7234 [30]) that may be present in the downloaded *@context*. This means that implementations shall periodically invalidate the "Cached" *@contexts* according to the headers mentioned above. Since origin servers do not always provide explicit expiration times, implementations should assign a heuristic (for instance according to IETF RFC 7234 [30] section 4.2.2) expiration time when an explicit time is not specified.

6.3.17 Distributed Operations Caching and Timeout Behaviour

The caching of response data received from forwarded HTTP GET requests is optionally supported by **Context Brokers**. For HTTP GET operations performed over the resources `/entities` and `/entities/{entity-id}`, where a **Context Source Registration** matches the request and a previous forwarded response has been cached and a subsequent request occurs before the **Context Source Registration** *cacheDuration* (as defined in Table 5.2.34-1) has been reached, the result may incorporate data cached from a previous response. To indicate that data from a cache has been included in the response, the HTTP header "NGSILD-Warning" shall be included. The value shall match the IANA Warning Code [32] 110 - Response is Stale.

"NGSILD-Warning" HTTP headers shall also be used to indicate instances of abnormal behaviour for distributed HTTP GET operations performed over the resources `/entities` and `/entities/{entity-id}`.

Table 6.3.17-1: NGSILD Warning Codes

IANA Warning Code	Remarks
110 - Response is Stale	No request was made to a specified registration endpoint - data from a cached value has been reused and it has been incorporated into the response.
111 - Revalidation Failed	Although data was received from the registration endpoint within the specified timeout period, the payload of the response was invalid. This could occur if the payload was corrupted or a non-NGSI payload was received.
199 - Miscellaneous Warning	No response was received from the registration endpoint within the specified timeout period or a registration loop has been detected.
299 - Miscellaneous Persistent Warning	An error response (such as 403 - Forbidden) was received from the registration endpoint within the specified timeout period. This could occur if the Context Broker has insufficient access rights to retrieve the data.

For distributed HTTP GET operations, registered context sources should always respond with a valid NGSILD payload. The **Context Broker** shall successfully parse this data and invalid non-NGSI-LD payloads shall be rejected and not incorporated into the overall response. It should be noted that a registration endpoint responding with no data and the HTTP status code **404 - Not Found** should not be considered as abnormal behaviour for distributed operations.

For all other operations, which correspond to HTTP Unsafe methods, the error response should be as informative as possible.

In the case of an **exclusive** or **redirect** registration, where all of the data is held outside of the **Context Broker** and held in a single registered source, the following errors shall be returned:

- 508 Loop Detected - if the single registered source and tenant is registered to redirect back on to the **Context Broker**.
- 504 Gateway Timeout - if the single registered source fails to respond in time.
- 404 Not Found - if resources not found within the single registered source.
- 502 Bad Gateway - if the single forwarded request fails for any other reason such as the **Context Broker** itself having insufficient access rights.

In the case of an **inclusive** or **redirect** registration, where an entity is distributed over multiple equally valid endpoints, but when updating the state of the distributed entity, an error response is returned from one or more registered sources:

- 207 Multi Status.

In the case of an **auxiliary** registration HTTP unsafe methods are not supported.

Whenever failures or timeouts occur, **Context Brokers** may optionally decline to make subsequent requests to the same registration endpoint until the cooldown period (as defined in Table 5.2.9-1) has been reached.

6.3.18 Limiting Distributed Operations

The parameter described in this clause limits the execution of an operation to a local *Context Source* or *Context Broker* (clause 5.5.13). It amends the matching *Context Source Registrations* behaviour as described in clause 5.12, in the case of the HTTP binding in order to avoid cascading distributed operations (see clause 4.3.6.4). For all operations the resources */entities/*, */types/*, */attributes/*, */entityOperations/*, */temporal/entities/* and */temporal/entityOperations/* (and their respective child resources) implementations shall support the HTTP query parameter specified in Table 6.3.18-1. The Registry API operations are always local and thus are not included here. Operations on a Snapshot (see clause 5.5.15) are always implicitly local scope, overriding setting the *local* parameter to *false*, i.e. such a setting is to be ignored by implementations.

Table 6.3.18-1: Limiting distributed operations: local parameter

Name	Data Type	Cardinality	Remarks
local	Boolean	0..1	If <i>local=true</i> then no <i>Context Source Registrations</i> shall be considered as matching to avoid cascading distributed operations (see clause 4.3.6.4).

Furthermore, to avoid infinite loops, for all operations, the resources */entities/*, */types/*, */attributes/*, */subscriptions/*, */csourceSubscriptions/*, */entityOperations/*, */temporal/entities/* and */temporal/entityOperations/* implementations shall fully support the HTTP Via Header (IETF RFC 7230 [27]) as specified in Table 6.3.18-2. Any *Context Broker* implementation passing a distributed operation request onward to another *Context Source* shall send an additional field value on the Via header field using its own unique *Context Source hostAlias* (see clause 5.2.40) as the pseudonym.

Table 6.3.18-2: Request Headers

Name	Data Type	Cardinality	Remarks
Via	String as defined in IETF RFC 7230 [27]	0..1	If present, the listing of previously encountered <i>Context Sources</i> supplied is used when determining matching registrations.

6.3.19 Extra information to provide when contacting Context Source

As specified in clauses 4.3.6.5 and 4.3.6.6, extra information to be provided when contacting a *Context Source* can be specified in the optional array (of *KeyValuePair* type, as defined by clause 5.2.22) *contextSourceInfo* of the *CSourceRegistration*.

In the HTTP binding, if the "jsonldContext" key is present, the URL value is placed in an HTTP Link Header as described by the JSON-LD specification [2], section 6.2 and, whenever a payload body is present in the request, the HTTP Content-Type Header is set to "application/json". For all other keys, a new custom HTTP header is added for each member named "key" of the key-value pairs that make up the array shall be generated and included in the HTTP list of headers. The content of each custom header shall be set equal to the content of the corresponding "value" member of the *KeyValuePair*, unless the special value "urn:ngsi-ld:request" has been set, in which case the value is to be taken from the triggering request, if present there. "Key" and "value" members shall adhere to IETF RFC 7230 [27] Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing definitions concerning HTTP headers.

Headers derived from other elements of the *CSourceRegistration*, e.g. "NGSILD-Tenant", take precedence and cannot be overridden using *contextSourceInfo*. The same applies for headers generally set by HTTP itself, e.g. Content-length.

6.3.20 Invalid parameters

If an HTTP request for an operation contains parameters that are incompatible with the operation, or it contains values of the *options* parameter that are not supported by the operation, an HTTP error response of type *InvalidRequest* should be returned.

6.3.21 Optional profile information regarding versioning and datatype conformance

In the HTTP Binding, if an HTTP Link header with a *profile* parameter in accordance with IETF RFC 6906 [33], is found set to <https://uri.etsi.org/ngsi-ld/profile/version> then implementations that are only partially capable of interpreting the datatypes conformant to the supplied NGSI-LD version may use this information to allow the payload to be accepted within the constraints of the current implementation (see clause 5.5.4) through amending the payload body and applying the fallbacks defined within clause 4.3.6.8.

6.3.22 Snapshot specification

If the system implementing the NGSI-LD API supports Snapshots as described in clause 4.3.7 and clause 5.5.15, the **Snapshot**, to which the NGSI-LD HTTP operation is targeted, is specified as the HTTP header "NGSILD-Snapshot", whose value is the identifier of the Snapshot. If the HTTP header "NGSILD-Snapshot" is present in the HTTP request, it shall also be present in HTTP response. This also applies to HTTP notifications sent as a result of subscriptions with an "NGSILD-Snapshot" HTTP header (clause 6.3.8).

6.4 Resource: entities/

6.4.1 Description

This resource represents the entities known to an NGSI-LD system.

6.4.2 Resource definition

Resource URI:

- /entities/

6.4.3 Resource methods

6.4.3.1 POST

This method is bound to the operation "Create Entity" and shall exhibit the behaviour defined by clause 5.6.1, taking the entity to be created from the HTTP request payload body. Figure 6.4.3.1-1 shows the Create Entity interaction and Table 6.4.3.1-1 describes the request body and possible responses.




Figure 6.4.3.1-1: Create Entity interaction

Table 6.4.3.1-1: Create Entity request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Entity	1	Payload body in the request contains a JSON-LD object which represents the entity that is to be created.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	201 Created	The HTTP response shall include a "Location" HTTP header that contains the relative path of the created entity.
	BatchOperationResult	1	207 Multi-Status	The HTTP response shall include a "Location" HTTP header that contains the relative path of the created entity. If the entity input data matches to a registration, the relevant parts of the request are forwarded as a distributed operation. In the case when an error response is received back from any distributed operation, a response body containing the result returned from each registration is returned in a <i>BatchOperationResult</i> structure. Errors can occur whenever a distributed operation is unsupported, fails or times out, see clause 6.3.17.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> member should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	409 Conflict	It is used to indicate that the entity or an exclusive or redirect registration defining the entity already exists, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	422 Unprocessable Entity	It is used to indicate that the operation is not available, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.4.3.2 GET

This method is associated to the operation "Query Entities" and shall exhibit the behaviour defined by clause 5.7.2, providing entities as part of the HTTP response payload body. In addition to this method, an alternative way to perform "Query Entities" operations via POST is defined in clause 6.23. Figure 6.4.3.2-1 shows the query entities interaction.




Figure 6.4.3.2-1: Query Entities interaction

The URL parameters that shall be supported by implementations are those defined in Table 6.4.3.2-1, the request headers that shall be supported by implementations are those defined in Table 6.4.3.2-2 and Table 6.4.3.2-3 describes the request body and possible responses.

Table 6.4.3.2-1: Query Entities URL parameters

Name	Data Type	Cardinality	Remarks
attrs	Comma separated list of strings	0..1 At least one among: <i>type</i> , <i>attrs</i> , <i>q</i> , or <i>georel</i> shall be present, unless the execution of the request is limited to local scope (see clause 5.5.13).	A synonym for a combination of the <i>pick</i> and <i>q</i> members. Deprecated Each String is an Attribute (Property or Relationship) name. List of Attributes to be matched by the Entities and also included in the response, i.e. only Entities that contain at least one of the Attributes in <i>attrs</i> are to be included in the response, and only the Attributes listed in <i>attrs</i> are to be included in each of the Entities of the response.
collation	String	0..1	An ICU collation (see IETF RFC 6067 [36]), When defined, the Entities returned in the payload shall be ordered according to the collation given. It is part of Entity ordering.
containedBy	Comma separated list of URIs	0..1	List of entity ids which have previously been encountered whilst retrieving the Entity Graph. Only applicable if <i>joinLevel</i> is present.
coordinates	String	0..1 It shall be one if <i>geometry</i> or <i>georel</i> are present.	Coordinates serialized as a string as per clause 4.10. It is part of geoquery.

Name	Data Type	Cardinality	Remarks
csf	String	0..1	Context Source filter as per clause 4.9.
datasetId	Comma separated list of strings	0..1	Shall be valid URIs, "@none" for including the default Attribute instances. Specifies the datasetIds of the Attribute instances to be selected for each matched Attribute as per clause 4.5.5.
entityMap	Boolean	0..1	If <i>true</i> , the location of the EntityMap used in the operation is returned in the response.
entityMapLifetime	String	0..1	Suggested duration, represented in ISO 8601 [17] format [17], for which the requester wants the EntityMap to exist. The actual expiresAt time of the EntityMap shall be set by the Context Broker or Context Source, possibly overriding the requested duration. Only applicable if <i>entityMap</i> is set to <i>true</i> .
expandValues	Comma separated list of strings	0..1	Each String is an Attribute (Property or Relationship) name. List of Attributes whose values shall be expanded into URIs according to the supplied <i>@context</i> prior to executing a query. It is part of query.
geometry	String	0..1 It shall be 1 if <i>georel</i> or <i>coordinates</i> are present. At least one among: <i>type</i> , <i>attrs</i> , <i>q</i> , or <i>geometry</i> shall be present, unless the execution of the request is limited to local scope (see clause 5.5.13).	Geometry as per clause 4.10. It is part of geoquery.
geometryProperty	String	0..1	It represents a Property name. In the case of GeoJSON Entity representation, this parameter indicates which GeoProperty to use for the top-level <i>geometry</i> field.
geoproperty	String	0..1 It shall be ignored unless a <i>geoquery</i> is present.	It represents the name of the <i>Property</i> that contains the geospatial data that will be used to resolve the geoquery. By default, will be <i>location</i> (see clause 4.7).
georel	String	0..1 It shall be 1 if <i>geometry</i> or <i>coordinates</i> are present.	Geo relationship as per clause 4.10. It is part of geoquery.
id	Comma separated list of strings	0..1	Each String shall be a valid URI. List of entity ids to be retrieved.
idPattern	Regular expression as defined by [11]	0..1	Regular expression that shall be matched by entity ids.
join	String	0..1	The type of Linked Entity retrieval to apply (see clause 4.5.23). Allowed values: "flat", "inline", "@none".
joinLevel	Positive Integer	0..1	Depth of Linked Entity retrieval to apply. Only applicable if <i>join</i> parameter is present.
jsonKeys	Comma separated list of strings	0..1	Each String is an Attribute (Property or Relationship) name. Values of the identified attributes are to be considered uninterpretable as JSON-LD and should not be expanded against the supplied <i>@context</i> using JSON-LD type coercion prior to executing the query.

Name	Data Type	Cardinality	Remarks
lang	String	0..1	It represents the preferred natural language of the response. It is used to reduce <i>languageMaps</i> to a string or string array property in a single preferred language.
omit	Comma separated list of strings	0..1	Each String is an Entity member ("id", "type", "scope" or a projected Attribute name). When defined, the listed Entity members are removed from each Entity within the payload.
orderBy	Comma separated list of strings	0..1	Each String is an Entity member ("id", "type", "scope" or an Attribute name) appended with an optional sorting style ("asc", "desc", "dist-asc", "dist-desc") as per clause 4.23. When defined, the Entities returned in the payload shall be ordered according to members defined. It is part of Entity ordering.
orderFrom	String	0..1 It shall be one if <i>orderBy</i> uses order by distance	Coordinates of a Geometry serialized as a string as per clause 4.10. It is part of Entity ordering.
orderGeometry	String	0..1	A Geometry type (with the exception of <i>GeometryCollection</i>) as defined by the GeoJSON specification (IETF RFC 7946 [8]). It is part of Entity ordering.
pick	Comma separated list of strings	0..1	Each String is an Entity member ("id", "type", "scope" or a projected Attribute name). When defined, every Entity within the payload body is reduced down to only contain the listed Entity members.
q	String	0..1 At least one among: <i>type</i> , <i>attrs</i> , <i>q</i> , or <i>georel</i> shall be present, unless the execution of the request is limited to local scope (see clause 5.5.13).	Query as per clause 4.9.
scopeQ	String	0..1	Scope query (see clause 4.19).
splitEntities	Boolean	0..1	If <i>true</i> it is assumed that single Entities are distributed between different Context Brokers and/or Context Sources and this has to be taken into account when applying any kind of filters (<i>q</i> , <i>geoQ</i> , <i>scopeQ</i> , Attributes etc.). If <i>false</i> it is expected that Context Broker and/or Context Source always have complete Entities, which allows applying filters locally. The parameter does not apply in case <i>local</i> is <i>true</i> . The default value should be decided by implementation; it should be configurable.
type	String	0..1 At least one among: <i>type</i> , <i>attrs</i> , <i>q</i> , or <i>georel</i> shall be present, unless the execution of the request is limited to local scope (see clause 5.5.13).	Selection of Entity Types as per clause 4.17. "*" is also allowed as a value and <i>local</i> is implicitly set to <i>true</i> and shall not be explicitly set to <i>false</i> .

Table 6.4.3.2-2: Query Entities request Headers

Name	Data Type	Cardinality	Remarks
NGSILD-EntityMap	URI	0..1	If present, the EntityMap supplied is used for determining the set of Entities requested during the query operation. The location of the EntityMap used in the query operation is returned in the response.

Table 6.4.3.2-3: Query Entities request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	Entity[]	1	200 OK 201 Created (in case an EntityMap has been (re)created)	A response body containing the query result as a list of entities, unless the Accept Header indicates that the Entities are to be rendered as GeoJSON. If an EntityMap has been requested, the HTTP response shall include an "NGSILD-EntityMap" HTTP header that contains the resource URI of the EntityMap resource used in the operation.
	Entity[]	1	203 Non-Authoritative Information	As above, but returning an altered response body , amended to conform to a specific version of the NGSI-LD specification as mandated in clause 4.3.6.8. The response shall also include a "Preference-Applied" HTTP header set to "ngsi-ld=<conformant-version>".
	GeoJSON FeatureCollection	1	200 OK 201 Created (in case an EntityMap has been (re)created)	If the Accept Header indicates that the Entities are to be rendered as GeoJSON, a response body containing the query result as GeoJSON FeatureCollection is returned. If an EntityMap has been requested, the HTTP response shall include an "NGSILD-EntityMap" HTTP header that contains the resource URI of the EntityMap resource used in the operation.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	501 Not Implemented	It is used by Registered Context Sources to indicate that the data format of the request is unsupported see clause 6.3.7.

6.4.3.3 DELETE

This method is associated to the operation "Purge Entities" and shall exhibit the behaviour defined by clause 5.6.21, providing entities as part of the HTTP response payload body. Figure 6.4.3.3-1 shows the query entities interaction.




Figure 6.4.3.3-1: Purge Entities interaction

The URL parameters that shall be supported by implementations are those defined in Table 6.4.3.3-1, the request headers that shall be supported by implementations are those defined in Table 6.4.3.3-2 and Table 6.4.3.3-3 describes the request body and possible responses.

Table 6.4.3.3-1: Purge Entities URL parameters

Name	Data Type	Cardinality	Remarks
coordinates	String	0..1 It shall be one if <i>geometry</i> or <i>georel</i> are present.	Coordinates serialized as a string as per clause 4.10. It is part of geoquery.
csf	String	0..1	Context Source filter as per clause 4.9.
drop	Comma separated list of strings	0..1	Each String is an Attribute name. When defined, every Entity within the payload body is reduced to not contain the listed Entity members.
geometry	String	0..1 It shall be 1 if <i>georel</i> or <i>coordinates</i> are present. At least one among: <i>type</i> , <i>attrs</i> , <i>q</i> , or <i>geometry</i> shall be present, unless the execution of the request is limited to local scope (see clause 5.5.13).	Geometry as per clause 4.10. It is part of geoquery.
geometryProperty	String	0..1	It represents a Property name. In the case of GeoJSON Entity representation, this parameter indicates which GeoProperty to use for the top-level <i>geometry</i> field.
geoproperty	String	0..1 It shall be ignored unless a <i>geoquery</i> is present.	It represents the name of the <i>Property</i> that contains the geospatial data that will be used to resolve the geoquery. By default, will be <i>location</i> (see clause 4.7).
georel	String	0..1 It shall be 1 if <i>geometry</i> or <i>coordinates</i> are present.	Geo relationship as per clause 4.10. It is part of geoquery.
id	Comma separated list of strings	0..1	Each String shall be a valid URI. List of entity ids to be retrieved.
idPattern	Regular expression as defined by [11]	0..1	Regular expression that shall be matched by entity ids.
keep	Comma separated list of strings	0..1	Each String is an Attribute name. When defined, every Entity within the payload body is reduced down to only contain the listed Entity members.

Name	Data Type	Cardinality	Remarks
q	String	0..1 At least one among: <i>type</i> , <i>attrs</i> , <i>q</i> , or <i>georeL</i> shall be present, unless the execution of the request is limited to local scope (see clause 5.5.13).	Query as per clause 4.9.
scopeQ	String	0..1	Scope query (see clause 4.19).
type	String	0..1 At least one among: <i>type</i> , <i>attrs</i> , <i>q</i> , or <i>georeL</i> shall be present, unless the execution of the request is limited to local scope (see clause 5.5.13).	Selection of Entity Types as per clause 4.17. "*" is also allowed as a value and <i>local</i> is implicitly set to <i>true</i> and shall not be explicitly set to <i>false</i> .

Table 6.4.3.3-2: Purge Entities request Headers

Name	Data Type	Cardinality	Remarks
NGSILD-EntityMap	URI	0..1	If present, the EntityMap supplied is used for determining the set of Entities requested during the purge operation. The location of the EntityMap used in the purge operation is returned in the response.

Table 6.4.3.3-3: Purge Entities request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
	BatchOperationResult	1	207 Multi-Status	If the entity input data matches to a registration, the relevant parts of the request are forwarded as a distributed operation. In the case when an error response is received back from any distributed operation, a response body containing the result returned from each registration is returned in a <i>BatchOperationResult</i> structure. Errors can occur whenever a distributed operation is unsupported, fails or times out, see clause 6.3.17.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.5 Resource: entities/{entityId}

6.5.1 Description

This resource represents an entity known to an NGSI-LD system.

6.5.2 Resource definition

Resource URI:

- /entities/{entityId}

Resource URI variables for this resource are defined in Table 6.5.2-1.

Table 6.5.2-1: URI variables

Name	Definition
entityId	Id (URI) of the entity to be retrieved

6.5.3 Resource methods

6.5.3.1 GET

This method is associated to the operation "Retrieve Entity" and shall exhibit the behaviour defined by clause 5.7.1. The Entity identifier is the value of the resource URI variable "entityId". Figure 6.5.3.1-1 shows the retrieve entity interaction.




Figure 6.5.3.1-1: Retrieve Entity interaction

The URL parameters that shall be supported are those defined in Table 6.5.3.1-1, the request headers that shall be supported by implementations are those defined in Table 6.5.3.1-2 and Table 6.5.3.1-3 describes the request body and possible responses.

Table 6.5.3.1-1: Retrieve Entity URL parameters

Name	Data Type	Cardinality	Remarks
attrs	Comma separated list of strings	0..1	A synonym for the <i>pick</i> parameter, except that <i>id</i> , <i>type</i> , <i>scope</i> are not allowed. Deprecated Each String is an Attribute (Property or Relationship) name. List of Attributes to be matched by the Entity and included in the response. If the Entity does not have any of the Attributes in <i>attrs</i> , then a <i>404 Not Found</i> shall be retrieved. If <i>attrs</i> is not specified, no matching is performed and all Attributes related to the Entity shall be retrieved.
containedBy	Comma separated list of URIs	0..1	List of entity ids which have previously been encountered whilst retrieving the Entity Graph. Only applicable if <i>joinLevel</i> is present.
datasetId	Comma separated list of strings	0..1	Shall be valid URIs, "@none" for including the default Attribute instances. Specifies the datasetIds of the Attribute instances to be selected for each matched Attribute as per clause 4.5.5.
entityMap	Boolean	0..1	If <i>true</i> , the location of the EntityMap used in the operation is returned in the response.
entityMapLifetime	String	0..1	Suggested duration, represented in ISO 8601 [17] format, for which the requester wants the EntityMap to exist. The actual expiresAt time of the EntityMap shall be set by the Context Broker or Context Source, possibly overriding the requested duration. Only applicable if <i>entityMap</i> is set to <i>true</i> .
geometryProperty	String	0..1	It represents a <i>GeoProperty</i> name. In the case of GeoJSON Entity representation, this parameter indicates which GeoProperty to use for the "geometry" element. By default, it shall be <i>location</i> .
join	String	0..1	The type of Linked Entity retrieval to apply (see clause 4.5.23). Allowed values: "flat", "inline", "@none".
joinLevel	Positive Integer	0..1	Depth of Linked Entity retrieval to apply. Default is 1 Only applicable if <i>join</i> parameter is: "flat" or "inline".
lang	String	0..1	It represents the preferred natural language of the response. It is used to reduce <i>languageMaps</i> to a string or string array property in a single preferred language.
omit	Comma separated list of strings	0..1	Each String is an Entity member ("id", "type", "scope" or a projected Attribute name). When defined, the listed Entity members are removed from the Entity (applies to all Entities in the payload in the case of Linked Entity retrieval).
pick	Comma separated list of strings	0..1	Each String is an Entity member ("id", "type", "scope" or a projected Attribute name). When defined, the Entity is reduced down to only contain the listed Entity members (applies to all Entities in the payload in the case of Linked Entity retrieval).
type	String	0..1	Selection of Entity Types as per clause 4.17.

Table 6.5.3.1-2: Retrieve Entity request Headers

Name	Data Type	Cardinality	Remarks
NGSILD-EntityMap	URI	0..1	If present, the EntityMap supplied is used for determining the extent of the Entity data requested during the retrieval operation. The location of the EntityMap used in the retrieval operation is returned in the response.

Table 6.5.3.1-3: Retrieve Entity request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	Entity	1	200 OK	A response body containing the JSON-LD representation of the target entity which consists only of the selected Attributes, unless the Accept Header indicates that the Entity is to be rendered as GeoJSON. If an EntityMap has been requested, the HTTP response shall include an "NGSILD-EntityMap" HTTP header that contains the resource URI of the EntityMap resource used in the operation.
	Entity	1	203 Non-Authoritative Information	As above, but returning an altered response body , amended to conform to a specific version of the NGSI-LD specification as mandated in clause 4.3.6.8. The response shall also include a "Preference-Applied" HTTP header set to "ngsi-ld=<conformant-version>".
	GeoJSON Feature	1	200 OK	If the Accept Header indicates that the Entity is to be rendered as GeoJSON, a GeoJSON <i>Feature</i> is returned. If an EntityMap has been requested, the HTTP response shall include an "NGSILD-EntityMap" HTTP header that contains the resource URI of the EntityMap resource used in the operation.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an Entity identifier (URI) not known to the system, see clause 6.3.2.
	ProblemDetails (see IETF RFC 7807 [10])	1	501 Not Implemented	It is used by Registered Context Sources to indicate that the data format of the request is unsupported see clause 6.3.7.

6.5.3.2 DELETE

This method is associated to the operation "Delete Entity" and shall exhibit the behaviour defined by clause 5.6.6. The Entity identifier is the value of the resource URI variable "entityId". Figure 6.5.3.2-1 shows the delete entity interaction.




Figure 6.5.3.2-1: Delete Entity interaction

The URL parameters that shall be supported are those defined in Table 6.5.3.2-1 and Table 6.5.3.2-2 describes the request body and possible responses.

Table 6.5.3.2-1: Delete Entity URL parameters

Name	Data Type	Cardinality	Remarks
type	String	0..1	Selection of Entity Types as per clause 4.17.

Table 6.5.3.2-2: Delete Entity request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
	BatchOperationResult	1	207 Multi-Status	If the entity input data matches to a registration, the relevant parts of the request are forwarded as a distributed operation. In the case when an error response is received back from any distributed operation, a response body containing the result returned from each registration is returned in a <i>BatchOperationResult</i> structure. Errors can occur whenever a distributed operation is unsupported, fails or times out, see clause 6.3.17.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an Entity identifier (URI) not known to the system, see clause 6.3.2.

6.5.3.3 PUT

This method is bound to the "Replace Entity" operation and shall exhibit the behaviour defined by clause 5.6.18. The Entity identifier is the value of the resource URI variable "entityId". The data to be updated shall be contained in the HTTP request payload body. Figure 6.5.3.3-1 shows the Replace Entity interaction.




Figure 6.5.3.3-1: Replace Entity interaction

The URL parameters that shall be supported are those defined in Table 6.5.3.3-1 and Table 6.5.3.3-2 describes the request body and possible responses.

Table 6.5.3.3-1: Replace Entity URL parameters

Name	Data Type	Cardinality	Remarks
type	String	0..1	Selection of Entity Types as per clause 4.17.

Table 6.5.3.3-2: Replace Entity request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Entity Fragment	1	Entity Fragment containing a complete representation of the Entity to be replaced.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No content	The entity was replaced successfully.
	BatchOperationResult	1	207 Multi-Status	If the entity input data matches to a registration, the relevant parts of the request are forwarded as a distributed operation. In the case when an error response is received back from any distributed operation, a response body containing the result returned from each registration is returned in a <i>BatchOperationResult</i> structure. Errors can occur whenever a distributed operation is unsupported, fails or times out, see clause 6.3.17.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an Entity identifier not known to the system, see clause 6.3.2.

6.5.3.4 PATCH

This method is bound to the "Merge Entity" operation and shall exhibit the behaviour defined by clause 5.6.17. The Entity identifier is the value of the resource URI variable "entityId". The data to be updated shall be contained in the HTTP request payload body. Figure 6.5.3.4-1 shows the Merge Entity interaction.




Figure 6.5.3.4-1: Merge Entity interaction

The URL parameters that shall be supported are those defined in Table 6.5.3.4-1 and Table 6.5.3.4-2 describes the request body and possible responses.

Table 6.5.3.4-1: Merge Entity URL parameters

Name	Data Type	Cardinality	Remarks
format	String	0..1	It shall be one of: "simplified" (or its synonym "keyVa lues"). Where present it indicates that a simplified representation of Entities has been provided as defined by clause 4.5.4. In this case, when a merge operation applies to an existing Attribute the <i>type</i> field of the Attribute shall remain unchanged (any attempt to modify the <i>type</i> of an Attribute shall result in a <i>BadRequest</i> error).
lang	String	0..1	It represents the natural language of data held in the request. When a merge operation applies to a pre-existing <i>LanguageProperty</i> and the value is supplied as a string or string array in the payload body, this query parameter shall be used to determine the key within the <i>languageMap</i> JSON Object to update.
observedAt	String	0..1	It shall be a <i>DateTime</i> (see clause 4.6.3). When a merge operation applies to a pre-existing Attribute which previously contained an <i>observedAt</i> sub-attribute, the value held in this query parameter shall be used if no specific <i>observedAt</i> sub-Attribute is found in the payload body.
options	Comma separated list of strings	0..1	An alternative mechanism to include the <i>format</i> parameter. Deprecated When its value includes the keyword "simplified" (or its synonym "keyVa lues"), this indicates that a simplified representation of Entities has been provided as defined by clause 4.5.4. If both <i>format</i> and <i>options</i> are present, the value of the <i>format</i> parameter shall take precedence.
type	String	0..1	Selection of Entity Types as per clause 4.17.

Table 6.5.3.4-2: Merge Entity request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Entity Fragment	1	Entity Fragment containing a complete representation of the Attributes to be merged.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No content	All the Attributes were merged successfully.
	BatchOperationResult	1	207 Multi-Status	If the entity input data matches to a registration, the relevant parts of the request are forwarded as a distributed operation. In the case when an error response is received back from any distributed operation, a response body containing the result returned from each registration is returned in a <i>BatchOperationResult</i> structure. Errors can occur whenever a distributed operation is unsupported, fails or times out, see clause 6.3.17.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an Entity identifier not known to the system, see clause 6.3.2.

6.6 Resource: entities/{entityId}/attrs/

6.6.1 Description

This resource represents all the Attributes (Properties or Relationships) of an NGSI-LD Entity.

6.6.2 Resource definition

Resource URI:

- /entities/{entityId}/attrs

Resource URI variables for this resource are defined in Table 6.6.2-1.

Table 6.6.2-1: URI variables

Name	Definition
entityId	Id (URI) of the concerned entity

6.6.3 Resource methods

6.6.3.1 POST

This method is bound to the "Append Attributes" operation and shall exhibit the behaviour defined by clause 5.6.3. The Entity identifier is the value of the resource URI variable "entityId". The data to be appended shall be contained in the HTTP request payload body. Figure 6.6.3.1-1 shows the Append Attributes interaction.




Figure 6.6.3.1-1: Append Attributes interaction

The URL parameters that shall be supported are those defined in Table 6.6.3.1-1 and Table 6.6.3.1-2 describes the request body and possible responses.

Table 6.6.3.1-1: Append Attributes URL parameters

Name	Data Type	Cardinality	Remarks
options	Comma separated list of strings	0..1	"noOverwrite" indicates that no attribute overwrite shall be performed.
type	String	0..1	Selection of Entity Types as per clause 4.17.

Table 6.6.3.1-2: Append Attributes request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Entity Fragment	1	Entity Fragment containing a complete representation of the Attributes to be added.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No content	All the Attributes were appended successfully.
	UpdateResult	1	207 Multi-Status	Only the Attributes included in the response payload body were successfully appended. If the entity input data matches to a registration, the relevant parts of the request are forwarded as a distributed operation. In the case when an error response is received back from any distributed operation, a response body containing the result returned from each registration is returned in a UpdateResult structure. Errors can occur whenever a distributed operation is unsupported, fails or times out, see clause 6.3.17.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an Entity identifier (URI) not known to the system, see clause 6.3.2.

6.6.3.2 PATCH

This method is bound to the "Update Attributes" operation and shall exhibit the behaviour defined by clause 5.6.2. The Entity identifier is the value of the resource URI variable "entityId". The data to be updated shall be contained in the HTTP request payload body. Figure 6.6.3.2-1 shows the Update Attributes interaction.




Figure 6.6.3.2-1: Update Attributes interaction

The URL parameters that shall be supported are those defined in Table 6.6.3.2-1 and Table 6.6.3.2-2 describes the request body and possible responses.

Table 6.6.3.2-1: Update Attributes URL parameters

Name	Data Type	Cardinality	Remarks
type	String	0..1	Selection of Entity Types as per clause 4.17.

Table 6.6.3.2-2: Update Attributes request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Entity Fragment	1	Entity Fragment containing a complete representation of the Attributes to be updated.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No content	All the Attributes were updated successfully.
	UpdateResult	1	207 Multi-Status	Only the Attributes included in the response payload body were successfully updated. If no Attributes were successfully updated the <i>updated</i> array of <i>UpdateResult</i> (see clause 5.2.18) will be empty. If the entity input data matches to a registration, the relevant parts of the request are forwarded as a distributed operation. In the case when an error response is received back from any distributed operation, a response body containing the result returned from each registration is returned in a <i>BatchOperationResult</i> structure. Errors can occur whenever a distributed operation is unsupported, fails or times out, see clause 6.3.17.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an Entity identifier not known to the system, see clause 6.3.2.

6.7 Resource: entities/{entityId}/attrs/{attrId}

6.7.1 Description

This resource represents an attribute (Property or Relationship) of an NGSI-LD Entity.

6.7.2 Resource definition

Resource URI:

- /entities/{entityId}/attrs/{attrId}

Resource URI variables for this resource are defined in Table 6.7.2-1.


Table 6.7.2-1: URI variables

Name	Definition
entityId	Id (URI) of the concerned entity
attrId	Attribute name (Property or Relationship)

6.7.3 Resource methods

6.7.3.1 PATCH

This method is bound to the "Partial Attribute Update" operation and shall exhibit the behaviour defined by clause 5.6.4. The Entity identifier is the value of the resource URI variable "entityId". The attribute name is the value of the resource URI variable "attrId". The Entity Fragment shall be contained in the HTTP request payload body. Figure 6.7.3.1-1 shows the Partial Attribute Update interaction.

**Figure 6.7.3.1-1: Partial Attribute Update interaction**

The URL parameters that shall be supported are those defined in Table 6.7.3.1-1 and Table 6.7.3.2-2 describes the request body and possible responses.

Table 6.7.3.1-1: Partial Attribute Update URL parameters

Name	Data Type	Cardinality	Remarks
type	String	0..1	Selection of Entity Types as per clause 4.17.

Table 6.7.3.1-2: Partial Attribute Update request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Entity Fragment	1	Entity Fragment containing the elements of the attribute to be updated.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	The attribute was updated successfully.
	UpdateResult	1	207 Multi-Status	If the entity input data matches to a registration, the relevant parts of the request are forwarded as a distributed operation. In the case when an error response is received back from any distributed operation, a response body containing the result returned from each registration is returned in a <i>UpdateResult</i> structure. Errors can occur whenever a distributed operation is unsupported, fails or times out, see clause 6.3.17.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an Entity identifier or attribute name not known to the system, see clause 6.3.2.

6.7.3.2 DELETE

This method is associated to the operation "Delete Attribute" and shall exhibit the behaviour defined by clause 5.6.5. The Entity identifier is the value of the resource URI variable "entityId". The attribute name is the value of the resource URI variable "attrId". Figure 6.7.3.2-1 shows the Delete Attribute interaction, Table 6.7.3.2-1 shows the delete parameters to be supported and Table 6.7.3.2-2 describes the request body and possible responses.




Figure 6.7.3.2-1: Delete Attribute interaction

Table 6.7.3.2-1: Delete Attribute URL parameters

Name	Data Type	Cardinality	Remarks
datasetId	String	0..1	Shall be a valid URI. Specifies the <i>datasetId</i> of the dataset to be deleted.
deleteAll	Boolean	0..1	If <i>true</i> , all attribute instances are deleted. Otherwise (default) only the Attribute instance specified by the <i>datasetId</i> is deleted. In case neither the deleteAll flag nor a <i>datasetId</i> is present, the default Attribute instance is deleted.
type	String	0..1	Selection of Entity Types as per clause 4.17.

Table 6.7.3.2-2: Delete Attribute request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
Response Body	UpdateResult	1	207 Multi-Status	If the entity input data matches to a registration, the relevant parts of the request are forwarded as a distributed operation. In the case when an error response is received back from any distributed operation, a response body containing the result returned from each registration is returned in a <i>UpdateResult</i> structure. Errors can occur whenever a distributed operation is unsupported, fails or times out, see clause 6.3.17.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an Entity identifier (URI) or attribute name not known to the system. see clause 6.3.2.

6.7.3.3 PUT

This method is bound to the "Replace Attribute" operation and shall exhibit the behaviour defined by clause 5.6.19. The Entity identifier is the value of the resource URI variable "entityId". The attribute name is the value of the resource URI variable "attrId". The Attribute Fragment shall be contained in the HTTP request payload body. Figure 6.7.3.3-1 shows the Replace Attribute interaction.




Figure 6.7.3.3-1: Replace Attribute interaction

The URL parameters that shall be supported are those defined in Table 6.7.3.3-1 and Table 6.7.3.3-2 describes the request body and possible responses.

Table 6.7.3.3-1: Replace Attribute URL parameters

Name	Data Type	Cardinality	Remarks
type	String	0..1	Selection of Entity Types as per clause 4.17.

Table 6.7.3.3-2: Replace Attribute request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Attribute Fragment	1	Attribute Fragment replacing the previous data.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	The attribute was replaced successfully.
	UpdateResult	1	207 Multi-Status	If the entity input data matches to a registration, the relevant parts of the request are forwarded as a distributed operation. In the case when an error response is received back from any distributed operation, a response body containing the result returned from each registration is returned in a <i>UpdateResult</i> structure. Errors can occur whenever a distributed operation is unsupported, fails or times out, see clause 6.3.17.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an Entity identifier or attribute name not known to the system, see clause 6.3.2.

6.8 Resource: csourceRegistrations/

6.8.1 Description

This resource represents the context source registrations known to an NGSI-LD system.

6.8.2 Resource definition

Resource URI:

- /csourceRegistrations/

6.8.3 Resource methods

6.8.3.1 POST

This method is bound to the operation "Register Context Source" and shall exhibit the behaviour defined by clause 5.9.2, taking the context source registration to be created from the HTTP request payload body. Figure 6.8.3.1-1 shows the Register Context Source interaction and Table 6.8.3.1-1 describes the request body and possible responses.





Figure 6.8.3.1-1: Register Context Source interaction

Table 6.8.3.1-1: Register Context Source request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	CSourceRegistration	1	Payload body in the request contains a JSON-LD object which represents the context source registration that is to be created.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	201 Created	The HTTP response shall include a "Location" HTTP header that contains the relative path of the created context source registration.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	409 Conflict	It is used to indicate that the context source registration already exists, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
Response Body	ProblemDetails (see IETF RFC 7807 [10])	1	422 Unprocessable Context Source Registration	It is used to indicate that the operation is not available see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.8.3.2 GET

This method is associated to the operation "Query Context Source Registrations" and shall exhibit the behaviour defined by clause 5.10.2, i.e. the parameters in the request describe entity related information, but instead of directly providing this entity information, the context source registration data, which describes context sources that can possibly provide the information, are returned as part of the HTTP response payload body. Figure 6.8.3.2-1 shows the Query Context Source Registrations interaction.

**Figure 6.8.3.2-1: Query Context Source Registrations interaction**

The URL parameters that shall be supported by implementations are those defined in Table 6.8.3.2-1 and Table 6.8.3.2-2 describes the request body and possible responses.

Table 6.8.3.2-1: Query Context Source Registrations URL parameters

Name	Data Type	Cardinality	Remarks
attrs	Comma separated list strings	0..1 At least one among: <i>type</i> , <i>attrs</i> , <i>q</i> , or <i>georel</i> shall be present.	A synonym for a combination of the <i>pick</i> and <i>q</i> members. Deprecated Each String is an Attribute (Property or Relationship) name. List of Attributes (Properties or Relationships) to be retrieved.
coordinates	String	0..1 It shall be one if <i>geometry</i> or <i>georel</i> are present.	Coordinates serialized as a string as per clause 4.10. It is part of geoquery.
csf	String	0..1	Context Source filter as per clause 4.9.
endTimeAt	String	0..1	It represents the <i>endTimeAt</i> parameter as defined by clause 4.1. It shall be a <i>DateTime</i> . Cardinality shall be 1 if <i>timerel</i> is equal to "between".
geometry	String	0..1 It shall be 1 if <i>georel</i> or <i>coordinates</i> are present. At least one among: <i>type</i> , <i>attrs</i> , <i>q</i> , or <i>georel</i> shall be present.	Geometry as per clause 4.10. It is part of geoquery.
geometryProperty	String	0..1	It represents a Property name. In the case of GeoJSON Entity representation, this parameter indicates which <i>GeoProperty</i> to use for the top-level <i>geometry</i> field.
geoproperty	String	0..1 It shall be ignored if no geoquery is present.	It represents the name of the Property that contains the geospatial data that will be used to resolve the geoquery.
georel	String	0..1 It shall be 1 if <i>geometry</i> or <i>coordinates</i> are present.	Geo relationship as per clause 4.10. It is part of geoquery.
id	Comma separated list of strings	0..1	Each String shall be a valid URI. List of entity ids to be retrieved.
idPattern	Regular expression as defined by [11]	0..1	Regular expression that shall be matched by entity ids satisfying the query
lang	String	0..1	It represents the preferred natural language of the response. It is used to reduce <i>languageMaps</i> to a string or string array property in a single preferred language.
omit	Comma separated list of strings	0..1	Each String is an Entity member ("id", "type", "scope" or a projected Attribute name).
pick	Comma separated list of strings	0..1	Each String is an Entity member ("id", "type", "scope" or a projected Attribute name).
q	String	0..1 At least one among: <i>type</i> , <i>attrs</i> , <i>q</i> , or <i>georel</i> shall be present.	Query as per clause 4.9.
scopeQ	String	0..1	Scope query (see clause 4.19).
timeAt	String	0..1	It represents the <i>timeAt</i> parameter as defined by clause 4.1. It shall be a <i>DateTime</i> . Cardinality shall be 1 if <i>timerel</i> is present.

Name	Data Type	Cardinality	Remarks
timeproperty	String	0..1 It shall be ignored if no temporal query is present.	It represents a Temporal Property name. Allowed values: "observedAt", "createdAt", "modifiedAt" and "deletedAt". If not specified, the default is "observedAt". (See clause 4.8).
timerel	String	0..1	It represents the temporal relationship as defined by clause 4.1. Allowed values: "before", "after", "between".
type	String	0..1 At least one among: <i>type</i> , <i>attrs</i> , <i>q</i> , or <i>georel</i> shall be present.	Selection of Entity Types as per clause 4.17.

Table 6.8.3.2-2: Retrieve Context Source Registrations request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	CSourceRegistration[]	1	200 OK	A response body containing the query result as an array of context source registrations.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.9 Resource: csourceRegistrations/{registrationId}

6.9.1 Description

This resource represents the context source registration, identified by *registrationId*, known to an NGSI-LD system.

6.9.2 Resource definition

Resource URI:

- /csourceRegistrations/{registrationId}

Resource URI variables for this resource are defined in Table 6.9.2-1.

Table 6.9.2-1: URI variables

Name	Definition
registrationId	Id (URI) of the context source registration

6.9.3 Resource methods

6.9.3.1 GET

This method is associated with the operation "Retrieve Context Source Registration" and shall exhibit the behaviour defined by clause 5.10.1. The registration identifier is the value of the resource URI variable "registrationId". Figure 6.9.3.1-1 shows the Retrieve Context Source Registration interaction and Table 6.9.3.1-1 describes the request body and possible responses.




Figure 6.9.3.1-1: Retrieve Context Source Registration interaction

Table 6.9.3.1-1: Retrieve Context Source Registration request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	CSourceRegistration	1	200 OK	A response body containing the JSON-LD representation of the target context source registration.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided a context source registration identifier (URI) not known to the system, see clause 6.3.2.

6.9.3.2 PATCH

This method is bound to the "Update Context Source Registration" operation and shall exhibit the behaviour defined by clause 5.9.3. The context source registration identifier is the value of the resource URI variable "registrationId". The context source registration to be updated shall be contained in the HTTP request payload body. Figure 6.9.3.2-1 shows the Update Context Source Registration interaction and Table 6.9.3.2-1 describes the request body and possible responses.




Figure 6.9.3.2-1: Update Context Source Registration interaction

Table 6.9.3.2-1: Update Context Source Registration request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	CSourceRegistration Fragment	1	Payload body in the request contains a JSON-LD object which represents the context source registration that is to be updated.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	The context source registration was successfully updated.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided a context source registration identifier not known to the system, see clause 6.3.2.

6.9.3.3 DELETE

This method is associated to the operation "Delete Context Source Registration" and shall exhibit the behaviour defined by clause 5.9.4. The context source registration identifier is the value of the resource URI variable "registrationId". Figure 6.9.3.3-1 shows the Delete Context Source Registration interaction and Table 6.9.3.3-1 describes the request body and possible responses.




Figure 6.9.3.3-1: Delete Context Source Registration interaction

Table 6.9.3.3-1: Delete Context Source Registration request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided a context source registration identifier (URI) not known to the system, see clause 6.3.2.

6.10 Resource: subscriptions/

6.10.1 Description

This resource represents the subscriptions known to an NGSI-LD system.

6.10.2 Resource definition

Resource URI:

- /subscriptions/

6.10.3 Resource methods

6.10.3.1 POST

This method is bound to the operation "Create Subscription" and shall exhibit the behaviour defined by clause 5.8.1, taking the subscription to be created from the HTTP request payload body. Figure 6.10.3.1-1 shows the Create Subscription interaction and Table 6.10.3.1-1 describes the request body and possible responses.





Figure 6.10.3.1-1: Create Subscription interaction

Table 6.10.3.1-1: Create Subscription request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Subscription	1	Payload body in the request contains a JSON-LD object which represents the subscription that is to be created.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	201 Created	The HTTP response shall include a "Location" HTTP header that contains the relative path of the created subscription.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	409 Conflict	It is used to indicate that the subscription already exists see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.10.3.2 GET

This method is associated to the operation "Query Subscriptions" and shall exhibit the behaviour defined by clause 5.8.4, providing the subscription data as part of the HTTP response payload body. Figure 6.10.3.2-1 shows the Query Subscriptions interaction.

**Figure 6.10.3.2-1: Query Subscriptions interaction**

The URL parameters that shall be supported by implementations are those defined in Table 6.10.3.2-1 and Table 6.10.3.2-2 describes the request body and possible responses.

Table 6.10.3.2-1: Query Subscriptions URL parameters

Name	Data Type	Cardinality	Remarks
limit	Number	0..1	Maximum number of subscriptions to be retrieved

Table 6.10.3.2-2: Query Subscriptions request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	Subscription[]	1	200 OK	A response body containing a list of subscriptions.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.11 Resource: subscriptions/{subscriptionId}

6.11.1 Description

This resource represents a subscription known to an NGSI-LD system.

6.11.2 Resource definition

Resource URI:

- /subscriptions/{subscriptionId}

Resource URI variables for this resource are defined in Table 6.11.2-1.

Table 6.11.2-1: URI variables

Name	Definition
subscriptionId	Id (URI) of the concerned subscription

6.11.3 Resource methods

6.11.3.1 GET

This method is associated to the operation "Retrieve Subscription" and shall exhibit the behaviour defined by clause 5.8.3. The subscription identifier is the value of the resource URI variable "subscriptionId". Figure 6.11.3.1-1 shows the Retrieve Subscription interaction and Table 6.11.3.1-1 describes the request body and possible responses.



**Figure 6.11.3.1-1: Retrieve Subscription interaction**

Table 6.11.3.1-1: Retrieve Subscription request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	Subscription	1	200 OK	A response body containing the JSON-LD representation of the target subscription.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided a subscription identifier (URI) not known to the system, see clause 6.3.2.

6.11.3.2 PATCH

This method is associated to the operation "Update Subscription" and shall exhibit the behaviour defined by clause 5.8.2. The subscription identifier is the value of the resource URI variable "subscriptionId". Figure 6.11.3.2-1 shows the Update Subscription interaction and Table 6.11.3.2-1 describes the request body and possible responses.

**Figure 6.11.3.2-1: Update Subscription interaction****Table 6.11.3.2-1: Update Subscription request body and possible responses**

Request Body	Data Type	Cardinality	Remarks	
	Subscription Fragment	1	Subscription Fragment including id, type and any other subscription field to be changed	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided a subscription identifier (URI) not known to the system, see clause 6.3.2.

6.11.3.3 DELETE

This method is associated to the operation "Delete Subscription" and shall exhibit the behaviour defined by clause 5.8.5. The subscription identifier is the value of the resource URI variable "subscriptionId". Figure 6.11.3.3-1 shows the Delete Subscription interaction and Table 6.11.3.3-1 describes the request body and possible responses.




Figure 6.11.3.3-1: Delete Subscription interaction

Table 6.11.3.3-1: Delete Subscription request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided a subscription identifier (URI) not known to the system, see clause 6.3.2.

6.12 Resource: csourceSubscriptions/

6.12.1 Description

This resource represents the context source registration subscriptions known to an NGSI-LD system.

6.12.2 Resource definition

Resource URI:

- /csourceSubscriptions/

6.12.3 Resource methods

6.12.3.1 POST

This method is bound to the operation "Create Context Source Registration Subscription" and shall exhibit the behaviour defined by clause 5.11.2, taking the context source registration subscription to be created from the HTTP request payload body. Figure 6.12.3.1-1 shows the Create Context Source Registration Subscription interaction and Table 6.12.3.1-1 describes the request body and possible responses.




Figure 6.12.3.1-1: Create Context Source Registration Subscription interaction

Table 6.12.3.1-1: Create Context Source Registration Subscription request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Subscription	1	Payload body in the request contains a JSON-LD object which represents the context source registration subscription that is to be created.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	201 Created	The HTTP response shall include a "Location" HTTP header that contains the relative path of the created context source registration subscription.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	409 Conflict	It is used to indicate that the context source registration subscription already exists, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.12.3.2 GET

This method is associated to the operation "Query Context Source Registration Subscriptions" and shall exhibit the behaviour defined by clause 5.11.5, providing the context source registration subscription data as part of the HTTP response payload body. Figure 6.12.3.2-1 shows the Query Context Source Registration Subscriptions interaction.




Figure 6.12.3.2-1: Query Context Source Registration Subscriptions interaction

The URL parameters that shall be supported by implementations are those defined in Table 6.12.3.2-1 and Table 6.12.3.2-2 describes the request body and possible responses.

Table 6.12.3.2-1: Query Context Source Registration Subscriptions URL parameters

Name	Data Type	Cardinality	Remarks
limit	Number	0..1	Maximum number of subscriptions to be retrieved

Table 6.12.3.2-2: Query Context Source Registration Subscriptions request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	Subscription[]	1	200 OK	A response body containing a list of context source registration subscriptions.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.13 Resource: csourceSubscriptions/{subscriptionId}

6.13.1 Description

This resource represents the context source registration subscription, identified by *subscriptionId*, known to an NGSI-LD system.

6.13.2 Resource definition

Resource URI:

- /csourceSubscriptions/{subscriptionId}

Resource URI variables for this resource are defined in Table 6.13.2-1.

Table 6.13.2-1: URI variables

Name	Definition
subscriptionId	Id (URI) of the concerned context source registration subscription.

6.13.3 Resource methods

6.13.3.1 GET

This method is associated to the operation "Retrieve Context Source Registration Subscription" and shall exhibit the behaviour defined by clause 5.11.4. The subscription identifier is the value of the resource URI variable "subscriptionId". Figure 6.13.3.1-1 shows the Retrieve Context Source Registration interaction and Table 6.13.3.1-1 describes the request body and possible responses.




Figure 6.13.3.1-1: Retrieve Context Source Registration Subscription interaction

Table 6.13.3.1-1: Retrieve Context Source Registration Subscription request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	Subscription	1	200 OK	A response body containing the JSON-LD representation of the target context source registration subscription.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided a subscription identifier (URI) not known to the system, see clause 6.3.2.

6.13.3.2 PATCH

This method is associated to the operation "Update Context Source Registration Subscription" and shall exhibit the behaviour defined by clause 5.11.3. The subscription identifier is the value of the resource URI variable "subscriptionId". Figure 6.13.3.2-1 shows the Update Context Source Registration Subscription interaction and Table 6.13.3.2-1 describes the request body and possible responses.




Figure 6.13.3.2-1: Update Context Source Registration Subscription interaction

Table 6.13.3.2-1: Update Context Source Registration Subscription request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Subscription Fragment	1	Subscription Fragment including id, type and any other context source registration subscription field to be changed.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided a subscription identifier (URI) not known to the system, see clause 6.3.2.

6.13.3.3 DELETE

This method is associated to the operation "Delete Context Source Registration Subscription" and shall exhibit the behaviour defined by clause 5.11.6. The subscription identifier is the value of the resource URI variable "subscriptionId". Figure 6.13.3.3-1 shows the Delete Context Source Registration Subscription interaction and Table 6.13.3.3-1 describes the request body and possible responses.




Figure 6.13.3.3-1: Delete Context Source Registration Subscription interaction

Table 6.13.3.3-1: Delete Context Source Registration Subscription request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided a subscription identifier (URI) not known to the system, see clause 6.3.2.

6.14 Resource: entityOperations/create

6.14.1 Description

A sub-resource, pertaining to the *entityOperations/* resource, intended to enable batch entity creation for the NGSI-LD API.

6.14.2 Resource definition

Resource URI:

- /entityOperations/create

6.14.3 Resource methods

6.14.3.1 POST

This method is associated to the operation "Batch Entity Creation" and shall exhibit the behaviour defined by clause 5.6.7. Figure 6.14.3.1-1 shows the operation interaction and Table 6.14.3.1-1 describes the request body and possible responses.




Figure 6.14.3.1-1: Batch Entity Creation interaction

Table 6.14.3.1-1: Batch Entity Creation request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Entity[]	1	Array of entities to be created.	
Response Body	Data Type	Cardinality	Response Code	Remarks
	String[]	1	201 Created	If all entities have been successfully created, an array of Strings containing URIs is returned in the response. Each URI represents the Entity ID of a created entity. There is no restriction as to the order of the Entity IDs.
	BatchOperationResult	1	207 Multi-Status	If only some or none of the entities have been successfully created, a response body containing the result of each operation contained in the batch is returned in a <i>BatchOperationResult</i> structure. It contains two arrays. The first array (<i>success</i>) contains the URIs of the successfully created entities, while the second array (<i>errors</i>) contains information about the error for each of the entities that could not be created. There is no restriction as to the order of the Entity IDs in the arrays. If any of the entities matches to a registration, the relevant parts of the request are forwarded as a distributed operation. In the case when an error response is received back from any distributed operation, a response body containing the result returned from each registration is returned in a <i>BatchOperationResult</i> structure. Errors can occur whenever a distributed operation is unsupported, fails or times out, see clause 6.3.17.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.15 Resource: entityOperations/upsert

6.15.1 Description

A sub-resource, pertaining to the *entityOperations/* resource, intended to enable batch entity creation or update for the NGSI-LD API.

6.15.2 Resource definition

Resource URI:

- /entityOperations/upsert

6.15.3 Resource methods

6.15.3.1 POST

This method is associated to the operation "Batch Entity Creation or Update (Upsert)" and shall exhibit the behaviour defined by clause 5.6.8. Figure 6.15.3.1-1 shows the operation interaction and Table 6.15.3.1-1 describes the request body and possible responses.

The *options* query parameter for this request can take the following values:

- "replace". Indicates that all the existing Entity content shall be replaced (default mode);
- "update". Indicates that existing Entity content shall be updated.




Figure 6.15.3.1-1: Batch Entity Creation or Update interaction

Table 6.15.3.1-1: Batch Entity Creation or Update request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Entity[]	1	Array of entities to be created/updated.	
Response Body			Response Code	Remarks
	String[]	1	201 Created	If all entities not existing prior to this request have been successfully created and the others have been successfully updated, an array of String (with the URIs representing the Entity IDs of the created entities only) is returned in the response. There is no restriction as to the order of the Entity IDs. The merely updated entities do not take part in the response (corresponding to 204 No Content returned in the case of updates).
	N/A	N/A	204 No Content	If all entities already existed and are successfully updated, there is no payload body in the response.
	BatchOperationResult	1	207 Multi-Status	If only some or none of the entities have been successfully created or updated, a response body containing the result of each operation contained in the batch is returned in a <i>BatchOperationResult</i> structure. It contains two arrays. The first array (<i>success</i>) contains the URIs of the successfully created or updated entities, while the second array (<i>errors</i>) contains information about the error for each of the entities that could not be created or updated. There is no restriction as to the order of the Entity IDs in the arrays. If any of the entities matches to a registration, the relevant parts of the request are forwarded as a distributed operation. In the case when an error response is received back from any distributed operation, a response body containing the result returned from each registration is returned in a <i>BatchOperationResult</i> structure. Errors can occur whenever a distributed operation is unsupported, fails or times out, see clause 6.3.17.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.16 Resource: entityOperations/update

6.16.1 Description

A sub-resource, pertaining to the *entityOperations/* resource, intended to enable batch entity update for the NGSI-LD API.

6.16.2 Resource definition

Resource URI:

- /entityOperations/update

6.16.3 Resource methods

6.16.3.1 POST

This method is associated to the operation "Batch Entity Update" and shall exhibit the behaviour defined by clause 5.6.9. Figure 6.16.3.1-1 shows the operation interaction and Table 6.16.3.1-1 describes the request body and possible responses.

The *options* query parameter for this request can take the following values:

- "noOverwrite". Indicates that no attribute overwrite shall be performed.




Figure 6.16.3.1-1: Batch Entity Update interaction

Table 6.16.3.1-1: Batch Entity Update request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Entity[]	1	Array of Entities to be updated	
Response Body	Data Type	Cardinality	Response Code	Remarks
	N/A	N/A	204 No Content	If all entities have been successfully updated, there is no payload body in the response.
	BatchOperationResult	1	207 Multi-Status	If only some or none of the entities have been successfully updated, a response body containing the result of each operation contained in the batch is returned in a <i>BatchOperationResult</i> structure. It contains two arrays. The first array (<i>success</i>) contains the URIs of the successfully updated entities, while the second array (<i>errors</i>) contains information about the error for each of the entities that could not be updated. There is no restriction as to the order of the Entity IDs in the arrays. If any of the entities matches to a registration, the relevant parts of the request are forwarded as a distributed operation. In the case when an error response is received back from any distributed operation, a response body containing the result returned from each registration is returned in a <i>BatchOperationResult</i> structure. Errors can occur whenever a distributed operation is unsupported, fails or times out, see clause 6.3.17.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.17 Resource: entityOperations/delete

6.17.1 Description

A sub-resource, pertaining to the *entityOperations/* resource, intended to enable batch entity deletion for the NGSI-LD API.

6.17.2 Resource definition

Resource URI:

- /entityOperations/delete

6.17.3 Resource methods

6.17.3.1 POST

This method is associated to the operation "Batch Entity Delete" and shall exhibit the behaviour defined by clause 5.6.10. Figure 6.17.3.1-1 shows the operation interaction and Table 6.17.3.1-1 describes the request body and possible responses.




Figure 6.17.3.1-1: Batch Entity Delete interaction

Table 6.17.3.1-1: Batch Entity Delete request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	String[]	1	Array of String (URLs representing Entity IDs) to be deleted	
Response Body	Data Type	Cardinality	Response Code	Remarks
	N/A	N/A	204 No Content	If all entities existed and have been successfully deleted, there is no payload body in the response.
	BatchOperationResult	1	207 Multi-Status	If some or all of the entities have not been successfully deleted, or did not exist, a response body containing the result of each operation contained in the batch is returned in a <i>BatchOperationResult</i> structure. It contains two arrays. The first array (<i>success</i>) contains the URLs of the successfully deleted entities, while the second array (<i>errors</i>) contains information about the error for each of the entities that could not be deleted. There is no restriction as to the order of the Entity IDs in the arrays. If any of the entities matches to a registration, the relevant parts of the request are forwarded as a distributed operation. In the case when an error response is received back from any distributed operation, a response body containing the result returned from each registration is returned in a <i>BatchOperationResult</i> structure. Errors can occur whenever a distributed operation is unsupported, fails or times out, see clause 6.3.17.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.18 Resource: temporal/entities/

6.18.1 Description

This resource represents the Temporal Evolution of Entities known to an NGSI-LD system.

6.18.2 Resource definition

Resource URI:

- /temporal/entities/

6.18.3 Resource methods

6.18.3.1 POST

This method is associated to the operation "Create or Update Temporal Evolution of an Entity" and shall exhibit the behaviour defined by clause 5.6.11, taking the temporal representation of Entity to be created from the HTTP request payload body. Figure 6.18.3.1-1 shows this interaction and Table 6.18.3.1-1 describes the request body and possible responses.




Figure 6.18.3.1-1: Create or Update Temporal Evolution of an Entity interaction

Table 6.18.3.1-1: Create or Update Temporal Evolution of an Entity request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	EntityTemporal	1	Payload body in the request contains a JSON-LD object which represents the temporal representation of the Entity that is to be created (or updated).	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	201 Created	Upon creation success, the HTTP response shall include a "Location" HTTP header that contains the relative path of the created entity.
	N/A	N/A	204 No Content	Upon update success.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> member should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	422 Unprocessable Entity	It is used to indicate that the operation is not available, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.18.3.2 GET

This method is associated to the operation "Query Temporal Evolution of Entities" and shall exhibit the behaviour defined by clause 5.7.4, providing the Temporal Evolution of the matching Entities as part of the HTTP response payload body. In addition to this method, an alternative way to perform "Query Temporal Evolution of Entities" operations via POST is defined in clause 6.24. Figure 6.18.3.2-1 shows this interaction.




Figure 6.18.3.2-1: Query Temporal Evolution of Entities interaction

The URL parameters that shall be supported by implementations are those defined in Table 6.18.3.2-1 and Table 6.18.3.2-2 describes the request body and possible responses.

Table 6.18.3.2-1: Query Temporal Evolution of Entities URL parameters

Name	Data Type	Cardinality	Remarks
aggrMethods	Comma separated list of strings	0..1 It shall be 1 if <i>aggregatedValues</i> is present in the <i>options</i> parameter	Each String represents an aggregation method, as defined by clause 4.5.19. Only applicable if "aggregatedValues" is present in the <i>format</i> or <i>options</i> parameter.
aggrPeriodDuration	String	0..1	It represents the duration of each period used for the aggregation as defined by clause 4.5.19. If not specified, it defaults to a duration of 0 seconds and is interpreted as a duration spanning the whole time-range specified by the temporal query. Only applicable if "aggregatedValues" is present in the <i>format</i> or <i>options</i> parameter.
attrs	Comma separated list of strings	0..1 It shall be 1 if <i>type</i> is not present, unless the execution of the request is limited to local scope (see clause 5.5.13).	A synonym for a combination of the <i>pick</i> and <i>q</i> parameters. Deprecated Each String is an Attribute (Property or Relationship) name. List of Attributes (Properties or Relationships) to be retrieved.
collation	String	0..1	An ICU collation (see IETF RFC 6067 [36]), When defined, the Entities returned in the payload shall be ordered according to the collation given. It is part of Entity ordering.
coordinates	String	0..1 It shall be one if <i>georel</i> or <i>geometry</i> are present	Coordinates serialized as a string as per clause 4.10. It is part of geoquery.
csf	String	0..1	Context Source filter as per clause 4.9.

Name	Data Type	Cardinality	Remarks
datasetId	Comma separated list of strings	0..1	Shall be valid URIs, "@none" for including the default Attribute instances. Specifies the datasetIds of the Attribute instances to be selected for each matched Attribute as per clause 4.5.5.
endTimeAt	String	0..1	It is representing the <i>endTimeAt</i> parameter as defined by clause 4.11. It shall be a <i>DateTime</i> . Cardinality shall be 1 if <i>timerel</i> is equal to "between".
entityMap	Boolean	0..1	If <i>true</i> , the location of the EntityMap used in the operation is returned in the response.
entityMapLifetime	String	0..1	Suggested duration, represented in ISO 8601 [17] format, for which the requester wants the EntityMap to exist. The actual expiresAt time of the EntityMap shall be set by the Context Broker or Context Source, possibly overriding the requested duration. Only applicable if <i>entityMap</i> is set to <i>true</i> .
expandValues	Comma separated list of strings	0..1	Each String is an Attribute (Property or Relationship) name. List of Attributes whose values shall be expanded into URIs according to the supplied <i>@context</i> prior to executing a query. It is part of query.
geometry	String	0..1 It shall be 1 if <i>georel</i> or <i>coordinates</i> are present	Geometry as per clause 4.10. It is part of geoquery.
geoproperty	String	0..1 It shall be ignored if no geoquery is present	The name of the <i>GeoProperty</i> that contains the geospatial data that will be used to resolve the geoquery. By default, will be <i>location</i> . (See clause 4.7).
georel	String	0..1 It shall be 1 if <i>geometry</i> or <i>coordinates</i> are present	Geo relationship as per clause 4.10. It is part of geoquery.
id	Comma separated list of strings	0..1	Each String shall be a valid URI. List of entity ids to be retrieved.
idPattern	Regular expression as defined by [11]	0..1	Regular expression that shall be matched by entity ids.
jsonKeys	Comma separated list of strings	0..1	Each String is an Attribute (Property or Relationship) name. Values of the identified attributes are to be considered uninterpretable as JSON-LD and should not be expanded against the supplied <i>@context</i> using JSON-LD type coercion prior to executing the query.
lang	String	0..1	It represents the preferred natural language of the response. It is used to reduce <i>languageMaps</i> to a string or string array property in a single preferred language.
lastN	Positive integer	0..1	Only the last n instances, per Attribute, per Entity (under the specified time interval) shall be retrieved.
omit	Comma separated list of strings	0..1	Each String is an Entity member ("id", "type", "scope" or an Attribute name). When defined, the listed Entity members are removed from each Entity within the payload.
orderBy	String	0..1	The Entity member ("id") appended with an optional sorting style ("asc", "desc") as per clause 4.23. When defined, the Entities returned in the payload shall be ordered according to members defined. It is part of Entity ordering.

Name	Data Type	Cardinality	Remarks
pick	Comma separated list of strings	0..1	Each String is an Entity member ("id", "type", "scope" or an Attribute name). When defined, every Entity within the payload body is reduced down to only contain the listed Entity members.
q	String	0..1	Query as per clause 4.9.
scopeQ	String	0..1	Scope query (see clause 4.19).
splitEntities	Boolean	0..1	If <i>true</i> it is assumed that single Entities are distributed between different Context Brokers and/or Context Sources and this has to be taken into account when applying any kind of filters (q, geoQ, scopeQ, Attributes etc.). If <i>false</i> it is expected that Context Broker and/or Context Source always have complete Entities, which allows applying filters locally. The parameter does not apply in case <i>Local</i> is <i>true</i> . The default value should be decided by implementation; it should be configurable.
timeAt	String	1	representing the <i>timeAt</i> parameter as defined by clause 4.11. It shall be a <i>DateTime</i> .
timeproperty	String	0..1	It represents a Temporal Property name. Allowed values: "observedAt", "createdAt", "modifiedAt" and "deletedAt". If not specified, the default is "observedAt". (See clause 4.8).
timerel	String	1	It represents the temporal relationship as defined by clause 4.11. Allowed values: "before", "after", "between".
type	String	0..1 It shall be 1 if <i>attrs</i> is not present, unless the execution of the request is limited to local scope (see clause 5.5.13).	Selection of Entity Types as per clause 4.17. "*" is also allowed as a value and <i>Local</i> is implicitly set to <i>true</i> and shall not be explicitly set to <i>false</i> .

Table 6.18.3.2-2: Query Temporal Evolution of Entities request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	EntityTemporal[]	1	200 OK 201 Created (in case an EntityMap has been (re)created)	A response body containing the query result as a list of temporal representation of Entities. If an EntityMap has been requested, the HTTP response shall include an "NGSILD-EntityMap" HTTP header that contains the resource URI of the EntityMap resource used in the operation.
	EntityTemporal[]	1	203 Non-Authoritative Information	As above, but returning an altered response body , amended to conform to a specific version of the NGSI-LD specification as mandated in clause 4.3.6.8. The response shall also include a "Preference-Applied" HTTP header set to "ngsi-ld=<conformant-version>".
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.19 Resource: temporal/entities/{entityId}

6.19.1 Description

This resource is associated to the Temporal Evolution of an Entity known to an NGSI-LD system.

6.19.2 Resource definition

Resource URI:

- /temporal/entities/{entityId}

Resource URI variables for this resource are defined in Table 6.19.2-1.

Table 6.19.2-1: URI variables

Name	Definition
entityId	Id (URI) of the entity to be retrieved

6.19.3 Resource methods

6.19.3.1 GET

This method is associated to the operation "Retrieve Temporal Evolution of an Entity" and shall exhibit the behaviour defined by clause 5.7.3. The Entity identifier is the value of the resource URI variable *entityId*. Figure 6.19.3.1-1 shows the retrieve temporal representation of an entity interaction.




Figure 6.19.3.1-1: Retrieve Temporal Evolution of an Entity interaction

The URL parameters that shall be supported are those defined in Table 6.19.3.1-1 and Table 6.19.3.1-2 describes the request body and possible responses.

Table 6.19.3.1-1: Retrieve Temporal Evolution of an Entity URL parameters

Name	Data Type	Cardinality	Remarks
aggrMethods	Comma separated list of strings	0..1 It shall be 1 if <i>aggregatedValues</i> is present in the <i>options</i> parameter	Each String represents the aggregation methods as defined by clause 4.5.19. Only applicable if "aggregatedValues" is present in the <i>format</i> or <i>options</i> parameter.
aggrPeriodDuration	String	0..1	It represents the duration of each period used for the aggregation as defined by clause 4.5.19. If not specified, it defaults to a duration of 0 seconds and is interpreted as a duration spanning the whole time-range specified by the temporal query. Only applicable if "aggregatedValues" is present in the <i>format</i> or <i>options</i> parameter.
attrs	Comma separated list of strings	0..1	A synonym for the <i>pick</i> parameter, except that <i>id</i> , <i>type</i> , <i>scope</i> are not allowed. Deprecated Each String is an Attribute (Property or Relationship) name. List of Attributes to be retrieved. If not specified, all Attributes related to the temporal representation of an Entity shall be retrieved.
datasetId	Comma separated list of strings	0..1	Shall be valid URIs, "@none" for including the default Attribute instances. Specifies the datasetIds of the Attribute instances to be selected for each matched Attribute as per clause 4.5.5.
endTimeAt	String	0..1 It shall be 1 if <i>timereI</i> is equal to "between"	It represents the <i>endTimeAt</i> parameter as defined by clause 4.11. It shall be a <i>DateTime</i> .
lang	String	0..1	It represents the preferred natural language of the response. It is used to reduce <i>languageMaps</i> to a string or string array property in a single preferred language.
lastN	Positive integer	0..1	Only the last n Attribute instances (under the concerned time interval) shall be retrieved.
omit	Comma separated list of strings	0..1	Each String is an Entity member ("id", "type", "scope" or an Attribute name). When defined, the listed Entity members are removed from each Entity within the payload.

Name	Data Type	Cardinality	Remarks
pick	Comma separated list of strings	0..1	Each String is an Entity member ("id", "type", "scope" or an Attribute name). When defined, every Entity within the payload body is reduced down to only contain the listed Entity members.
timeAt	String	0..1 It shall be 1 if <i>timereI</i> is present	It represents the <i>timeAt</i> parameter as defined by clause 4.11. It shall be a <i>DateTime</i> .
timeproperty	String	0..1	It represents a Temporal Property name. Allowed values: "observedAt", "createdAt", "modifiedAt" and "deletedAt". If not specified, the default is "observedAt". (See clause 4.8).
timereI	String	0..1 It shall be 1 if <i>timeAt</i> is present	It represents the Temporal Relationship as defined by clause 4.11. Allowed values: "before", "after", "between".

Table 6.19.3.1-2: Get Temporal Evolution of an Entity request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	EntityTemporal	1	200 OK	A response body containing the JSON-LD temporal representation of the target Entity containing the selected Attributes.
	EntityTemporal	1	203 Non-Authoritative Information	As above, but returning an altered response body , amended to conform to a specific version of the NGSI-LD specification as mandated in clause 4.3.6.8. The response shall also include a "Preference-Applied" HTTP header set to "ngsi-ld=<conformant-version>".
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an Entity identifier (URI) not known to the system, see clause 6.3.2.

6.19.3.2 DELETE

This method is associated to the operation "Delete Temporal Evolution of an Entity" and shall exhibit the behaviour defined by clause 5.6.16. The Entity identifier is the value of the resource URI variable *entityId*. Figure 6.19.3.2-1 shows the delete entity interaction and Table 6.19.3.2-1 describes the request body and possible responses.




Figure 6.19.3.2-1: Delete Temporal Evolution of an Entity interaction

Table 6.19.3.2-1: Delete Temporal Evolution of an Entity request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an Entity identifier (URI) not known to the system, see clause 6.3.2.

6.20 Resource: temporal/entities/{entityId}/attrs/

6.20.1 Description

This resource represents all the Attributes (Properties or Relationships) of a Temporal Evolution of an Entity.

6.20.2 Resource definition

Resource URI:

- /temporal/entities/{entityId}/attrs/

Resource URI variables for this resource are defined in Table 6.20.2-1.

Table 6.20.2-1: URI variables

Name	Definition
entityId	Id (URI) of the concerned entity

6.20.3 Resource methods

6.20.3.1 POST

This method is bound to the "Add Attributes to Temporal Evolution of an Entity" operation and shall exhibit the behaviour defined by clause 5.6.12. The Entity identifier is the value of the resource URI variable *entityId*. The data to be added shall be contained in the HTTP request payload body. Figure 6.20.3.1-1 shows the Add Attributes interaction and Table 6.20.3.1-1 describes the request body and possible responses.




Figure 6.20.3.1-1: Add Attributes to Temporal Evolution of an Entity interaction

Table 6.20.3.1-1: Add Attributes to Temporal Evolution of an Entity request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	EntityTemporal Fragment	1	EntityTemporal Fragment containing a complete representation of the Attribute instances to be added.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No content	All the Attributes were added successfully.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an Entity identifier (URI) not known to the system, see clause 6.3.2.

6.21 Resource: temporal/entities/{entityId}/attrs/{attrId}

6.21.1 Description

This resource represents an Attribute (Property or Relationship) of a Temporal Evolution of an Entity.

6.21.2 Resource definition

Resource URI:

- /temporal/entities/{entityId}/attrs/{attrId}

Resource URI variables for this resource are defined in Table 6.21.2-1.

Table 6.21.2-1: URI variables

Name	Definition
entityId	Id (URI) of the concerned entity
attrId	Attribute name (Property or Relationship)

6.21.3 Resource methods

6.21.3.1 DELETE

This method is associated to the operation "Delete Attribute from Temporal Evolution of an Entity" and shall exhibit the behaviour defined by clause 5.6.13. The Entity identifier is the value of the resource URI variable *entityId*. The Attribute name is the value of the resource URI variable *attrId*. Figure 6.21.3.1-1 shows the "Delete Attribute from Temporal Evolution of an Entity" interaction, Table 6.21.3.1-1 shows the delete parameters to be supported and Table 6.21.3.1-2 describes the request body and possible responses.




Figure 6.21.3.1-1: Delete Attribute from Temporal Evolution of an Entity interaction

Table 6.21.3.1-1: Delete Attribute from Temporal Evolution of an Entity URL parameters

name	Data Type	Cardinality	Remarks
datasetId	String	0..1	Shall be a valid URI. Specifies the <i>datasetId</i> of the dataset to be deleted.
deleteAll	Boolean	0..1	If <i>true</i> , all attribute instances are deleted. Otherwise (default) only the Attribute instance specified by the <i>datasetId</i> is deleted. In case neither the deleteAll flag nor a <i>datasetId</i> is present, the default Attribute instance is deleted.

Table 6.21.3.1-2: Delete Attribute from Temporal Evolution of an Entity request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an Entity identifier (URI) or Attribute name not known to the system. See clause 6.3.2.

6.22 Resource: temporal/entities/{entityId}/attrs/{attrId}/{instanceId}

6.22.1 Description

This resource represents an Attribute (Property or Relationship) instance of a Temporal Evolution of an Entity.

6.22.2 Resource definition

Resource URI:

- /temporal/entities/{entityId}/attrs/{attrId}/{instanceId}

Resource URI variables for this resource are defined in Table 6.22.2-1.

Table 6.22.2-1: URI variables

Name	Definition
entityId	Id (URI) of the concerned entity
attrId	Attribute name (Property or Relationship)
instanceId	Id (URI) identifying a particular Attribute instance

6.22.3 Resource methods

6.22.3.1 PATCH

This method is associated to the operation "Modify attribute instance from Temporal Evolution of an Entity" and shall exhibit the behaviour defined by clause 5.6.14. The Entity identifier is the value of the resource URI variable *entityId*. The attribute name is the value of the resource URI variable *attrId*. The instance identifier is the value of the resource URI variable *instanceId*. Figure 6.22.3.1-1 shows the Modify Attribute instance interaction and Table 6.22.3.1-1 describes the request body and possible responses.




Figure 6.22.3.1-1: Modify Attribute instance from Temporal Evolution interaction

Table 6.22.3.1-1: Modify Attribute instance from Temporal Evolution request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	EntityTemporal Fragment	1	EntityTemporal Fragment containing a complete representation of the Attribute instance to be replaced.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an Entity identifier (URI), attribute name or instance identifier not known to the system. See clause 6.3.2.

6.22.3.2 DELETE

This method is associated to the operation "Delete Attribute instance from Temporal Evolution of an Entity" and shall exhibit the behaviour defined by clause 5.6.15. The Entity identifier is the value of the resource URI variable *entityId*. The Attribute name is the value of the resource URI variable *attrId*. The instance identifier is the value of the resource URI variable *instanceId*. Figure 6.22.3.2-1 shows the Delete Attribute instance interaction and Table 6.22.3.2-1 describes the request body and possible responses.




Figure 6.22.3.2-1: Delete Attribute instance from Temporal Evolution of an Entity interaction

Table 6.22.3.2-1: Delete Attribute instance from Temporal Evolution of an Entity request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an Entity identifier (URI), attribute name or instance identifier not known to the system. See clause 6.3.2.

6.23 Resource: entityOperations/query

6.23.1 Description

A sub-resource, pertaining to the *entityOperations/* resource, intended to enable querying for entities by means of a POST method. The behaviour of this clause mirrors the one in clause 6.4.3.2, which performs the "Query Entity" operation (defined by clause 5.7.2) by means of a GET method. The reason to provide an alternative via POST is that, using GET:

- The client may end up assembling very long URLs, due to the URI parameters for *id*, *q*, *type*, *attrs*, etc., being included in the URL. Problems with too long URLs may arise with some applications that cut URLs to a maximum length.
- There is a need to URL-encode the resulting URL. By using POST, there is no need to url-encode.

6.23.2 Resource definition

Resource URI:

- /entityOperations/query

6.23.3 Resource methods

6.23.3.1 POST

This method is associated to the operation "Query Entities" and shall exhibit the behaviour defined by clause 5.7.2. Figure 6.23.3.1-1 shows the operation interaction and Table 6.23.3.1-1 describes the request body and possible responses.




Figure 6.23.3.1-1: Query Entity via POST interaction

Table 6.23.3.1-1: Query Entity via POST request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Query	1	Payload body in the request contains a JSON-LD object which represents the query to be performed.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	Entity[]	1	200 OK 201 Created (in case an EntityMap has been (re)created)	A response body containing the query result as a list of Entities. The HTTP response shall include an "NGSILD-EntityMap" HTTP header that contains the resource URI of the EntityMap resource created in the operation.
	GeoJSON FeatureCollection	1	200 OK 201 Created (in case an EntityMap has been (re)created)	If the Accept Header indicates that the Entities are to be rendered as GeoJSON, a response body containing the query result as GeoJSON FeatureCollection is returned. The HTTP response shall include an "NGSILD-EntityMap" HTTP header that contains the resource URI of the EntityMap resource created in the operation.
	Entity[]	1	203 Non-Authoritative Information	As above, but returning an altered response body , amended to conform to a specific version of the NGSI-LD specification as mandated in clause 4.3.6.8. The response shall also include a "Preference-Applied" HTTP header set to "ngsi-ld=<conformant-version>".
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.24 Resource: temporal/entityOperations/query

6.24.1 Description

A sub-resource, pertaining to the *temporal/entityOperations/* resource, intended to enable temporal querying for entities by means of a POST method. The behaviour of this clause mirrors the one in clause 6.18.3.2, which performs the "Query Temporal Evolution of Entities" (defined by clause 5.7.4) operation by means of a GET method. The reason to provide an alternative via POST is that, using GET:

- The client may end up assembling very long URLs, due to the URI parameters for *id*, *q*, *type*, *attrs*, etc., being included in the URL. Problems with too long URLs may arise with some applications that cut URLs to a maximum length.
- There is a need to URL-encode the resulting URL. By using POST, there is no need to url-encode.

6.24.2 Resource definition

Resource URI:

- /temporal/entityOperations/query

6.24.3 Resource methods

6.24.3.1 POST

This method is associated to the operation "Query Temporal Evolution of Entities" and shall exhibit the behaviour defined by clause 5.7.4. Figure 6.24.3.1-1 shows the operation interaction and Table 6.24.3.1-1 describes the request body and possible responses.




Figure 6.24.3.1-1: Temporal Query Entity via POST interaction

Table 6.24.3.1-1: Temporal Query Entity via POST request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Query	1	Payload body in the request contains a JSON-LD object which represents the query to be performed.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	EntityTemporal[]	1	200 OK 201 Created (in case an EntityMap has been (re)created)	A response body containing the query result as a list of Entities. The HTTP response shall include an "NGSI-LD-EntityMap" HTTP header that contains the resource URI of the EntityMap resource created in the operation.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.25 Resource: types/

6.25.1 Description

This resource represents the entity types available in an NGSI-LD system.

6.25.2 Resource definition

Resource URI:

- /types/

6.25.3 Resource methods

6.25.3.1 GET

This method is associated to the operations "Retrieve Available Entity Types" and "Retrieve Details of Available Entity Types" (if the *details* parameter is set to *true*) and shall exhibit the behaviour defined by clauses 5.7.5 and 5.7.6 respectively.




Figure 6.25.3.1-1: Retrieve Available Entity Types interaction

The request parameters that shall be supported are those defined in Table 6.25.3.1-1 and Table 6.25.3.1-2 describes the request body and possible responses.

Table 6.25.3.1-1: Retrieve Available Entity Types URL parameters

Name	Data Type	Cardinality	Remarks
details	Boolean	0..1	If <i>true</i> , then detailed entity type information represented as an array with elements of the Entity Type data structure (clause 5.2.25) is to be returned.

Table 6.25.3.1-2: Retrieve Available Entity Types request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	EntityTypeInfoList	1	200 OK	A response body containing the JSON-LD representation of the EntityTypeInfoList (clause 5.2.24) is to be returned, unless <i>details=true</i> is specified.
	EntityTypeInfo[]	1	200 OK	If <i>details=true</i> is specified, a response body containing a JSON-LD array with elements of the EntityTypeInfo data structure (clause 5.2.25) is to be returned.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.26 Resource: types/{type}

6.26.1 Description

This resource represents the specified entity type for which entity instances are available in an NGSI-LD system.

6.26.2 Resource definition

Resource URI:

- /types/{type}

Resource URI variables for this resource are defined in Table 6.26.2-1.

Table 6.26.2-1: URI variables

Name	Definition
type	Name of the entity type for which detailed information is to be retrieved. The Fully Qualified Name (FQN) as well as the short name can be used, given that the latter is part of the JSON-LD <i>@context</i> provided.

6.26.3 Resource methods

6.26.3.1 GET

This method is associated to the operation "Retrieve Available Entity Type Information" and shall exhibit the behaviour defined by clause 5.7.7. The entity type is the value of the resource URI variable "type". Figure 6.26.3.1-1 shows the retrieve available entity type interaction.




Figure 6.26.3.1-1: Retrieve Available Entity Type interaction

Table 6.26.3.1-1 describes the request body and possible responses.

Table 6.26.3.1-1: Retrieve Available Entity Type request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	EntityTypeInfo	1	200 OK	A response body containing the JSON-LD representation of the detailed information about the available entity type.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an entity type not known to the system, see clause 6.3.2.

6.27 Resource: attributes/

6.27.1 Description

This resource represents the attributes available in an NGSI-LD system.

6.27.2 Resource definition

Resource URI:

- /attributes/

6.27.3 Resource methods

6.27.3.1 GET

This method is associated to the operations "Retrieve Available Attributes" and "Retrieve Details of Available Attributes" (if the *details* parameter is set to *true*) and shall exhibit the behaviour defined by clauses 5.7.8 and 5.7.9 respectively.




Figure 6.27.3.1-1: Retrieve Available Attributes interaction

The request parameters that shall be supported are those defined in Table 6.27.3.1-1 and Table 6.27.3.1-2 describes the request body and possible responses.

Table 6.27.3.1-1: Retrieve Available Attributes URL parameters

Name	Data Type	Cardinality	Remarks
details	Boolean	0..1	If <i>true</i> , then detailed attribute information represented as an array with elements of the Attribute data structure (clause 5.2.28) is to be returned.

Table 6.27.3.1-2: Retrieve Available Attributes request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	AttributeList	1	200 OK	A response body containing the JSON-LD representation of the AttributeList (clause 5.2.27) is to be returned, unless <i>details=true</i> is specified.
	Attribute[]	1	200 OK	If <i>details=true</i> is specified, a response body containing a JSON-LD array with elements of the Attribute data structure (clause 5.2.28) is to be returned.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.28 Resource: attributes/{attrId}

6.28.1 Description

This resource represents the specified attribute that belongs to entity instances existing within the NGSI-LD system.

6.28.2 Resource definition

Resource URI:

- /attributes/{attrId}

Resource URI variables for this resource are defined in Table 6.28.2-1.

Table 6.28.2-1: URI variables

Name	Definition
attrId	Name of the attribute for which detailed information is to be retrieved. The Fully Qualified Name (FQN) as well as the short name can be used, given that the latter is part of the JSON-LD <i>@context</i> provided.

6.28.3 Resource methods

6.28.3.1 GET

This method is associated to the operation "Retrieve Available Attribute Information" and shall exhibit the behaviour defined by clause 5.7.10. The attribute is the value of the resource URI variable "attrId". Figure 6.28.3.1-1 shows the retrieve available attribute information interaction.




Figure 6.28.3.1-1: Retrieve Available Attribute Information interaction

Table 6.28.3.1-1 describes the request body and possible responses.

Table 6.28.3.1-1: Retrieve Available Attribute Information request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	Attribute	1	200 OK	A response body containing the JSON-LD representation of the detailed information about the available attribute.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an attribute name not known to the system, see clause 6.3.2.

6.29 Resource: jsonldContexts/

6.29.1 Description

This resource represents the *@contexts* known to an NGSI-LD system.

6.29.2 Resource definition

Resource URI:

- /jsonldContexts/

6.29.3 Resource methods

6.29.3.1 POST

This method is bound to the operation "Add *@context*" and shall exhibit the behaviour defined by clause 5.13.2, taking the *@context* to be added from the HTTP request payload body. Figure 6.29.3.1-1 shows the Add *@context* interaction and Table 6.29.3.1-1 describes the request body and possible responses.




Figure 6.29.3.1-1: Add @context interaction

Table 6.29.3.1-1: Add @context request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	JSON Object	1	Payload body in the request contains a JSON object that has a root node named <i>@context</i> , which represents a JSON-LD "local context".	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	201 Created	The HTTP response shall include a "Location" HTTP header that contains the local relative path of the added <i>@context</i> .
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.29.3.2 GET

This method is associated to the operation "List @contexts" and shall exhibit the behaviour defined by clause 5.13.3, and it provides information about stored @contexts as part of the HTTP response payload body. Figure 6.29.3.2-1 shows the List @contexts interaction.




Figure 6.29.3.2-1: List @contexts interaction

The request parameters that shall be supported by implementations are those defined in Table 6.29.3.2-1 and Table 6.29.3.2-2 describes the request body and possible responses.

Table 6.29.3.2-1: List @contexts URL parameters

Name	Data Type	Cardinality	Remarks
details	Boolean	0..1	Whether a list of URLs or a more detailed list of JSON Objects is requested.
kind	String	0..1	Can be either "Cached", "Hosted", or "ImplicitlyCreated".

Table 6.29.3.2-2: List @contexts request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	String[] or JSON Object[]	1	200 OK	A response body containing a list of URLs or a list of JSON Objects, as defined in clause 5.13.3.5, representing metadata about stored @contexts.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.30 Resource: jsonldContexts/{contextId}

6.30.1 Description

This resource represents a JSON-LD @context stored in the broker's internal @context storage.

6.30.2 Resource definition

Resource URI:

- /jsonldContexts/{contextId}

Resource URI variables for this resource are defined in Table 6.30.2-1.

Table 6.30.2-1: URI variables

Name	Definition
contextId	Local identifier of the @context to be managed (served or deleted). For @contexts of kind "Cached" this can also be the original URL the broker downloaded the @context from.

6.30.3 Resource methods

6.30.3.1 GET

This method is associated to the operation "Serve @context" and shall exhibit the behaviour defined by clause 5.13.4. The @context identifier is the value of the resource URI variable "contextId". Figure 6.30.3.1-1 shows the HTTP Serve @context interaction.




Figure 6.30.3.1-1: Serve @context interaction

The request parameters that shall be supported by implementations are those defined in Table 6.30.3.1-1 and Table 6.30.3.1-2 describes the request body and possible responses.

Table 6.30.3.1-1: Serve @contexts URL parameters

Name	Data Type	Cardinality	Remarks
details	Boolean	0..1	Whether the content of the @context or its metadata is requested.

Table 6.30.3.1-2: Serve @context request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	JSON Object	1	200 OK	If the parameter details are <i>false</i> or missing, response body contains a JSON object that has a root node named @context, which represents a JSON-LD "local context". If the parameter details are <i>true</i> , response body contains a JSON object as defined in clause 5.13.4.5, which metadata of a JSON-LD "local context".
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned ProblemDetails structure, the detail attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an @context identifier not known to the system, see clause 6.3.2.
	ProblemDetails (see IETF RFC 7807 [10])	1	422 Unprocessable	It is used when a client indicated an @context of type "Cached", see clause 6.3.2.

6.30.3.2 DELETE

This method is associated to the operation "Delete and Reload @context" and shall exhibit the behaviour defined by clause 5.13.5. The Entity identifier is the value of the resource URI variable "contextId". Figure 6.30.3.2-1 shows the delete entity interaction. The request parameters that shall be supported are those defined in Table 6.30.3.2-1 and Table 6.30.3.2-2 describes the request body and possible responses.

Table 6.30.3.2-1: Delete and Reload @context URL parameters

Name	Data Type	Cardinality	Remarks
reload	Boolean	0..1	indicates to perform a download and replace of the @context, as specified in clause 5.13.5.4.




Figure 6.30.3.2-1: Delete and Reload @context interaction

Table 6.30.3.2-2: Delete and Reload @context request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an @context identifier not known to the system, see clause 6.3.2.
	ProblemDetails (see IETF RFC 7807 [10])	1	504 Gateway Timeout	It is used when re-downloading fails.

6.31 Resource: entityOperations/merge

6.31.1 Description

A sub-resource, pertaining to the *entityOperations/* resource, intended to enable batch entity merge for the NGSI-LD API.

6.31.2 Resource definition

Resource URI:

- /entityOperations/merge

6.31.3 Resource methods

6.31.3.1 POST

This method is associated to the operation "Batch Entity Merge" and shall exhibit the behaviour defined by clause 5.6.20. Figure 6.31.3.1-1 shows the operation interaction and Table 6.31.3.1-1 describes the request body and possible responses.




Figure 6.31.3.1-1: Batch Entity Merge interaction

Table 6.31.3.1-1: Batch Entity Merge request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Entity[]	1	Array of Entities to be merged.	
Response Body	Data Type	Cardinality	Response Code	Remarks
	N/A	N/A	204 No Content	If all entities have been successfully merged, there is no payload body in the response.
	BatchOperationResult	1	207 Multi-Status	If only some or none of the entities have been successfully merged, a response body containing the result of each operation contained in the batch is returned in a <i>BatchOperationResult</i> structure. It contains two arrays. The first array (<i>success</i>) contains the URIs of the successfully merged entities, while the second array (<i>errors</i>) contains information about the error for each of the entities that could not be merged-patched. There is no restriction as to the order of the Entity IDs in the arrays. If any of the entities matches to a registration, the relevant parts of the request are forwarded as a distributed operation. In the case when an error response is received back from any distributed operation, a response body containing the result returned from each registration is returned in a <i>BatchOperationResult</i> structure. Errors can occur whenever a distributed operation is unsupported, fails or times out, see clause 6.3.17.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.32 Resource: entityMaps/{entityMapId}

6.32.1 Description

This resource represents an EntityMap available in the broker's internal storage or memory.

6.32.2 Resource definition

Resource URI:

- /entityMaps/{entityMapId}

Resource URI variables for this resource are defined in Table 6.32.2-1.

Table 6.32.2-1: URI variables

Name	Definition
entityMapId	Id (URI) of the EntityMap to be retrieved, updated or deleted.

6.32.3 Resource methods

6.32.3.1 GET

This method is associated to the operation "Retrieve EntityMap" and shall exhibit the behaviour defined by clause 5.14.1. The EntityMap identifier is the value of the resource URI variable "entityMapId". Figure 6.32.3.1-1 shows the Retrieve EntityMap interaction and Table 6.32.3.1-1 describes the request body and possible responses.




Figure 6.32.3.1-1: Retrieve EntityMap

Table 6.32.3.1-1: Retrieve EntityMap request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	EntityMap	1	200 OK	A response body containing the JSON-LD representation of the target entity.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an EntityMap identifier not known to the system, see clause 6.3.2.

6.32.3.2 PATCH

This method is associated to the operation "Update EntityMap" and shall exhibit the behaviour defined by clause 5.14.2. The EntityMap identifier is the value of the resource URI variable "entityMapId". Figure 6.32.3.2-1 shows the Update EntityMap interaction and Table 6.32.3.2-1 describes the request body and possible responses.




Figure 6.32.3.2-1: Update EntityMap

Table 6.32.3.2-1: Update EntityMap request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	EntityMap Fragment	1	Payload body in the request contains a JSON-LD object which represents the EntityMap fragment with which the EntityMap is to be updated.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an EntityMap identifier not known to the system, see clause 6.3.2.

6.32.3.3 DELETE

This method is associated to the operation "Delete EntityMap" and shall exhibit the behaviour defined by clause 5.14.3. The Entity identifier is the value of the resource URI variable "contextId". Figure 6.32.3.3-1 shows the delete entity interaction and Table 6.32.3.3-1 describes the request body and possible responses.




Figure 6.32.3.3-1: Delete and Reload @context interaction

Table 6.32.3.3-1: Delete EntityMap request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an <i>@context</i> identifier not known to the system, see clause 6.3.2.

6.33 Resource: info/sourceIdentity

6.33.1 Description

This resource represents identity information about the Context Source itself.

6.33.2 Resource definition

Resource URI:

- /info/sourceIdentity

6.33.3 Resource methods

6.33.3.1 GET

This method is associated to the operation "Retrieve Context Source Identity Information" and shall exhibit the behaviour defined by clause 5.15.1. Figure 6.33.3.1-1 shows the Retrieve Context Source Identity Information interaction and Table 6.33.3.1-1 describes the request body and possible responses.




Figure 6.33.3.1-1: Retrieve Context Source Identity Information

Table 6.33.3.1-1: Retrieve Context Source Identity Information request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	ContextSourceIdentity	1	200 No Content	A response body containing the JSON-LD representation of the Context Source Identity Information.
	ProblemDetails (see IETF RFC 7807 [10])	1	501 Not Implemented	It is used by Context Sources to indicate that retrieval of the Context Source information is unsupported see clause 5.15.1.4.

6.34 Resource: entityMaps

6.34.1 Description

This resource represents the Entity maps in an NGSI-LD system.

6.34.2 Resource definition


Resource URI:

- /entityMaps/

6.34.3 Resource methods

6.34.3.1 GET

This method is associated to the operation "Create EntityMap for Query Entities" and shall exhibit the behaviour defined by clause 5.14.4, providing an EntityMap as part of the HTTP response payload body. In addition to this method, an alternative way to perform "Create EntityMap for Query Entities " operations via POST is defined in clause 6.34.3.2. Figure 6.34.3.1-1 shows the Create EntityMap for Query Entities interaction.

**Figure 6.34.3.1-1: Create EntityMap for Query Entities interaction**

The URL parameters that shall be supported by implementations are the same as those for Query Entities and can be found in Table 6.4.3.2-1. Table 6.34.3.1-1 describes the request body and possible responses.

Table 6.34.3.1-1: Create EntityMap for Query Entities request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	EntityMap	1	201 Created	A response body containing the entityMap with the identifiers of the Entities matching the query. The HTTP response shall include an "NGSILD-EntityMap" HTTP header that contains the resource URI of the EntityMap resource created in the operation.
	EntityMap	1	203 Non-Authoritative Information	As above, but returning an altered response body , amended to conform to a specific version of the NGSI-LD specification as mandated in clause 4.3.6.8. The response shall also include a "Preference-Applied" HTTP header set to "ngsi-ld=<conformant-version>".
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	501 Not Implemented	It is used by Registered Context Sources to indicate that the data format of the request is unsupported see clause 6.3.7.

6.34.3.2 POST

This method is associated to the operation "Create EntityMap for Query Entities" and shall exhibit the behaviour defined by clause 5.14.4. Figure 6.34.3.2-1 shows the operation interaction and Table 6.34.3.2-1 describes the request body and possible responses.




Figure 6.34.3.2-1: Create EntityMap for Query Entities via POST interaction

Table 6.34.3.2-1: Create EntityMap for Query Entities via POST request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Query	1	Payload body in the request contains a JSON-LD object which represents the query to be performed.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	EntityMap	1	201 Created	A response body containing the entityMap with the identifiers of the Entities matching the query. The HTTP response shall include an "NGSILD-EntityMap" HTTP header that contains the resource URI of the EntityMap resource created in the operation.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.35 Resource: temporal/entityMaps

6.35.1 Description

This resource is used for creating entityMaps based on temporal queries in an NGSI-LD system.

6.35.2 Resource definition


Resource URI:

- /temporal/entityMaps/

6.35.3 Resource methods

6.35.3.1 GET

This method is associated to the operation "Create EntityMap for Query Temporal Evolution of Entities" and shall exhibit the behaviour defined by clause 5.14.5, an EntityMap as part of the HTTP response payload body. In addition to this method, an alternative way to perform "Create EntityMap for Query Temporal Evolution of Entities" operations via POST is defined in clause 6.35.3.2. Figure 6.35.3.1-1 shows this interaction.

**Figure 6.35.3.1-1: Create EntityMap for Query Temporal Evolution of Entities interaction**

The URL parameters that shall be supported by implementations are the same as those for Query Temporal Evolution of Entities and can be found in Table 6.18.3.2-1. Table 6.35.3.1-1 describes the request body and possible responses.

Table 6.35.3.1-1: Create EntityMap for Query Temporal Evolution of Entities request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	EntityMap	1	201 Created	A response body containing the entityMap with the identifiers of the Entities matching the query. The HTTP response shall include an "NGSILD-EntityMap" HTTP header that contains the resource URI of the EntityMap resource created in the operation.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.35.3.2 POST

This method is associated to the operation "Create EntityMap for Query Temporal Evolution of Entities" and shall exhibit the behaviour defined by clause 5.14.5. Figure 6.35.3.2-1 shows the operation interaction and Table 6.35.3.2-1 describes the request body and possible responses.




Figure 6.35.3.2-1: Create EntityMap for Query Temporal Evolution of Entities via POST interaction

Table 6.35.3.2-1: Create EntityMap for Query Temporal Evolution of Entities via POST request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Query	1	Payload body in the request contains a JSON-LD object which represents the query to be performed.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	EntityMap	1	201 Created	A response body containing the entityMap with the identifiers of the Entities matching the query. The HTTP response shall include an "NGSILD-EntityMap" HTTP header that contains the resource URI of the EntityMap resource created in the operation.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.36 Resource: snapshots

6.36.1 Description

This resource represents Snapshots available in an NGSI-LD system.

6.36.2 Resource definition

Resource URI:

- /snapshots/

6.36.3 Resource methods

6.36.3.1 POST

This method is associated to the operation "Create Snapshot" and shall exhibit the behaviour defined by clause 5.16.1. Figure 6.36.3.1-1 shows the operation interaction and Table 6.36.3.1-1 describes the request body and possible responses.




Figure 6.36.3.1-1: Create Snapshot interaction

Table 6.36.3.1-1: Create Snapshot request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Snapshot	1	Payload body in the request contains a JSON-LD object which represents parameters for the snapshot to be created.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	201 Created	The HTTP response shall include a "Location" HTTP header that contains the relative path of the created snapshot.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	409 Conflict	It is used to indicate that a Snapshot having the Snapshot identifier included in the request body already exists, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

6.36.3.2 DELETE

This method is associated to the operation "Purge Snapshots" and shall exhibit the behaviour defined by clause 5.16.7. Figure 6.36.3.2-1 shows the Purge Snapshot interaction.




Figure 6.36.3.2-1: Purge Snapshots interaction

The URL parameters that shall be supported by implementations are those defined in Table 6.36.3.2-1 and Table 6.36.3.2-2 describes the request body and possible responses.

Table 6.36.3.2-1: Purge Snapshots URL parameters

Name	Data Type	Cardinality	Remarks
q	String	1	Query as per clause 4.9, restricted to members of the Snapshot data type.

Table 6.36.3.2-2: Purge Snapshots request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
	BatchOperationResult	1	207 Multi-Status	If some or all of the Snapshots have not been successfully deleted, or did not exist, a response body containing the result of each operation contained in the batch is returned in a <i>BatchOperationResult</i> structure. It contains two arrays. The first array (<i>success</i>) contains the URIs of the successfully deleted Snapshots, while the second array (<i>errors</i>) contains information about the error for each of the Snapshots that could not be deleted. There is no restriction as to the order of the Snapshots IDs in the arrays.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided a filter not matching any Snapshots, see clause 6.3.2.

6.37 Resource: snapshots/{snapshotId}

6.37.1 Description

This resource represents a snapshot in an NGSI-LD system.

6.37.2 Resource definition

Resource URI:

- /snapshots/{snapshotId}

Resource URI variables for this resource are defined in Table 6.37.2-1.

Table 6.37.2-1: URI variables

Name	Definition
snapshotId	Id (URI) of the Snapshot whose status is to be retrieved or updated, or the Snapshot to be deleted.

6.37.3 Resource methods

6.37.3.1 GET

This method is associated to the operation "Retrieve Snapshot Status" and shall exhibit the behaviour defined by clause 5.16.3. The Snapshot identifier is the value of the resource URI variable "snapshotId". Figure 6.37.3.1-1 shows the Retrieve Snapshot Status interaction and Table 6.37.3.1-1 describes the request body and possible responses.




Figure 6.37.3.1-1: Retrieve Snapshot Status interaction

Table 6.37.3.1-1: Retrieve Snapshot Status request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	Snapshot	1	200 OK	A response body containing the JSON-LD representation of the target Snapshot status.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an EntityMap identifier not known to the system, see clause 6.3.2.

6.37.3.2 PATCH

This method is associated to the operation "Update Snapshot Status" and shall exhibit the behaviour defined by clause 5.16.4. The Snapshot identifier is the value of the resource URI variable "snapshotId". Figure 6.37.3.2-1 shows the Update Snapshot Status interaction and Table 6.37.3.2-1 describes the request body and possible responses.




Figure 6.37.3.2-1: Update Snapshot Status interaction

Table 6.37.3.2-1: Update Snapshot Status request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Snapshot Fragment	1	Payload body in the request contains a JSON-LD object which represents the Snapshot fragment with which the Snapshot status is to be updated.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	Snapshot	1	200 OK	A response body containing the JSON-LD representation of the updated Snapshot status.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided an EntityMap identifier not known to the system, see clause 6.3.2.

6.37.3.3 DELETE

This method is associated to the operation "Delete Snapshot" and shall exhibit the behaviour defined by clause 5.16.5. The Snapshot identifier is the value of the resource URI variable "snapshotId". Figure 6.37.3.3-1 shows the Delete Snapshot interaction and Table 6.37.3.3-1 describes the request body and possible responses.




Figure 6.37.3.3-1: Delete Snapshot interaction

Table 6.37.3.3-1: Delete Snapshot request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	N/A	N/A		
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	204 No Content	
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided a snapshot identifier not known to the system, see clause 6.3.2.

6.38 Resource: snapshots/{snapshotId}/clone

6.38.1 Description

This resource enables the cloning of a snapshot in an NGSI-LD system.

6.38.2 Resource definition

Resource URI:

- /snapshots/{snapshotId}/clone

Resource URI variables for this resource are defined in Table 6.38.2-1.

Table 6.38.2-1: URI variables

Name	Definition
snapshotId	Id (URI) of the Snapshot to be queried or deleted.

6.38.3 Resource methods

6.38.3.1 POST

This method is associated to the operation "Clone Snapshot" and shall exhibit the behaviour defined by clause 5.16.2. The Snapshot identifier is the value of the resource URI variable "snapshotId". Figure 6.38.3.1-1 shows the Clone Snapshot interaction and Table 6.38.3.1-1 describes the request body and possible responses.


**Figure 6.38.3.1-1: Clone Snapshot interaction**

Table 6.38.3.1-1: Clone Snapshot request body and possible responses

Request Body	Data Type	Cardinality	Remarks	
	Snapshot	1	Payload body in the request contains a JSON-LD object which represents parameters for the cloned snapshot.	
Response Body	Data Type	Cardinality	Response Codes	Remarks
	N/A	N/A	201 Created	The HTTP response shall include a "Location" HTTP header that contains the relative path of the cloned snapshot.
	ProblemDetails (see IETF RFC 7807 [10])	1	400 Bad Request	It is used to indicate that the request or its content is incorrect, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.
	ProblemDetails (see IETF RFC 7807 [10])	1	404 Not Found	It is used when a client provided a Snapshot identifier (URI) not known to the system, see clause 6.3.2.
	ProblemDetails (see IETF RFC 7807 [10])	1	409 Conflict	It is used to indicate that the Snapshot having the Snapshot identifier included in the request body already exists, see clause 6.3.2. In the returned <i>ProblemDetails</i> structure, the <i>detail</i> attribute should convey more information about the error.

7 API MQTT Notification Binding

7.1 Introduction

This clause defines the optional support of the NGSI-LD API for sending notifications via the MQTT protocol [24] and [25]. The subscriptions are handled using the HTTP binding as described in clause 6, but instead of an HTTP endpoint, an MQTT endpoint is provided.

7.2 Notification behaviour

In case a subscription received via HTTP specifies an MQTT endpoint in the "notification.endpoint.uri" member of the subscription structure (defined by clauses 5.2.12, 5.2.14 and 5.2.15), and the MQTT notification binding is supported by the NGSI-LD implementation, notifications related to this subscription shall be sent via the MQTT protocol.

The syntax of an MQTT endpoint URI is `mqtt[s]://[<username>][:<password>]@<host>[:<port>]/<topic>[/<subtopic>]*` and follows an existing convention for representing an MQTT endpoint as a URI [i.19].

Username and password can be optionally specified as part of the endpoint URI. If the port is not explicitly specified, the default MQTT port is 1883 for MQTT over TCP and 8883 for mqtts, i.e. Secure MQTT over TLS. MQTT supports the structuring of topics as a hierarchy with any number of subtopic levels, which can be specified as part of the endpoint URI.

In MQTT, all non-protocol information has to be included into the MQTT message. This means that the actual notification as specified in clause 5.3.1, as well as additional information like MIME type, possibly the link to the *@context* and additional user-specified information, which in the HTTP case is provided as headers, has to be included into the MQTT message. The MQTT notification message shall be provided as a JSON Object with the two elements "metadata" and "body". The actual notification, as specified in clause 5.3.1 is the value of "body", whereas any additional information is provided as key-value pairs in "metadata".

For the MQTT protocol, there are currently two versions supported, MQTTv3.1.1 [24] and MQTTv5.0 [25]. Also, there are three levels of quality of service:

- at most once (0);
- at least once (1); and
- exactly once (2).

These can be specified in the subscription as part of the optional array of KeyValuePair type (defined by clause 5.2.22) "notification.endpoint.notifierInfo". The MQTT protocol parameters can be found in Table 7.2-1. If not present, the given default value is used.

Table 7.2-1: Protocol parameters for MQTT in notifierInfo

Key	Possible Values	Default	Source	Description
MQTT-QoS	0, 1, 2	0	Subscription's <i>notification.endpoint.notifierInfo</i>	MQTT Quality of service, at most once (0), at least once (1) and exactly once (2)
MQTT-Version	mqtt3.1.1, mqtt5.0	mqtt5.0	Subscription's <i>notification.endpoint.notifierInfo</i>	Version of MQTT protocol

The MIME type associated with the notification shall be "application/json" by default. However, this can be changed to "application/ld+json" by means of the "endpoint.accept" member. The MIME type is specified as Content-Type in the "metadata" element of the MQTT message. If the target MIME type is "application/json" then the reference to the JSON-LD @context is provided as Link in the "metadata" element of the MQTT message, following the specification of the HTTP Link header as mandated by the JSON-LD specification [2], section 6.2 (to the default JSON-LD @context if none available). Table 7.2-2 lists these "receiver side" metadata parameters.

Table 7.2-2: Parameters for MQTT in "metadata"

Key	Possible Values	Default	Source	Description
Content-Type	"application/json", "application/ld+json"	"application/json"	Subscription's <i>notification.endpoint.accept</i>	MIME type of the notification included in the "body" element of the MQTT message
Link	Same format as specified in JSON-LD specification [2], section 6.2 for the HTTP Link header		Link Header provided in Subscription	Contains the reference to the @context in case Content-Type is "application/json". Example: <http://myhost.org/mycontext>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

Additionally, if the optional array of KeyValuePair type (defined by clause 5.2.22) *notification.endpoint.receiverInfo* of the subscription is present, then a new entry for each member named "key" of the key, value pairs that make up the array shall be generated and added to the "metadata" element of the MQTT message. The content of each entry shall be set equal to the content of the corresponding "value" member of the *KeyValuePair*.

Annex A (normative): NGSI-LD identifier considerations

A.1 Introduction

The purpose of identifiers is to allow uniquely identifying NGSI-LD elements (Entities, Context Subscriptions or Context Source Registrations) within an NGSI-LD system. This annex is intended to clarify the different issues around the design of identifiers in NGSI-LD.

A.2 Entity identifiers

In order to enable the participation of NGSI-LD in linked data scenarios, all Entities are identified by **URIs**. If those URIs are expected to participate in external linked data relationships they **should** be dereferenceable.

It is noteworthy that the identifier from the point of view of NGSI-LD is different from the inherent identifier that a specific Entity may have. For instance, an NGSI-LD Entity of Type "Vehicle" may have a Property named *licencePlateNumber*, which it is actually a unique identifier from the point of view of the Entity domain, as it uniquely identifies the specific vehicle instance. However, from the point of view of the NGSI-LD system, it may have another identifier which might or might not include such licence plate number identifier.

A.3 NGSI-LD namespace

NGSI-LD defines a specific URN [9] namespace intended to help API users to design readable, clean and simple identifiers. As it is based on URNs, the usage of this identification approach is not recommended when dereferenceable URIs are needed (fully-fledged linked data scenarios).

The referred namespace is defined as follows (to be registered with IANA):

- namespace identifier: NID = "ngsi-ld"
- namespace specific string: NSS = EntityType " : " EntityIdentificationString

EntityType shall be an Entity Type name which can be expanded to a URI as per the *@context*.

EntityIdentificationString shall be a string that allows uniquely identifying the subject Entity in combination with the other items being part of the NSS.

EXAMPLE: urn:ngsi-ld:Person:28976543.

It is recommended that applications use this URN namespace when applicable.

In general, the URN specification defines namespace equivalence in a case-insensitive manner, however it is assumed that context-broker implementations shall always use lowercase letters in namespaces where they have a choice in case, unless there is a strong reason otherwise. Restricting the namespace prefix to lower case `urn:ngsi-ld:` can improve caching and retrieval, since this ensures since alphabetic characters within the namespace specific string are always consistent.

Annex B (normative):

Core NGSI-LD @context definition

Below is the definition of the Core NGSI-LD *@context* which shall be supported by implementations.

Such definition has been tested using [i.7].

```
{
  "@context": {
    "@version": 1.1,
    "@protected": true,
    "ngsi-ld": "https://uri.etsi.org/ngsi-ld/",
    "geojson": "https://purl.org/geojson/vocab#",
    "id": "@id",
    "type": "@type",
    "Attribute": "ngsi-ld:Attribute",
    "AttributeList": "ngsi-ld:AttributeList",
    "ContextSourceIdentity": "ngsi-ld:ContextSourceIdentity",
    "ContextSourceNotification": "ngsi-ld:ContextSourceNotification",
    "ContextSourceRegistration": "ngsi-ld:ContextSourceRegistration",
    "Date": "ngsi-ld:Date",
    "DateTime": "ngsi-ld:DateTime",
    "EntityType": "ngsi-ld:EntityType",
    "EntityTypeInfo": "ngsi-ld:EntityTypeInfo",
    "EntityTypeList": "ngsi-ld:EntityTypeList",
    "ExecutionResultDetails": "ngsi-ld:ExecutionResultDetails",
    "Feature": "geojson:Feature",
    "FeatureCollection": "geojson:FeatureCollection",
    "GeoProperty": "ngsi-ld:GeoProperty",
    "GeometryCollection": "geojson:GeometryCollection",
    "JsonProperty": "ngsi-ld:JsonProperty",
    "LanguageProperty": "ngsi-ld:LanguageProperty",
    "LineString": "geojson:LineString",
    "ListProperty": "ngsi-ld:ListProperty",
    "ListRelationship": "ngsi-ld:ListRelationship",
    "MultiLineString": "geojson:MultiLineString",
    "MultiPoint": "geojson:MultiPoint",
    "MultiPolygon": "geojson:MultiPolygon",
    "Notification": "ngsi-ld:Notification",
    "Point": "geojson:Point",
    "Polygon": "geojson:Polygon",
    "Property": "ngsi-ld:Property",
    "Relationship": "ngsi-ld:Relationship",
    "Snapshot": "ngsi-ld:Snapshot",
    "SnapshotNotification": "ngsi-ld:SnapshotNotification",
    "Subscription": "ngsi-ld:Subscription",
    "TemporalProperty": "ngsi-ld:TemporalProperty",
    "Time": "ngsi-ld:Time",
    "VocabProperty": "ngsi-ld:VocabProperty",
    "accept": "ngsi-ld:accept",
    "aggrParams": "ngsi-ld:aggrParams",
    "aggrMethods": "ngsi-ld:aggrMethods",
    "aggrPeriodDuration": "ngsi-ld:aggrPeriodDuration",
    "attributeCount": "attributeCount",
    "attributeDetails": "attributeDetails",
    "attributeList": {
      "@id": "ngsi-ld:attributeList",
      "@type": "@vocab"
    },
    "attributeName": {
      "@id": "ngsi-ld:attributeName",
      "@type": "@vocab"
    },
    "attributeNames": {
      "@id": "ngsi-ld:attributeNames",
      "@type": "@vocab"
    },
    "attributeTypes": {
      "@id": "ngsi-ld:attributeTypes",
      "@type": "@vocab"
    },
    "attributes": {
      "@id": "ngsi-ld:attributes",
      "@type": "@vocab"
    }
  },
}
```

```

"attrs": "ngsi-ld:attrs",
"avg": {
  "@id": "ngsi-ld:avg",
  "@container": "@list"
},
"bbox": {
  "@container": "@list",
  "@id": "geojson:bbox"
},
"cacheDuration": "ngsi-ld:cacheDuration",
"collation": "ngsi-ld:collation",
"containedBy": "ngsi-ld:isContainedBy",
"contextSourceAlias": "ngsi-ld:contextSourceAlias",
"contextSourceExtras": {
  "@id": "ngsi-ld:contextSourceExtras",
  "@type": "@json"
},
"contextSourceInfo": "ngsi-ld:contextSourceInfo",
"contextSourceTimeAt": {
  "@id": "ngsi-ld:contextSourceTimeAt",
  "@type": "DateTime"
},
"contextSourceUptime": "ngsi-ld:contextSourceUptime",
"cooldown": "ngsi-ld:cooldown",
"coordinates": {
  "@container": "@list",
  "@id": "geojson:coordinates"
},
"createdAt": {
  "@id": "ngsi-ld:createdAt",
  "@type": "DateTime"
},
"csf": "ngsi-ld:csf",
"data": "ngsi-ld:data",
"dataset": {
  "@id": "ngsi-ld:hasDataset",
  "@container": "@index"
},
"datasetId": {
  "@id": "ngsi-ld:datasetId",
  "@type": "@id"
},
"deletedAt": {
  "@id": "ngsi-ld:deletedAt",
  "@type": "DateTime"
},
"description": "http://purl.org/dc/terms/description",
"detail": "ngsi-ld:detail",
"distinctCount": {
  "@id": "ngsi-ld:distinctCount",
  "@container": "@list"
},
"endAt": {
  "@id": "ngsi-ld:endAt",
  "@type": "DateTime"
},
"endTimeAt": {
  "@id": "ngsi-ld:endTimeAt",
  "@type": "DateTime"
},
"endpoint": "ngsi-ld:endpoint",
"entities": "ngsi-ld:entities",
"entity": "ngsi-ld:entity",
"entityCount": "ngsi-ld:entityCount",
"entityId": {
  "@id": "ngsi-ld:entityId",
  "@type": "@id"
},
"entityList": {
  "@id": "ngsi-ld:entityList",
  "@container": "@list"
},
"entityMap": "ngsi-ld:hasEntityMap",
"error": "ngsi-ld:error",
"errors": "ngsi-ld:errors",
"expandValues": "ngsi-ld:expandValues",
"expiresAt": {
  "@id": "ngsi-ld:expiresAt",

```

```

    "@type": "DateTime"
  },
  "features": {
    "@container": "@set",
    "@id": "geojson:features"
  },
  "format": "ngsi-ld:format",
  "geoQ": "ngsi-ld:geoQ",
  "geometry": "geojson:geometry",
  "geoproperty": "ngsi-ld:geoproperty",
  "georel": "ngsi-ld:georel",
  "idPattern": "ngsi-ld:idPattern",
  "information": "ngsi-ld:information",
  "instanceId": {
    "@id": "ngsi-ld:instanceId",
    "@type": "@id"
  },
  "isActive": "ngsi-ld:isActive",
  "join": "ngsi-ld:join",
  "joinLevel": "ngsi-ld:hasJoinLevel",
  "json": {
    "@id": "ngsi-ld:hasJSON", "@type": "@json"
  },
  "jsonKeys": "ngsi-ld:jsonKeys",
  "jsons": {
    "@id": "ngsi-ld:jsons",
    "@container": "@list"
  },
  "key": "ngsi-ld:hasKey",
  "lang": "ngsi-ld:lang",
  "languageMap": {
    "@id": "ngsi-ld:hasLanguageMap",
    "@container": "@language"
  },
  "languageMaps": {
    "@id": "ngsi-ld:hasLanguageMaps",
    "@container": "@list"
  },
  "langString": "http://www.w3.org/1999/02/22-rdf-syntax-ns#langString",
  "lastFailure": {
    "@id": "ngsi-ld:lastFailure",
    "@type": "DateTime"
  },
  "lastNotification": {
    "@id": "ngsi-ld:lastNotification",
    "@type": "DateTime"
  },
  "lastSuccess": {
    "@id": "ngsi-ld:lastSuccess",
    "@type": "DateTime"
  },
  "lastUsedAt": {
    "@id": "ngsi-ld:lastUsedAt",
    "@type": "DateTime"
  },
  "linkedMaps": "ngsi-ld:linkedMaps",
  "localOnly": "ngsi-ld:localOnly",
  "location": "ngsi-ld:location",
  "management": "ngsi-ld:management",
  "managementInterval": "ngsi-ld:managementInterval",
  "max": {
    "@id": "ngsi-ld:max",
    "@container": "@list"
  },
  "min": {
    "@id": "ngsi-ld:min",
    "@container": "@list"
  },
  "mode": "ngsi-ld:mode",
  "modifiedAt": {
    "@id": "ngsi-ld:modifiedAt",
    "@type": "DateTime"
  },
  "ngsildproof": {
    "@id": "ngsi-ld:ngsildproof",
    "@context": [{
      "entityIdSealed": "ngsi-ld:entityIdSealed",
      "entityTypeSealed": {

```

```

        "@id": "ngsi-ld:entityTypeSealed",
        "@type": "@vocab"
    },
    ],
    "https://w3id.org/security/data-integrity/v2"
]
},
"notification": "ngsi-ld:notification",
"notificationTrigger": "ngsi-ld:notificationTrigger",
"notifiedAt": {
    "@id": "ngsi-ld:notifiedAt",
    "@type": "DateTime"
},
"notifierInfo": "ngsi-ld:notifierInfo",
"notUpdated": "ngsi-ld:notUpdated",
"object": {
    "@id": "ngsi-ld:hasObject",
    "@type": "@id"
},
"objectList": {
    "@id": "ngsi-ld:hasObjectList",
    "@container": "@list"
},
"objectLists": {
    "@id": "ngsi-ld:hasObjectLists",
    "@container": "@list"
},
"objects": {
    "@id": "ngsi-ld:hasObjects",
    "@container": "@list"
},
"objectType": {
    "@id": "ngsi-ld:hasObjectType",
    "@type": "@vocab"
},
"observationInterval": "ngsi-ld:observationInterval",
"observationSpace": "ngsi-ld:observationSpace",
"observedAt": {
    "@id": "ngsi-ld:observedAt",
    "@type": "DateTime"
},
"omit": "ngsi-ld:omit",
"operations": "ngsi-ld:operations",
"operationSpace": "ngsi-ld:operationSpace",
"orderBy": {
    "@container": "@list",
    "@id": "ngsi-ld:orderBy"
},
"ordering": "ngsi-ld:ordering",
"pick": "ngsi-ld:pick",
"previousJson": {
    "@id": "ngsi-ld:hasPreviousJson",
    "@type": "@json"
},
"previousLanguageMap": {
    "@id": "ngsi-ld:hasPreviousLanguageMap",
    "@container": "@language"
},
"previousObject": {
    "@id": "ngsi-ld:hasPreviousObject",
    "@type": "@id"
},
"previousObjectList": {
    "@id": "ngsi-ld:hasPreviousObjectList",
    "@container": "@list"
},
"previousValue": "ngsi-ld:hasPreviousValue",
"previousValueList": {
    "@id": "ngsi-ld:hasPreviousValueList",
    "@container": "@list"
},
"previousVocab": {
    "@id": "ngsi-ld:hasPreviousVocab",
    "@type": "@vocab"
},
"problemDetails": {
    "@id": "ngsi-ld:problemDetails",
    "@type": "@json"
}

```



```

},
"properties": "geojson:properties",
"propertyNames": {
  "@id": "ngsi-ld:propertyNames",
  "@type": "@vocab"
},
"q": "ngsi-ld:q",
"reason": "ngsi-ld:reason",
"receiverInfo": "ngsi-ld:receiverInfo",
"refreshRate": "ngsi-ld:refreshRate",
"registrationId": "ngsi-ld:registrationId",
"registrationName": "ngsi-ld:registrationName",
"relationshipNames": {
  "@id": "ngsi-ld:relationshipNames",
  "@type": "@vocab"
},
"resultStatus": "ngsi-ld:resultStatus",
"scope": "ngsi-ld:scope",
"scopeQ": "ngsi-ld:scopeQ",
"showChanges": "ngsi-ld:showChanges",
"snapshotId": "ngsi-ld:snapshotId",
"snapshotLifetime": "ngsi-ld:snapshotLifetime",
"snapshotPriority": "ngsi-ld:snapshotPriority",
"snapshotQueries": {
  "@id": "ngsi-ld:snapshotQueries",
  "@container": "@list"
},
"snapshotQueriesDetails": {
  "@id": "ngsi-ld:snapshotQueriesDetails",
  "@container": "@list"
},
"snapshotStatus": "ngsi-ld:snapshotStatus",
"snapshotTemporalQueries": {
  "@id": "ngsi-ld:snapshotTemporalQueries",
  "@container": "@list"
},
"snapshotTemporalQueriesDetails": {
  "@id": "ngsi-ld:snapshotTemporalQueriesDetails",
  "@container": "@list"
},
"startAt": {
  "@id": "ngsi-ld:startAt",
  "@type": "DateTime"
},
"status": "ngsi-ld:status",
"stddev": {
  "@id": "ngsi-ld:stddev",
  "@container": "@list"
},
"subscriptionId": {
  "@id": "ngsi-ld:subscriptionId",
  "@type": "@id"
},
"subscriptionName": "ngsi-ld:subscriptionName",
"success": {
  "@id": "ngsi-ld:success",
  "@type": "@id"
},
"sum": {
  "@id": "ngsi-ld:sum",
  "@container": "@list"
},
"sumsq": {
  "@id": "ngsi-ld:sumsq",
  "@container": "@list"
},
"sysAttrs": "ngsi-ld:sysAttrs",
"temporalQ": "ngsi-ld:temporalQ",
"tenant": {
  "@id": "ngsi-ld:tenant",
  "@type": "@id"
},
"throttling": "ngsi-ld:throttling",
"timeAt": {
  "@id": "ngsi-ld:timeAt",
  "@type": "DateTime"
},
"timeInterval": "ngsi-ld:timeInterval",

```

```

"timeout": "ngsi-ld:timeout",
"timeproperty": "ngsi-ld:timeproperty",
"timereel": "ngsi-ld:timereel",
"timesFailed": "ngsi-ld:timesFailed",
"timesSent": "ngsi-ld:timesSent",
"title": "http://purl.org/dc/terms/title",
"totalCount": {
  "@id": "ngsi-ld:totalCount",
  "@container": "@list"
},
"triggerReason": "ngsi-ld:triggerReason",
"typeList": {
  "@id": "ngsi-ld:typeList",
  "@type": "@vocab"
},
"typeName": {
  "@id": "ngsi-ld:typeName",
  "@type": "@vocab"
},
"typeNames": {
  "@id": "ngsi-ld:typeNames",
  "@type": "@vocab"
},
"unchanged": "ngsi-ld:unchanged",
"unitCode": "ngsi-ld:unitCode",
"updated": "ngsi-ld:updated",
"uri": "ngsi-ld:uri",
"value": "ngsi-ld:hasValue",
"valueList": {
  "@id": "ngsi-ld:hasValueList",
  "@container": "@list"
},
"valueLists": {
  "@id": "ngsi-ld:hasValueLists",
  "@container": "@list"
},
"values": {
  "@id": "ngsi-ld:hasValues",
  "@container": "@list"
},
"valueType": {
  "@id": "ngsi-ld:hasValueType",
  "@type": "@vocab"
},
"vocab": {
  "@id": "ngsi-ld:hasVocab",
  "@type": "@vocab"
},
"vocabs": {
  "@id": "ngsi-ld:hasVocabs",
  "@container": "@list"
},
"watchedAttributes": {
  "@id": "ngsi-ld:watchedAttributes",
  "@type": "@vocab"
},
"@vocab": "https://uri.etsi.org/ngsi-ld/default-context/"
}
}

```

NOTE: Implementers can take advantage of prefixed terms, i.e. in the form `ngsi-ld:term`, to provide a terser representation of the Core *@context*.

Annex C (informative): Examples of using the API

C.1 Introduction

This annex is informative and is intended to show in action the JSON-LD representation defined by NGSI-LD.

JSON representations of the examples shown in this annex can be found at [i.15].

C.2 Entity Representation

C.2.1 Property Graph

Figure C.2.1-1 shows a diagram representing a property graph to be used for the examples discussed in this clause.




Figure C.2.1-1: Reference example

As per the algorithms described above and as per the rules for generating the JSON-LD representation of NGSI-LD entities the above graph will result in the following JSON-LD representations. The syntax has been checked using the JSON-LD Playground tool [i.5].

C.2.2 Vehicle Entity

Normalized Representation

The normalized representation is a lossless representation of an Entity, where every *Property* is defined by a *type* and a *value* and every *Relationship* is defined by a *type* and an *object*.

Below there is a representation of an Entity of Type "Vehicle". It can be observed that the *@context* is composed of different parts, namely the Core *@context* and several vocabulary-specific *@contexts*.

It is noteworthy that the *@context* corresponding to the Parking domain is included as it is referenced through the *isParked* Relationship.

```
{
  "id": "urn:ngsi-ld:Vehicle:A4567",
  "type": "Vehicle",
  "brandName": {
    "type": "Property",
    "value": "Mercedes"
  },
  "street": {
    "type": "LanguageProperty",
    "languageMap": {
      "fr": "Grand Place",
      "nl": "Grote Markt"
    }
  },
  "isParked": {
    "type": "Relationship",
    "objectType": "OffStreetParking",
    "object": "urn:ngsi-ld:OffStreetParking:Downtown1",
    "observedAt": "2017-07-29T12:00:04Z",
    "providedBy": {
      "type": "Relationship",
      "object": "urn:ngsi-ld:Person:Bob"
    }
  },
  "category": {
    "type": "VocabProperty",
    "vocab": "non-commercial"
  },
  "tyreTreadDepths": {
    "type": "ListProperty",
    "valueList": [300, 300, 120, 6],
    "valueType": "Integer",
    "unitCode": "MMT"
  },
  "passengers": {
    "type": "Relationship",
    "objectType": "Person",
    "object": ["urn:ngsi-ld:Person:Alice", "urn:ngsi-ld:Person:Bob"]
  },
  "route": {
    "type": "ListRelationship",
    "objectType": "City",
    "objectList": [
      {"object": "urn:ngsi-ld:City:Antwerp"},
      {"object": "urn:ngsi-ld:City:Rotterdam"},
      {"object": "urn:ngsi-ld:City:Amsterdam"}
    ]
  },
  "@context": [
    "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
    "http://example.org/ngsi-ld/latest/vehicle.jsonld",
    "http://example.org/ngsi-ld/latest/parking.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}
```

Normalized Representation when inline Linked Entity retrieval is used

When inline Linked Entity retrieval (see clause 4.5.23.2) is specified, any *Relationships* which target Entities stored locally or include an *objectType* Attribute are returned in an expanded format. Attributes of type "Relationship" are returned with an additional *entity* sub-Attribute, which holds the normalized Linked Entity data corresponding to the *Relationship's* target *object* URI. Attributes of type "ListRelationship" are returned with an additional *entityList* sub-Attribute which in turn holds an ordered array of the normalized Linked Entities corresponding to the target "objectList" URIs.

```
{
  "id": "urn:ngsi-ld:Vehicle:A4567",
  "type": "Vehicle",
  "brandName": {
    "type": "Property",
    "value": "Mercedes"
  },
  "street": {
    "type": "LanguageProperty",
    "languageMap": {
      "fr": "Grand Place",
      "nl": "Grote Markt"
    }
  },
  "isParked": {
    "type": "Relationship",
    "objectType": "OffStreetParking",
    "object": "urn:ngsi-ld:OffStreetParking:Downtown1",
    "entity": {
      "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
      "type": "OffStreetParking",
      "name": {
        "value": "Top Parking",
        "type": "Property"
      },
      "operatedBy": {
        "object": "urn:ngsi-ld:Company:BigParkingCorp",
        "type": "Relationship"
      },
    },
    "observedAt": "2017-07-29T12:00:04Z",
    "providedBy": {
      "object": "urn:ngsi-ld:Person:Bob"
    }
  },
  "category": {
    "type": "VocabProperty",
    "vocab": "non-commercial"
  },
  "tyreTreadDepths": {
    "type": "ListProperty",
    "valueList": [300, 300, 120, 6],
    "valueType": "Integer",
    "unitCode": "MMT"
  },
  "passengers": {
    "type": "Relationship",
    "objectType": "Person",
    "object": ["urn:ngsi-ld:Person:Alice", "urn:ngsi-ld:Person:Bob"],
    "entity": [
      {
        "id": "urn:ngsi-ld:Person:Alice",
        "type": "Person",
        "name": {
          "value": "Alice",
          "type": "Property"
        }
      },
      {
        "id": "urn:ngsi-ld:Person:Bob",
        "type": "Person",
        "name": {
          "value": "Bob",
          "type": "Property"
        }
      }
    ]
  }
}
```

```

    },
    "route": {
      "type": "ListRelationship",
      "objectType": "City",
      "objectList": [
        { "object": "urn:ngsi-ld:City:Antwerp" },
        { "object": "urn:ngsi-ld:City:Rotterdam" },
        { "object": "urn:ngsi-ld:City:Amsterdam" }
      ],
      "entityList": [
        {
          "id": "urn:ngsi-ld:City:Antwerp",
          "type": "City",
          "name": {
            "value": "Antwerp",
            "type": "Property"
          }
        },
        {
          "id": "urn:ngsi-ld:City:Rotterdam",
          "type": "City",
          "name": {
            "value": "Rotterdam",
            "type": "Property"
          }
        },
        {
          "id": "urn:ngsi-ld:City:Amsterdam",
          "type": "City",
          "name": {
            "value": "Amsterdam",
            "type": "Property"
          }
        }
      ]
    },
    "@context": [
      "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "http://example.org/ngsi-ld/latest/parking.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  }
}

```

Normalized Representation when flattened Linked Entity retrieval is used

When flattened Linked Entity retrieval (see clause 4.5.23.3) is specified, an array of normalized Entities is returned. Whenever a *Relationship* Attribute targets an Entity stored locally or includes an *objectType*, an additional normalized Linked Entity holding data corresponding to the *Relationship's* target *object* URI is appended to the array. For Attributes of type "ListRelationship", an array of normalized Linked Entities is appended, which hold the data corresponding to each of the target URIs found within its *objectList*.

```

[
  {
    "id": "urn:ngsi-ld:Vehicle:A4567",
    "type": "Vehicle",
    "brandName": {
      "type": "Property",
      "value": "Mercedes"
    },
    "street": {
      "type": "LanguageProperty",
      "languageMap": {
        "fr": "Grand Place",
        "nl": "Grote Markt"
      }
    },
    "isParked": {
      "type": "Relationship",
      "objectType": "OffStreetParking",
      "object": "urn:ngsi-ld:OffStreetParking:Downtown1",
      "observedAt": "2017-07-29T12:00:04Z",
      "providedBy": {
        "type": "Relationship",
        "object": "urn:ngsi-ld:Person:Bob"
      }
    }
  }
]

```

```

    },
    "category": {
      "type": "VocabProperty",
      "vocab": "non-commercial"
    },
    "tyreTreadDepths": {
      "type": "ListProperty",
      "valueList": [300, 300, 120, 6],
      "valueType": "Integer",
      "unitCode": "MMT"
    },
    "passengers": {
      "type": "Relationship",
      "objectType": "Person",
      "object": ["urn:ngsi-ld:Person:Alice", "urn:ngsi-ld:Person:Bob"]
    },
    "route": {
      "type": "ListRelationship",
      "objectType": "City",
      "objectList": [
        {"object": "urn:ngsi-ld:City:Antwerp"},
        {"object": "urn:ngsi-ld:City:Rotterdam"},
        {"object": "urn:ngsi-ld:City:Amsterdam"}
      ]
    },
    "@context": [
      "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "http://example.org/ngsi-ld/latest/parking.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  },
  {
    "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
    "type": "OffStreetParking",
    "name": {
      "value": "Top Parking",
      "type": "Property"
    },
    "operatedBy": {
      "object": "urn:ngsi-ld:Company:BigParkingCorp",
      "type": "Relationship"
    },
    "observedAt": "2017-07-29T12:00:04Z",
    "providedBy": {
      "object": "urn:ngsi-ld:Person:Bob"
    },
    "@context": [
      "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "http://example.org/ngsi-ld/latest/parking.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  },
  {
    "id": "urn:ngsi-ld:Person:Alice",
    "type": "Person",
    "name": {
      "value": "Alice",
      "type": "Property"
    },
    "@context": [
      "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "http://example.org/ngsi-ld/latest/parking.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  },
  {
    "id": "urn:ngsi-ld:Person:Bob",
    "type": "Person",
    "name": {
      "value": "Bob",
      "type": "Property"
    },
    "@context": [
      "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",

```

```

        "http://example.org/ngsi-ld/latest/parking.jsonld",
        "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
},
{
    "id": "urn:ngsi-ld:City:Antwerp",
    "type": "City",
    "name": {
        "value": "Antwerp",
        "type": "Property"
    },
    "@context": [
        "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
        "http://example.org/ngsi-ld/latest/vehicle.jsonld",
        "http://example.org/ngsi-ld/latest/parking.jsonld",
        "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
},
{
    "id": "urn:ngsi-ld:City:Rotterdam",
    "type": "City",
    "name": {
        "value": "Rotterdam",
        "type": "Property"
    },
    "@context": [
        "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
        "http://example.org/ngsi-ld/latest/vehicle.jsonld",
        "http://example.org/ngsi-ld/latest/parking.jsonld",
        "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
},
{
    "id": "urn:ngsi-ld:City:Amsterdam",
    "type": "City",
    "name": {
        "value": "Amsterdam",
        "type": "Property"
    },
    "@context": [
        "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
        "http://example.org/ngsi-ld/latest/vehicle.jsonld",
        "http://example.org/ngsi-ld/latest/parking.jsonld",
        "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
}
]

```

Normalized Representation when Language Filter is used

When the Language Filter (see clause 4.15) is used, Properties of type "LanguageProperty" are returned as type "Property", and their *languageMaps* are reduced to simple strings. For example if the language filter `lang=fr` is specified, only the value for French language is present.

```

{
    "id": "urn:ngsi-ld:Vehicle:A4567",
    "type": "Vehicle",
    "brandName": {
        "type": "Property",
        "value": "Mercedes"
    },
    "street": {
        "type": "Property",
        "value": "Grand Place",
        "lang": "fr"
    },
    "isParked": {
        "type": "Relationship",
        "object": "urn:ngsi-ld:OffStreetParking:Downtown1",
        "observedAt": "2017-07-29T12:00:04Z",
        "providedBy": {
            "type": "Relationship",
            "object": "urn:ngsi-ld:Person:Bob"
        }
    },
    "category": {

```



```

    "type": "VocabProperty",
    "vocab": "non-commercial"
  },
  "tyreTreadDepths": {
    "type": "ListProperty",
    "valueList": [300, 300, 120, 6],
    "valueType": "Integer",
    "unitCode": "MMT"
  },
  "passengers": {
    "type": "Relationship",
    "objectType": "Person",
    "object": ["urn:ngsi-ld:Person:Alice", "urn:ngsi-ld:Person:Bob"]
  },
  "route": {
    "type": "ListRelationship",
    "objectType": "City",
    "objectList": [
      {"object": "urn:ngsi-ld:City:Antwerp"},
      {"object": "urn:ngsi-ld:City:Rotterdam"},
      {"object": "urn:ngsi-ld:City:Amsterdam"}
    ]
  },
  "@context": [
    "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
    "http://example.org/ngsi-ld/latest/vehicle.jsonld",
    "http://example.org/ngsi-ld/latest/parking.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}

```

Concise Representation

The concise representation is a terser, lossless form of the normalized representation, where redundant Attribute *type* members are omitted and the following rules are applied:

- Every *Property* without further sub-attributes is represented directly by the *Property* value only.
- Every *Property* that includes further sub-attributes is represented by a value key-value pair.
- Every *GeoProperty* without further sub-attributes is represented by the *GeoProperty*'s GeoJSON representation only.
- Every *GeoProperty* that includes further sub-attributes is represented by a value key-value pair.
- Every *LanguageProperty* is represented by a *languageMap* key-value pair.
- Every *ListProperty* is represented directly by the array of *Property* values.
- Every *JsonProperty* is represented by a *json* the value of which is raw JSON which is not available for JSON-LD representation.
- Every *VocabProperty* is represented by a *vocab* the value of which is a compacted URI.
- Every *Relationship* is represented by an object key-value pair.
- Every *ListRelationship* is represented by an array of URIs.

```

{
  "id": "urn:ngsi-ld:Vehicle:A4567",
  "type": "Vehicle",
  "brandName": "Mercedes",
  "street": {
    "languageMap": {
      "fr": "Grand Place",
      "nl": "Grote Markt"
    }
  },
  "isParked": {
    "object": "urn:ngsi-ld:OffStreetParking:Downtown1",
    "objectType": "OffStreetParking",
    "observedAt": "2017-07-29T12:00:04Z",
    "providedBy": {
      "object": "urn:ngsi-ld:Person:Bob"
    }
  }
}

```

```

    },
    "category": {
      "vocab": "non-commercial"
    },
    "tyreTreadDepths": {
      "valueList": [300, 300, 120, 6],
      "valueType": "Integer",
      "unitCode": "MMT"
    },
    "passengers": {
      "objectType": "Person",
      "object": ["urn:ngsi-ld:Person:Alice", "urn:ngsi-ld:Person:Bob"]
    },
    "route": {
      "objectType": "City",
      "objectList": [
        "urn:ngsi-ld:City:Antwerp",
        "urn:ngsi-ld:City:Rotterdam",
        "urn:ngsi-ld:City:Amsterdam"
      ]
    },
    "@context": [
      "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "http://example.org/ngsi-ld/latest/parking.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  }
}

```

Concise Representation when inline Linked Entity retrieval is used

When inline Linked Entity retrieval (see clause 4.5.23.2) is specified, any *Relationships* which target Entities stored locally or include an *objectType* Attribute are returned in an expanded format. The concise Linked Entity data corresponding to the *Relationship's* target *object* URI is returned within an *entity* sub-Attribute. Attributes of type "ListRelationship" are returned within an *entityList* sub-Attribute which in turn holds an ordered array of the Linked Entities in the concise format corresponding to each of the target *objectList* URIs.

```

{
  "id": "urn:ngsi-ld:Vehicle:A4567",
  "type": "Vehicle",
  "brandName": "Mercedes",
  "street": {
    "languageMap": {
      "fr": "Grand Place",
      "nl": "Grote Markt"
    }
  },
  "isParked": {
    "object": "urn:ngsi-ld:OffStreetParking:Downtown1",
    "objectType": "OffStreetParking",
    "entity": {
      "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
      "type": "OffStreetParking",
      "name": "Top Parking",
      "operatedBy": {
        "object": "urn:ngsi-ld:Company:BigParkingCorp"
      }
    }
  },
  "observedAt": "2017-07-29T12:00:04Z",
  "providedBy": {
    "object": "urn:ngsi-ld:Person:Bob"
  }
},
"category": {
  "vocab": "non-commercial"
},
"tyreTreadDepths": {
  "valueList": [300, 300, 120, 6],
  "valueType": "Integer",
  "unitCode": "MMT"
},
"passengers": {
  "objectType": "Person",
  "object": ["urn:ngsi-ld:Person:Alice", "urn:ngsi-ld:Person:Bob"],
  "entity": [

```

```

    {
      "id": "urn:ngsi-ld:Person:Alice",
      "type": "Person",
      "name": "Alice"
    },
    {
      "id": "urn:ngsi-ld:Person:Bob",
      "type": "Person",
      "name": "Bob"
    }
  ],
  "route": {
    "objectType": "City",
    "objectList": [
      "urn:ngsi-ld:City:Antwerp",
      "urn:ngsi-ld:City:Rotterdam",
      "urn:ngsi-ld:City:Amsterdam"
    ],
    "entityList": [
      {
        "id": "urn:ngsi-ld:City:Antwerp",
        "type": "City",
        "name": "Antwerp"
      },
      {
        "id": "urn:ngsi-ld:City:Rotterdam",
        "type": "City",
        "name": "Rotterdam"
      },
      {
        "id": "urn:ngsi-ld:City:Amsterdam",
        "type": "City",
        "name": "Amsterdam"
      }
    ]
  },
  "@context": [
    "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
    "http://example.org/ngsi-ld/latest/vehicle.jsonld",
    "http://example.org/ngsi-ld/latest/parking.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}

```

Concise Representation when flattened Linked Entity retrieval is used

When flattened Linked Entity retrieval (see clause 4.5.23.3) is specified, an array of concise Entities is returned. Whenever a *Relationship* Attribute targets an Entity stored locally or includes an *objectType*, an additional concise Linked Entity holding data corresponding to the *Relationship's* target *object* URI is appended to the response. For Attributes of type "ListRelationship", an array of concise Linked Entities is appended to the response which hold the data corresponding to each of the target URIs found within its *objectList*.

```

[
  {
    "id": "urn:ngsi-ld:Vehicle:A4567",
    "type": "Vehicle",
    "brandName": "Mercedes",
    "street": {
      "languageMap": {
        "fr": "Grand Place",
        "nl": "Grote Markt"
      }
    },
    "isParked": {
      "object": "urn:ngsi-ld:OffStreetParking:Downtown1",
      "objectType": "OffStreetParking",
      "observedAt": "2017-07-29T12:00:04Z",
      "providedBy": {
        "object": "urn:ngsi-ld:Person:Bob"
      }
    },
    "category": {
      "vocab": "non-commercial"
    }
  }
]

```

```

    },
    "tyreTreadDepths": {
      "valueList": [300, 300, 120, 6],
      "valueType": "Integer",

      "unitCode": "MMT"
    },
    "passengers": {
      "objectType": "Person",
      "object": ["urn:ngsi-ld:Person:Alice", "urn:ngsi-ld:Person:Bob"]
    },
    "route": {
      "objectType": "City",
      "objectList": [
        "urn:ngsi-ld:City:Antwerp",
        "urn:ngsi-ld:City:Rotterdam",
        "urn:ngsi-ld:City:Amsterdam"
      ]
    },
    "@context": [
      "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "http://example.org/ngsi-ld/latest/parking.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  },
  {
    "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
    "type": "OffStreetParking",
    "name": "Top Parking",
    "operatedBy": {
      "object": "urn:ngsi-ld:Company:BigParkingCorp",
    },
    "@context": [
      "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "http://example.org/ngsi-ld/latest/parking.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  },
  {
    "id": "urn:ngsi-ld:Person:Alice",
    "type": "Person",
    "name": "Alice",
    "@context": [
      "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "http://example.org/ngsi-ld/latest/parking.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  },
  {
    "id": "urn:ngsi-ld:Person:Bob",
    "type": "Person",
    "name": "Bob",
    "@context": [
      "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "http://example.org/ngsi-ld/latest/parking.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  },
  {
    "id": "urn:ngsi-ld:City:Antwerp",
    "type": "City",
    "name": "Antwerp",
    "@context": [
      "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "http://example.org/ngsi-ld/latest/parking.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  },
  {
    "id": "urn:ngsi-ld:City:Rotterdam",
    "type": "City",
    "name": "Rotterdam",
    "@context": [

```

```

        "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
        "http://example.org/ngsi-ld/latest/vehicle.jsonld",
        "http://example.org/ngsi-ld/latest/parking.jsonld",
        "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
},
{
    "id": "urn:ngsi-ld:City:Amsterdam",
    "type": "City",
    "name": "Amsterdam",
    "@context": [
        "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
        "http://example.org/ngsi-ld/latest/vehicle.jsonld",
        "http://example.org/ngsi-ld/latest/parking.jsonld",
        "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
}
]

```

Concise representation when Language Filter is used

The rules apply as defined in the previous examples. For example if the language filter `lang=fr` is specified.

```

{
    "id": "urn:ngsi-ld:Vehicle:A4567",
    "type": "Vehicle",
    "brandName": "Mercedes"
},
{
    "street": {
        "value": "Grand Place",
        "lang": "fr"
    },
    "isParked": {
        "objectType": "OffStreetParking",
        "object": "urn:ngsi-ld:OffStreetParking:Downtown1",
        "observedAt": "2017-07-29T12:00:04Z",
        "providedBy": {
            "object": "urn:ngsi-ld:Person:Bob"
        }
    },
    "category": {
        "vocab": "non-commercial"
    },
    "tyreTreadDepths": {
        "valueList": [300, 300, 120, 6],
        "valueType": "Integer",
        "unitCode": "MMT"
    },
    "passengers": {
        "objectType": "Person",
        "objectList": ["urn:ngsi-ld:Person:Alice", "urn:ngsi-ld:Person:Bob"]
    },
    "route": {
        "objectType": "City",
        "objectList": [
            "urn:ngsi-ld:City:Antwerp",
            "urn:ngsi-ld:City:Rotterdam",
            "urn:ngsi-ld:City:Amsterdam"
        ]
    },
    "@context": [
        "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
        "http://example.org/ngsi-ld/latest/vehicle.jsonld",
        "http://example.org/ngsi-ld/latest/parking.jsonld",
        "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
}

```

Simplified Representation

The simplified representation is a collapsed, lossy representation of an Entity, which focuses on Property Values and Relationship objects present at the first level of the graph only.

```

{
    "id": "urn:ngsi-ld:Vehicle:A4567",
    "type": "Vehicle",
    "brandName": "Mercedes",

```

```

    "street": {
      "languageMap": {
        "fr": "Grand Place",
        "nl": "Grote Markt"
      }
    },
    "isParked": "urn:ngsi-ld:OffStreetParking:Downtown1",
    "category": {
      "vocab": "non-commercial"
    },
    "tyreTreadDepths": [300, 300, 120, 6],
    "passengers": ["urn:ngsi-ld:Person:Alice", "urn:ngsi-ld:Person:Bob"],
    "route": [
      "urn:ngsi-ld:City:Antwerp",
      "urn:ngsi-ld:City:Rotterdam",
      "urn:ngsi-ld:City:Amsterdam"
    ],
    "@context": [
      "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "http://example.org/ngsi-ld/latest/parking.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  }
}

```

Simplified Representation when inline Linked Entity retrieval is used

When inline Linked Entity retrieval (see clause 4.5.23.2) is specified, any *Relationships* which target Entities stored locally or include an *objectType* Attribute are returned as a JSON object holding key-value pairs corresponding to the data from the *Relationship's* target *object* URI in simplified format. Attributes of type "ListRelationship" are returned as array of JSON objects each holding key-value pairs corresponding to the data obtained from the target *objectList* URIs.

```

{
  "id": "urn:ngsi-ld:Vehicle:A4567",
  "type": "Vehicle",
  "brandName": "Mercedes",
  "street": {
    "languageMap": {
      "fr": "Grand Place",
      "nl": "Grote Markt"
    }
  },
  "isParked": {
    "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
    "type": "OffStreetParking",
    "name": "Top Parking",
    "operatedBy": "urn:ngsi-ld:Company:BigParkingCorp"
  },
  "category": {
    "vocab": "non-commercial"
  },
  "tyreTreadDepths": [300, 300, 120, 6],
  "passengers": [
    {
      "id": "urn:ngsi-ld:Person:Alice",
      "type": "Person",
      "name": "Alice"
    },
    {
      "id": "urn:ngsi-ld:Person:Bob",
      "type": "Person",
      "name": "Bob"
    }
  ],
  "route": [
    {
      "id": "urn:ngsi-ld:City:Antwerp",
      "type": "City",
      "name": "Antwerp"
    },
    {
      "id": "urn:ngsi-ld:City:Rotterdam",
      "type": "City",
      "name": "Rotterdam"
    }
  ]
}

```

```

        "id": "urn:ngsi-ld:City:Amsterdam",
        "type": "City",
        "name": "Amsterdam"
    }
},
"@context": [
    "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
    "http://example.org/ngsi-ld/latest/vehicle.jsonld",
    "http://example.org/ngsi-ld/latest/parking.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
]
}

```

Simplified Representation when flattened Linked Entity retrieval is used

When flattened Linked Entity retrieval (see clause 4.5.23.3) is specified, an array of JSON Objects is returned. Whenever a *Relationship* Attribute targets an Entity stored locally or includes an *objectType*, an additional JSON Object of key-value pairs holding data corresponding to the *Relationship's* target *object* URI is appended to the response. For Attributes of type "ListRelationship", an array of JSON Objects each holding key-value pairs corresponding to the data obtained from the target *objectList* URIs is appended to the response.

```

[
    {
        "id": "urn:ngsi-ld:Vehicle:A4567",
        "type": "Vehicle",
        "brandName": "Mercedes",
        "street": {
            "languageMap": {
                "fr": "Grand Place",
                "nl": "Grote Markt"
            }
        }
        "isParked": "urn:ngsi-ld:OffStreetParking:Downtown1",
        "category": {
            "vocab": "non-commercial"
        },
        "tyreTreadDepths": [300, 300, 120, 6],
        "passengers": ["urn:ngsi-ld:Person:Alice", "urn:ngsi-ld:Person:Bob"],
        "route": [
            "urn:ngsi-ld:City:Antwerp",
            "urn:ngsi-ld:City:Rotterdam",
            "urn:ngsi-ld:City:Amsterdam"
        ],
        "@context": [
            "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
            "http://example.org/ngsi-ld/latest/vehicle.jsonld",
            "http://example.org/ngsi-ld/latest/parking.jsonld",
            "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
        ]
    },
    {
        "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
        "type": "OffStreetParking",
        "name": "Top Parking",
        "operatedBy": "urn:ngsi-ld:Company:BigParkingCorp"
    },
    {
        "id": "urn:ngsi-ld:Person:Alice",
        "type": "Person",
        "name": "Alice"
    },
    {
        "id": "urn:ngsi-ld:Person:Bob",
        "type": "Person",
        "name": "Bob"
    },
    {
        "id": "urn:ngsi-ld:City:Antwerp",
        "type": "City",
        "name": "Antwerp"
    },
    {
        "id": "urn:ngsi-ld:City:Rotterdam",
        "type": "City",
        "name": "Rotterdam"
    }
],

```

```
{
  "id": "urn:ngsi-ld:City:Amsterdam",
  "type": "City",
  "name": "Amsterdam"
}
```

Simplified Representation of Natural Language Attributes

The simplified natural language representation is a collapsed representation of an Entity, which focuses on Property Values and Relationship objects present at the first level of the graph, and where *languageMaps* are reduced to simple string properties. For example if the language filter `lang=fr` is specified.

```
{
  "id": "urn:ngsi-ld:Vehicle:A4567",
  "type": "Vehicle",
  "brandName": "Mercedes",
  "street": "Grand Place",
  "isParked": "urn:ngsi-ld:OffStreetParking:Downtown1",
  "category": {
    "vocab": "non-commercial"
  },
  "tyreTreadDepths": [300, 300, 120, 6],
  "passengers": ["urn:ngsi-ld:Person:Alice", "urn:ngsi-ld:Person:Bob"],
  "route": [
    "urn:ngsi-ld:City:Antwerp",
    "urn:ngsi-ld:City:Rotterdam",
    "urn:ngsi-ld:City:Amsterdam"
  ],
  "@context": [
    "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
    "http://example.org/ngsi-ld/latest/vehicle.jsonld",
    "http://example.org/ngsi-ld/latest/parking.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}
```

Multiple attribute example

Below is an example, where the speed of the car is provided by two different sources. As both may be relevant at the same time, there are two individual attribute instances for speed; each is identified by a *datasetId* and both instances are represented in an array. The *datasetId* enables individually creating, updating and deleting a particular instance without affecting the instance from another source.

```
{
  "id": "urn:ngsi-ld:Vehicle:A4567",
  "type": "Vehicle",
  "speed": [{
    "type": "Property",
    "value": 55,
    "source": {
      "type": "Property",
      "value": "Speedometer"
    }
  },
  "datasetId": "urn:ngsi-ld:Property:speedometerA4567-speed"
],
{
  "type": "Property",
  "value": 54.5,
  "source": {
    "type": "Property",
    "value": "GPS"
  },
  "datasetId": "urn:ngsi-ld:Property:gpsBxyz123-speed"
}],
  "@context": [
    {
      "Vehicle": "http://example.org/Vehicle",
      "speed": "http://example.org/speed",
      "source": "http://example.org/hasSource"
    },
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}
```


Simplified Representation of a multi-attribute

The simplified representation is a collapsed, lossy representation of an Entity, which focuses on Property Values and Relationship objects present at the first level of the graph only.

```
{
  "id": "urn:ngsi-ld:Vehicle:A4567",
  "type": "Vehicle",
  "speed": {
    "dataset": {
      "urn:ngsi-ld:Property:speedometerA4567-speed": 55,
      "urn:ngsi-ld:Property:gpsBxyz123-speed": 54.5
    }
  },
  "@context": [
    {
      "Vehicle": "http://example.org/Vehicle",
      "speed": "http://example.org/speed"
    },
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}
```

C.2.3 Parking Entity

Normalized Representation

The normalized representation is a lossless representation of an Entity, where every *Property* is defined by a *type* and a *value* and every *Relationship* is defined by a *type* and an *object*.

Below there is a representation of an Entity of Type "OffStreetParking". It can be observed that the *@context* is composed of two different elements, the Core one and the vocabulary-specific one.

```
{
  "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
  "type": "OffStreetParking",
  "name": {
    "type": "Property",
    "value": "Downtown One"
  },
  "availableSpotNumber": {
    "type": "Property",
    "value": 121,
    "observedAt": "2017-07-29T12:05:02Z",
    "reliability": {
      "type": "Property",
      "value": 0.7
    }
  },
  "providedBy": {
    "type": "Relationship",
    "object": "urn:ngsi-ld:Camera:C1"
  },
  "totalSpotNumber": {
    "type": "Property",
    "value": 200
  },
  "location": {
    "type": "GeoProperty",
    "value": {
      "type": "Point",
      "coordinates": [-8.5, 41.2]
    }
  },
  "@context": [
    "http://example.org/ngsi-ld/latest/parking.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}
```

Concise Representation

The concise representation is a terser, lossless form of the normalized representation, where redundant Attribute *type* members are omitted and the following rules are applied:

- Every *Property* without further sub-attributes is represented by the *Property* value only.
- Every *Property* that includes further sub-attributes is represented by a value key-value pair.
- Every *GeoProperty* without further sub-attributes is represented by the *GeoProperty*'s GeoJSON representation only.
- Every *GeoProperty* that includes further sub-attributes is represented by a value key-value pair.
- Every *LanguageProperty* is defined by a *languageMap* key-value pair.
- Every *VocabProperty* is represented by a *vocab* the value of which is a compacted URI.
- Every *Relationship* is defined by an object key-value pair.

```
{
  "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
  "type": "OffStreetParking",
  "name": "Downtown One",
  "availableSpotNumber": {
    "value": 121,
    "observedAt": "2017-07-29T12:05:02Z",
    "reliability": 0.7,
    "providedBy": {
      "object": "urn:ngsi-ld:Camera:C1"
    }
  },
  "totalSpotNumber": 200,
  "location": {
    "type": "Point",
    "coordinates": [
      -8.5,
      41.2
    ]
  },
  "@context": [
    "http://example.org/ngsi-ld/latest/parking.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}
```

Simplified representation

The Simplified Representation (also known as "*keyValues*") is a lossy, collapsed representation of an Entity, which focuses on Property Values and Relationship objects present at the first level of the graph only.

```
{
  "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
  "type": "OffStreetParking",
  "name": "Downtown One",
  "availableSpotNumber": 121,
  "totalSpotNumber": 200,
  "location": {
    "type": "Point",
    "coordinates": [-8.5, 41.2]
  },
  "@context": [
    "http://example.org/ngsi-ld/latest/parking.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}
```

Normalized GeoJSON Representation

The normalized GeoJSON representation of a single Entity is defined as a single GeoJSON *Feature* object as follows:

```
{
  "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
```

```

"type": "Feature",
"geometry": {
  "type": "Point",
  "coordinates": [-8.51, 41.1]
},
"properties": {
  "type": "OffStreetParking",
  "name": {
    "type": "Property",
    "value": "Downtown One"
  },
  "availableSpotNumber": {
    "type": "Property",
    "value": 121,
    "observedAt": "2017-07-29T12:05:02Z",
    "reliability": {
      "type": "Property",
      "value": 0.7
    },
  },
  "providedBy": {
    "type": "Relationship",
    "object": "urn:ngsi-ld:Camera:C1"
  }
},
"location": {
  "type": "GeoProperty",
  "value": {
    "type": "Point",
    "coordinates": [-8.51, 41.1]
  }
},
"totalSpotNumber": {
  "type": "Property",
  "value": 200
},
"@context": [
  "http://example.org/ngsi-ld/latest/parking.jsonld",
  "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
]
}
}

```

The GeoJSON representation of multiple Entities is defined as a GeoJSON *FeatureCollection* object containing an array of GeoJSON features corresponding to the individual Entity representations.

```

{
  "type": "FeatureCollection",
  "features": [
    {
      "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
      "type": "Feature",
      "geometry": {
        "type": "Point",
        "coordinates": [-8.5, 41.1]
      },
      "properties": {
        "type": "OffStreetParking",
        "name": {
          "type": "Property",
          "value": "Downtown One"
        },
        "availableSpotNumber": {
          "type": "Property",
          "value": 121,
          "observedAt": "2017-07-29T12:05:02Z",
          "reliability": {
            "type": "Property",
            "value": 0.7
          },
        },
        "providedBy": {
          "type": "Relationship",
          "object": "urn:ngsi-ld:Camera:C1"
        }
      },
      "totalSpotNumber": {
        "type": "Property",
        "value": 200
      },
    },
  ],
}

```

```

    "location": {
      "type": "GeoProperty",
      "value": {
        "type": "Point",
        "coordinates": [-8.51, 41.1]
      }
    }
  },
  {
    "id": "urn:ngsi-ld:OffStreetParking:Downtown2",
    "type": "Feature",
    "geometry": {
      "type": "Point",
      "coordinates": [-8.51, 41.1]
    },
    "properties": {
      "type": "OffStreetParking",
      "name": {
        "type": "Property",
        "value": "Downtown Two"
      },
      "availableSpotNumber": {
        "type": "Property",
        "value": 99,
        "observedAt": "2017-07-29T12:05:02Z",
        "reliability": {
          "type": "Property",
          "value": 0.8
        }
      },
      "providedBy": {
        "type": "Relationship",
        "object": "urn:ngsi-ld:Camera:C2"
      }
    },
    "totalSpotNumber": {
      "type": "Property",
      "value": 100
    },
    "location": {
      "type": "GeoProperty",
      "value": {
        "type": "Point",
        "coordinates": [-8.51, 41.1]
      }
    }
  }
}
],
"@context": [
  "http://example.org/ngsi-ld/latest/parking.jsonld",
  "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
]
}

```

Concise GeoJSON Representation

The concise GeoJSON representation of a single Entity is defined as a single GeoJSON *Feature* object as follows:

```

{
  "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
  "type": "Feature",
  "geometry": {
    "type": "Point",
    "coordinates": [
      -8.51,
      41.1
    ]
  },
  "properties": {
    "type": "OffStreetParking",
    "name": "Downtown One",
    "availableSpotNumber": {
      "value": 121,
      "observedAt": "2017-07-29T12:05:02Z",
      "reliability": 0.7,
      "providedBy": {
        "object": "urn:ngsi-ld:Camera:C1"
      }
    }
  }
}

```

```

    }
  },
  "location": {
    "type": "Point",
    "coordinates": [
      -8.51,
      41.1
    ]
  },
  "totalSpotNumber": 200,
  "@context": [
    "http://example.org/ngsi-ld/latest/parking.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}
}

```

The concise GeoJSON representation of multiple Entities is defined as a GeoJSON *FeatureCollection* object containing an array of GeoJSON features corresponding to the individual Entity representations in concise GeoJSON format.

```

{
  "type": "FeatureCollection",
  "features": [
    {
      "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
      "type": "Feature",
      "geometry": {
        "type": "Point",
        "coordinates": [
          -8.5,
          41.1
        ]
      },
      "properties": {
        "type": "OffStreetParking",
        "name": "Downtown One",
        "availableSpotNumber": {
          "value": 121,
          "observedAt": "2017-07-29T12:05:02Z",
          "reliability": 0.7,
          "providedBy": {
            "object": "urn:ngsi-ld:Camera:C1"
          }
        },
        "totalSpotNumber": 200,
        "location": {
          "type": "Point",
          "coordinates": [
            -8.51,
            41.1
          ]
        }
      }
    },
    {
      "id": "urn:ngsi-ld:OffStreetParking:Downtown2",
      "type": "Feature",
      "geometry": {
        "type": "Point",
        "coordinates": [
          -8.51,
          41.1
        ]
      },
      "properties": {
        "type": "OffStreetParking",
        "name": "Downtown Two",
        "availableSpotNumber": {
          "value": 99,
          "observedAt": "2017-07-29T12:05:02Z",
          "reliability": 0.8,
          "providedBy": {
            "object": "urn:ngsi-ld:Camera:C2"
          }
        },
        "totalSpotNumber": 100,
        "location": {
          "type": "Point",

```

```

        "coordinates": [
            -8.51,
            41.1
        ]
    }
}
],
"@context": [
    "http://example.org/ngsi-ld/latest/parking.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
]
}

```

Simplified GeoJSON Representation

The simplified GeoJSON representation of a single Entity is defined as a single GeoJSON *Feature* object where the properties represent a collapsed representation of the Entity, which focuses on Property Values and Relationship objects present at the first level of the graph.

```

{
  "id": "urn:ngsi-ld:offstreetparking:Downtown1",
  "type": "Feature",
  "geometry": {
    "type": "Point",
    "coordinates": [-8.51, 41.1]
  },
  "properties": {
    "type": "OffStreetParking",
    "name": "Downtown One",
    "availableSpotNumber": 121,
    "totalSpotNumber": 200,
    "location": {
      "type": "Point",
      "coordinates": [-8.51, 41.1]
    }
  },
  "@context": [
    "http://example.org/ngsi-ld/latest/parking.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}

```

The simplified GeoJSON representation of multiple Entities is defined as a GeoJSON *FeatureCollection* object containing an array of GeoJSON features corresponding to the individual Entity representations in simplified GeoJSON format.

```

{
  "type": "FeatureCollection",
  "features": [
    {
      "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
      "type": "Feature",
      "geometry": {
        "type": "Point",
        "coordinates": [-8.5, 41.2]
      },
      "properties": {
        "type": "OffStreetParking",
        "name": "Downtown One",
        "availableSpotNumber": 121,
        "totalSpotNumber": 200,
        "location": {
          "type": "Point",
          "coordinates": [-8.5, 41.2]
        }
      }
    },
    {
      "id": "urn:ngsi-ld:OffStreetParking:Downtown2",
      "type": "Feature",
      "geometry": {
        "type": "Point",
        "coordinates": [-8.51, 41.1]
      },
      "properties": {

```

```

    "type": "OffStreetParking",
    "name": "Downtown Two",
    "availableSpotNumber": 99,
    "totalSpotNumber": 100,
    "location": {
      "type": "Point",
      "coordinates": [-8.51, 41.1]
    }
  }
},
"@context": [
  "http://example.org/ngsi-ld/latest/parking.jsonld",
  "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
]
}

```

C.2.4 @context

The disposition of the *@context* can be as an inline JSON object, as a dereferenceable URI or as a (multiple) combination of both. In the examples above the *@context* is provided through several dereferenceable URIs. The resulting *@context* (obtained by merging the content of the resource referenced by the referred URIs) is shown below.

NOTE 1: For brevity reasons the *@context* does not contain the API terms defined by clause 5.2.

NOTE 2: Some extra terms are defined because they will be used in examples later presented.

```

{
  "id": "@id",
  "type": "@type",
  "Property": "https://uri.etsi.org/ngsi-ld/Property",
  "Relationship": "https://uri.etsi.org/ngsi-ld/Relationship",
  "value": "https://uri.etsi.org/ngsi-ld/hasValue",
  "object": {
    "@type": "@id",
    "@id": "https://uri.etsi.org/ngsi-ld/hasObject"
  },
  "observedAt": {
    "@type": "https://uri.etsi.org/ngsi-ld/DateTime",
    "@id": "https://uri.etsi.org/ngsi-ld/observedAt"
  },
  "datasetId": {
    "@id": "https://uri.etsi.org/ngsi-ld/datasetId",
    "@type": "@id"
  },
  "location": "https://uri.etsi.org/ngsi-ld/location",
  "GeoProperty": "https://uri.etsi.org/ngsi-ld/GeoProperty",
  "Vehicle": "http://example.org/vehicle/Vehicle",
  "street": "http://example.org/vehicle/street",
  "brandName": "http://example.org/vehicle/brandName",
  "category": "http://example.org/vehicle/category",
  "tyreTreadDepths": "http://example.org/vehicle/tyreTreadDepths",
  "passengers": "http://example.org/vehicle/passengers",
  "speed": "http://example.org/vehicle/speed",
  "isParked": {
    "@type": "@id",
    "@id": "http://example.org/common/isParked"
  },
  "OffStreetParking": "http://example.org/parking/OffStreetParking",
  "availableSpotNumber": "http://example.org/parking/availableSpotNumber",
  "totalSpotNumber": "http://example.org/parking/totalSpotNumber",
  "isNextToBuilding": {
    "@type": "@id",
    "@id": "http://example.org/common/isNextToBuilding"
  },
  "reliability": "http://example.org/common/reliability",
  "providedBy": {
    "@type": "@id",
    "@id": "http://example.org/common/providedBy"
  },
  "name": "http://example.org/common/name",
  "commercial": "http://example.org/vehicle/commercial",
  "non-commercial": "http://example.org/vehicle/non-commercial",
  "Integer": "http://www.w3.org/2001/XMLSchema#Integer"
}

```

C.3 Context Source Registration

Below there is an example representation of a Context Source Registration. It makes use of the *@context* formerly described.

```
{
  "id": "urn:ngsi-ld:ContextSourceRegistration:csr1a3456",
  "type": "ContextSourceRegistration",
  "information": [
    {
      "entities": [
        {
          "id": "urn:ngsi-ld:Vehicle:A456",
          "type": "Vehicle"
        }
      ],
      "propertyNames": ["brandName", "speed"],
      "relationshipNames": ["isParked"]
    },
    {
      "entities": [
        {
          "idPattern": ".*downtown$",
          "type": "OffStreetParking"
        },
        {
          "idPattern": ".*47$",
          "type": "OffStreetParking"
        }
      ],
      "propertyNames": ["availableSpotNumber", "totalSpotNumber"],
      "relationshipNames": ["isNextToBuilding"]
    }
  ],
  "endpoint": "http://my.csource.org:1026",
  "location": {
    "type": "Polygon",
    "coordinates": [
      [ [100.0, 0.0], [101.0, 0.0], [101.0, 1.0],
        [100.0, 1.0], [100.0, 0.0] ] ]
    ],
    "managementInterval": {
      "startAt": "2017-11-29T14:53:15Z"
    }
  },
  "@context": [
    "http://example.org/ngsi-ld/latest/commonTerms.jsonld",
    "http://example.org/ngsi-ld/latest/vehicle.jsonld",
    "http://example.org/ngsi-ld/latest/parking.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}
```

The Registration is referring to a Temporal Context Source capable of providing temporal information from Entities of type "Vehicle" and "OffStreetParking", meeting certain id requirements. More concretely, it can only provide the referenced Properties and Relationships. Temporal information is provided for the given *managementInterval*, i.e. related to *createdAt* and *modifiedAt* Temporal Properties. The *managementInterval* is specified as an open interval, so only a starting point is given. In addition, the Registration example covers a particular geographical area.

C.4 Context Subscription

Below there is an example of a Context Subscription. It makes use of the *@context* formerly described.

```
{
  "id": "urn:ngsi-ld:Subscription:mySubscription",
  "type": "Subscription",
  "entities": [
    {
```



```

        "type": "Vehicle"
      }
    ],
    "watchedAttributes": ["speed"],
    "q": "speed>50",
    "geoQ": {
      "georel": "near;maxDistance==2000",
      "geometry": "Point",
      "coordinates": [-1,100]
    },
    "notification": {
      "attributes": ["speed"],
      "format": "keyValues",
      "endpoint": {
        "uri": "http://my.endpoint.org/notify",
        "accept": "application/json"
      }
    },
    "@context": [
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  }
}

```

The subject of the Context Subscription is Entities of Type *Vehicle* which *speed* is greater than 50, and located close to a certain area defined by a reference spatial point. Every time the *speed* (watched Attribute) of a concerned vehicle, changes, a new notification (including the new speed value) will be received in the specified endpoint.

C.5 HTTP REST API Examples

C.5.1 Introduction

This clause introduces some simple usage examples of the NGSI-LD API (HTTP REST binding). They are not intended to be exhaustive but just a sample for helping readers to understand better the present document. ETSI ISG CIM published a Developer's Primer with many more examples, see ETSI GR CIM 008 [i.21].

C.5.2 Create Entity of Type Vehicle

C.5.2.1 HTTP Request

POST /ngsi-ld/v1/entities/

Content-Type: application/ld+json

Content-Length: 556

<Insert Here the JSON-LD representation of a Vehicle as described by clause C.2.2 Vehicle Entity>

C.5.2.2 HTTP Response

201 Created

Location: /ngsi-ld/v1/entities/urn:ngsi-ld:Vehicle:A4567

C.5.3 Query Entities

C.5.3.1 Introduction

EXAMPLE: Give back all the Entities of type "Vehicle" whose *brandName* attribute is not "Mercedes". Only give back the *brandName* attribute and provide the data in the NGSI-LD Simplified Format.

C.5.3.2 HTTP Request

GET /ngsi-ld/v1/entities/?type=Vehicle&q=brandName!="Mercedes"&format=**simplified**

Accept: application/ld+json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

C.5.3.3 HTTP Response

200 OK

Content-Type: application/ld+json

```
[
  {
    "id": "urn:ngsi-ld:Vehicle:B9211",
    "type": "Vehicle",
    "brandName": "Volvo",
    "@context": [
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  }
]
```

C.5.4 Query Entities (Pagination)

C.5.4.1 Introduction

EXAMPLE: Give back all the Entities of type "Vehicle". Only give back the *brandName* attribute and provide the data in the NGSI-LD Simplified Format. Limit the number of entities retrieved to 2.

C.5.4.2 HTTP Request

GET /ngsi-ld/v1/entities/?type= Vehicle&format=**simplified**&limit=2

Accept: application/ld+json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

C.5.4.3 HTTP Response

200 OK

Content-Type: application/ld+json

Link: </ngsi-ld/v1/entities/?type= Vehicle&format=**simplified**&limit=2&offset=2>; rel="next"; type="application/ld+json"

```
[
  {
    "id": "urn:ngsi-ld:Vehicle:B9211",
    "type": "Vehicle",
    "brandName": "Volvo",
    "@context": [
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  },
  {
    "id": "urn:ngsi-ld:Vehicle:A456",
    "type": "Vehicle",
    "brandName": "Mercedes",
  }
]
```

```

    "@context": [
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  }
]

```

C.5.5 Temporal Query

C.5.5.1 Introduction

EXAMPLE 1: Give back the temporal evolution of the attribute *speed* of Entities of type "Vehicle" whose *brandName* attribute is not "Mercedes" between the 1st of August at noon and the 1st of August at 01 PM.

EXAMPLE 2: Give back the temporal evolution of the attribute *speed* and *brandName* of Entities of type "Vehicle" whose *brandName* attribute is not "Mercedes" between the 1st of August at noon and the 1st of August at 01 PM. As *brandName* attribute does not have any temporal evolution, *brandName* attribute is omitted in the response.

C.5.5.2 HTTP Request #1

GET /ngsi-ld/v1/temporal/entities/?type=Vehicle&q=brandName!=Mercedes&attrs=speed&timerel=between&timeAt=2018-08-01T12:00:00Z&endTimeAt=2018-08-01T13:00:00Z

Accept: application/ld+json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

C.5.5.3 HTTP Response #1

200 OK

Content-Type: application/ld+json

```

[
  {
    "id": "urn:ngsi-ld:Vehicle:B9211",
    "type": "Vehicle",
    "speed": [
      {
        "type": "Property",
        "value": 120,
        "observedAt": "2018-08-01T12:03:00Z"
      },
      {
        "type": "Property",
        "value": 80,
        "observedAt": "2018-08-01T12:05:00Z"
      },
      {
        "type": "Property",
        "value": 100,
        "observedAt": "2018-08-01T12:07:00Z"
      }
    ],
    "@context": [
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  }
]

```

C.5.5.4 HTTP Request #2

GET /ngsi-ld/v1/temporal/entities/?type=Vehicle&q=brandName!=Mercedes&attrs=speed,brandName&timerel=between&timeAt=2018-08-01T12:00:00Z&endTimeAt=2018-08-01T13:00:00Z

Accept: application/ld+json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

C.5.5.5 HTTP Response #2

200 OK

Content-Type: application/ld+json

```
[
  {
    "id": "urn:ngsi-ld:Vehicle:B9211",
    "type": "Vehicle",
    "speed": [
      {
        "type": "Property",
        "value": 120,
        "observedAt": "2018-08-01T12:03:00Z"
      },
      {
        "type": "Property",
        "value": 80,
        "observedAt": "2018-08-01T12:05:00Z"
      },
      {
        "type": "Property",
        "value": 100,
        "observedAt": "2018-08-01T12:07:00Z"
      }
    ],
    "@context": [
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  }
]
```

C.5.6 Temporal Query (Simplified Representation)

C.5.6.1 Introduction

EXAMPLE: Give back the temporal evolution of the *speed* attribute for Entities of type "Vehicle" whose *brandName* attribute is not "Mercedes" between the 1st of August at noon and the 1st of August at 01 PM. Simplified representation is required.

C.5.6.2 HTTP Request

GET /ngsi-ld/v1/temporal/entities/?type=Vehicle&q=brandName!=Mercedes&attrs=speed&timerel=between&timeAt=2018-08-01T12:00:00Z&endTimeAt=2018-08-01T13:00:00Z&format=**temporalValues**

Accept: application/ld+json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

C.5.6.3 HTTP Response

200 OK

Content-Type: application/ld+json

```
[
  {
    "id": "urn:ngsi-ld:Vehicle:B9211",
    "type": "Vehicle",
    "speed": {
      "type": "Property",
      "values": [
        [
          120,
          "2018-08-01T12:03:00Z"
        ],
        [
          80,
          "2018-08-01T12:05:00Z"
        ],
        [
          100,
          "2018-08-01T12:07:00Z"
        ]
      ]
    },
    "@context": [
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  }
]
```

C.5.7 Retrieve Available Entity Types

C.5.7.1 Introduction

EXAMPLE: Give back all entity types for which entity instances are currently available in the NGSI-LD system.

C.5.7.2 HTTP Request

GET /ngsi-ld/v1/types

Accept: application/json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

C.5.7.3 HTTP Response

200 OK

Content-Type: application/json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

```
{
  "id": "urn:ngsi-ld:EntityTypeList:34534657",
  "type": "EntityTypeList",
  "typeList": [
    "Vehicle",
    "OffStreetParking",
    "http://example.org/parking/ParkingSpot"
  ]
}
```

NOTE: All entity types that can be found in the provided *@context* are given as short names, the others as Fully Qualified Names (FQNs).

C.5.8 Retrieve Details of Available Entity Types

C.5.8.1 Introduction

EXAMPLE: Give back the details of all entity types for which entity instances are currently available in the NGSI-LD system.

C.5.8.2 HTTP Request

GET /ngsi-ld/v1/types?details=true

Accept: application/json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

C.5.8.3 HTTP Response

200 OK

Content-Type: application/json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

```
[
  {
    "id": "http://example.org/vehicle/Vehicle",
    "type": "EntityType",
    "typeName": "Vehicle",
    "attributeNames": [
      "brandName",
      "isParked",
      "location",
      "speed"
    ]
  },
  {
    "id": "http://example.org/parking/OffStreetParking",
    "type": "EntityType",
    "typeName": "OffStreetParking",
    "attributeNames": [
      "availableSpotNumber",
      "isNextToBuilding",
      "location",
      "totalSpotNumber"
    ]
  },
  {
    "id": "http://example.org/parking/ParkingSpot",
    "type": "EntityType",
    "typeName": "http://example.org/parking/ParkingSpot",
    "attributeNames": [
      "location",
      "http://example.org/parking/status"
    ]
  }
]
```

NOTE: The type name of all entity types and all attribute names that can be found in the provided *@context* are given as short names, the others as Fully Qualified Names (FQNs). The id is always an FQN.

C.5.9 Retrieve Available Entity Type Information

C.5.9.1 Introduction

EXAMPLE: Give back the details of entity type "Vehicle" (for which entity instances are currently available in the NGSI-LD system).

C.5.9.2 HTTP Request

GET /ngsi-ld/v1/types/Vehicle

[Alternative with FQN: **GET** /ngsi-ld/v1/attributes/http%3A%2F%2Fexample.org%2Fvehicle%2FVehicle]

Accept: application/json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

C.5.9.3 HTTP Response

200 OK

Content-Type: application/json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

```
{
  "id": "http://example.org/vehicle/Vehicle",
  "type": "EntityTypeInfo",
  "typeName": "Vehicle",
  "entityCount": 2,
  "attributeDetails": [
    {
      "id": "http://example.org/vehicle/brandName",
      "type": "Attribute",
      "attributeName": "brandName",
      "attributeTypes": [
        "Property"
      ]
    },
    {
      "id": "http://example.org/vehicle/isParked",
      "type": "Attribute",
      "attributeName": "isParked",
      "attributeTypes": [
        "Relationship"
      ]
    },
    {
      "id": "https://uri.etsi.org/ngsi-ld/location",
      "type": "Attribute",
      "attributeName": "location",
      "attributeTypes": [
        "GeoProperty"
      ]
    },
    {
      "id": "http://example.org/vehicle/speed",
      "type": "Attribute",
      "attributeName": "speed",
      "attributeTypes": [
        "Property"
      ]
    }
  ]
}
```

C.5.10 Retrieve Available Attributes

C.5.10.1 Introduction

EXAMPLE: Give back all attribute names for which entity instances are currently available in the NGSI-LD system that have an attribute with the respective name.

C.5.10.2 HTTP Request

GET /ngsi-ld/v1/attributes

Accept: application/json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

C.5.10.3 HTTP Response

200 OK

Content-Type: application/json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

```
{
  "id": "urn:ngsi-ld:AttributeList:56534657",
  "type": "AttributeList",
  "attributeList": [
    "brandName",
    "isParked",
    "location",
    "speed",
    "http://example.org/parking/status"
  ]
}
```

NOTE: The attribute names that can be found in the provided *@context* are given as short names, the others as Fully Qualified Names (FQNs).

C.5.11 Retrieve Details of Available Attributes

C.5.11.1 Introduction

EXAMPLE: Give back the details of all attributes for which entity instances are currently available in the NGSI-LD system to which an attribute with the respective attribute name belongs.

C.5.11.2 HTTP Request

GET /ngsi-ld/v1/attributes?details=true

Accept: application/json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

C.5.11.3 HTTP Response

200 OK

Content-Type: application/json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

```
[
  {
    "id": "http://example.org/vehicle/brandName",
    "type": "Attribute",
    "attributeName": "brandName",
    "typeNames": [
      "Vehicle"
    ]
  },
  {
    "id": "http://example.org/vehicle/isParked",
    "type": "Attribute",
    "attributeName": "isParked",
    "typeNames": [
      "Vehicle"
    ]
  },
  {
    "id": "https://uri.etsi.org/ngsi-ld/location",
    "type": "Attribute",
    "attributeName": "location",
    "typeNames": [
      "Vehicle",
      "OffStreetParking",
      "http://example.org/parking/ParkingSpot"
    ]
  },
  {
    "id": "http://example.org/vehicle/speed",
    "type": "Attribute",
    "attributeName": "speed",
    "typeNames": [
      "Vehicle"
    ]
  },
  {
    "id": "http://example.org/parking/status",
    "type": "Attribute",
    "attributeName": "http://example.org/parking/status",
    "typeNames": [
      "http://example.org/parking/ParkingSpot"
    ]
  }
]
```

NOTE: The attribute name and all type names that can be found in the provided *@context* are given as short names, the others as Fully Qualified Names (FQNs). The *id* is always an FQN.

C.5.12 Retrieve Available Attribute Information

C.5.12.1 Introduction

EXAMPLE: Give back the details of the attribute named *brandName* (for which entity instances with an attribute of this name are currently available in the NGSI-LD system).

C.5.12.2 HTTP Request

GET /ngsi-ld/v1/attributes/brandName

[Alternative with FQN: **GET** /ngsi-ld/v1/attributes/http%3A%2F%2Fexample.org%2Fvehicle%2FbrandName]

Accept: application/json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

C.5.12.3 HTTP Response

200 OK

Content-Type: application/json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

```
{
  "id": "http://example.org/vehicle/brandName",
  "type": "Attribute",
  "attributeName": "brandName",
  "attributeTypes": ["Property"],
  "typeNames": ["Vehicle"],
  "attributeCount": 2
}
```

C.5.13 Query Entities (Natural Language Filtering)

C.5.13.1 Introduction

EXAMPLE: Give back all the Entities of type "Vehicle" where the *marque* attribute in British English is "Vauxhall Viva". Only give back the *marque* attribute and provide the data in the NGSI-LD Simplified Format and only return language strings in German.

C.5.13.2 HTTP Request

GET /ngsi-ld/v1/entities/?type=Vehicle&attrs=marque&q=marque[en-GB]== "Vauxhall Viva"&format=simplified&lang=de

Accept: application/ld+json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

C.5.13.3 HTTP Response

200 OK

Content-Type: application/ld+json

```
[
  {
    "id": "urn:ngsi-ld:Vehicle:A4567",
    "type": "Vehicle",
    "marque": "Opel Karl",
    "@context": [
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  }
]
```

C.5.14 Temporal Query (Aggregated Representation)

C.5.14.1 Introduction

EXAMPLE: Give back the maximum and average speed of Entities of type "Vehicle" whose *brandName* attribute is not "Mercedes" between the 1st of August at noon and the 1st of August at 01 PM, aggregated by periods of 4 minutes.

C.5.14.2 HTTP Request

GET /ngsi-ld/v1/temporal/entities/?type=Vehicle&q=brandName!=Mercedes&attrs=speed&timerel=between&timeAt=2018-08-01T12:00:00Z&endTimeAt=2018-08-01T13:00:00Z&aggrMethods=max,avg&aggrPeriodDuration=PT4M&format=**aggregatedValues**

Accept: application/ld+json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

C.5.14.3 HTTP Response

200 OK

Content-Type: application/ld+json

```
[
  {
    "id": "urn:ngsi-ld:Vehicle:B9211",
    "type": "Vehicle",
    "speed": {
      "type": "Property",
      "max": [
        [
          120,
          "2018-08-01T12:00:00Z",
          "2018-08-01T12:04:00Z"
        ],
        [
          100,
          "2018-08-01T12:04:00Z",
          "2018-08-01T12:08:00Z"
        ]
      ],
      "avg": [
        [
          120,
          "2018-08-01T12:00:00Z",
          "2018-08-01T12:04:00Z"
        ],
        [
          90,
          "2018-08-01T12:04:00Z",
          "2018-08-01T12:08:00Z"
        ]
      ]
    },
    "@context": [
      "http://example.org/ngsi-ld/latest/vehicle.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  }
]
```

C.5.15 Scope Queries

C.5.15.1 Introduction

EXAMPLE: Give back all the Entities of type "OffStreetParking" that are within the Scope /Madrid/Centro or /Madrid/Cortes.

C.5.15.2 HTTP Request

GET /ngsi-ld/v1/entities/?type=OffStreetParking&scopeQ="/Madrid/Centro,/Madrid/Cortes"

Accept: application/ld+json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

C.5.15.3 HTTP Response

200 OK

Content-Type: application/ld+json

```
[
  {
    "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
    "type": "OffStreetParking",
    "scope": "/Madrid/Centro",
    "name": {
      "type": "Property",
      "value": "Downtown One"
    },
    "availableSpotNumber": {
      "type": "Property",
      "value": 121,
      "observedAt": "2017-07-29T12:05:02Z",
      "reliability": {
        "type": "Property",
        "value": 0.7
      },
      "providedBy": {
        "type": "Relationship",
        "object": "urn:ngsi-ld:Camera:C1"
      }
    },
    "totalSpotNumber": {
      "type": "Property",
      "value": 200
    },
    "location": {
      "type": "GeoProperty",
      "value": {
        "type": "Point",
        "coordinates": [
          -8.5,
          41.2
        ]
      }
    },
    "@context": [
      "http://example.org/ngsi-ld/latest/parking.jsonld",
      "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
  },
  {
    "id": "urn:ngsi-ld:OffStreetParking:Corte4",
    "type": "OffStreetParking",
    "scope": [
      "/Madrid/Cortes",
      "/Company894/UnitC"
    ],
    "name": {
```

```

        "type": "Property",
        "value": "Corte4"
    },
    "availableSpotNumber": {
        "type": "Property",
        "value": 121,
        "observedAt": "2017-07-29T12:05:02Z",
        "reliability": {
            "type": "Property",
            "value": 0.7
        },
        "providedBy": {
            "type": "Relationship",
            "object": "urn:ngsi-ld:Camera:C1"
        }
    },
    "totalSpotNumber": {
        "type": "Property",
        "value": 100
    },
    "location": {
        "type": "GeoProperty",
        "value": {
            "type": "Point",
            "coordinates": [
                -8.6,
                41.3
            ]
        }
    },
    "@context": [
        "http://example.org/ngsi-ld/latest/parking.jsonld",
        "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
    ]
}
]

```

C.5.16 Temporal Scope Queries

C.5.16.1 Introduction

EXAMPLE: Give back the *speed* of all the Entities of type "Vehicle" that have been within the Scope /Madrid/Centro between the 1st of August 2018 at noon and the 1st of August 2018 at 01 PM. Note that the value of the Scope has to match for the given timeframe, which means it is possible that it has been set before, e.g. on 1st of August 2018 at 11 AM.

C.5.16.2 HTTP Request

GET /ngsi-ld/v1/temporal/entities/?type=Vehicle&attrs=speed,scope&timerel=between&timeAt=2018-08-01T12:00:00Z&endTimeAt=2018-08-01T13:00:00Z&scopeQ="/Madrid/Centro"

Accept: application/ld+json

Link: <http://example.org/ngsi-ld/latest/aggregatedContext.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

C.5.16.3 HTTP Response

200 OK

Content-Type: application/ld+json

```

[
  {
    "id": "urn:ngsi-ld:Vehicle:B9211",
    "type": "Vehicle",
    "scope": {
      "type": "Property",
      "values": [

```

```

        [
          "/Madrid/Centro",
          "2018-08-01T11:00:00Z"
        ]
      ],
      "speed": {
        "type": "Property",
        "values": [
          [
            30,
            "2018-08-01T12:03:00Z"
          ],
          [
            60,
            "2018-08-01T12:05:00Z"
          ],
          [
            50,
            "2018-08-01T12:07:00Z"
          ]
        ]
      },
      "@context": [
        "http://example.org/ngsi-ld/latest/vehicle.jsonld",
        "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
      ]
    },
    {
      "id": "urn:ngsi-ld:Vehicle:A8311",
      "type": "Vehicle",
      "scope": {
        "type": "Property",
        "values": [
          [
            [
              "/Madrid/Centro",
              "/Company123/UnitA"
            ],
            "2018-08-01T12:10:00Z"
          ]
        ]
      },
      "speed": {
        "type": "Property",
        "values": [
          [
            40,
            "2018-08-01T12:12:00Z"
          ],
          [
            60,
            "2018-08-01T12:14:00Z"
          ],
          [
            50,
            "2018-08-01T12:16:00Z"
          ]
        ]
      },
      "@context": [
        "http://example.org/ngsi-ld/latest/vehicle.jsonld",
        "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
      ]
    }
  ]
}

```

Vehicle B9211 has already been within the Scope `/Madrid/Centro` before the beginning of the request interval, whereas Vehicle A8311 only entered the Scope within the request interval. Thus in the latter case only Property values are included that have been observed after the Scope has become valid.

C.6 Date Representation

In NGSI-LD, a *TemporalProperty* is represented only by its value, i.e. no sub-Properties of *TemporalProperty* nor sub-Relationships of *TemporalProperty* can be conveyed. In more formal language, a *TemporalProperty* does not allow reification. The term *TemporalProperty* has been reserved for non-reified structural timestamps (*observedAt*, *createdAt*, *modifiedAt*, *deletedAt*), which capture the temporal evolution of Attributes. Only such structural timestamps can be used as *timeproperty* in Temporal Queries as mandated by clause 4.11.

The following examples show how time values (*Date*, *Time*, or *DateTime*) can be represented in NGSI-LD as reified Properties. For a reified Property whose value is assigned the JSON type *Date*, *DateTime* or *Time*, one mechanism is to use the Property's *valueType* to hold the datatype ("Date", "Datetime" or "Time"), as shown below:

```
{
  "id": "urn:ngsi-ld:Vehicle:B9211",
  "type": "Vehicle",
  "testedAt": {
    "type": "Property",
    "value": "2018-12-04T12:00:00Z"
    "valueType": "DateTime"
  },
  "@context": [
    "http://example.org/ngsi-ld/latest/vehicle.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}
```

Alternatively the data can be structured to use the JSON-LD *@value* syntax structure, as shown below:

```
{
  "id": "urn:ngsi-ld:Vehicle:B9211",
  "type": "Vehicle",
  "testedAt": {
    "type": "Property",
    "value": {
      "@type": "DateTime",
      "@value": "2018-12-04T12:00:00Z"
    }
  },
  "@context": [
    "http://example.org/ngsi-ld/latest/vehicle.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}
```

A third alternative to achieve the same result would be to use JSON-LD "type coercion". With type coercion, values with a special data type are defined with *@type* in the *@context*. This enforces the correct type for any occurrence. Such an *@context* fragment is shown below:

```
"testedAt": {
  "@type": "https://uri.etsi.org/ngsi-ld/DateTime",
  "@id": "http://example.org/test/testedAt"
}
```

The above does not work, when using the *@context* to perform compaction, in the normalized and compact representation of NGSI-LD, due to reification of the Property, because in this case *testedAt* is a complex JSON object, which cannot be compacted to a *DateTime* type as the *@context* specifies. Thus, the full URI *http://example.org/test/testedAt* is kept, instead of the short name *testedAt*. In summary, user *@contexts* used for the normalized and compact NGSI-LD representation cannot use the JSON-LD type coercion feature.

However, in the simplified (keyValue) representation case, such an *@context* with the specification of *testedAt* could be used, as there is no reification.

As a side note, when using the above *@value* + *@type* approach, since *type* is mapped to *@type* in the NGSI-LD core *@context*, JSON-LD compaction will result in the following compacted value, instead of the one shown above, because *@type* is compacted to *type*:

```
"value": {
  "type": "DateTime",
  "@value": "2018-12-04T12:00:00Z"
}
```

C.7 @context utilization clarifications

When expanding or compacting JSON-LD terms, the JSON-LD *@context* to be used is always the one provided in the current API request. For the benefit of users and implementers the following examples illustrate this concept.

The scenario starts with the creation of an Entity using a JSON-LD *@context* as follows:

```
POST /ngsi-ld/v1/entities/

Content-Type: application/ld+json

Content-Length: 200

{
  "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
  "type": "OffStreetParking",
  "name": {
    "type": "Property",
    "value": "Downtown One"
  },
  "availableSpotNumber": {
    "type": "Property",
    "value": 121,
    "observedAt": "2017-07-29T12:05:02Z"
  },
  "totalSpotNumber": {
    "type": "Property",
    "value": 200
  },
  "@context": [
    "http://example.org/ngsi-ld/latest/parking.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}
```

The content of the *@context* utilized for the referred Entity creation (at <http://example.org/ngsi-ld/latest/parking.jsonld>) is as follows:

```
{
  "OffStreetParking": "http://example.org/parking/OffStreetParking",
  "availableSpotNumber": "http://example.org/parking/availableSpotNumber",
  "totalSpotNumber": "http://example.org/parking/totalSpotNumber",
  "name": "http://example.org/parking/name"
}
```

At Entity creation time the implementation will perform the expansion of terms using the JSON-LD *@context* depicted above.

Now it is needed to retrieve our initial Entity. For retrieving such Entity, this time, a different JSON-LD *@context* is going to be utilized, as follows:

```
{
  "offP": "http://example.org/parking/OffStreetParking",
  "ava": "http://example.org/parking/availableSpotNumber",
  "total": "http://example.org/parking/totalSpotNumber"
}
```

This new *@context*, even though it makes use of the same set of Fully Qualified Names, is defining new short strings as terms. The reasons for that could be multiple: to facilitate data consumption by clients, to save some bandwidth, to enable a more (or less) human-readable response payload body in a language different than English, etc.

In this particular case, the result of the Entity retrieval will be as depicted below. It can be observed that the terms defined by the JSON-LD *@context* **provided at retrieval time** are used to render the Entity content (compaction), and **not** the terms that were provided at creation time (which may be no longer known by the broker).

It is also interesting to note that the *@context* array of the response payload body contains, indeed, our header-supplied *@context*:

```
GET /ngsi-ld/v1/entities/urn:ngsi-ld:OffStreetParking:Downtown1

Accept: application/ld+json
```


Link: <http://example.org/ngsi-ld/latest/parking-abbreviated.jsonld>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

```
{
  "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
  "type": "OffP",
  "name": {
    "type": "Property",
    "value": "Downtown One"
  },
  "ava": {
    "type": "Property",
    "value": 121,
    "observedAt": "2017-07-29T12:05:02Z"
  },
  "total": {
    "type": "Property",
    "value": 200
  },
  "@context": [
    "http://example.org/ngsi-ld/latest/parking-abbreviated.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}
```

Another interesting case to note is the one when an *@context* with no matching terms or no *@context* at all is supplied. See the following example:

GET /ngsi-ld/v1/entities/urn:ngsi-ld:OffStreetParking:Downtown1

Accept: application/ld+json

```
{
  "id": "urn:ngsi-ld:OffStreetParking:Downtown1",
  "type": "http://example.org/parking/OffStreetParking",
  "http://example.org/parking/name": {
    "type": "Property",
    "value": "Downtown One"
  },
  "http://example.org/parking/availableSpotNumber": {
    "type": "Property",
    "value": 121,
    "observedAt": "2017-07-29T12:05:02Z"
  },
  "http://example.org/parking/totalSpotNumber": {
    "type": "Property",
    "value": 200
  },
  "@context": "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
}
```

In this particular case it can be observed that the user names (Entity Type, Attributes) in the response payload body have not been compacted, and as a result the Fully Qualified Names are included. However, the core API terms have been compacted, as the Core *@context* is always considered to be implicitly present if not specified explicitly (and that is why it is included in the JSON-LD response, as mandated by the specification).

C.8 Link header utilization clarifications

The JSON-LD Specification [2] states clearly that **only one HTTP Link header** with the link relationship <http://www.w3.org/ns/json-ld#context> is required to appear. Such statement has implications in terms of providing the JSON-LD *@context* when using the NGSI-LD API. The main implication is that if the *@context* is a compound one, i.e. an *@context* which references multiple individual *@context*, served by resources behind different URIs, then a **wrapper** *@context* has to be created and hosted. The final aim is that only one *@context* is referenced from the JSON-LD Link header. This can be illustrated with an example:

Imagine that it is desired to create an Entity providing *@context* terms which are defined in two different JSON-LD *@context* resources:

- http://example.org/vehicle/v1/vehicle-context.jsonld
- https://schema.org

If a developer wants to reference these two *@context* resources from a Link header, a wrapper *@context* can be easily created as follows:

```
{
  "@context": [
    "http://example.org/vehicle/v1/vehicle-context.jsonld",
    "https://schema.org",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}
```

As such wrapper *@context* needs to be referenced from a Link header by using a URI, then it will have to be hosted at some place on the Web. Usually, developers will host *@context* using popular and simple solutions such as GitHub or GitLab pages. As a result, developers will be able to use *@context* in queries or when using "application/json" as main content type managed by their applications.

It is a **good practice to include the Core @context** in the wrapper *@context* so it can be used, off-the-shelf, by external JSON-LD processing tools. However, it should be noted this is not necessary for NGSI-LD, as the Core *@context* is always implicitly included.

Then, using such wrapper *@context*, (in our example hosted at <https://hosting.example.com/ngsi-ld/v1/wrapper-context.jsonld>), the developer will be able to issue requests like:

POST /ngsi-ld/v1/entities/

Content-Type: application/json

Content-Length: 200

Link: <<https://hosting.example.com/ngsi-ld/v1/wrapper-context.jsonld>>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

```
{
  "id": "urn:ngsi-ld:Vehicle:V1",
  "type": "Vehicle",
  "builtYear": {
    "type": "Property",
    "value": "2014"
  },
  "speed": {
    "type": "Property",
    "value": 121,
    "observedAt": "2017-07-29T12:05:02Z"
  }
}
```

201 Created

Location: /ngsi-ld/v1/entities/urn:ngsi-ld:Vehicle:V1

Link: <<https://hosting.example.com/ngsi-ld/v1/wrapper-context.jsonld>>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

GET /ngsi-ld/v1/entities/urn:ngsi-ld:Vehicle:V1

Accept: application/ld+json

Link: <<https://hosting.example.com/ngsi-ld/v1/wrapper-context.jsonld>>; rel="http://www.w3.org/ns/json-ld#context"; type="application/ld+json"

200 OK

Content-Type: application/ld+json

```
{
  "id": "urn:ngsi-ld:Vehicle:V1",
  "type": "Vehicle",
  "builtYear": {
    "type": "Property",
    "value": "2014"
  },
}
```

```

    "speed": {
      "type": "Property",
      "value": 121,
      "observedAt": "2017-07-29T12:05:02Z"
    },
    "@context": "https://hosting.example.com/ngsi-ld/v1/wrapper-context.jsonld"
  }

```

Observe that in this case the broker is responding with the same wrapper *@context* in the Link header of the HTTP Response or within the JSON-LD response payload body (when MIME type accepted is *"application/ld+json"*). However, that could not be always the case, as there could be situations where the broker could need to provide a wrapper *@context* hosted by itself, for instance, when there are inline *@context* terms or when the Core *@context* has not been previously included by the wrapper *@context* (not recommended) provided within developer's requests.

C.9 @context processing clarifications

JSON-LD Specification [2] says that "If a term is redefined within a context, all previous rules associated with the previous definition are removed". In addition, it is stated that "Multiple contexts may be combined using an array, which is processed in order".

In contrast to the JSON-LD Specification, the NGSI-LD specification states that the Core *@context* is protected and has to remain immutable. This essentially means that the Core *@context* has final precedence and, therefore, is always to be processed as the last one in the *@context* array. For clarity, data providers should place the Core *@context* in the final position. From the point of view of Data providers, care has to be taken so that there are no unexpected or undesired term expansions. See the following example:

```

{
  "id": "urn:ngsi-ld:Building:B1",
  "type": "Building",
  "name": {
    "type": "Property",
    "value": "Empire State"
  },
  "location": {
    "type": "Property",
    "value": "20 West 34th Street, New York City, NY 10001"
  },
  "@context": [
    "https://schema.org/version/latest/schemaorg-current-https.jsonld",
    "http://example.org/ngsi-ld/latest/parking.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}

```

The problem of the example above is that the term *"location"* is defined in both the Core *@context* and the schema.org user *@context* and the Core *@context* takes precedence for the expansion. In these situations, if one wanted to refer to the schema.org *"location"*, the solution is to prefix the conflicting terms, so that there cannot be any clashing. Therefore, if the intent is to refer to <https://schema.org/location> throughout, the example above can be modified as shown below:

```

{
  "id": "urn:ngsi-ld:Building:B1",
  "type": "Building",
  "name": {
    "type": "Property",
    "value": "Empire State"
  },
  "schema:location": {
    "type": "Property",
    "value": "20 West 34th Street, New York City, NY 10001"
  },
  "@context": [
    "https://schema.org/version/latest/schemaorg-current-https.jsonld",
    "http://example.org/ngsi-ld/latest/parking.jsonld",
    "https://uri.etsi.org/ngsi-ld/v1/ngsi-ld-core-context-v1.9.jsonld"
  ]
}

```

Note that the Core *@context* should be placed in the last position of the *@context* array. NGSI-LD implementations are required to render content following this approach, which has been undertaken in order to maximize compatibility with JSON-LD processing tools. This example works because the "schema:" prefix has already been defined by the schema.org *@context*.

C.10 ValueType datatype utilization clarifications

Using JSON-LD [2] syntax, typed values can be expressed using the JSON-LD *@type* keyword when defining a term, where *@type* value holds a URI which indicates the value's datatype. However, it can be desirable for a **Context Broker** to be able to hold simpler untyped values within a Property's *value* attribute and separately use a Property's *valueType* to hold the value's datatype. This format can be used to accommodate multiple acceptable datatype formats for the data to be held within a single Property and still hold sufficient meta data to be able to distinguish between them.

For example, consider an ontology for an Entity of type "Place" which has an *address* Property, where the *address* Property can **either** be an unstructured address in the form of a "String" **or** a structured "PostalAddress" JSON Object with *street*, *city* and *country* attributes - the following JSON-LD schema can be defined:

Example JSON-LD schema

```
{
  "example": "http://example.org/",
  "xsd": "http://www.w3.org/2001/XMLSchema#",
  "address": "example:address",
  "city": "example:city",
  "country": "example:country",
  "street": "example:street",
  "Place": "example:Place",
  "PostalAddress": "example:PostalAddress",
  "String": "xsd:String"
}
```

Then the following two Entities of type "Place" can be created using the *address.valueType* Property to distinguish between the two formats:

```
[
  {
    "id": "urn:ngsi-ld:Place:27182",
    "type": "Place",
    "address": {
      "type": "Property",
      "value": "Pariser Platz, Berlin, Germany",
      "valueType": "String"
    }
  },
  {
    "id": "urn:ngsi-ld:Place:31415",
    "type": "Place",
    "address": {
      "type": "Property",
      "value": {
        "street": "Straße des 17. Juni",
        "city": "Berlin",
        "country": "Germany"
      },
      "valueType": "PostalAddress"
    }
  }
]
```

C.11 Entity with digital signature for a Property

As specified in [35], the atomic piece of information that creators can digitally sign in an NGSI-LD ecosystem is each single Attribute of an Entity. In the following example, an Entity of type "Store" with two Properties, "address" and "location" is presented. The "address" Property is digitally signed. The signature is created using one Ed25519 instantiation of the Edwards-Curve Digital Signature Algorithm (EdDSA). The used crypto suite is "eddsa-rdfc-2022".

EXAMPLE: Entity of type "Store" with two Properties. The "address" Property is digitally signed.

```
{
  "id": "urn:ngsi-ld:Store:002",
  "type": "Store",
  "address": {
    "type": "Property",
    "value": {
      "streetAddress": ["Tiger Street 4", "al"],
      "addressRegion": "Metropolis",
      "addressLocality": "Cat City",
      "postalCode": "42420"
    }
  },
  "ngsildproof": {
    "type": "Property",
    "entityIdSealed": "urn:ngsi-ld:Store:002",
    "entityTypeSealed": "Store",
    "value": {
      "type": "DataIntegrityProof",
      "created": "2025-01-27T21:02:24Z",
      "verificationMethod": "https://example.edu/issuers/565049#z6MkwXG2WjeQnN...Hc6SaVWoT",
      "cryptosuite": "eddsa-rdfc-2022",
      "proofPurpose": "assertionMethod",
      "proofValue": "z3XrH3diVCqpVHXkE7WbnictqyQCKJBGTx...NRTzmuowU1Y2FyqGfSV9eS"
    }
  },
  "location": {
    "type": "GeoProperty",
    "value": {
      "type": "Point",
      "coordinates": [57.5522, -20.3484]
    }
  },
  "@context": "https://uri.etsi.org/ngsi-ld/primer/store-context.jsonld"
}
```

Annex D (informative): Transformation Algorithms

D.1 Introduction

These algorithms are informative but NGSI-LD implementations should aim at either implementing them as they are described here or devising similar algorithms which take exactly the same input and provides exactly the same output (or an equivalent one as per the JSON-LD specification [2]).

D.2 Algorithm for transforming an NGSI-LD Entity into a JSON-LD document (ALG1)

This algorithm takes as input an NGSI-LD graph which top level node is a particular Entity and returns as output a JSON-LD document which represents all the data associated to the entity. The JSON-LD document (and its associated *@context*) corresponds to a representation of the Entity in JSON-LD as per the NGSI-LD Information Model.

NOTE: An early implementation of this algorithm can be found at [i.5].

Let:

- **G** be a graph defined as follows:
 - Let **N** be G's top level node.
 - **N** is an Entity instance of type **T** or types **Ts**. Type name is "AliasT" or there is an Array of Type names ["AliasT1", ..., "AliasTn"], **N**'s identifier is **I**.
 - **N** has 0 or more associated Property. Each Property (**Psi**) is defined as follows:
 - Property type identifier is **Pi**.
 - Property name is "AliasPi".
 - Property Value is **Vi**.
 - Property Value's associated data type is **Di**.
 - **N** is the subject of 0 or more Relationship. Each Relationship is defined as follows:
 - Relationship type identifier is **Ri**.
 - Relationship name is "AliasRi".
 - Relationship target object identifier is **Robji**.
- **O** be a JSON object initialized to the empty object ({}).
- **C** be a JSON-LD *@context* initialized as described by annex B.

The algorithm should run as follows, provided all the preconditions defined above are satisfied:

- 1) Add to **C** a new member <"AliasT", T> or new members <"AliasT1", T1> ... <"AliasTn", Tn>.
- 2) Add to **O** two new members:
 - a) <"id", I>.
 - b) <"type", "AliasT"> or <"type", ["AliasT1", ..., "AliasTn"]> .>.

- 3) For each Property P_i (P_i , "AliasP", V_i , D_i) associated to N :
 - a) Run Algorithm *ALG1.1* taking the following inputs:
 - $P_s \rightarrow P_i$.
 - $O \rightarrow O$.
 - $C \rightarrow C$.
- 4) For each Relationship R_s (R_i , Alias R_i , Robj i) associated to N :
 - a) Run Algorithm *ALG1.2* taking the following inputs:
 - $R_s \rightarrow R_i$.
 - $O \rightarrow O$.
 - $C \rightarrow C$.
- 5) Return (O , C) and end of the algorithm.

D.3 Algorithm for transforming an NGSI-LD Property into JSON-LD (ALG1.1)

Let **P**_s be the Property that has to be transformed. It is defined by (P , "AliasP", V , D), where **P** denotes a Property Type Id, "AliasP" is the Property name, **V** is the Property Value and **D** is the Property Value's data type.

P_s might be associated to extra Properties or Relationships.

Let O be the output JSON-LD object and C the associated JSON-LD context:

- 1) Execute the following steps:
 - a) If no member with "AliasP" is present in O , add a new member to O with key "AliasP" and value an object structure, let it be named **Op** as defined in the following. Otherwise, add all existing members with "AliasP" to a JSON-LD array and in addition put the object structure Op as defined in the following:
 - $\langle \text{"type"}, \text{"Property"} \rangle$.
 - If D is not a native JSON data type add a new member to Op with name "value" and which value has to be an object structure as follows:
 - 1) $\langle \text{"@type"}, D \rangle$.
 - 2) $\langle \text{"@value"}, V \rangle$.
 - Else If D is a native JSON data type add a new member to Op as follows:
 - 1) $\langle \text{"value"}, V \rangle$.
 - b) Add a new member to C as follows:
 - $\langle \text{"AliasP"}, P \rangle$.
 - c) For each Property associated to P_s (P_{ss}) recursively run the present algorithm (*ALG1.1*) taking the following inputs:
 - $P_s \rightarrow P_{ss}$.
 - $O \rightarrow Op$.
 - $C \rightarrow C$.

- d) For each Relationship associated to Ps (Rss) run algorithm *ALG1.2* taking the following inputs:
 - $R_s \rightarrow R_{ss}$.
 - $O \rightarrow O_p$.
 - $C \rightarrow C$.
- 2) Return (O,C) and end of the algorithm.

D.4 Algorithm for transforming an NGSI-LD Relationship into JSON-LD (ALG1.2)

Let **Rs** be the Relationship that has to be transformed. It is defined by (R, "AliasR", Robj), where **R** denotes a Relationship Type Id, "**AliasR**" is the Relationship's name and **Robj** is the identifier of the target object of the Relationship.

Rs might be associated to extra Properties or Relationships.

Let O be the output JSON-LD object and C the current JSON-LD context:

- 1) Execute the following statements:
 - a) If no member with "AliasR" is present in O, add a new member to O with key "AliasR" and value an object structure, let it be named **Or**, and defined as in the following. Otherwise, add all existing members with "AliasR" to a JSON-LD array and in addition put the object structure Or as defined in the following:
 - $\langle \text{"object"}, \text{Robj} \rangle$.
 - $\langle \text{"type"}, \text{"Relationship"} \rangle$.
 - b) For each Property associated to Rs (Pss) run the algorithm *ALG1.1* taking the following inputs:
 - $P_s \rightarrow P_{ss}$.
 - $O \rightarrow O_r$.
 - $C \rightarrow C$.
 - c) For each Relationship associated to Rs (Rss) recursively run the present algorithm *ALG1.2* taking the following inputs:
 - $R_s \rightarrow R_{ss}$.
 - $O \rightarrow O_r$.
 - $C \rightarrow C$.
- 2) Return (O,C) and end of the algorithm.

Annex E (informative): RDF-compatible specification of NGSI-LD meta-model

The content of this annex is now in ETSI GS CIM 006 [i.8].

Annex F (informative): Conventions and syntax guidelines

When new terms are defined, they are marked in bold, and terms are capitalized thereafter.

EXAMPLE 1: **NGSI-LD Linked Entity**, **Linked Entity**.

API Parameter names are always in lowercase.

EXAMPLE 2: *options*.

Entity Types are defined using lowercase but with a starting capital letter.

EXAMPLE 3: **Vehicle**, **Building**, **ParkingSpace**.

JSON-LD nodes and terms are always defined using camel case notation starting with lower case.

EXAMPLE 4: *createdAt*, *value*, *unitCode*.

When referring to special terms, data types or words defined previously in the present document or by other referenced specifications, italics format is used.

EXAMPLE 5: *ListRelationship*, *GeoProperty*, *Geometry*, *Second*, *Number*.

When referring to literal strings double quotes are used.

EXAMPLE 6: "application/json", "Subscription".

When referring to the JSON-LD Context the mnemonic text string *@context* is used as a placeholder.

All the dates and times are given in UTC format.

EXAMPLE 7: 2018-02-09T11:00:00Z.

The measurement units used in the API are those defined by the International System of Units.

EXAMPLE 8: The distance in geo-queries is provided in meters.

When defining application-specific elements or API extensions the same conventions and syntax guidelines should be followed.

Annex G (informative): Localization and Internationalization Support

G.0 Foreword

These algorithms described below are informative, but NGSI-LD implementations should aim at either implementing them as they are described here or providing equivalent *@context* elements for their payloads to provide interoperability with their internationalized context entities.

G.1 Introduction

G.1.0 Foreword

Since Internationalization is not core to context information management, any direct support within NGSI-LD systems is limited. Annex G proposes a series of best practices for maintaining, querying and displaying interoperable internationalized data.

The content of the *@context* utilized for the referred Entities within these examples uses pre-existing URNs used for internationalization and is as follows:

```
{
  "inLanguage": "http://schema.org/inLanguage",
  "sameAs": "http://schema.org/sameAs"
}
```

G.1.1 Associating an Entity with a Natural Language

Where a context Entity is associated with a single natural language, include a well-defined *Property* indicating the natural language of the content. For example an Event taking place in French may be defined as follows:

```
{
  "type": "Event",
  "id": "urn:ngsi-ld:Event:bonjourLeMonde",
  "name": {
    "type": "Property",
    "value": "Bonjour le Monde"
  },
  "description": {
    "type": "Property",
    "value": "«Bonjour le monde» sont les mots traditionnellement écrits par un programme informatique simple"
  },
  "inLanguage": {
    "type": "Property",
    "value": "fr"
  }
}
```

G.1.2 Associating a Property with a Natural Language

Where a *Property* of a context entity can be associated to one more natural language, include additional metadata as a sub-Property of that Property. For example, a Hotel with booking forms available in English, French and German may be defined as follows:

```
{
  "type": "Hotel",
  "id": "urn:ngsi-ld:Hotel:XXXXX",
  "name": {
    "type": "Property",
    "value": "Grand Hotel"
  }
}
```

```

    },
    "bookingUrl": {
      "type": "Property",
      "value": [
        "http://example.com/booking-in-french/",
        "http://example.com/booking-in-english/",
        "http://example.com/booking-in-german/"
      ],
      "inLanguage": {
        "type": "Property",
        "value": ["fr", "en", "de"]
      }
    }
  }
}

```

G.1.3 Associating as equivalent entity

Where equivalent context entities in multiple natural languages exist, they may be associated with each other through the use of a one-to-many relationship, where each relationship holds an additional sub-Property indicating the natural language of the equivalent entities.

For example, three Events (such as a walking tour which is available in English, French and German) may be associated to each other as follows:

```

{
  "type": "Event",
  "id": "urn:ngsi-ld:Event:bonjourLeMonde",
  "name": {
    "type": "Property",
    "value": "Bonjour le Monde"
  },
  "sameAs": [
    {
      "type": "Relationship",
      "datasetId": "urn:ngsi-ld:Relationship:1",
      "object": "urn:ngsi-ld:Event:helloWorld",
      "inLanguage": {
        "type": "Property",
        "value": "en"
      }
    },
    {
      "type": "Relationship",
      "object": "urn:ngsi-ld:Event:halloWelt",
      "inLanguage": {
        "type": "Property",
        "value": "de"
      }
    }
  ]
}

```

G.2 Natural Language Collation Support

G.2.0 Foreword

All strings within an NGSI-LD system are defined and sorted as a sequence of Unicode characters. As such there is no simple collation mechanism to query entities ignoring case, diacritic marks or matching diphthong single letters such as the German "ö" to also match with "oe".

Many databases support a degree of natural language support, in general collation support will always depend upon the underlying database and as such will vary from implementation to implementation. This therefore cannot be standardized and exposed as part of the context information management API. Furthermore, collation is slow and processor intensive, and for massive systems is better achieved using a separate index.

For systems that require it, this clause proposes a mechanism as an extension to a NGSi-LD Context Broker which can be modified and used to offer collation support to the natural language attributes found within context entities where necessary through creating, querying and maintaining an additional property of a property for collated attributes.

G.2.1 Maintain collations as metadata

- Create a subscription on the attribute (e.g. *name*)
- Create a simple microservice to add/upsert a *name.collate* property-of-a-property using a simple function to strip all diacritic marks - for example:

```
str.normalize("NFD").replace(/[^\u0300-\u036f]/g, "").toLowerCase()
```

Other substitutions could be made where local spelling rules vary (for example different for German ö = oe).

G.2.2 Route language sensitive queries via a proxy

Create a simple forwarding proxy around the NGSi-LD system. For any urls with a *q* param (and a *collate* flag) run a clean-up of the *q* param and amend the query string:

The following request on the proxy:

```
GET /ngsi-ld/v1/entities/?type=Building&q=name==%22Schöne%20Grüße%22&collate=name
```

is altered on the fly and is sent to the NGSi-LD system as shown:

```
GET /ngsi-ld/v1/entities/?type=Building&q=name.collate==%22schoene%20gruesse%22
```

Once again, the substitutions to make to the query string will depend on the rules of the natural language to be supported.

G.3 Localization of Dates, Currency formats, etc.

G.3.0 Foreword

Context data entities are designed to be interoperable and therefore all dates are held as UTC dates, all currency amounts are held as JSON numbers (with the unitCode property-of-a-property available to hold the currency), etc. Localization should not occur within the context data entities themselves. Offering fully localized responses is not a concern of the NGSi-LD API.

If localization support is necessary, a simple proxying a conversion mechanism could be used to amend the context data received from the NGSi-LD system before being passed to a third party system for display.

G.3.1 Localizing Dates

For example, if a system needs to display *DateTime* data in Islamic Date format

The following request on the proxy:

```
GET /ngsi-ld/v1/entities/urn:ngsi-ld:Event:XXX?attrs=date&format=simplified
```

is forwarded unaltered and is sent to the NGSi-LD system as shown:

```
GET /ngsi-ld/v1/entities/urn:ngsi-ld:Event:XXX?attrs=date&format=simplified
```

The response from the NGSi-LD system is always in UTC format:

```
{"date": "2020-09-28T17:13:39+02:00"}
```

And the proxy can be used to update this to the desired format:

```
{"date": "11 Safar, 1442 1:13:39PM"}
```

Using an internationalization script such as the following:

```
new Intl.DateTimeFormat("en-u-ca-islamic", {day: 'numeric', month: 'long', weekday: 'long', year :  
'numeric'}).format(date);
```

It should be noted that post-localization, the transformed date is no longer valid NGSI-LD.

Annex H (informative): Suggested actuation workflows

H.1 Actuators and feedback to the consumer

Actuators are things that can change their state (light on/off) or execute actions (move forward, detect face, etc.).

There is currently no explicitly and precisely specified support for actuation in the NGSI-LD API. Thus, this clause describes some conventions that represent a proposed best-practice about how NGSI-LD API and data models can be used for the interaction between applications and actuators represented by NGSI-LD Entities.

The conventions and approach described in this clause are not powerful enough to implement complex actuation jobs that depend on each other and, for instance, make actuation decisions conditional on the outcome of other actuations, unless that behaviour is implemented in a custom way within the application logic. The concept of a more evolved service execution logic, being a first-class citizen of the NGSI-LD API and able to offer more structured building blocks for actuation, is outside the scope of this annex.

An NGSI-LD system that comprises an actuator and supports actuation workflows is represented as one or more NGSI-LD Entities, plus a `Context Broker`, a `Context Source` or a `Context Producer`, and a `Context Consumer`, which collaborate.

The interaction between actuator and `Context Consumer` needs to be bidirectional. Thus, actuators are triggered by the reception of actuation-specific commands (e.g. "set the on state of the lamp to false", to turn the light off) that are encoded as NGSI-LD data, following a suggested data model. They respond with feedback, similarly encoded as NGSI-LD data.

Command feedbacks may serve to control the maximum operations rate a controlling application needs to achieve, and different levels of feedback can be requested, by associating a specific Quality of Service value to the command:

- Some applications need high operation rate but no feedback. For this case a $QoS = 0$ can be used. The typical example is to control the arms of a robot with a joystick.
- Some applications need to be sure that the actuators actually received the command request or need to get back a payload in response to the command. In this case a $QoS = 1$ can be used. The typical case is switching on a light with confirmation, or request face-detection with consequent notification of matching events.
- Commands can either require a short or a long execution time. For commands with long execution time, the application may require a continuous status feedback. In this case a $QoS = 2$ can be used. The typical example is that of a door opening, where feedback continuously reports the current level (10 % open, 50 % open, etc.).

H.2 Architecture for actuation

In this architecture, the application acts as `Context Consumer`, and the terms are used interchangeably.

Commands are sent to the `Context Broker` by the `Context Consumer`, using the standard NGSI-LD API and a suggested convention for representing them. In turn, feedback about command execution is received by the `Context Consumer`, both as continuous status updates and/or a final command result.

More specifically, the component that handles direct communication with the actuator is the `Context Source` or the `Context Producer`: it uses an actuator-specific protocol to control the actuator and get responses and updates from it, i.e. from the real system. Such `Context Source/Consumer` or `Context Producer/Storage` acting as a proxy or intermediary to the actuator is referred to, in the following, as `Context Adapter`.

Thus, the `Context Adapter` is able to use the NGSI-LD API to receive NGSI-LD command requests from the NGSI-LD `Context Broker` and send back command status and result to it, as well as using an actuator-specific protocol to communicate with the actuator.

The NGSI-LD `Context Broker` is responsible for handling direct communication with the `Context Consumer`.

Thus, to support actuation, there is a need to specify:

- Additional NGSI-LD Properties the NGSI-LD system has to have, in order to represent and manage command Request, Status, Result.
- A communication model that allows commands to flow in forward direction and feedback to flow in reverse. This communication model has to comprise a mapping, to be held within the NGSI-LD system, that is able to route the command requests to the appropriate handler within the Context Adapter and vice-versa.




Figure H.2-1: Architecture for actuation

H.3 Structure of Commands and additional Properties

H.3.0 Introduction

The NGSI-LD system has, in addition to the usual NGSI-LD Properties representing the actuator's status, a set of additional, dedicated NGSI-LD Properties associated with:

- the list of available commands, i.e. the list of commands supported by the actuator;
- command endpoints, one for each command, that are used to send and receive command related messages and optionally hold state for the ongoing commands.

The structure of the commands needs to be specified, but not the internal format of their payloads. By using commands with a custom payload, one can support all kinds of operations, for example:

- "set-on": "true"
- "detect-face": {"face-features": "..."}
- "move": "forward"

The data model for command requests, status and responses has to include metadata such as the QoS of the command, its identifier, and the custom payload itself.

Both the requests/responses and the list of commands the NGSI-LD system is able to support can be represented with additional NGSI-LD Properties, as follows.

H.3.1 Property for listing available commands

The additional Property dedicated to the list of available commands is as follows:

```
"commands": {
  "type": "Property",
  "value": ["<cmd_name1>", "<cmd_name2>", ..., "<cmd_nameN>"]
}
```

It is a *Property* whose value is an array of Strings, each string representing the unique name of a supported command.

H.3.2 Properties for command endpoints

For each available command, a set of three endpoints is to be additionally created within the NGSI-LD system, by means of three dedicated Properties per command. The first endpoint will manage that command's requests, the second endpoint will manage its status, and the third endpoint will manage command's results.

This convention dictates that:

- The NGSI-LD Property that manages requests has the same name as the command, e.g. "<cmd_name1>".
- The NGSI-LD Property that manages results has the same name as the command plus the "-RESULT" suffix.
- The NGSI-LD Property that manages status has the same name as the command plus the "-STATUS" suffix.

Each endpoint can receive multiple requests or responses, and it supports queueing of messages. For example, the command *moveToLocation* may receive a sequence of requests that are to be stored in an array and orderly processed depending on the arrival timestamps. A number of respective responses may be produced, as well. Thus, each endpoint, represented by its dedicated NGSI-LD Property, exploits the multi-Attribute feature (see clause 4.5.5), as follows:

Command Request endpoint

```
"<cmd_name>": {
  "datasetId": a URI uniquely identifying the specific command request
                (optional, if the use case does not need command queueing),
  "type":      "Property",
  "qos":       an Integer, representing the desired QoS (optional, default=0),
  "value":     custom parameters of the command (mandatory)
}
```

Command Status endpoint

```
"<cmd_name>-STATUS": {
  "datasetId": a URI uniquely identifying the specific status feedback message
                (optional, if the use case does not need queueing them),
  "type":      "Property",
  "value":     custom status of the command (mandatory)
}
```

Command Result endpoint

```
"<cmd_name>-RESULT": {
  "datasetId": a URI uniquely identifying the specific result feedback message
                (optional, if the use case does not need queueing them),
  "type":      "Property",
  "value":     custom result of the command (mandatory)
}
```

Usually, the Context Adapter (or the actuator behind it), upon receiving a command request with a specific *datasetId*, will then generate status and result with the same *datasetId*, so that, when the status/result is received by the application, it can link it back to the corresponding command that is generating the received feedback. The value of the request, status and result is generic, and it is up to the specific application to define useful values. As an example, the PackML language for the control of packaging machines defines a set of possible values for statuses during an actuation workflow.

EXAMPLE 1: An example follows, where the NGSI-LD system represents a simple actuator. The example shows the NGSI-LD Entity representing a light that can change colour by manipulation of its brightness, hue and saturation values; further, it is possible to turn the lamp on and off. Apart from the *id* and the *type*, the actuator entity has a set of regular properties that represent the current status of the lamp. In the example these are *colorRGB* and *is-on*. Then it has the conventional Property named *commands*, signalling that it supports four commands: "turn-on", "set-saturation", "set-brightness", "set-hue". Further, it has four (times three) additional properties serving the purpose of command endpoints.

```
{
  "id": "urn:ngsi-ld:pHueActuator:light1",
  "type": "Lamp",

  REGULAR PROPERTIES
  "colorRGB": {"type": "Property", "value": "0xABABAB"},
  "is-on": {"type": "Property", "value": true},

  AVAILABLE COMMANDS
  "commands": {
    "type": "Property",
    "value": ["turn-on", "set-saturation", "set-hue", "set-brightness"]
  }

  COMMAND ENDPOINTS
  "turn-on": {"type": "Property", "value": <custom request>}
  "turn-on-STATUS": {"type": "Property", "value": <custom status>}
  "turn-on-RESULT": {"type": "Property", "value": <custom response>}
  "set-hue": ...
  "set-hue-STATUS": ...
  "set-hue-RESULT": ...
  ...
}
```

EXAMPLE 2: The following example shows an NGSI-LD Entity Fragment that can be used as a command request to request that the lamp be turned off.

```
{
  "id": "urn:ngsi-ld:pHueActuator:light1",
  "type": "Lamp",
  "turn-on": {
    "type": "Property",
    "qos": {
      "type": "Property",
      "value": 1
    },
    "value": false
  }
}
```

EXAMPLE 3: In the following example, the value of the command request is a more complex JSON Object, to show that complex actions can be conveyed by just one request. Further, the request has an identifier that makes it possible to enqueue it, together with other request that may arrive to the same command endpoint within a timespan.

```
{
  "id": "urn:ngsi-ld:pHueActuator:light1",
  "type": "Lamp",
  "set-hue": {
    "type": "Property",
    "qos": {
      "type": "Property",
      "value": 1
    },
    "datasetId": "myapp:mycommand:1342",
    "value": {"red": "1 seconds", "green": "2 seconds"}
  }
}
```

In summary, the suggested convention prescribes a *commands* property that contains a list of commands supported by the actuator. For each of these commands, the convention requires a command endpoint consisting of three properties, the name of the command, e.g. "turn-on", the status property, which is the name of the command with "-STATUS" as suffix, and the result, which is the name of the command with "-RESULT" as suffix. Nevertheless, it is noted that such suffixes are just a convention to distinguish the endpoints. So far, two practical implementations exist, see clauses H.5 and H.6, that adopt the general scheme of this convention, with minor deviations. In fact, this convention is derived as a generalization that leverages the full potential of NGSI-LD sub-Attributes and multi-Attributes.

H.4 Communication model

H.4.1 Possible communication models

This convention can be leveraged by two different communication models:

- Subscription/notification, where both the application and the Context Adapter use NGSI-LD Subscriptions to have the command requests delivered to the appropriate handler within the Context Adapter and vice-versa. In this case the Context Adapter acts as a **Context Source** as well as a **Context Consumer**.
- Forwarding, which uses the NGSI-LD Registry and a Context Adapter able to federate itself with the **Context Broker** holding the actuator's Entity, as a means to deliver the commands. In this case the Context Adapter acts as a **Context Storage** as well as a **Context Producer**.

H.4.2 Subscription/notification model

For the interaction to work, the Context Adapter, acting as a proxy to the actuator, subscribes to all command properties; in example 1 of clause H.3.2, these are "set-brightness", "set-saturation", "set-hue" and "turn-on". When the application, acting as the actuation client, updates the value of a command property, the Context Adapter will receive the notification with the new value. This will be translated into the proprietary format and forwarded to the actuator using the actuator-specific protocol. The application in turn can subscribe to the command status and the result. The Context Adapter updates the status of the actuation during the execution of the command, which is primarily relevant in the case of longer-lasting actuations, and finally updates the result once the actuation has been completed. If the application has subscribed to the status and result, it will receive the corresponding notifications. Independent of the command-related properties, the status of the actuator, held within its regular properties, will be updated.

The detailed workflow is depicted in Figure H.4.2-1, and can be interpreted as follows:

- 1) Application updates *turn-on* command Property with "value": true
- 2) Context Adapter gets notification of the new value true
- 3) Context Adapter updates *turn-on-STATUS* command Property with "value": "PENDING"
- 4) Application gets notification of the new value "PENDING"
- 5) Context Adapter updates *is-on* regular Property with "value": true
- 6) Application gets notification with value: true
- 7) Context Adapter updates *turn-on-RESULT* command Property with "value": "OK"
- 8) Application gets notification with of the new value "OK"




Figure H.4.2-1: Steps of the actuation workflow using subscription/notification

H.4.3 Forwarding model

The forwarding model uses registrations and forwarding of requests. Actuation of commands is provisioned via registration(s) to the NGSI-LD Registry done by the Context Adapter that states "I am responsible for command property <X>". When the Application changes the value of a command property, first the NGSI-LD Context Broker asks to the NGSI-LD Registry whether the property is delegated to some other component. The NGSI-LD Registry knows that property <X> of the Entity is delegated to the Context Adapter. Hence, the request is forwarded to the Context Adapter. Similar to the other communication model, the request will then be translated into the proprietary format and forwarded to the actuator using the actuator-specific protocol.

In this model, the NGSI-LD Entity is distributed over two different components, because some of its properties live in the Context Brokers and other properties live in the Context Adapter, as indicated in Figure H.4.3-1 with a dotted rectangle.

The rest of the workflow, i.e. delivery of status and result messages to the application, is done similarly to the subscription/notification model. The detailed workflow is depicted in Figure H.4.3-1, and can be interpreted as follows:

- 1) Application updates *turn-on* command Property with "value": true
- 2a) Context Broker ask Registry where to forward the request
- 2b) Context Broker forwards request to Context Adapter
- 3) Context Adapter updates *turn-on-STATUS* command Property with "value": "PENDING"
- 4) Application gets notification of the new value "PENDING"
- 5) Context Adapter updates *is-on* regular Property with "value": true
- 6) Application gets notification with value: true
- 7) Context Adapter updates *turn-on-RESULT* command Property with "value": "OK"
- 8) Application gets notification with of the new value "OK"




Figure H.4.3-1: Steps of the actuation workflow using forwarding

H.5 Implementation of the subscription-based actuation workflow

The Fed4IoT project (<https://fed4iot.org>) leverages the NGSI-LD architecture and the subscription/notification workflow for actuation, in order to implement the concept of a Cloud of Things. It enables virtualization of existing IoT sensors/actuators through Virtual Things and IoT Brokers. IoT application developers can simply rent the Virtual Things and the Brokers their applications need.

The Fed4IoT's Cloud of Things is named VirIoT (<https://github.com/fed4iot/VirIoT>), and it is based on the concept of Virtual Silos as-a-service: isolated and secure IoT environments made of Virtual Things whose data can be accessed through standard IoT Brokers (oneM2M, NGSI, NGSI-LD, etc.).

In Figure H.5-1 a diagram shows how VirIoT implements the concept of a large-scale and distribute NGSI-LD system that leverages the architecture and the workflow convention described in clause H.4.2.




Figure H.5-1: VirIoT implementation of the NGSI-LD system and actuation workflow

All components encapsulate requests in a neutral-format message that leverages NGSI-LD Entities at its core. But, since VirIoT uses MQTT as its internal data distribution system, all information and actuation commands are encoded as NGSI-LD entities, plus an additional "meta header" that is used by the MQTT to publish and subscribe in a broadcast fashion to multiple vThings, because the same virtual sensor can be used by multiple applications at the same time.

For the actuation workflow, the "data" part of this message contains the command request, as specified in clause H.3, but with an additional value key that is the "command notification uri" (`cmd-nuri`), representing a location where feedback (status, result) should be sent by the ThingVisor. For example, the `cmd-nuri` contains the "data_in" MQTT topic of the issuing vSilo, so that command feedbacks (status and results) are sent to it, only, instead of being broadcasted to all subscribing applications.

VirIoT is agnostic to access control issues to a virtual actuator, since the relevant policies are implemented in the specific ThingVisor, which can grant tokens to execute actuation-commands to a subset of vSilos only, through preliminary exchange of specific actuation-commands (a kind of log-in).

Fed4IoT has developed several different ThingVisors (Context Adapters for different sensors and hardware): for example, lamp systems and robot devices are virtualized through specific ThingVisors, and applications can control the lighting system of a rented conference room or control camera and position of a bot by adding related virtual actuators to their vSilo.

H.6 Implementation of the registration-based actuation workflow

The IoT Agent node library [i.22] introduces the concept of an IoT Agent, which is a component that lets a group of devices send their data to and be managed from a **Context Broker** using their own native protocols. Thus, it corresponds to the Context Adapter, and wires up the IoT devices so that measurements can be read and commands can be sent using NGSI-LD requests sent to an NGSI-LD compliant context **Context Broker**.

IoT Agents already exist or are in development for many IoT communication protocols and data models. Examples include the following:

- IoTAgent-JSON - a bridge between HTTP/MQTT messaging (with a JSON payload) and NGSI-LD
- IoTAgent-LWM2M - a bridge between the Lightweight M2M protocol and NGSI-LD
- IoTAgent-UL - a bridge between HTTP/MQTT messaging (with an UltraLight2.0 payload) and NGSI-LD
- IoTAgent-LoRaWAN - a bridge between the LoRaWAN protocol and NGSI-LD

This implementation follows the communication model described in clause H.4.3, as explained in Figure H.6-1. In this workflow:

- Requests between User and Context Broker use NGSI-LD
- Requests between Context Broker and IoT Agent use NGSI-LD
- Requests between IoT Agent and IoT Device use native protocols
- Requests between IoT Device and IoT Agent use native protocols
- Requests between IoT Agent and Context Broker use NGSI-LD




Figure H.6-1: IoT Agent node library implementation of the NGSI-LD system and actuation workflow

Provisioning of the devices will be carried out (via REST API) through IoT Agents, as well. This provisioning implies that, on the one hand, the corresponding entities (with their commands), that represent the devices, are generated in the **Context Broker** and, on the other hand, that the corresponding IoT Agent is configured for communication with the corresponding device, all in one provisioning step. Below, an example how to provision a device which supports start and stop commands is presented.

```
{
  "devices": [
    {
      "device_id": "device001",
      "entity_name": "urn:ngsi-ld:Device:001",
      "entity_type": "Device",
      "attributes": [
        { "object_id": "s", "name": "isOpen", "type": "boolean" }
      ],
      "commands": [
        { "name": "start", "type": "command" },
        { "name": "stop", "type": "command" }
      ],
      "static_attributes": [
        { "name": "provision", "type": "Text", "value": "Device:001 provision" }
      ]
    }
  ]
}
```


Annex I (informative): Change history

Date	Version	Information about changes
February 2020	1.2.10	Early draft copied from API version 1.2.1
February 2020	1.2.11	Unicode characters. Query Language syntax changes to Attribute path, and extension to accept specifying just Query or Geoquery when Querying Entities Acknowledgements to EU projects. Lightweight Figures
March 2020	1.2.12	Extending to other interactions the above changes to query entities interaction Changes to ABNF Query Language syntax to access complex objects value of properties more easily Generalized Notification Headers, in order to carry authentication etc., info Novel &option=count and associated Header to indicate number of Entities in response to a query Novel/entityOperations/query and/temporal/entityOperations/query endpoints to perform query via POST Clarified attrs URL parameter behaviour Support for Multiple Attributes Support for Multiple Tenants
May 2020	Candidate 1.2.13	from 101r1: Multi-Attribute-Support-fix-in-4.5.5 from 102r1: Batch_Operation_Error_Codes from 110r1: JSON-LD Validation clause from 112r1: IRI Support for International Characters from 115r2: More Core Context Changes from 130: Entity Types MQTT Notifications GeoJSON Representation
9 July 2020	1.3.1	Technical Officer verifications for submission to editHelp! publication pre-processing
August 2020	1.3.2	New baseline towards v1.4.1
November 2020	1.3.3	From 272r1: Support for natural languages via LanguageProperty; annex G
December 2020	1.3.4	From 319: Align Table 6.8.3.2-1 with clause 5.10.2 for query via attrs
December 2020	1.3.4	From 310r2: Dot vs. comma in DateTime
December 2020	1.3.4	From 309r1: Remove sentences referring to old multi attributes representation
December 2020	1.3.4	From 308r: id and type for JSON-LD compliance
December 2020	1.3.4	From 313r1: FIXES to Cross domain data model for LanguageProperties Bug fixes and errata
December 2020	1.3.5	From 275r3: Temporal Aggregation Functions
December 2020	1.4.0	1.3.5 with small typos corrected, approved as 1.4.0
January 2021	1.4.1	ETSI Technical Officer review for ETSI EditHelp publication pre-processing
March 2021	1.4.2	Editorial Changes, clarifications added, better references, figures replacements and corrections, figures merged, typos, code indentation
April 2021	1.4.2	Temporal Pagination
April 2021	1.4.2	Clarified behaviour when multiple instances of the same Entity are in an input array
July 2021	1.4.3	From 130r6: NGSI-LD Scope
July 2021	1.4.3	From 143r6: Storing, managing and serving @contexts
July 2021	1.4.3	From 120r4: API structuring
October 2021	1.4.4	From 156: Remove static elements from temporal representations
October 2021	1.4.4	From 155: Existence query
October 2021	1.4.4	From 152: Remove null value deletion
October 2021	1.4.4	From 151: attrs missing in core context
October 2021	1.5.1	ETSI Technical Officer review for ETSI EditHelp publication pre-processing
November 2021	1.5.2	First early draft after EditHelp publication of V1.5.1 to prepare next V1.6.1 publication
January 2022	1.5.3	Concise representation
May 2022	1.5.4	PUT/PATCH Entity
May 2022	1.5.4	Distributed operations
July 2022	1.5.5	From 99r6: Deletions and advanced notifications
July 2022	1.5.5	From 106r1: Workflow for actuation
July 2022	1.5.5	From 105r1: Context Source Info in Context Source Registration
July 2022	1.5.5	From 93r1: default context clarification
July 2022	1.5.5	From 91r1: CR_on_Scope_ABNF_format
July 2022	1.6.1	Final Technical Official check for EditHelp publication

Date	Version	Information about changes
October 2022	1.6.2	New early draft: corrected Annex C6 date representation from 149r3: generalized description of @context in bullet lists Fixed usage of NGSI-LD Null in Attributes' definitions
December 2022	1.6.4	From 188r2_Registration_Clarifications
December 2022	1.6.4	From 182r2_Add_NGSI-LD_Roles_for_Context_Registration
December 2022	1.6.4	From 156r2_VocabProperty/URI_type_coercion
December 2022	1.6.4	177r2 Clarify usage of Accept, Content-Type and Linked @context when forwarding to partial Context Brokers
December 2022	1.6.4	178 Add Batch Query to Federation Ops
December 2022	1.6.4	183r1 Clarify Temporal query behaviour
December 2022	1.6.4	149r4 Forbid scoped and nested @contexts
December 2022	1.6.4	023006 Fixing CSource registration example in C.3
December 2022	1.6.4	Fix: Tenants URI becomes String
December 2022	1.6.4	Fix: POST-QUERY-COUNT-PAGINATION
December 2022	1.6.4	Fix: CHECK-URI-PARAM
December 2022	1.6.4	Updated examples and text to context v1.7.jsonld
March 2023	1.6.6	CIM(23)000006_Adding_previousValue_to_GeoProperty_type_definition
March 2023	1.6.6	cSource -> CSource; "implementations shall do the following"
March 2023	1.6.7	000013r4_Updated_Distributed_Execution_Behaviour
March 2023	1.6.8	CIM(22)000195r3_type_passing_for_M2M_callReviewed
April 2023	1.6.9	Fixing URI→String datatypes
June 2023	1.7.2	CIM(23)000053r1_Entity_Graph_Retrieval (for FIWARE SUMMIT)
June 2023	1.7.3	000056r2_APIv172_towards_v18.docx (for FIWARE SUMMIT)
October 2023	1.7.4	From 25023r2: Use Temporal Evolution instead of Temporal Representation + Updated figures in clause 5 and 6
November 2023	1.7.5	From 68r5: Additional id only format and attribute projection via pick and omit From 121r1: Relationship as Array From 123r1: URN Namespace Definitions From 149r3: Allow Broader Local Requests From 153r1: JsonProperty From 159: Bug fixed in CIM 009: GeoJSON responses in figures From 160: Replace Attribute fix
December 2023	1.7.6	From 164: Host alias /info Endpoint From 154r2: EntityMap Updated figure in clause 6.2
January 2024	1.7.7	From 168r1: 504 error instead of 503 in JSON-LD context endpoints From 169r1: Allow forbidden characters From 170: Remove Scope from PATCH /attrs operations From 1005r2: URI for value of several attributes From 173r2: Clarify match in distributed operations From 174: Protect core context From 25002r2: API Issue #10 Filter on value with specific datasetId
January 2024	1.7.8	From 164r5: fix Tenant in Host Alias (164) and /info/sourceldentity Endpoint + figure Updated figure in clause 6.2 Updated figures in clause 4.2
January 2024	1.7.9	CIM(24)000007r2_Query_Language_Extension_for_Linked_Entity_Retrieval
January 2024	1.7.10	Internal revision, cleanup
January 2024	1.7.10	FIX: CIM(24)000015_Projection_attributes_error_messaging
January 2024	1.7.10	FIX: CIM(24)000014r1_POST_Query_Parameters
January 2024	1.7.11	TO revision
January 2024	1.7.12	ISG CIM revision and cleanup after TO revision. New keywords
February 2024	1.7.13	Footnotes in Tables
February 2024	1.7.15	Change of NGSI-LD Term style to HTML Keyboard
February 2024	1.7.16	Added expiresAt to @context serving
March 2024	1.8.1	Published
April 2024	1.8.2	Clone of published
April 2024	1.8.3	Minor typos and style cleanup
April 2024	1.8.3	000048_Fix_operations_allowing_sysAttrs.docx
April 2024	1.8.3	000047r2_Accept_header_in_case_of_partial_success_207_.docx
April 2024	1.8.3	000049r1_Clarify_options_allowed_for_POST_queries.docx
April 2024	1.8.4	Track changes removed
May 2024	1.8.5	TooLargeResponse
June 2024	1.8.6	CIM(24)000036r1_Loop_Detection

Date	Version	Information about changes
June 2024	1.8.7	CIM(24)000033r6 Backwards Compatibility Indicators
June 2024	1.8.8	CIM(24)000027r5 Value Type
June 2024	1.8.9	CIM(24)000029 Purge Entities
June 2024	1.8.10	CIM(24)000028r4 Transient Entity bugfixed
June 2024	1.8.11	CIM(24)000070 Additional Updates to ExpiresAt Conformance etc
October 2024	1.8.12	CIM(24)000076r8 Entity Maps and Split Entities
October 2024	1.8.13	CIM(24)000108 Adding Missing Elements in Core Context and Data Types.zip
October 2024	1.8.14	Updated figures and new baseline and created Stable Draft out of this
November 2024	1.8.15	CIM(24)000128r1 Signed Attributes and Scoped Context
December 2024	1.8.16	VocabProperty instead of VocabularyProperty in C2.2 and Table 5.2.35-1
January 2025	1.8.17	Reordering table rows in alpha (almost) order
January 2025	1.8.18	Ngsildproof example and @context. Switching to new @context URI for v1.9
February 2025	1.8.19	CIM(25)000005 temporal bounds clarifications and misc fixes
March 2025	1.8.20	CIM(24)000111r5 Ordered Entites
March 2025	1.8.21	CIM(24)000138r4 Snapshot
March 2025	1.8.22	CLEAN. Removed all track changes. Comments to be still addressed
March 2025	1.8.23	CIM(25)000011r1 Updated figures and text for Snapshots (clauses 5 and 6)
March 2025	1.8.24	CIM(25)000012r1 Harmonize output and 203
April 2025	1.8.25	CIM(25)000014 GS CIM 009 v1825 fixSubscriptions
April 2025	1.8.26	CIM(25)000016 Explanation of valueType as data type
April 2025	1.8.27	CIM(25)000015 Precision Clarification + Harmonize captions in clause 6 + addressing fixes requested in the comments
May 2025	1.8.28	CIM(25)000023 Order Table rows + small editorial fixes + aggregation parameters for POST query
May 2025	1.8.29	Final review prior to EditHelp
June 2025	1.8.30	TO review: fuzzy figures, table numbering. Editorial: fixed core context and value in ngsildproof example

History

Version	Date	Status
V1.1.1	January 2019	Publication
V1.2.1	October 2019	Publication
V1.2.2	February 2020	Publication
V1.3.1	August 2020	Publication
V1.4.1	February 2021	Publication
V1.4.2	April 2021	Publication
V1.5.1	November 2021	Publication
V1.6.1	August 2022	Publication
V1.7.1	June 2023	Publication
V1.8.1	March 2024	Publication
V1.9.1	July 2025	Publication