ETSI TS 104 178 V1.1.1 (2026-07)



TECHNICAL SPECIFICATION

Data Solutions (DATA);
NGSI-LD Information Model

Sample Document

get full document from standards.iteh.ai

Reference
DTS/DATA-00104178

Keywords
information model, interoperability, smart cities

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 of this document may be reproduced in any form, by any means and in any media, without the prior written authorization of ETSI and except as expressly permitted below.
By way of exception and when the document is a normative deliverable (European Standard (EN), Technical Specification (TS), Group Specification (GS) or ETSI Standard (ES)), ETSI authorizes to reproduce and incorporate into products, services and technical documentation only those extracts (e.g. templates) that are strictly necessary for the technical implementation of the normative deliverable, to ensure compliance with the latter.
Nothing in this notice shall be construed as limiting any mandatory exceptions to copyright provided by applicable law.

© ETSI 2026.
All rights reserved.

ETSI TS 104 178 V1.1.1 (2026-07) • https://standards.iteh.ai/catalog/standards/etsi/dd123511-cac3-4ef1-b028-bbab27139033/etsi-ts-104-178-v1-1-1-2026-07

Contents

Intellectual Property Rights 5
Foreword 5
Modal verbs terminology 5
1 Scope 6
2 References 6
2.1 Normative references 6
2.2 Informative references 6
3 Definition of terms, symbols and abbreviations 7
3.1 Terms 7
3.2 Symbols 8
3.3 Abbreviations 8
4 Rationale for a multi-layered and multi-scale graph-based context information model 9
4.1 Why use a graph-based model? 9
4.2 Separating semantic referencing from structural descriptions 10
4.3 Graph Examples used in the present document 10
5 NGSI-LD meta-model 12
5.0 Introduction 12
5.1 Fundamentals of property graphs and graph databases 12
5.2 Reification with blank nodes 13
5.3 Formal definition 15
5.3.0 Introduction 15
5.3.1 Entity Types 17
5.3.2 Attribute Types 17
5.4 Serialization with JSON-LD 18
6 ETSI NGSI-LD cross-domain ontology 19
6.1 Rationale 19
6.2 NGSI-LD API compatibility 20
6.3 Formal specification 20
6.3.0 Comparison with other approaches 20
6.3.1 Mobility (of Entities) 24
6.3.2 Properties 24
6.3.3 Location (Property or Relationship) 25
6.3.4 Values 26
6.3.5 Temporal Properties and Values 26
6.3.6 Systems Composition 27
6.3.6.0 Introduction 27
6.3.6.1 Top-down system composition 27
6.3.6.2 Bottom-up system composition and clustering 28
Annex A (informative): Guidelines for Entity Typing 30
A.0 Introduction 30
A.1 Additional implementation requirements 30
A.2 Modelling recommendations 31
A.3 Using OWL/RDFS/RDF modelling 31
A.3.0 Introduction 31
A.3.1 OWL/RDFS/RDF modelling 31
A.3.2 Object-Oriented modelling 32
Annex B (informative): Distinguishing real-word entities from NGSI-LD Entities 33
B.0 Introduction 33

Sample Document
get full document from standards.iteh.ai

ETSI TS 104 178 V1.1.1 (2026-07) • https://standards.iteh.ai/catalog/standards/etsi/dd123511-cac3-4ef1-b028-bbab27139033/etsi-ts-104-178-v1-1-1-2026-07

B.1 W3C View 33
B.2 NGSI-LD View 33
Annex C (informative): OWL-DL representation of the Information Model 35
Annex D (informative): Change history 39
History 40

Sample Document
get full document from standards.iteh.ai

ETSI TS 104 178 V1.1.1 (2026-07) • https://standards.iteh.ai/catalog/standards/etsi/dd123511-cac3-4ef1-b028-bbab27139033/etsi-ts-104-178-v1-1-1-2026-07

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to normative deliverables (European Standard (EN), Technical Specification (TS), Group Specification (GS) or ETSI Standard (ES)) 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.

Sample Document
get full document from standards.iteh.ai

Foreword

This Technical Specification (TS) has been produced by ETSI Technical Committee Data Solutions (DATA).

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.

ETSI TS 104 178 V1.1.1 (2026-07) • https://standards.iteh.ai/catalog/standards/etsi/dd123511-cac3-4ef1-b028-bbab27139033/etsi-ts-104-178-v1-1-1-2026-07

1 Scope

The purpose of the present document is to give property graphs a formal semantic grounding based on RDF/RDFS/OWL, with blank nodes reification, geared to JSON-LD serialization. On top of it, a set of core cross-domain ontology classes have been defined, based on this meta-model. This whole information model is meant to be used by many applications as a basis for data representations. It is compatible with the NGSI-LD API defined in ETSI GS CIM 009 [2].

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.

Sample Document

[1] W3C® Recommendation 16 July 2020: "JSON-LD 1.1 - A JSON-based Serialization for Linked Data".

[2] ETSI GS CIM 009 (V1.7.1): "Context Information Management (CIM); NGSI-LD API".

[3] W3C® Candidate Recommendation Draft 15 November 2022: "Time Ontology in OWL".

get full document from standards.iteh.ai

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] Guinard, D., & Trifa, V. (2016): "Building the web of things", Shelter Island: Manning.

[i.2] Tim Berners-Lee (2006-07-27): "Linked Data", Design Issues W3C.

[i.3] J. Frey, K. Müller, S. Hellmann, E. Rahm and M.-E. Vidal (2017): Semantic Web - Interoperability, Usability, Applicability an IOS Press Journal: "Evaluation of Metadata Representations in RDF stores".

[i.4] Cassandras, C. G., & Lafortune, S. (2009) Springer Science & Business Media: "Introduction to discrete event systems".

[i.5] W3C® Editor's Draft 11 August 2005: "Simple part-whole relations in OWL Ontologies".

[i.6] W3C® Working Group Note 9 March 2006: "A Semantic Web Primer for Object-Oriented Software Developers".

[i.7] W3C®: "HttpRange14Webography".

ETSI TS 104 178 V1.1.1 (2026-07) • https://standards.iteh.ai/catalog/standards/etsi/dd123511-cac3-4ef1-b028-bbab27139033/etsi-ts-104-178-v1-1-1-2026-07

[i.8] W3C® Interest Group Note 3 December 2008: "Cool URIs for the Semantic Web", Leo Sauermann and Richard Cyganiak.

3 Definition of terms, symbols and abbreviations

3.1 Terms

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

cross-domain ontology: part of the information model that defines generic classes (formal concepts and constructs, with associated constraints) that serve as common denominators between domain specific models, addressing the temporal and structural description of physical systems

domain-specific ontologies: information models that define base classes and their constraints, within specific technical domains (e.g. buildings, transportation, agriculture) and define their structure and vocabulary

meta-model: part of the information model that formally defines the NGSI-LD foundational classes (Entities, Attributes, Relationships, Properties and reification constructs) on the basis of RDF/RDFS/OWL

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 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 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 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. 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 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

Sample Document
get full document from standards.iteh.ai

ETSI TS 104 178 V1.1.1 (2026-07) • https://standards.iteh.ai/catalog/standards/etsi/dd123511-cac3-4ef1-b028-bbab27139033/etsi-ts-104-178-v1-1-1-2026-07

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 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 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 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 type coercion into the IRI the <http://example.com/non-commercial>.

3.2 Symbols

Void.

3.3 Abbreviations

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

API	Application Programming Interface
CIM	Context Information Management
CNIT	Consorzio Nazionale Interuniversitario per le Telecomunicazioni
DATA	Data Solutions
DBMS	DataBase Management System
GSMA™	GSM Association
HTTP	HyperText Transfer Protocol
IoT	Internet of Things
IRI	Internationalized Resource Identifier
ISG	Industry Specification Group

Sample Document

ETSI TS 104 178 V1.1.1 (2026-07) • https://standards.iteh.ai/catalog/standards/etsi/dd123511-cac3-4ef1-b028-bbab27139033/etsi-ts-104-178-v1-1-1-2026-07

JSON	JavaScript Object Notation
JSON-LD	JavaScript Object Notation for Linked Data
LD	Linked Data
LOD	Linked Open Data
NGSI	Next Generation Service Interfaces
NIR	Non-Informational Resource
OGC	Open Geospatial Consortium
OWL	Web Ontology Language
OWL-DL	Web Ontology Language - Description Logic
PG	Property Graph
RDF	Resource Description Framework
RDFS	Resource Description Framework Schema
SAREF	Smart Applications REFerence ontology
SAS	Société par Actions Simplifiée
SPARQL	SPARQL Protocol and RDF Query Language
URI	Uniform Resource Identifier
WoT	Web of Things
XML	eXtended Markup language
XSD	eXtensible markup language Schema Definition language

4 Rationale for a multi-layered and multi-scale graph-based context information model

4.1 Why use a graph-based model?

Sample Document
Get full document from standards.iteh.ai

Systems and environments about which context information is stored and managed encompass many physical and non-physical entities. Context comprises all characteristics of these entities, as well as their states and other dynamic properties, together with relationships that stand for real-world connections, i.e. physical, virtual, or even abstract connections between them. This context information may be consolidated on the basis of data obtained from many different primary sources and infrastructures. Typical examples of such systems would be smart homes, buildings, or cities. Such systems, due to the wide range of requirements and granularities, are complex from the semantic, structural and behavioural viewpoints.

The expressivity and versatility of graph-based models allows to bring the whole corpus of graph theory to bear and to capture key information about such complex environments, in a directly usable way, as the graph matches all kinds of real-world connections between different physical and non-physical entities.

Graph models bring a fresh view on the definition of context information. In the first wave of context-awareness research dating back to the early 2000s, context used to be mostly, and implicitly, user-centric, typically capturing e.g. the activity or location of mobile users to adapt services offered to them. A widely publicized definition of context, dating from this stage of research, was "any information that can be used to characterize the situation of an entity". In a broader view of context where the very notion of context is de-centred and relative, this definition may in fact remain valid if entities are represented as the nodes of a graph. Rather than through a vague notion of situation, context is defined in the present document as the set of properties characterizing these nodes, together with the set of relationships that enmesh them together, and the properties of these relationships. In this perspective, the primary data of one application may be the context of another, and vice versa. Context is, though decentred and broadly defined, the graph itself. NGSI-LD thus maintains and exposes context information as a graph of matching links between the informational units corresponding to real-world entities of these environments.

Traditional (mono-centred) context fitted rather well a classical object-oriented or key-value description, with a set of more or less detailed context features attached to a single entity. The multi-centred notion of context address in the present document requires breaking this rigid hierarchical model by using a more expressive, flexible and adaptable information model. Graphs are the only model adapted to capture the complex structure of inter-entity relationships that make up context information in the sense in which it is defined it here. This information need not be semantically defined from the outset: it may be natively structural information, capturing e.g. containment or adjacency relationships. The semantics of this context may be added in a later stage of graph enrichment. This model fits the natively distributed nature of context data sources.

ETSI TS 104 178 V1.1.1 (2026-07) • https://standards.iteh.ai/catalog/standards/etsi/dd123511-cac3-4ef1-b028-bbab27139033/etsi-ts-104-178-v1-1-1-2026-07

The Web of Things (WoT) [i.1] does also involve a graph of sorts, but it dispenses with maintaining it explicitly inside a database. The WoT graph is, like the graph of the original web, an implicit 100 % distributed graph of hyperlinks, not between web pages but between resources corresponding to connected devices that expose an interface on the web (e.g. using HTTP, WebSockets, etc.).

This view also aligns well with the grand evolution of the Web towards Linked Data, an evolution proposed by W3C from the Semantic Web project [i.2] that is currently supported by RDF-derived graph models. Linked Data provides a method of publishing structured data so that it can be interlinked and support semantic queries.

4.2 Separating semantic referencing from structural descriptions

The NGSI-LD information model separates semantic referencing, used in the classical sense of the Semantic Web, from the structural description proper. The structural description may itself be decomposed into a basis structural graph whose nodes are physically-matched entities, and an overlay layer used to capture the way in which these entities are clustered into subgraphs.

Semantic referencing used by NGSI-LD is based on standard RDF/RDFS/OWL typing and public ontologies, as shared by all other semantic information models. All nodes and edges of the structural graph are thus matched to several relevant classes/categories of these ontologies that jointly characterize the features shared by all instances of these classes.

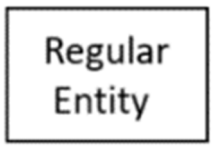
A structural graph is a model of the structural description of an environment, capturing the relationships between the different subsystems that make up this environment. This description is, to some extent, independent of the overlaying semantic referencing, and it could be considered to "stand on its own", even without this referencing. A structural graph does in fact have a different kind of semantics of its own, such as e.g. when a graph captures and matches the structure of a physical network like a power grid or a water distribution network. These semantics apply to the graphs as a whole and are not reducible to the kind of "per-resource" semantics, which RDF is meant to describe.

Sample Document
get full document from standards.iteh.ai

4.3 Graph Examples used in the present document

Two examples of structural representations of city environments will be used as lead examples throughout the present document and are presented in Figure 1 Property graph example (1) and Figure 2 Property graph example (2).

The following graphical conventions are used throughout the present document:



Regular (physically-matched) entities are represented as black rectangles.

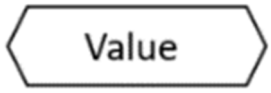


Relationships between these entities are represented as diamonds (rhombuses) overlaid on the corresponding arc of the graph, a convention borrowed from "entity-relationship" diagrams.



Properties are represented by ovals that are on an arc between their entity or relationship and the property value, but often the arc is shortened to zero length for compactness.

ETSI TS 104 178 V1.1.1 (2026-07) • https://standards.iteh.ai/catalog/standards/etsi/dd123511-cac3-4ef1-b028-bbab27139033/etsi-ts-104-178-v1-1-1-2026-07



Values are represented as hexagons that may about the oval of the property of which they are the target, omitting an arc between the two.

Figure 1 describes a parking scenario, adjacent to two different streets. Information about the streets, parking places, and the sensors that monitor are attached to entities as shown in Figure 1. This example is intended to illustrate the full expressivity of a property graph as used to capture not only pure semantics, as an RDF graph would, but also structural and behavioural (in this case, the real-time state) information.

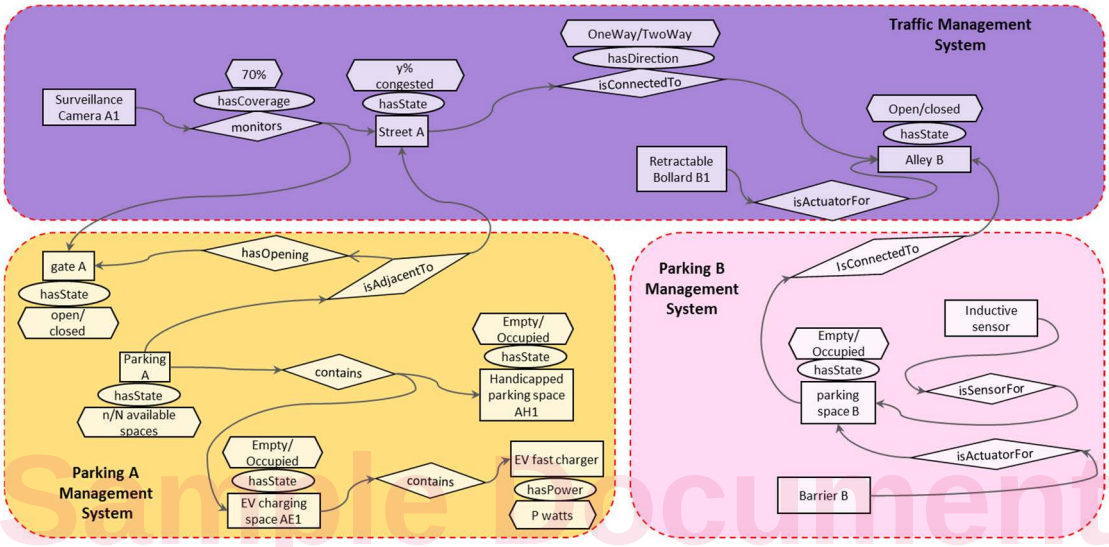


Figure 1: Property graph example (1)

Sample Document
get full document from standards.iteh.ai

Figure 2 example (2) is a more complex example used to illustrate intersecting domains and intertwined technical systems. The example consists of a building and its parts (using "hasPart" relationships) forming the structure of the building, in addition to other technical systems that are included in the building. The building is comprised of a garage and apartments (only one instance is represented below). A parking place within the garage belongs to the apartment, thus forming one system together. The building is equipped with a security system containing security devices. Additionally, there is a separate public parking that also appears in the example.

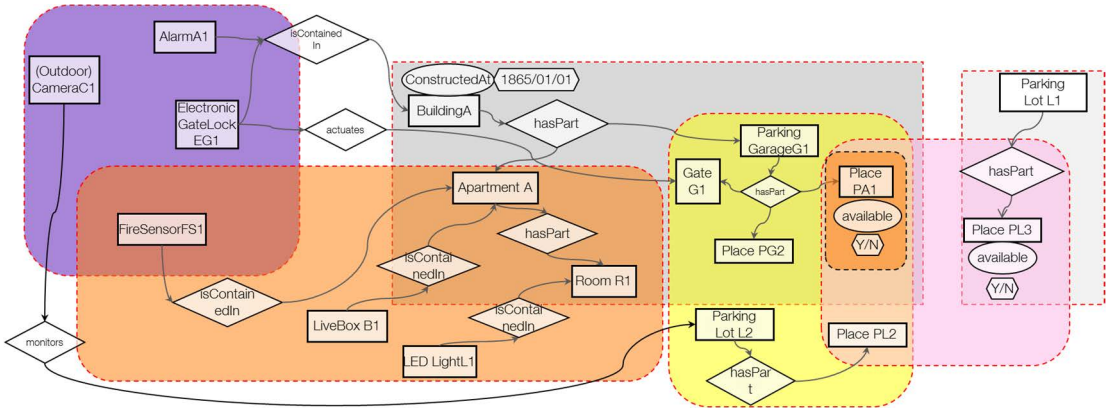


Figure 2: Property graph example (2)

ETSI TS 104 178 V1.1.1 (2026-07) • https://standards.iteh.ai/catalog/standards/etsi/dd123511-cac3-4ef1-b028-bbab27139033/etsi-ts-104-178-v1-1-1-2026-07

5 NGSI-LD meta-model

5.0 Introduction

The NGSI-LD meta-model provides a formal basis for representing "property graphs" using RDF/RDFS/OWL. It makes it possible to perform back and forth conversion between datasets based on the property graph model on the one hand and linked data datasets which rely on the RDF framework, on the other hand. This may be seen as raising the semantic expressivity of RDF triples to the level of property graphs. Property graphs may, contrary to RDF, use predicates as subjects of other predicates (properties of properties and properties of relationships).

5.1 Fundamentals of property graphs and graph databases

Property graphs are the implicit semi-formal data models underlying most present-day graph databases. They have gained widespread following, more in industry than in academia. They make it possible to attach properties (defined as key-value pairs) to relationships, a feature which RDF does not directly support, but they lack the standardization and formal underpinnings of RDF and do not interoperate directly with linked data and other RDF datasets. Also they do not lend themselves to reasoning with RDF-based reasoning tools or querying with standard query languages such as SPARQL.

Property graphs are usually defined (informally) as follows:

- A property graph is made up of nodes (vertices), relationships, and properties.
- Nodes may have properties in the form of arbitrary key-value pairs. Keys are strings and values are arbitrary data types.
- A relationship is an arc (uni-directional, i.e. directed edge) of the graph proper, which always has an identifier, a start node and an end node. Like nodes, relationships can have properties attached to them.

Sample Document

There are several key differences between property graphs (PG in the following) and RDF graphs:

get full document from standards.iteh.ai

- RDF properties are expressed as regular triples, i.e. arcs of the graph with start node and end node, and their target can be either a literal, an IRI or a blank node [1], whereas the target of a PG property always corresponds to an RDF literal.
- PG relationships (i.e. primary graph links between PG vertices) are first-class citizens of the PG model and have an internal structure similar to that of a vertex, inherited from object-oriented modelling object, with an optional set of properties defined by key-value pairs.
- The distinction between relationships and properties in the PG model is similar to the distinction between object properties and datatype properties in OWL, but stronger.
- PG properties are, for simplicity and avoiding clutter in diagrams, usually not represented as additional arcs of an underlying graph, but are represented as attached to vertices or relationships.
- Identifiers in a Property Graph need only be unique within the scope of a given graph (typically as internal identifiers assigned by the Graph DBMS), and need not be universally unique like URIs/IRIs.
- Property graphs can be queried with graph-specific query languages that may use graph patterns (complete subgraphs) as query terms, i.e. are not limited to only nodes identified with specific key/values.
- PG properties and relationships are individually identified when instantiated, whereas RDF properties are not instantiated nor identified as individual resources, being only defined by their property type.
- Properties cannot be directly attached to the arc (predicate) of an RDF triple, but RDF reification makes it possible, in several different ways, to circumvent this limitation by turning a triple into a resource.

ETSI TS 104 178 V1.1.1 (2026-07) • https://standards.iteh.ai/catalog/standards/etsi/dd123511-cac3-4ef1-b028-bbab27139033/etsi-ts-104-178-v1-1-1-2026-07