

**UN/CEFACT**

UNITED NATIONS

Centre for Trade Facilitation and Electronic Business

(UN/CEFACT)

1 METHODOLOGY AND TECHNOLOGY PROGRAMME DEVELOPMENT AREA

2 SPECIFICATIONS DOMAIN

3 **OPENAPI NAMING AND DESIGN RULES**4 **TECHNICAL SPECIFICATION****SOURCE:** API TechSpec Project Team**ACTION:** Ready for publication**DATE:** 13 September 2022**STATUS:** **v1.0****Disclaimer (Updated UN/CEFACT Intellectual Property Rights Policy – ECE/TRADE/C/CEFACT/ 2010/20/Rev.2)**

ECE draws attention to the possibility that the practice or implementation of its outputs (which include but are not limited to Recommendations, norms, standards, guidelines and technical specifications) may involve the use of a claimed intellectual property right.

Each output is based on the contributions of participants in the UN/CEFACT process, who have agreed to waive enforcement of their intellectual property rights pursuant to the UN/CEFACT IPR Policy (document ECE/TRADE/C/CEFACT/2010/20/Rev.2 available at [http://www.unece.org/cefact/cf\\_docs.html](http://www.unece.org/cefact/cf_docs.html) or from the ECE secretariat). ECE takes no position concerning the evidence, validity or applicability of any claimed intellectual property right or any other right that might be claimed by any third parties related to the implementation of its outputs. ECE makes no representation that it has made any investigation or effort to evaluate any such rights.

Implementers of UN/CEFACT outputs are cautioned that any third-party intellectual property rights claims related to their use of a UN/CEFACT output will be their responsibility and are urged to ensure that their use of UN/CEFACT outputs does not infringe on an intellectual property right of a third party.

5 ECE does not accept any liability for any possible infringement of a claimed intellectual property right or any other right that might be  
6 claimed to relate to the implementation of any of its outputs.

## 7 **Abstract**

8 This OpenAPI Naming and Design Rules technical specification defines an architecture and  
9 a set of rules necessary to specify, describe and implement APIs based on an OpenAPI  
10 specification to consistently express business information. It is based on the OpenAPI  
11 specification and the UN/CEFACT Core Components Technical Specification. This  
12 specification describes the requirements that UN/CEFACT compliant APIs should fulfil. It  
13 will be used by other organisations who are interested in maximizing inter- and intra-  
14 industry interoperability.  
15

16	<i>Abstract</i> .....	2
	1.1 DOCUMENT HISTORY .....	5
	1.2 CHANGE LOG.....	5
	1.3 OPENAPI NAMING AND DESIGN RULES PROJECT TEAM .....	6
	1.4 ACKNOWLEDGEMENTS.....	6
	1.5 CONTACT INFORMATION .....	6
	1.6 NOTATION .....	6
	1.7 AUDIENCE.....	7
17	<b>2 INTRODUCTION</b> .....	<b>8</b>
	2.1 OBJECTIVES .....	8
	2.2 REQUIREMENTS .....	8
	2.3 DEPENDENCIES .....	8
	2.4 CAVEATS AND ASSUMPTIONS .....	8
	2.5 GUIDING PRINCIPLES.....	9
	2.6 INTEROPERABILITY.....	10
18	<b>3 API NAMING AND DESIGN RULES</b> .....	<b>13</b>
	3.1 CONFORMANCE AND COMPLIANCE.....	13
	3.2 DESIGN RULES.....	14
19	3.2.1 <i>Media type for structured data exchange</i> .....	14
20	3.2.2 <i>Endpoints</i> .....	14
21	3.2.3 <i>Discoverability</i> .....	16
22	3.2.4 <i>Date and Time</i> .....	16
23	3.2.5 <i>Using the UN/CEFACT semantics</i> .....	16
24	3.2.6 <i>Operations</i> .....	21
25	3.2.7 <i>Pagination</i> .....	23
26	3.2.8 <i>Filtering</i> .....	26
27	3.2.9 <i>Sorting</i> .....	27
28	3.2.10 <i>API Responses and error handling</i> .....	28
29	3.2.11 <i>Error Response Payload</i> .....	33
30	3.2.12 <i>Design rule examples</i> .....	35
31	<b>4 WELL-DOCUMENTED APIS</b> .....	<b>36</b>
	4.1 GENERAL CONSIDERATIONS .....	36

4.2 API VERSIONING .....36

32 4.2.1 Versioning Scheme..... 36

33 4.2.2 URI Versioning ..... 37

34 4.2.3 Providing version information..... 37

35 4.2.4 Robustness..... 39

36 4.2.5 Deprecation and End of Life Policy..... 40

4.3 HYPERMEDIA .....43

37 4.3.1 Hypermedia - Linked Data..... 43

38 4.3.2 HATEOAS ..... 43

39 4.3.3 Hypermedia Compliant API..... 44

40 **5 API SECURITY ..... 46**

41 **6 EVENT DRIVEN DATA EXCHANGE..... 48**

6.1 CALLBACKS .....48

6.2 WEBHOOKS.....49

6.3 SECURITY GUIDELINE FOR CALLBACKS (INFORMATIVE).....49

42 6.3.1 Subscription setup (informative)..... 50

43 6.3.2 Performing a subscription call (informative) ..... 51

44 **7 APPENDIX A: EXAMPLES ..... 53**

45 **8 APPENDIX B: NAMING AND DESIGN RULES LIST (NORMATIVE) ..... 54**

46 **9 APPENDIX C: GLOSSARY..... 58**

47

48 **1.1 Document History**

49

<b>Phase</b>	<b>Status</b>	<b>Date Last Modified</b>
Draft development	First draft	06 September 2022
Ready for approval	First version	13 September 2022

50

**Table 1 – Document history**51 **1.2 Change Log**

52 The change log is designed to alert users about significant changes that occurred during the  
53 development of this document.

54

<b>Date of Change</b>	<b>Version</b>	<b>Paragraph Changed</b>	<b>Summary of Changes</b>
30 May 2022	0.3		First draft TOC
07 June 2022	0.4		Drafted up to chapter 3.2.7
	0.5		Drafted up to chapter 3.2.9
20 June 2022	0.6		Completion up to chapter 6
05 Sept 2022	0.7	1.6 2.6 R 1 R 16 6.3 7 Appendix A Appendix B Appendix C	Considering public review comments
13 Sept 2022	1.0		Minor corrections

55

**Table 2 - Document change log**

### 56 **1.3 OpenAPI Naming and Design Rules Project Team**

57 We would like to recognize the following for their significant participation in the  
58 development of this Unites Nations Centre for Trade Facilitation and Electronic Business  
59 (UN/CEFACT) OpenAPI Naming and Design Rules technical specification.

#### **ATG2 Chair**

Marek Laskowski

#### **Project Lead**

Jörg Walther

#### **Lead editors**

Andreas Pelekies

Gerhard Heemskerk

### 60 **1.4 Acknowledgements**

61 This version of UN/CEFACT OpenAPI Naming and Design Rules Technical Specification  
62 has been created to foster convergence among Standards Development Organisations  
63 (SDOs). It has been developed in close coordination with these organisations:

- 64 • Digital Container Shipping Association
- 65 • GS1
- 66 • Odette

### 67 **1.5 Contact information**

68 ATG2 – Marek Laskowski, [Marek.laskowski@gmail.com](mailto:Marek.laskowski@gmail.com)

69 NDR Project Lead – Jörg Walther, [jwalther@odette.org](mailto:jwalther@odette.org)

70 Editor – Andreas Pelekies, [Andreas@pelekies.de](mailto:Andreas@pelekies.de)

71 Editor – Gerhard Heemskerk, [Gerhard.heemskerk@kpnmail.nl](mailto:Gerhard.heemskerk@kpnmail.nl)

### 72 **1.6 Notation**

73 The keywords MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD,  
74 SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL, when they appear in this  
75 specification, are to be interpreted as described in Internet Engineering Task Force (IETF)  
76 Request For Comments (RFC) 2119<sup>1</sup>.

77 **Example** A representation of a definition or a rule. Examples are informative.

---

<sup>1</sup> Key words for use in RFCs to Indicate Requirement Levels - Internet Engineering Task Force, Request For Comments 2119, March 1997, <http://www.ietf.org/rfc/rfc2119.txt?number=2119>

- 78 `[Note]` Explanatory information. Notes are informative.
- 79 `[R n|c]` Identification of a rule that requires conformance. Rules are normative. In  
80 order to ensure continuity across versions of the specification, rule numbers  
81 “n” are randomly generated. The number of a rule that is deleted will not be  
82 re-issued. Rules that are added will be assigned a previously unused random  
83 number.  
84 The second number “c” after the pipe symbol `|` identifies the conformance  
85 category of the given rule as defined in section 3.1. A `[+Inf]` may be added to  
86 identify rules that are informative and not normative.
- 87 `Courier` All words appearing in **bolded courier font** are values, objects or keywords.  
88 Representation of non-printable characters like white space are surrounded  
89 by double-quotes, e.g. `" "`.
- 90 `<<var>>` All placeholders are surrounded by double less-than and greater-than  
91 characters. The meaning of the placeholder is described in the text.

## 92 **1.7 Audience**

93 The audience for this UN/CEFACT OpenAPI Schema Naming and Design Rules Technical  
94 Specification is:

- 95 • Members of the UN/CEFACT Applied Technologies Groups who are responsible for  
96 development and maintenance of UN/CEFACT OpenAPI specifications and  
97 recommendations.
- 98 • The wider membership of the other UN/CEFACT Groups who participate in the  
99 process of creating and maintaining UN/CEFACT OpenAPI specifications.
- 100 • Designers of tools who need to design OpenAPI specifications adhering to the rules  
101 defined in this document.
- 102 • Designers of OpenAPI specifications outside of the UN/CEFACT Forum community.  
103 These include designers from other organisations that have found these rules suitable  
104 for their own organisations.

## 105 **2 Introduction**

### 106 **2.1 Objectives**

107 This OpenAPI NDR technical specification document forms part of a suite of documents  
108 that aim to support modern web developers to make use of UN/CEFACT semantics.

109 Taking any layer of the UN/CEFACT Reference Data Models to create conformant  
110 OpenAPI specifications in accordance with the UN/CEFACT Core Components Technical  
111 Specification Version 2.01. This includes comprehensive RDMs like Buy-Ship-Pay, or  
112 Accounting as well as their contextualization like the Supply-Chain-Reference-Data-Model  
113 (SC-RDM), Multi-Modal-Transport-Reference-Data-Model (MMT-RDM) down to single  
114 message implementation like the Road Consignment Note (eCMR) or the certificate of  
115 origin (COO).

### 116 **2.2 Requirements**

117 Users of this specification should have an understanding of basic data modelling concepts,  
118 basic business information exchange concepts and basic (REST) API concepts.

### 119 **2.3 Dependencies**

120 This document depends on

- 121 1. UN/CEFACT Core Components Technical Specification Version 2.01.
- 122 2. JSON Schema Naming and Design Rules Technical Specification.

### 123 **2.4 Caveats and Assumptions**

124 Specifications created as a result of employing this specification should be made publicly  
125 available as OpenAPI specification documents in a universally free, accessible, and  
126 searchable library. UN/CEFACT will make its contents freely available to any government,  
127 individual or organisation who wishes access.

128 Although this specification defines the data structures used in an OpenAPI specification as  
129 expressions of Reference Data Models, non-CCTS developers can also use it for other  
130 logical data models and information exchanges.

131 This specification does not address transformations via scripts or any other means. It does  
132 not address any other representation of CCTS artefacts – such as XML, JSON-LD, OWL,  
133 and XMI.

134 Standards foster interoperability. In the creation of this specification and definition of  
135 design principles, several sources were taken into account in the following order:

- 136 1. The OpenAPI 3.1.0 specification



- 137 2. Standards defined by internet standard organisations as RFCs
- 138 3. The DCSA API Design Principles 1.0
- 139 4. The json:api specification
- 140 5. Experts experience

## 141 **2.5 Guiding Principles**

- 142 3. OpenAPI Creation
- 143 UN/CEFACT OpenAPI design rules support OpenAPI specification creation through
- 144 handcrafting as well as automatic generation.
- 145 4. Tool Use and Support
- 146 The design of UN/CEFACT OpenAPI will not make any assumptions about
- 147 sophisticated tools for creation, management, storage, or presentation being available.
- 148 5. Technical Specifications
- 149 UN/CEFACT OpenAPI Naming and Design Rules will be based on technical
- 150 specifications holding the equivalent of OpenAPI recommendation status.
- 151 6. OpenAPI Specification
- 152 UN/CEFACT OpenAPI Naming and Design Rules will be fully conformant with the
- 153 OpenAPI specification recommendation.
- 154 7. Interoperability
- 155 The number of ways to express the same information in a UN/CEFACT OpenAPI
- 156 specification is to be kept as close to one as possible.
- 157 8. Maintenance
- 158 The design of an UN/CEFACT OpenAPI specification must facilitate maintenance.
- 159 9. Context Sensitivity
- 160 The design of an UN/CEFACT OpenAPI specification must ensure that context-
- 161 sensitive document types are not precluded.
- 162 10. Ease of implementation
- 163 An UN/CEFACT OpenAPI specification should be intuitive and reasonably clear in the
- 164 context for which they are designed. They should allow an intuitive implementation in
- 165 REST APIs, a.k.a. RESTful API, as well as other interchange appliances.
- 166

## 167 **2.6 Interoperability**

168 Decades of cross-industry and cross-national harmonisation of B2B and B2A processes  
169 have gone into the development of the semantic UN/CEFACT reference data models by  
170 thousands of experts. This tremendous achievement does not exist a second time in this  
171 scope and depth. The clear path from semantic definition to syntax - and not vice versa -  
172 means that these semantic data models are syntax-neutral and can thus be used not only  
173 with current syntaxes but also with future ones. For this purpose, either they are mapped  
174 directly into a (UN/CEFACT) syntax via NDR specifications, or they can be mapped to data  
175 models and syntaxes of other sectors.

176 The ideal of a REST API envisages the fully automatable connection of an API consumer to  
177 an API provider. In practice, this is often not the case today, as the corresponding standards  
178 for the design of an API, the scope and depth of the documentation and the modelling of  
179 processes and data in B2B and B2A communication via WebAPIs are still in their infancy.  
180 The keyword here is interoperability.

181 In classic EDI implementations (e.g. EDIFACT or XML), a variety of industry standards  
182 exist. With their help, the following dimensions of interoperability are promoted:

- 183 1. Business process interoperability: the business partners have the same understanding  
184 of the basic process flow, for example in the Order2Cash - process.
- 185 2. Semantic interoperability: the business partners have the same understanding of the  
186 technical terms. For example, the definition of consignment and shipment is the same  
187 for all business partners.
- 188 3. Syntax interoperability: a uniform syntax (e.g. UN/CEFACT XML) is used.
- 189 4. Contextualisation interoperability: Industry standards define how individual  
190 requirements are to be handled. Ideally, it is agreed that as few different  
191 contextualisations (consideration of individual requirements) as possible should take  
192 place. This means that information that is only required by some recipients will be  
193 read over by the remaining recipients instead of carrying out an individual  
194 implementation.
- 195 5. Interoperability of transmission: Business partners agree on uniform transmission  
196 methods as well as associated security measurements such as SFTP, OFTP2 or AS2.  
197 This dimension often plays a lesser role in classic EDI implementations, as the  
198 transmission of data usually takes place from one sender to one recipient at a time.  
199 EDI is usually optimised for mass data.

200 When implementing a WebAPI, the same requirements for interoperability exist in  
201 principle. An essential difference of previous WebAPIs is the approach to connect mass  
202 users to an API. For example, a map, route or booking service should be used by as many  
203 users as possible at the same time. The REST principle of composability also means that

204 different services (possibly from different providers) are often combined into an overall  
 205 solution for processing with WebAPIs. For example, in a flight booking service, the  
 206 capacities, conditions and tickets are allocated by the airlines, payment service providers are  
 207 connected, and often a specialised billing service that correctly calculates the different tax  
 208 constellations for cross-border flights. The aspect that many consumers have to use one API  
 209 (billing service) as well as one consumer has to use many APIs with the same processes  
 210 (contingent with airlines) extends the interoperability requirements for WebAPIs.

211 6. Interoperability of API design: This specification deals with the aspect of API design  
 212 interoperability. Uniform methods and rules in API design simplify the  
 213 understanding of APIs, errors during implementation are minimised, the handling of  
 214 error messages is standardised and the implication of similar APIs in a cross-  
 215 organisational (B2B) network is promoted.

216 7. Service interoperability: uniform endpoints in mapping the same process  
 217 requirements promote B2B communication via WebAPIs.

218 The following table shows how the seven dimensions of interoperability can be achieved in  
 219 WebAPIs:

Dimension of interoperability	Guideline
<b>Business process interoperability</b>	Within UN/CEFACT, business process interoperability is achieved by implementing the harmonised business requirement specifications (BRS).
<b>Semantic interoperability</b>	The CCTS and its derived semantic Reference Data Models (RDMs) are the basis for this dimension for UN/CEFACT users. The UN/CEFACT Vocabulary, the JSON Schema artefacts, and the UN/CEFACT XML standards implement these semantic requirements in the respective syntax.
<b>Syntax interoperability</b>	When a user group agrees on the use of a uniform data exchange syntax, this dimension is achieved. When creating an OpenAPI specification, it should be noted that the syntax to be used must always be modelled as a JSON schema, even if the later exchange syntax is an XML format, for example. It is defined in an OpenAPI specification.
<b>Contextualisation interoperability</b>	An implementation guideline (for example, of a particular industry) defines how contextualisations are to be applied to the data or message structures to be exchanged.
<b>Interoperability of transmission</b>	This dimension is also specified in an implementation guideline. In particular, it also includes security aspects including authorisation and authentication.
<b>Interoperability of API design</b>	This NDR specification defines the interoperability of API design. Among others, it includes rules for filtering, pagination and error handling.
<b>Service interoperability</b>	A good OpenAPI specification especially focuses on service interoperability. The interoperability of APIs designed to be implemented by several business partners can be fostered if the services are well designed. For instance, a user group agrees on a

set of services with a minimum subset. If a provider does not support a specific service it is still implemented, but always responds with a `501 Method not implemented` HTTP response code that includes a `HTTP Link Header` to the corresponding documentation

220

**Table 3: Interoperability of WebAPIs**

221

## 222 3 API Naming and Design Rules

### 223 3.1 Conformance and Compliance

224 Designers of OpenAPI specifications in governments, private sector, and other standards  
 225 organisations external to the UN/CEFACT community have found this specification  
 226 suitable for adoption. To maximize reuse and interoperability across this wide user  
 227 community, the rules in this specification have been categorised to allow these other  
 228 organisations to create conformant OpenAPI specifications while allowing for discretion or  
 229 extensibility in areas that have minimal impact on overall interoperability.

230 Accordingly, applications will be considered to be in full conformance with this technical  
 231 specification if they comply with the content of normative sections, rules and definitions.

232 [R 1|1]

233 Compliance and conformance SHALL be determined through adherence to the content of  
 234 the normative sections and rules. Furthermore, each rule is categorised to indicate the  
 235 intended audience for the rule by the following:

236

Category	Description
<b>1</b>	Rules, which must not be violated. Else, compliance and interoperability are lost.
<b>2</b>	Rules, which may be modified, while still conformant to the NDR structure. If all rules of categories 1 and 2 are followed, the API is fully compliant. If rules of category 2 are modified the API is not compliant anymore, but still conformant.
<b>Inf</b>	Rules that are informative only. If a different implementation is chosen this does not have any impact on the compliance and conformance of the implementation towards this specification.

237 **Table 4 - Conformance categories**

238 [R 2|1]

239 All API specifications based on this OpenAPI Naming and Design Rules technical  
 240 specification SHALL be compliant to the OpenAPI 3.1.x specification.

241

242 [R 3|1]

243 An API specification claiming conformance to this specification SHALL define schema  
 244 components as described in the JSON Schema Naming and Design Rules Technical  
 245 Specification.

246

247 **3.2 Design Rules**248 **3.2.1 Media type for structured data exchange**

249 [R 4|1]

250 Request body content and Response content used to transfer structured data information  
 251 SHALL use the `application/json` media type for JavaScript Object Notation (JSON). This  
 252 rule MAY only be deviated from, if the API implements a conversion service from or to  
 253 JSON in another media type.

254  
 255 Additional media types (e.g. `text/xml`) to transfer structured data information MAY be  
 256 used. If non-structured information is transferred any valid media type MAY be used.

257

258 [R 5|1]

259 Encoding SHALL be UTF-8.

260 **3.2.2 Endpoints**

261 [R 6|2]

262 The structure of the paths defined within APIs SHOULD be meaningful to the consumers.  
 263 Paths SHOULD follow a predictable, hierarchical structure to enhance understandability  
 264 and therefore usability.

265

266 [R 7|1]

267 The API URLs SHOULD follow the standard naming convention as described below:

268

```
269 https://{env}.api.{dnsdomain}/v{m}/{service}/{resource}/{id}/{sub-
270 resource}?{query}
```

271 The components are described as follows. If a rule is mandatory for a specific component of  
 272 the URL it SHALL be applied to any conformant API specification, even if the basic URL  
 273 structure is different from the one described above (e.g. if `api` is not used as a prefix to the  
 274 `dnsdomain`).

- 275 • `https://` SHALL be used as the web protocol.
- 276 • `{env}` indicates the environment (e.g. `test`, `sandbox` or `dev`) and is usually omitted for  
 277 production environment.
- 278 • `{dnsdomain}` is the DNS domain of the API implementer (e.g. `unece.org`)
- 279 • `{service}` is a logical grouping of API functions that represent a business service domain  
 280 (e.g. `transport`). The `{service}` component is optional.
- 281 • `v{m}` is the major version number of the API specification. This component SHALL be  
 282 stated in the URL. It MAY be provided at a different place in the URL (e.g. as a prefix  
 283 to the domain).
- 284 • `{resource}` is the plural noun representing an API resource (e.g. `consignments`)

- 285
- 286
- 287
- 288
- 289
- 290
- 291
- 292
- {id} is the unique identifier for the resource defined as a path parameter. Path parameters SHALL be used to identify a resource. This component is not part of the path if an operation is performed on a collection of the resource.
  - {sub-resource} is an optional sub-resource. Only used when there are contained collections or actions on a main resource (e.g. `consignmentItem`).
  - {query} is a list of additional parameters like filters that determine the results of a search (e.g. `consignments?loadingPort=AUSYD`).

293

294 [R 8|1]

295 The total number of characters in the URL, including the path and the query, SHALL NOT  
296 exceed 2000 characters in length including any formatting codes such as commas,  
297 underscores, question marks, hyphens, plus or slashes.

298

299 [R 9|1]

300 Endpoints SHALL NOT be actions. Services and resources SHALL consist of nouns. HTTP  
301 verbs SHALL be used for actions (See chapter 3.2.6).

302

303 [R 10|1]

304 Kebab-case<sup>2</sup> SHALL be used in services.

305

306 [R 11|1]

307 Lower camelCase<sup>3</sup> SHALL be used in resources, path parameters and query parameters.

308

309 [R 12|1]

310 Path parameters and query parameters with a relation to property names SHALL be  
311 consistent with property names.

312

313

314

315 [R 13|1]

---

<sup>2</sup> Kebab-case is a naming rule for a technical representation of identifiers consisting of several words. Hyphens are used to connect words. Example: `this identifier` is written as `this-identifier` in kebab-case.

<sup>3</sup> CamelCase is a naming rule for a technical representation of identifiers consisting of several words. White spaces are removed and every new word begins with a capital letter. Example: `this identifier` is written as `thisIdentifier` in camelCase. Lower camelCase means that the identifier must start with a small letter.

316 Query parameters SHALL be URL safe<sup>4</sup>.

317

318 [R 14|1]

319 Resource names SHALL be pluralised. Resource names SHOULD be consistent with  
320 schemas. If a schema is defined in singular, nevertheless the resource SHALL be pluralized.  
321 If the plural of a resource is non-standard, you MAY choose a more appropriate noun in its  
322 plural form.

323

324 Examples for good endpoints:

- 325 • /employees
- 326 • /customers
- 327 • /products

### 328 3.2.3 Discoverability

329 One of the REST design principles is service discoverability. The OpenAPI specification  
330 supports them via links. They SHALL be implemented via HTTP headers.

### 331 3.2.4 Date and Time

332 The date and time representation in the CCL supports an ISO8601 subset with only a few  
333 exceptions. Those exceptions may be present in the content body of a request or a response.

334 [R 15|1]

335 Query parameters SHALL use ISO8601 compliant date and time representations that are  
336 defined in `UNTDID 2379 json` as defined in the JSON schema NDR technical specification.  
337 To represent a specific date, time or date-time the format SHALL comply with the JSON  
338 schema definition for date, time or date-time.

### 339 3.2.5 Using the UN/CEFACT semantics

340 Decades of harmonisation and standardisation of business requirements resulted in the  
341 UN/CEFACT reference data models (RDM). These exist across different domains like Buy-  
342 Ship-Pay, Agriculture, Regulatory or Audit and Accounting.

343 As one example the Buy-Ship-Pay RDM contains subsets e.g. for multimodal transport  
344 (MMT-RDM) and the supply chain (SC-RDM). Over time, hundreds of business document  
345 structures were harmonised and standardised on a semantic model level. Different Syntax  
346 Naming and Design rules allow an automated creation and mapping of those semantic  
347 models to certain syntaxes such as XML.

---

<sup>4</sup> See [https://www.w3schools.com/tags/ref\\_urlencode.ASP](https://www.w3schools.com/tags/ref_urlencode.ASP) Example: `https://unece.org/this url` is invalid because of the space. Correct it looks like `https://unece.org/this%20url`



348 In the world of web APIs, the transmission of document structures is considered obsolete. If  
349 the limitations of REST principles are to be applied to a web API, business document  
350 structures are unsuitable for a RESTful implementation. These structures contradict the  
351 basic principle of loose coupling of resources. Instead, the exchange of information should  
352 be resource-based, where resources are information blocks leading in their combination to  
353 the complete information (e.g. business document).

354 Nevertheless, there are often limitations in B2B information exchange that make it difficult  
355 to completely move away from document structures. This includes technical reasons,  
356 procedural reasons, but also legal reasons. If the basic processes of communication between  
357 organisations are not changed, a shift purely to resource-based information exchange leads  
358 to a new level of media disruption and consistency challenges. If both the sending and  
359 receiving systems work on the basis of document structures (e.g. an invoice), then an  
360 intermediate, purely resource-based transmission leads to a number of challenges, such as  
361 the archiving obligation of such documents that exists in many countries to ensure  
362 subsequent verification.

363 On the other hand, if networks of platforms (e.g. for logistics) are established, a resource-  
364 based exchange can still be useful for certain purposes. For example, a platform could exist  
365 for a marketplace where free delivery capacities by carriers can be offered and booked. The  
366 division by resources usually leads to the need for identity providers and the clarification of  
367 the question of the single source of trust for individual resources.

368 At UN/CEFACT, there are two basic JSON-based publications of semantic data models: the  
369 UN/CEFACT vocabulary, and the UN/CEFACT JSON schema publication.

### 370 **3.2.5.1 Using the UN/CEFACT JSON schema publication**

371 JSON schema is the natural partner of an OpenAPI specification, as OpenAPI relies on  
372 JSON schema. The UN/CEFACT JSON schemas are published in two variants:

- 373 8. Streamlined stand-alone JSON schemas for the individual business documents.  
374 Those schemas contain every definition relevant for a specific business document  
375 and its applied contextualisation.
- 376 9. A JSON schema library of the different RDMs and their related business document  
377 structures. This variant uses an inheritance and validation technique supported by  
378 JSON schema. The basic data structures define the information blocks needed  
379 together in the reference data model. Subsets and contextualisation for the individual  
380 applications (e.g. MMT-RDM, SC-RDM, Invoice ...) are then formed on this basis.

381 The JSON schemas are published in the official UN/CEFACT repository. They can be used  
382 in two different ways:

383 First by referencing the needed data types directly from the repository. This leads to a  
384 maximum on interoperability. In an OpenAPI specification, it is easily possible to further  
385 contextualise (including extension) the JSON subschemas to the needed requirements of the  
386 specific process. This explicitly lets the users "tick off" unneeded optional attributes or  
387 supplementary components, restrict code lists or add user defined properties in a  
388 standardised and transparent way.

389 Additionally, maintenance becomes quite easy. If the API is to be updated to a newer  
390 version of the JSON schema publication, only the reference needs to be updated.

391 Alternatively, the JSON schemas can be downloaded to a local system or repository. In that  
392 case it is needed to update or remove the "\$id" properties of the schemas, as they link to the  
393 official UN/CEFACT repository.

394 The way in which the JSON schemas are defined allow a very simple transmission from  
395 using document-based structures to resource-based structures. On the RDM level, all ABIEs  
396 (data classes) are defined. For every RDM exists a master document structure. All of the  
397 business documents are derived from this. The hierarchic structure connects the different  
398 ABIEs through ASBIEs including cardinality information. At every single ASBIE node, the  
399 JSON schema publication allows to replace the provision of a substructure by the URN of  
400 the corresponding resource:

401 Let us assume you want to define an API to manage transport capacity booking. In a classic  
402 message-based scenario, you would define how those messages are interchanged. In many  
403 case you would design a `POST` and `GET` or `POST`, subscribe and `GET` scenario. Those scenarios  
404 need envelope-information around the message information in order to tell the API who the  
405 ultimate receiver is, who the sender is etc. In addition the message is quite complex and  
406 contains many sub-resources with details. Those include for instance "requester", "shipTo",  
407 "receiver", "carrier", "consignment-items" etc. If this scenario is planned to move towards a  
408 (more) resource-based information exchange it is very easy to do so. First, you have to  
409 identify which of your sub-resources should become stand-alone. Let us assume you want to  
410 manage trade party information master data as a single resource. In that case, you can  
411 specify a schema under `components/schemas` named `tradePartyType` and simply define it as  
412 a reference to the contextualised data type of the corresponding RDM or even the  
413 corresponding business document structure. The following example shows, how the  
414 document structure can be restricted to resource usage as well.

415 Example for a `tradePartyType` under `components/schemas`:

```
416 "tradePartyType": {  
417   "description": "Trade party definition according to MMT-RDM",  
418   "$ref": "https://raw.githubusercontent.com/uncefact/spec-  
419   JSONschema/main/JSONschema2020-12/library/BuyShipPay/D22A/UNECE-
```

```
420 MMTContextCCL.json#/$defs/tradePartyType"
421 }
422 "tradePartyType": {
423   "description": "Trade party definition according to the Multimodal
424     Transport Booking Recipient",
425   "$ref": "https://raw.githubusercontent.com/uncefact/spec-
426 JSONschema/main/JSONschema2020-12/library/BuyShipPay/D22A/UNECE-
427 MultimodalTransportBooking.json#/exchangedDocument/recipient"
428 }
429 "multimodalTransportBooking": {
430   "title": "Multimodal Transport Booking",
431   "description": "Restrict business document to resource usage for
432     recipient",
433   "allOf": [
434     { "$ref": "https://raw.githubusercontent.com/uncefact/spec-
435 JSONschema/main/JSONschema2020-12/library/BuyShipPay/D22A/UNECE-
436 MultimodalTransportBooking.json/#" },
437     {
438       "properties": {
439         "exchangedDocument": {
440           "properties": {
441             "recipient": { "type": "string", "format": "uri" }
442           }
443         }
444       }
445     }
446   ]
447 }
```

### 448 3.2.5.2 Using the UN/CEFACT vocabulary

449 The UN/CEFACT vocabulary uses the JSON-LD format in order to be conformant with the  
450 publication on schema.org.

451 The publication in JSON-LD follows a different approach. JSON-LD is a graph  
452 representation of context-enhanced semantic ABIE-representations derived from the  
453 combination of the corresponding RDMs. By applying the appropriate context, the subset of  
454 the defined graph can be used.

455 JSON-LD cannot directly be used and linked to in an OpenAPI specification. According to  
456 the maintenance body of the OpenAPI specification, this is not intended to change in the  
457 near future. In addition, the JSON-LD does not specify the cardinalities and subsets for the  
458 different contexts of business document structure definitions. Therefore, a web developer  
459 implementing an API for business related intra-organisational information exchange needs a  
460 reasonable knowledge of the underlying processes. On the other hand, JSON-LD unfolds

461 immense power wherever (publicly) available data is to be automatically crawled, filtered  
462 and evaluated. Examples of this are applications such as flight-radar, online search for  
463 recipes or searches for goods over the boundaries of online shops with specific criteria. In  
464 those scenarios, the individual resources get into focus, as well as their relationships (links)  
465 to other resources. The business-related-interdependencies are not part of the definitions  
466 themselves. Adding state machines in definitions could help with this. Unfortunately,  
467 currently there does not exist a widely supported exchange format for this kind of  
468 information<sup>5</sup>.

469 In order to use the JSON-LD vocabulary, additional tooling must be used, as there does not  
470 exist a direct support in OpenAPI specifications. As a proof-of-concept, in the JSON-LD  
471 vocabulary publication, a sample implementation is included to import the vocabulary into a  
472 UML design tool. Here the first conversion from JSON-LD to UML is performed. Now the  
473 designing of the API can be performed within the UML-Tool. Some assumptions are made  
474 how to define which operations should be supported for each of the specified endpoints.  
475 Having defined this a second conversion from the UML-Tool to the OpenAPI specification  
476 format is performed.

### 477 **3.2.5.3 Using other (standardised) data structures**

478 In chapter 2.6 seven dimension of interoperability for WebAPIs are defined. From a global  
479 cross-industry perspective, full interoperability can only be achieved if for all of the  
480 dimensions the implementation rules are clearly defined. In the context of UN/CEFACT,  
481 this means that the UN/CEFACT semantic definitions as well as the UN/CEFACT syntaxes  
482 must be used to be fully compliant.

483 However, this NDR specification is syntax-neutral, as it defines basic requirements for the  
484 design of an OpenAPI specification in a B2B context. The stipulations in this specification  
485 can thus also promote interoperability between APIs that use a different syntax or divergent  
486 semantic specifications within a (closed) user group. Therefore, the following rule is  
487 defined as a conformance criterion:

488 [R 16 1]
489 A prerequisite for an OpenAPI specification and its implementation to be fully compliant 490 with this NDR TS is the use of UN/CEFACT semantics and UN/CEFACT syntax (e.g. 491 UN/CEFACT XML, UN/CEFACT JSON Schema, and UN/CEFACT Vocabulary).
492 An OpenAPI specification that does not use UN/CEFACT syntax or UN/CEFACT 493 semantics may still be conformant to this NDR TS if it meets the criteria specified in [R 494 1 1].

---

<sup>5</sup> See for example the JSON Finite State Machine in JSON schema format at <https://github.com/ryankurte/jfsm>

495 **3.2.6 Operations**

496 [R 17|1]

497 Endpoints are RECOMMENDED to support CRUD operations. (Create, Read, Update,  
498 Delete). If an endpoint is not intended to support e.g. a delete operation, it SHALL return  
499 the HTTP response codes as defined in chapter 3.2.10.

HTTP Method	Description
<i>GET</i>	To <i>retrieve/read</i> a resource.
<i>POST</i>	To <i>create</i> a new resource or to <i>execute</i> an operation on a resource that changes the state of the system e.g. send a message.
<i>PUT</i>	To <i>replace</i> a resource with another supplied in the request.
<i>PATCH</i>	To perform a <i>partial update</i> to a resource.
<i>DELETE</i>	To <i>delete</i> a resource.
<i>HEAD</i>	For retrieving metadata about the request, e.g. how many results <i>would</i> a query return? (Without actually performing the query). This can be used to follow a link-chain in an HATEOS implementation as well. An example is shown in chapter 4.3.2.
<i>OPTIONS</i>	Used to determine if a CORS (cross-origin resource sharing) request can be made. This is primarily used in front-end web applications to determine if they can use APIs directly.

500 **3.2.6.1 Collection of Resources**

501 The following operations are applicable for a collection of resources:

HTTP method	Resource Path	Operation	Examples
GET	<i>/resources</i>	Get a collection of the resource	GET <i>/employees</i> or GET <i>/employees?status=open</i>
HEAD	<i>/resources</i>	Get header and link information of the resource collection, e.g. for pagination	HEAD <i>/employees</i> or HEAD <i>/employees?birthday=2022-04-16</i>

502

**Note**

503

Creating or updating multiple resource instances in the same request is not standardised and thus should be avoided. There are factors such as receipt acknowledgement and how to handle partial success in a set of batches that must be considered on a case-by-case basis.

504

505

506 **3.2.6.2 Single Resource**

507 The following operations are applicable for a single resource:

<b>HTTP method</b>	<b>Resource Path</b>	<b>Operation</b>
GET	<i>/resources/{id}</i>	Get the instance corresponding to the resource ID
PUT	<i>/resources/{id}</i>	To update a resource instance by replacing it – <i>"Take this new thing and <b>_put_</b> it there"</i>
DELETE	<i>/resources/{id}</i>	To delete the resource instance based on the resource e.g. id
HEAD	<i>/resources/{id}</i>	Get header and link information of the resource.
PATCH	<i>/resources/{id}</i>	Perform changes such as add, update, and delete to the specified attribute(s). Is used often to perform partial updates on a resource

508 **3.2.6.3 Idempotency**

509 An idempotent HTTP method is an HTTP method that can be called many times without  
510 different outcomes. In some cases, secondary calls will result in a different response code,  
511 but there will be no change of state of the resource.

512 As an example, when you invoke N similar DELETE requests, the first request will delete  
513 the resource and the response will be 200 (OK) or 204 (No Content). Further requests will  
514 return 404 (Not Found). Clearly, the response is different from first request, but there is no  
515 change of state for any resource on server side because the original resource is already  
516 deleted.

<b>HTTP Method</b>	<b>Is Idempotent</b>
<i>GET</i>	True
<i>POST</i>	False

HTTP Method	Is Idempotent
<i>PUT</i>	True
<i>PATCH</i>	False
<i>DELETE</i>	True
<i>HEAD</i>	True
<i>OPTIONS</i>	True

517 **Table 5 – Idempotency of operations**

518

519 [R 18|1]

520 APIs SHALL adhere to the idempotency of operations specified in the list above.

521

522 [R 19|1]

523 APIs SHOULD implement the `Idempotency-Key`<sup>6</sup> HTTP header field and the corresponding  
 524 implementation advice in order to make non-idempotent operations like POST and PATCH  
 525 fault-tolerant.

### 526 3.2.7 Pagination

527 Querying an API with a GET can theoretically result in a huge return collection. Image  
 528 querying the API of one of the big internet search engines without pagination. Hundreds of  
 529 millions of results would have to be downloaded and displayed on a single page. That API  
 530 would be unusable. Pagination helps to keep the data load to a reasonable amount and at the  
 531 same time supports security aspects.

532 Historically, many APIs use offset pagination. A maximum page size (e.g. 20) is specified  
 533 and the clients requests the starting record or the page number. However, this approach  
 534 leads to fuzzy results: Suppose an API is supposed to return a list of all planned transport  
 535 movements of a certain carrier ordered by destination. The first page of results is returned  
 536 accurately. Before the client requests the next page or set of records, three possible things  
 537 can happen.

- 538 • The databank does not change at all. Then the next page of records is accurate.

<sup>6</sup> <https://www.ietf.org/archive/id/draft-ietf-httpapi-idempotency-key-header-01.txt>

- 539 • A record is added to the database, which falls under the result list of the first page, which  
540 the client already received. In that case, the last result of the previous page is returned as  
541 the first result of the second page. The list therefore contains a duplicate.
- 542 • In the opposite case, a planned transport movement that has already been returned to the  
543 client on the first page is deleted. The first data record of the second page therefore moves  
544 to the previous page. If the client now queries the next page, this data record is not  
545 transmitted at all.

546 As an inter-organisational data exchange cannot accept this type of results, an alternative  
547 solution for pagination is needed. The solution to this problem is the so called keyset-based  
548 or cursor-pagination<sup>7</sup>. In addition, cursor-pagination is much more time-efficient on large  
549 datasets than offset-pagination.

550 [R 20|1]

551 If pagination is used in an API, keyset-based pagination (cursor-pagination) SHALL be  
552 used. This means that the consumer cannot request a specific page, instead the consumer  
553 has to select a page-link provided by the server. The server SHALL provide links in the  
554 HTTP response header to the previous and next page and SHOULD provide links to the  
555 first and last page. More links MAY be provided.

556 The cursor-value is a string, created by the server using whatever method it likes. It  
557 identifies a point in a list of results for a query containing filters and sorting parameters for  
558 a specific moment in time. Therefore, it divides the list into those that fall before the cursor  
559 and those that fall after the cursor. There may optionally be one result that falls "on" the  
560 cursor.

561 Cursor-pagination assures a consistent data set for a query with filtering/sorting criteria at a  
562 specific moment in time. If another consumer performs the same query a moment later, he  
563 may get a different data set.

564 [R 21|1]

565 GET requests on collection results SHOULD implement pagination. The default and  
566 maximum page size SHOULD be 100, if not specified on the endpoint. If SHOULD be  
567 smaller, if the resulting page load is large. The default page size MAY be changed per  
568 endpoint. A consumer SHOULD be able to override the default page size.

569 If the filter, sorting and/or page size used is changed when getting a result, the pagination  
570 SHALL BE reset to the first page.

571 The query parameters described in the following table SHALL be used, rules SHALL be  
572 applied.

---

<sup>7</sup> <https://jsonapi.org/profiles/ethanresnick/cursor-pagination/>, <https://medium.com/swlh/how-to-implement-cursor-pagination-like-a-pro-513140b65f32>



Type	Explanation	Example
<i>Page size</i>	Overrides the default page size defined by the server / specification.	Example for the first query:  GET /transportMovements? carrier=ABC &status=PLANNED &sort=estimatedTimeOfArrival &pageSize=50
<i>Current page</i>	A link to the current page.	Link: <https://api.unece.org/ transportMovements? cursor=XXX>; rel="current"
<i>First page</i>	A link to the first page. If it is the first page the link MAY be omitted.	Link: <https://api.unece.org/ transportMovements? cursor=XXX>; rel="first"
<i>Next page</i>	A link to the next page. If it is the last page, the link to the next page MAY be omitted. Otherwise, a <code>null</code> link shall be provided.	Link: <https://api.unece.org/ transportMovements? cursor=XXX>; rel="next"  Link: <null>; rel="next"
<i>Previous page</i>	A link to the previous page. If it is the first page, the link to the previous page MAY be omitted. Otherwise, a <code>null</code> link shall be provided.	Link: <https://api.unece.org/ transportMovements? cursor=XXX>; rel="prev"
<i>Last page</i>	A link to the last page. If it is the last page, the link to the last page MAY be omitted. Otherwise, a <code>null</code> link shall be provided.	Link: <https://api.unece.org/ transportMovements? cursor=XXX>; rel="last"

573 When multiple links are given, they are separated by comma.

574 Example for a combination of Links:

575 Link:  
576 <https://api.unece.org/transportMovements?cursor=XXX>; rel="current",  
577 <https://api.unece.org/transportMovements?cursor=YYY>; rel="first",

```
578 <https://api.unece.org/transportMovements?cursor=ZZZ>; rel="next",  
579 <https://api.unece.org/transportMovements?cursor=LLL>; rel="last"
```

580

### 581 3.2.8 Filtering

582 Providing the ability to filter and sort collections in an API allows your consumers greater  
583 flexibility and controls on how they choose to consume a conformant API.

584 [R 22|1]

585 Sorting and filtering SHALL be done using query parameters. Using a path parameter is  
586 only allowed to identify a specific resource.

#### 587 3.2.8.1 Output Selection

588 Consumers can specify the attributes they wish to return in the response payload by  
589 specifying the attributes in the query parameters

590 Example that returns only the *first\_name* and *last\_name* fields in the response:  
591 `?attributes=first_name,last_name`

#### 592 3.2.8.2 Simple Filtering

593 Attributes can be used to filter a collection of resources.

594 `?last_name=Citizen` will filter out the collection of resources with the  
595 attribute `last_name` that matches `citizen`.

596 `?last_name=Citizen&date_of_birth=1999-12-31` will filter out the  
597 collection of resources with the attribute `last_name` that  
598 matches `citizen` and `date_of_birth` that matches 31<sup>st</sup> of December 1999.

599 [R 23|1]

600 As a general guide, filtering SHOULD be done with case insensitivity. Whether you choose  
601 to filter with case insensitivity or not SHALL be clearly documented.

602 The equal = operator is the only supported operator when used in this technique. For other  
603 operators and conditions next section.

#### 604 3.2.8.3 Advanced filtering with LHS Operators

605 There are situations where simple filtering does not meet the needs and a more  
606 comprehensive approach is required. Use the reserved keyword filters to define a more  
607 complex filtering logic. The general pattern is

608 `/path?property[operator]=value&property[operator]=value`

609 The = sign in this case is there to maintain URL query string compatibility with RFC 3986.  
610 However, the API service will use the operator inside the brackets for the actual  
611 comparison. A logical AND combines all query conditions.

612 The following operators are supported:

- 613 • [gte] Greater than or equalled to
- 614 • [egt] Equalled to or greater than
- 615 • [gt] Greater than
- 616 • [lt] Less than
- 617 • [lte] Less than or equalled to
- 618 • [elt] Equalled to or less than
- 619 • [ne] Not equalled

620 Example for filtering with LHS attributes:

621 `/path?creation_date[gt]=2020-11-30`

### 622 3.2.8.4 Rich Query with Lucene Syntax

623 [R 24|1]

624 If an application needs to support a richer search and filter capability that includes logical  
625 operators, fuzzy search, grouping, and so on, API MAY apply a query string according to  
626 lucene query syntax<sup>8</sup>. In that case, the filtering and query parameters normally are  
627 transmitted in the request body.

### 628 3.2.8.5 GraphQL

629 When API implementers would like to allow their clients rich flexibility to define response  
630 data sets that might include data from multiple APIs with rich filtering capability then a  
631 GraphQL query interface could be provided. GraphQL is a different architecture to  
632 RESTful APIs, is especially tailored to queries across multiple entities, and allows clients to  
633 specify exactly which data elements they would like in the response. If you find yourself  
634 building very complex RESTful queries then you should consider GraphQL as an  
635 alternative.

636 GraphQL is not discussed further in this RESTful API design guide.

### 637 3.2.9 Sorting

638 Providing data in specific order is often the requirement from client applications and hence  
639 it is important to provide the flexibility for clients to retrieve the data in the order they need  
640 it.

641 [R 25|1]

---

<sup>8</sup> [https://lucene.apache.org/core/2\\_9\\_4/queryparsersyntax.html](https://lucene.apache.org/core/2_9_4/queryparsersyntax.html)

642 Sorting SHOULD be limited to specified fields. The sort direction MAY be omitted. The  
 643 default sort direction is ascending. A colon `:` is used to separate the field name and the sort  
 644 direction. Multiple sort fields are separated by comma `,`.

Query Parameter	Description
<i>sort=name</i> <i>sort=name:asc</i>	Sort by the name field in ascending order.
<i>sort=name:desc</i>	Sort by the name field in descending order.
<i>sort=yearOfBirth,name:dec</i>	Sort by year of birth in ascending order. If two equal years exist, sort the names by birth year in descending order.

645 **Table 6: Sort examples**

### 646 3.2.10 API Responses and error handling

647 [R 26|1]  
 648 HTTP response codes SHALL be used.  
 649 The following table defines HTTP response codes supported by conformant APIs. The  
 650 column **Response** indicates whether an additional error response payload is  
 651 RECOMMENDED to be returned as described in chapter 0.

652

Code	Status	Response	When to use
200	OK	No	The request was successfully processed
201	Created	No	The resource was created. The <b>Location</b> HTTP response header SHALL be returned to indicate where the newly created resource is accessible.
202	Accepted	No	The request was accepted, and is processed asynchronously.
204	No content	No	The server successfully processed the request and is not returning any content. There is no need for the client to move to a different location.
400	Bad Request	Yes	The server cannot process the request (such as malformed request syntax, size too large, invalid request message)

Code Status		Response When to use	
			framing, or deceptive request routing, invalid values in the request). For sensitive information, a code <b>404 Not found</b> MAY be returned instead.
401	Unauthorised	Yes	The request could not be authenticated. For sensitive information, a code <b>404 Not found</b> MAY be returned instead.
403	Forbidden	Yes	The request was authenticated but is not authorised to access the resource. For sensitive information, a code <b>404 Not found</b> MAY be returned instead.
404	Not found	Yes	The resource was not found.
405	Not Allowed		The method is not implemented for this resource. The response MAY include an <b>allow</b> HTTP response header containing a list of valid methods for the resource.
408	Request Timeout	No	The request timed out before a response was received. A <b>Retry-After</b> HTTP response header is RECOMMENDED to be returned.
415	Unsupported Media Type	Yes	This status code indicates that the server refuses to accept the request because the content type specified in the request is not supported by the server
422	Unprocessable Entity		This status code indicates that the server understands the content type of the request entity, and the syntax of the request entity is correct, but it was unable to process the contained instructions.
429	Too Many Requests		There have been too many requests (by the consumer). A <b>Retry-After</b> HTTP response header is RECOMMENDED to be returned. A response body MAY be returned containing information about the reason for the response code. A possible reason may be if a quota of requests for the day / hour / month etc. was exceeded.

Code Status	Response When to use
500 Internal Server error	An internal server error. The response body may contain error messages. The response body SHALL not reveal any server configuration information (e.g. version, paths, database used, etc.).
501 Method Not Implemented	It indicates that the request method is not supported by the server and cannot be handled for the requested resource. Implementing this response code allows a higher interoperability between API implementations based on the same specification, if a specific server does not support one of the specified methods (yet). A <code>Link</code> HTTP response header is RECOMMENDED to point to the specific documentation.
503 Service unavailable	It indicates that the service is unavailable (e.g. due to maintenance reasons). A <code>Retry-After</code> HTTP response header is RECOMMENDED to be returned.

653

Table 7: HTTP response codes

654

[R 27|1]

655

The following table defines which HTTP response codes SHALL be supported for a specific HTTP request method by conformant APIs. Column `use` indicates how a conformant API supports the specified http response code:

656

657

658

- **M** the code SHALL be supported

659

- **MA** SHALL be supported for requests where the response is handled asynchronous, for instance due to forwarding or processing time. In that case, a `Location` HTTP response header SHALL be gives that points to the respective resource. In addition, a `Retry-After` HTTP response header is RECOMMENDED to be returned.

660

661

662

663

- **R** the code is recommended to be supported.

664

The default response code for a positive response is marked in **bold**.

665

## HTTP

Request method	Code	Status	Use
<b>GET</b>	<b>200</b>	<b>OK</b>	<b>M</b>
	202	Accepted	MA
	400	Bad Request	R

**HTTP**

<b>Request method</b>	<b>Code</b>	<b>Status</b>	<b>Use</b>
	401	Unauthorised	M
	403	Forbidden	M
	404	Not found	M
	405	Not Allowed	M
	408	Request Timeout	R
	415	Unsupported Media Type	M
	429	Too Many Requests	R
	500	Internal Server error	M
	503	Service unavailable	R
<b>POST</b>	<b>201</b>	<b>Created</b>	M
	202	Accepted	MA
	400	Bad Request	M
	401	Unauthorised	M
	403	Forbidden	M
	408	Request Timeout	R
	415	Unsupported Media Type	M
	422	Unprocessable Entity	R
	429	Too Many Requests	R
	500	Internal Server error	M
	503	Service unavailable	R
<b>PATCH</b>	202	Accepted	MA
	<b>204</b>	<b>No content</b>	M
	400	Bad Request	M
	401	Unauthorised	M
	403	Forbidden	M
	404	Not found	M
	405	Not Allowed	M
	408	Request timeout	R
	415	Unsupported Media Type	M
	422	Unprocessable Entity	M
	429	Too Many Requests	R
	500	Internal Server error	M
	503	Service unavailable	R
<b>PUT</b>	202	Accepted	MA
	<b>204</b>	<b>No content</b>	M
	400	Bad Request	M
	401	Unauthorised	M
	403	Forbidden	M

**HTTP**

<b>Request method</b>	<b>Code</b>	<b>Status</b>	<b>Use</b>
	404	Not found	M
	405	Not Allowed	M
	408	Request Timeout	R
	415	Unsupported Media Type	M
	422	Unprocessable Entity	M
	429	Too Many Requests	R
	500	Internal Server error	M
	503	Service unavailable	R
	202	Accepted	MA
<b>DELETE</b>	<b>204</b>	<b>No content</b>	M
	400	Bad Request	M
	401	Unauthorised	M
	403	Forbidden	M
	404	Not found	M
	405	Not Allowed	M
	408	Request timeout	R
	415	Unsupported Media Type	M
	422	Unprocessable Entity	M
	429	Too Many Requests	R
	500	Internal Server error	M
	503	Service unavailable	R



667 **3.2.11 Error Response Payload**

668 For some errors, returning the HTTP status code is enough to convey the response.  
669 Additional error information can be supplemented in the response body. For example;  
670 HTTP 400 Bad request is considered too generic for a validation error and more information  
671 must be provided in the response body.

672 [R 28|1]

673 An API SHALL implement an error response schema to allow a standardised error  
674 handling. The response SHALL use the following JSON Schema. The JSON Schema MAY  
675 be extended.

```
676 {  
677   "$schema": "https://json-schema.org/draft/2020-12/schema",  
678   "type": "object",  
679   "properties": {  
680     "errors": {  
681       "type": "array",  
682       "items": {  
683         "type": "object",  
684         "properties": {  
685           "id": { "type": "string",  
686             "format": "uuid" },  
687           "code": { "type": "string" },  
688           "detail": { "type": "string" },  
689           "source": {  
690             "type": "object",  
691             "properties": {  
692               "parameter": { "type": "string" },  
693               "pointer": { "type": "string",  
694                 "format": "json-pointer" }  
695             },  
696             "unevaluatedProperties": false  
697           },  
698           "sourcePointer": { "type": "string",  
699             "format": "json-pointer" }  
700         },  
701         "required": ["code", "detail"],  
702         "patternProperties": { "^x-": true },  
703         "unevaluatedProperties": false  
704       },  
705       "minItems": 1  
706     },  
707   },  
708   "required": [ "errors" ],  
709   "patternProperties": { "^x-": true },  
710   "unevaluatedProperties": false  
711 }
```

712 The following definitions are applied:

### Error response

attributes	Description
<i>id</i>	Identifier of the specific error
<i>detail</i>	A human-readable explanation specific to this occurrence of the problem.
<i>code</i>	An application-specific error code
<i>source</i>	An object containing computer processable information about the origin of the error.
<i>parameter</i>	The (query) parameter where the error was caused.
<i>pointer</i>	JSON Pointer [RFC6901] to the associated entity in the request document [e.g. <code>"/data"</code> for a primary data object, or <code>"/data/attributes/title"</code> for a specific attribute].

713 **Table 8: Error response attributes**

714 Example for a 400 Bad Request error response:

```
715 {
716   "errors": [
717     {
718       "id": "86032cbe-a804-4c3b-86ce-ec3041e3effc",
719       "code": "19283",
720       "detail": "Invalid value(s) in request input",
721       "source": {
722         "parameter": "id"
723       }
724     }
725   ]
726 }
```

727 Example for a 503 Service unavailable error response:

```
728 Retry-After: Sat, 16 Apr 2022 15:00:00 GMT
729 {
730   "errors": [
731     {
732       "id": "45786a8f-452e-492f-a779-801b5d0bd0a7",
733       "code": "19284",
```

```
734     "detail": "The service is unavailable due to maintenance. Come back
735 at 15:00 GMT.",
736     "source": {
737         "pointer": "#/resources/12345"
738     }
739 }
740 ]
741 }
```

### 742 **3.2.12 Design rule examples**

#### 743 Good examples

744 Get a list of voyages:

745 *GET* https://api.logistics.io/v1/transport/voyages

746 Filtering in a query:

747 *GET* https://api.logistics.io/v1/transport/voyages?departure\_location=AUBN  
748 E&date=2022-04-16

749 Get a single voyage:

750 *GET* https://api.logistics.io/v1/transport/voyages/N234

751 Create a new voyage:

752 *POST* https://api.logistics.io/v1/transport/voyages

753 {content body with voyage data in JSON format}

754 Update a voyage status:

755 *PATCH* https://api.logistics.io/v1/transport/voyages/N234/status

756 {content body status data in JSON format}

## 757 4 Well-documented APIs

### 758 4.1 General considerations

759 [R 29|1]

760 The following rules are RECOMMENDED:

- 761 - The definitions in a conformant OpenAPI specification **SHALL** be considered as
- 762 technical contracts between designers and developers and between consumers and
- 763 providers.
- 764 - Mock APIs **SHOULD** be created using the API description to allow early code
- 765 integration for development.
- 766 - The behaviour and intent of the API **SHOULD** be described with as much information
- 767 as possible.
- 768 - Operations **SHOULD** provide examples for request and response bodies.
- 769 - Expected response codes and error messages **SHOULD** be provided in full.
- 770 - Known issues or limitations **SHOULD** be clearly documented.
- 771 - Expected performance, uptime and SLA/OLA **SHOULD** be clearly documented.
- 772 - Although YAML is a supported file format of an OpenAPI specification, the JSON
- 773 format **SHOULD** be used as the OpenAPI specification format.

### 774 4.2 API Versioning

#### 775 4.2.1 Versioning Scheme

776 [R 30|1]

777 All APIs **SHALL** apply Semantic versioning 2.0.0<sup>9</sup>:

778 MAJOR . MINOR . PATCH

779 The first version of an API **SHALL** start with a **MAJOR** version of 1.

780 Pre-release version<sup>10</sup> information and build metadata<sup>11</sup> version information **SHALL NOT** be  
781 used in API versioning.

782 Use the following guidelines when incrementing the API version number:

<sup>9</sup> <https://semver.org/spec/v2.0.0.html>

<sup>10</sup> <https://semver.org/spec/v2.0.0.html#spec-item-9>

<sup>11</sup> <https://semver.org/spec/v2.0.0.html#spec-item-10>

- 783 • **MAJOR** version when you make API changes that break backwards-compatibility,
- 784 • **MINOR** version when you add functionality in a backwards-compatible manner,
- 785 and
- 786 • **PATCH** version when you make backwards-compatible bug fixes. A PATCH does
- 787 not include new functionality.

## 788 4.2.2 URI Versioning

789 [R 31|1]

790 All APIs **SHALL** use URI versioning. They **SHALL** include the **MAJOR** version as part of  
791 the URI in the format of '**v**{**MAJOR**}'

792 Example:

793 `https://api.logistics.io/transport/v1/voyages`

794 The minor and patch version **SHALL NOT** be used in the URI.

## 795 4.2.3 Providing version information

796 APIs conforming to this technical specification are intended to be used with REST  
797 principles. Those mandate HATEOS (see chapter 4.3.2) support. On major aspect is the  
798 self-descriptiveness of an API. Although a support of HATEOS is not required, providing  
799 basic metadata about the called API including version information is useful even in not  
800 RESTful scenarios.

801 [R 32|1]

802 A custom header named **API-Version** **SHALL** be added to any response of the API. It  
803 **SHALL** be aligned with the URI version and **SHALL** state all three levels:

804 `API-Version: 1.21.5`

805

806 [R 33|1]

807 An **API-Version** custom header **MAY** be added to a request. If added, it **SHALL** only  
808 contain the **MAJOR** version.

809 `API-Version: 1`

810 In order to easily provide information about an API in a standardised way, the following  
811 information can be retrieved from any conformant API:

812 [R 34|1]

813 An API **SHALL** implement a response to a GET request to the base URI of the API. The  
814 response **SHALL** use the following JSON Schema:

```
815 {
816   "$schema": "https://json-schema.org/draft/2020-12/schema",
817   "type": "object",
818   "properties": {
819     "title": { "type": "string" },
```

```

820     "version": {
821         "type": "string",
822         "pattern": "^\\d+(-.+)?\\.\\.\\d+(-.+)?\\.\\.\\d+(-.+)?$"
823     },
824     "status": {
825         "type": "string",
826         "enum": ["DRAFT", "ACTIVE", "DEPRECATED", "RETIRED"]
827     },
828     "effective": {
829         "type": "string",
830         "format": "date-time"
831     },
832     "specification": {
833         "type": "string",
834         "format": "uri"
835     }
836 },
837 "required": [
838     "title", "version", "status", "effective", "specification"
839 ],
840 "$comment" : "Allow extensions to the API metadata",
841 "patternProperties": {
842     "^x-": true
843 },
844 "unevaluatedProperties": false
845 }

```

846 The following definitions are applied:

- 847 • **title**: The name of the API. It SHALL be identical to the API title defined in the
- 848 OpenAPI specification
- 849 • **version**: The API version
- 850 • **status**: The operation status of the API. The following values are used:
  - 851 ○ **ACTIVE**: The API is in its productive phase. Maintenance or deprecation of
  - 852 specific services SHALL be indicated at the service level. The **effective**
  - 853 defines the moment in the past since when API is in its productive phase.
  - 854 ○ **DEPRECATED**: The complete API is going to its end-of-life phase. The
  - 855 **effective** defines the moment in the future when the API is intended to
  - 856 switch to **RETIRED**. The rules of deprecation (see chapter 4.2.5) are applied
  - 857 additionally.
  - 858 ○ **RETIRED**: The complete API is to its end-of-life phase. The **effective** defines
  - 859 the moment in the past when the API was set to **RETIRED**. The rules of
  - 860 deprecation (see chapter 4.2.5) are applied additionally.
- 861 • **effective**: The moment in time corresponding to the **status**.
- 862 • **specification**: A valid URI to the OpenAPI specification of the current API. This
- 863 way the available services and data types become self-descriptive from their basic
- 864 structure. The OpenAPI specification SHOULD be public where possible and easily
- 865 accessible to those that require it.

866 Additional metadata can be added to the response if required.

```
867 Example:
868 GET https://api.uncefact.unece.org/v1/
869 HTTP 200 OK
870 content-type: application/json; charset=utf-8
871 API-Version: 1.0.0
872 {
873   "title": "UN/CEFACT Demo API",
874   "version": "1.0.0",
875   "status": "ACTIVE",
876   "effective": "2022-06-02T23:00:00Z",
877   "specification": "https://service.unece.org/demo/demoAPI.json",
878   "x-info" : "Additional information"
879 }
```

880 During the draft, development or testing phase of an API sandbox environments are used to  
881 validate the intended functionality. For those kinds of APIs in development no additional  
882 state like **DRAFT** is provided.

883 [R 35|2]  
884 APIs that are still in a **DRAFT** status **SHOULD** be placed in a sandbox environment. This  
885 could be done by changing the basis URL accordingly.

886 Example for a productive base URL:

887 <https://api.uncefact.unece.org/v1/>

888 Examples for a development base URL:

889 <https://sandbox.api.uncefact.unece.org/v1/>

890 <https://staging.api.uncefact.unece.org/v1/>

#### 891 **4.2.4 Robustness<sup>12</sup>**

892 It is critical that APIs are developed with loose coupling in mind to ensure backwards  
893 compatibility for consumers.

894 [R 36|1]  
895 Within a major release backward compatibility **SHALL NOT** be broken.

896 The following changes are deemed backwards compatible:

- 897 • Addition of a new optional field to a representation
- 898 • Addition of a new link to the `_links` array of a representation
- 899 • Addition of a new endpoint to an API
- 900 • Additional support of a new media type (e.g. `Accept: application/pdf`)

---

<sup>12</sup> [https://en.wikipedia.org/wiki/Robustness\\_principle](https://en.wikipedia.org/wiki/Robustness_principle)

901 The following changes are **NOT** deemed backwards compatible:

- 902 • Removal of fields from representations
- 903 • Changes of data types on fields (e.g. `string` to `boolean`)
- 904 • Changing semantic definitions
- 905 • Removal of endpoints or functions
- 906 • Removal of media type support

907

908 [R 37|1]

909 API clients and subscribers **SHOULD** be robust:

- 910 - Be conservative with API requests and data passed as input.
- 911 - Be tolerant with unknown fields in the payload, but do not eliminate them from payload
- 912 if needed for subsequent **PUT** requests.

## 913 4.2.5 Deprecation and End of Life Policy

914 When designing new APIs one of the most important dates to consider is when the API will  
915 be retired. APIs are not intended to last forever. Some APIs are retired after a short time as  
916 they may be proving a use-case; others may be removed when better options are available  
917 for users.

918 The End-of-Life (EOL) policy determines the process that APIs go through to move  
919 through their workflow from **ACTIVE** to the **RETIRED** state. The EOL policy is designed to  
920 ensure a consistent and reasonable transition period for API customers who need to migrate  
921 from the old API version to the new API version while enabling a healthy process to retire  
922 technical debt.

### 923 Major API Version EOL

924 Major API versions **MAY** be backwards compatible with preceding major versions. The  
925 following rules apply when retiring a major API version.

926 [R 38|1]

927 An API **SHALL NOT** be set to **DEPRECATED** until a replacement service is running with  
928 status **ACTIVE**.

929 The root service of the API **SHALL** provide the **Deprecation Header Field**<sup>13</sup> and the **Sunset**  
930 **HTTP Response Header Field**<sup>14</sup>.

931 A **Link header** **SHALL** be added in combination with the **Deprecation header**. It **SHALL**  
932 provide a link to the documentation. A second **Link header** **SHALL** be added linking to the  
933 replacement version of the API.

934 Additionally, the following thoughts should be considered:

---

<sup>13</sup> <https://tools.ietf.org/html/draft-dalal-deprecation-header-02>

<sup>14</sup> <https://tools.ietf.org/html/rfc8594#section-3>



- 935 1. A minimum transition period of 60 days should be planned to give users adequate  
936 notice to migrate.
- 937 2. Deprecation of API versions with external users should be considered on a case-by-  
938 case basis and may require additional deprecation time and/or constraints to  
939 minimise impact to users.
- 940 3. If a versioned API is **ACTIVE** or **DEPRECATED** state has no registered users, it may move  
941 to the **RETIRED** state immediately.

942  
943 [R 39|1]

944 Deprecated endpoints **SHALL** be documented in the OpenAPI specification using the  
945 **DEPRECATED** property introduced since OpenAPI 3.0.0.

946 Deprecated endpoints **SHOULD** provide the Deprecation Header Field and the Sunset  
947 HTTP Response Header Field.

948 A Link header **SHALL** be added in combination with the Deprecation header. It **SHALL**  
949 provide a link to the documentation.

950 Where possible, communication **SHOULD** be sent to consumers of deprecated endpoints.

951

952 [R 40|1]

953 The introduction of a major version **SHOULD** be avoided, whenever possible. This **MAY**  
954 be achieved as follows:

- 955 - Create a new service endpoint, if the process is changed.

956 Duplicate and Deprecate: add a **Deprecation Header** to the old service including a **Link**  
957 **Header** to documentation and to the new service. Eventually add a **Sunset Header**.

- 958 - Create a new resource (a variant of the old) in addition to the old.

## 959 **Minor API Version EOL**

960 Due to the specified URL versioning the URL does not change if the minor version of an  
961 API changes. Minor API versions are backwards compatible with preceding minor versions  
962 within the same major version.

963 Therefore, the status before, during or after a minor API version update does not change.  
964 The change should have no impact on existing subscribers so there is no need to transition  
965 through a **DEPRECATED** state to facilitate client migration.

966 [R 41|2]

967 New resources or service endpoints can be added during a minor release. In order to support  
968 the implementation of those new services a sandbox environment **SHOULD** be provided to  
969 the interested or affected consumers.

970

971 [R 42|1]

972 It is RECOMMENDED that no more than 3 parallel MAJOR versions are available.  
973 Implementers of the API SHALL NOT be more than 1 major version behind the latest  
974 version.

975 Example

976 Version 1 is **RETIRED**

977 Version 2 is **DEPRECATED**

978 Version 3 is **ACTIVE**

## 979 **4.3 Hypermedia**

### 980 **4.3.1 Hypermedia - Linked Data**

981 An API becomes RESTful by meeting the requirements of the REST principles. A key  
982 principle is the discoverability of the API. Ideally, this is achieved by an API being  
983 completely self-describing. According to the inventor of REST, Roy Fielding<sup>15</sup>, the use of  
984 hypermedia is a prerequisite for designing a RESTful API.

985 Hypermedia means that links are provided together with the response payload. They inform  
986 the consumers what options are available according to their original request. Though simple  
987 in concept hypermedia links in APIs, allow consumers to locate resource without the need  
988 to have an upfront understanding of the resource and its relationship.

989 This is similar to the navigation of a web page. The user is not expected to know the  
990 structure of the web page prior to visiting. They can simply browse to the home page and  
991 the navigation lets them browse the site as required.

992 APIs that do not provide links are more difficult to use and expect the consumer to refer to  
993 the documentation.

### 994 **4.3.2 HATEOAS**

995 *Hypermedia As The Engine Of Application State* is the concept of representing allowable  
996 actions as hyperlinks associated with resource. Similar to Hypermedia Linked Data concept  
997 the links defined in the response data represents state transitions that are available from that  
998 current state to adjacent states.

999 Example:

```
1000 HEAD /v1/accounts/4711
1001 HTTP/1.1 200 OK
1002 Link: <https://api.unece.org/v1/accounts/4711>; rel="self",
1003       <https://api.unece.org/v1/accounts/4711/deposit>; rel="deposit",
1004       <https://api.unece.org/v1/accounts/4711/withdraw>; rel="withdraw",
1005       <https://api.unece.org/v1/accounts/4711/transfer>; rel="transfer"
```

1006 If the same account is overdrawn, the only allowed action could be to deposit:  
1007

---

<sup>15</sup> [https://www.ics.uci.edu/~fielding/pubs/dissertation/fielding\\_dissertation\\_2up.pdf](https://www.ics.uci.edu/~fielding/pubs/dissertation/fielding_dissertation_2up.pdf)

```
1008 Example:
1009 GET /v1/accounts/4711
1010
1011 HTTP/1.1 200 OK
1012 Link: <https://api.unece.org/v1/accounts/4711>; rel="self",
1013       <https://api.unece.org/v1/accounts/4711/deposit>; rel="deposit"
1014 Content-Type: application/json
1015 Content-Length: ...
1016 {
1017   "accountId": "4711",
1018   "balance": {
1019     "currency": "EUR",
1020     "value": -25
1021   }
1022 }
```

### 1022 4.3.3 Hypermedia Compliant API

1023 In APIs, request methods such as *DELETE*, *PATCH*, *POST* and *PUT* initiate a transition in  
1024 the state of a resource. A *GET* request never changes the state of the resource that is  
1025 retrieved.

1026 [R 43|1]

1027 In order to provide a better experience for API consumers, APIs SHOULD provide a list of  
1028 state transitions that are available for each resource. As possible values for link relation  
1029 types the official IANA registry list<sup>16</sup> SHALL be used. It MAY be extended. Any extension  
1030 SHALL be documented in the API specification.

1031 An example of an API that exposes a set of operations to manage a user account lifecycle  
1032 and implements the HATEOAS interface constraint is as follows:

1033 A client starts their interaction with a service through the URI */users*. This fixed URI  
1034 supports both *GET* and *POST* operations. The client decides to do a *POST* operation to  
1035 create a user in the system.

1036 Request

```
1037 POST https://api.unece.org/v1/v1/users
1038 {
1039   "firstName": "John",
1040   "lastName" : "Smith",
```

---

<sup>16</sup> <https://www.iana.org/assignments/link-relations/link-relations.xhtml>

```
1041     ...
1042 }
```

1043 The API creates a new user from the input and returns the following links to the client in the  
1044 response.

- 1045 • A link to the created resource in the *Location* header (to comply with the 201 response  
1046 spec)
- 1047 • A link to retrieve the complete representation of the user (a.k.a. *self*-link) (*GET*).
- 1048 • A link to update the user (*PUT*).
- 1049 • A link to partially update the user (*PATCH*).
- 1050 • A link to delete the user (*DELETE*).

```
1051 HTTP/1.1 201 CREATED
1052 Location: https://api.unece.org/v1/users/JFWXHGUV7VI
1053 Link: <https://api.unece.org/v1/users/JFWXHGUV7VI>, rel="self",
1054       <https://api.unece.org/v1/users/JFWXHGUV7VI>, rel="delete",
1055       <https://api.unece.org/v1/users/JFWXHGUV7VI>, rel="replace",
1056       <https://api.unece.org/v1/users/JFWXHGUV7VI>, rel="edit"
```

1057 A client can store these links in its database for later use.

1058 In summary:

- 1059 • There is a well-defined index or navigation entry point for every API, which a client  
1060 navigates to in order to access all other resources.
- 1061 • The client does not need to build the logic of composing URIs to execute different  
1062 requests or code any kind of business rule by looking into the response details that  
1063 may be associated with the URIs and state changes.
- 1064 • The client acknowledges the fact that the process of creating URIs belongs to the  
1065 server.
- 1066 • Client treats URIs as opaque identifiers.
- 1067 • APIs using hypermedia in representations could be extended seamlessly. As new  
1068 methods are, introduced responses could be extended with relevant HATEOAS  
1069 links. These way clients could take advantage of the functionality in incremental  
1070 fashion. For example, if the API starts supporting a new *PATCH* operation then  
1071 clients could use it to do partial updates.

1072 The mere presence of links does not decouple a client from having to learn the data required  
1073 making requests for a transition and all associated link semantics particularly  
1074 for *POST/PUT/PATCH* operations.

## 1075 5 API Security

1076 [R 44|1]

1077 All API endpoints SHALL be secured. HTTPS SHALL be used. The OAuth2 security  
1078 scheme is RECOMMENDED. Other security schemes MAY be used.

1079 The receivers' endpoints of subscription callbacks MAY be designed with different security  
1080 measures like those described in chapter 6.3.

1081 The following aspects of API security are RECOMMENDED to be implemented:

### 1082 **Rate Limiting**

1083 Rate limiting and throttling policies are introduced to prevent abuse of your API.

1084 Appropriate alerts should be implemented and respond with informative errors when

1085 thresholds are nearing or have been exceeded. See [https://greenbytes.de/tech/webdav/draft-](https://greenbytes.de/tech/webdav/draft-ietf-httpapi-ratelimit-headers-latest.html)

1086 [ietf-httpapi-ratelimit-headers-latest.html](https://greenbytes.de/tech/webdav/draft-ietf-httpapi-ratelimit-headers-latest.html) for implementation details.

### 1087 **Error Handling**

1088 When your application displays error messages, it should not expose information that could

1089 be used to attack your system. You should establish the following controls when providing

1090 error messages:

- 1091 • Your API MUST mask any system related errors behind standard HTTP status  
1092 responses and error messages e.g. do not expose system level information in your  
1093 error response
- 1094 • Your API MUST NOT pass technical details (e.g. call stacks or other internal hints)  
1095 to the client

### 1096 **Audit Logs**

1097 An important aspect of security is to be notified when something wrong occurs, and to be

1098 able to investigate it. It is RECOMMENDED to implement logging.

- 1099 • Write audit logs before and after security related events which can trigger the alerts
- 1100 • Sanitizing the log data to prevent log injection attacks

### 1101 **Input Validation**

1102 Input validation is performed to ensure only properly formed data is received by your

1103 system, this helps to prevent malicious attacks

- 1104 • Input validation should happen as early as possible, preferably as soon as the data is  
1105 received from the external party
- 1106 • Define an appropriate request size limit and reject requests exceeding the limit
- 1107 • Validate input: e.g. length / range / format and type
- 1108 • Consider logging input validation failures. Assume that someone who is performing  
1109 hundreds of failed input validations per second has a malicious intent.
- 1110 • Constrain string inputs with regular expression where appropriate

### 1111 **Content Type Validation**

1112 Honour the specified content-type. Reject requests containing unexpected or missing  
1113 content type headers with HTTP response status *415 Unsupported Media Type*.

1114 **Gateway Security Features**

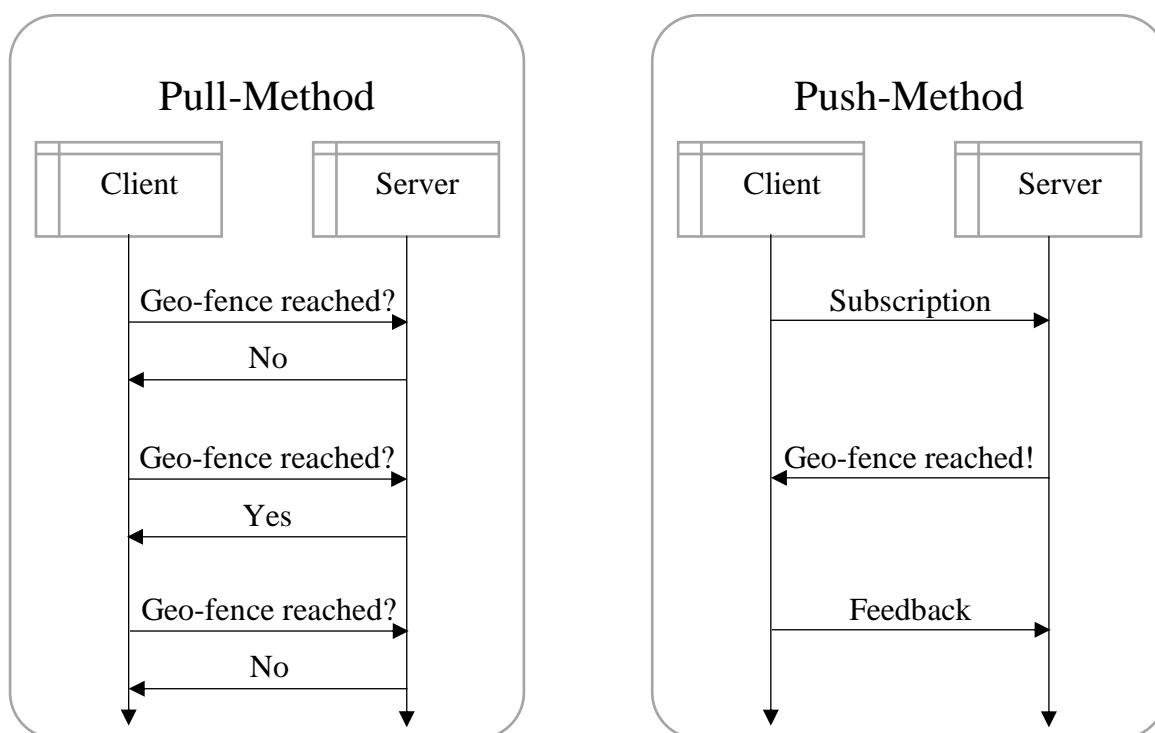
1115 It is RECOMMENDED to use the security policy features available in the gateway rather  
1116 than to implement the policies in your back-end API.

1117

## 1118 6 Event driven data exchange

1119 Classic B2B data exchange scenarios reach their limits especially when it comes to  
 1120 processing real-time data. For example, one of the most important pieces of information in  
 1121 just-in-time production is the expected arrival time (ETA) at the factory. PULL scenarios  
 1122 are often implemented, where the consumer periodically asks the data sender for the current  
 1123 status of the delivery. Alternatively, the carrier sends a status message at regular but short  
 1124 intervals on the current status of the delivery with detailed information for each  
 1125 consignment item. This leads to tremendous amounts of data, so that in practice the  
 1126 minimum interval of such updates is about 15 minutes. Thus, in such scenarios, real-time  
 1127 information is a long way off.

1128 One approach to solving this problem is now to define events when they occur and  
 1129 exchange the data instead of constantly exchanging (less relevant) information. This could  
 1130 be the case, for example, if a geo-fence is crossed, a temperature is exceeded or not reached,  
 1131 or a clearance takes longer than it is intended. In the consumer space, such scenarios are  
 1132 already familiar, for example, when the buyer of an online delivery is notified that the  
 1133 package is only 10 stops away from delivery.



1134

1135

Figure 1: Event driven data exchange – pull versus push method

### 1136 6.1 Callbacks

1137 In OpenAPI, you can define callbacks. Those are asynchronous requests to a consumer  
 1138 specified URL that are called in response to a specific event. An example is that a carrier is  
 1139 informed if a specific vessel approaches a port.



1140 In order to be able to receive this information, the receiver first needs to subscribe to this  
1141 event information in the API. When subscribing, he may pass filter criteria that define the  
1142 conditions under which the consumer will be informed. Examples are a specific journey  
1143 where the consumer wants to get informed if it approaches a specific port.

1144 The basic principle is that a consumer subscribes for an event, supplies a (callback) URL  
1145 and stands by for incoming HTTP requests to that URL.

## 1146 **6.2 Webhooks**

1147 Since OpenAPI 3.1, webhooks are supported as well. The main difference between  
1148 callbacks and webhooks is that webhooks are synchronous to the process flow handled by  
1149 the APIs. This means that a consumer can directly hook into the process and thus, if  
1150 necessary, change the processed information before it is further processed. A webhook is  
1151 used to extend the functionality of the API.

1152 A webhook defines a clear point in the process where the consumer is enabled to react on,  
1153 for example based on some external event. An example is if you want to react immediately  
1154 on any incoming order/payment etc. The payload itself is given with the webhook and often  
1155 allows modifications. Examples are the option to link to a GitHub push event or to define a  
1156 plugin for the WordPress content management system. The latter modifies for example the  
1157 displayed HTML page directly by adding new functionalities like images, tables, videos or  
1158 similar to the HTML page. Such modifications would not be possible with an asynchronous  
1159 callback.

## 1160 **6.3 Security guideline for callbacks (informative)**

1161 Since webhooks work synchronously, the same security rules apply to them as to the entire  
1162 API. In contrast, the call direction is reversed for asynchronous callbacks. This makes it  
1163 important to ensure that the callback URL is only called from the authorized API.

1164 The following rules are based on the current approach of the DCSA. They are in the trial  
1165 phase at the time of publication of this document. As soon as sufficient practicality has been  
1166 demonstrated, this specification will be updated accordingly. Against this background, the  
1167 following rules are purely informative and not normative.

1168

1169 [R 45|1+Inf]

1170 All event subscriptions SHALL be secured via a Shared Secret that is used to sign every  
1171 callback message as described in this section. The secret SHALL be provided BASE64  
1172 encoded. The provider SHALL NOT expose the `secret` in any endpoint. It is write-only.  
1173 The provider SHALL assure that the secret fulfils the security requirements of the applied  
1174 algorithm.

1175 [R 46|2+Inf]  
 1176 A sha256 signature SHALL be used computed as an HMAC-SHA246 over the request  
 1177 body<sup>17</sup>. The subscriber provided Shared Secret SHALL be of at least 32-byte length. It  
 1178 SHOULD not be longer than 64 byte, as longer keys do not provide additional security to  
 1179 that algorithm.  
 1180 To improve security, it is RECOMMENDED to update the `secret` (and together with it the  
 1181 `callbackURL`) on a regular basis.

1182 [R 47|1+Inf]  
 1183 The publisher API SHALL provide the following endpoints for subscriptions:  
 1184 • `POST .../subscriptions` to create a new subscription  
 1185 • `GET .../subscriptions` to list all subscriptions the subscriber has access to  
 1186 • `GET .../subscriptions/{subscriptionId}` to get details about a specific subscription  
 1187 • `PUT .../subscriptions/{subscriptionId}` to update a specific subscription  
 1188 • `PUT .../subscriptions/{subscriptionId}/secret` to update the secret of a specific  
 1189 subscription  
 1190 • `DELETE .../subscriptions/{subscriptionId}` to cancel a specific subscription

### 1191 6.3.1 Subscription setup (informative)

1192 The setup of a subscription follows the following steps:

- 1193 1. The subscriber defines a Shared Secret and registers with the `secret` and a  
 1194 `callbackURL` in the publisher's system. It is recommended to use a not-easy-to-guess<sup>18</sup>  
 1195 callback URL and to update it when the secret is changed.
- 1196 2. The publisher confirms the subscription and returns the `subscriptionId` to the  
 1197 subscriber.
- 1198 3. The subscriber records the `subscriptionId` associated with the shared `secret`.

1199

#### 1200 Example for a subscription setup

##### 1201 1. Initiating the subscription

```
1202 POST https://api.unece.org/v1/events/subscribe
1203 Content-Type: application/json
1204 Content-Length: ...
1205 {
1206   "callbackURL" : "https://callback.example.com/callback/Ujh4kkQ9A",
```

<sup>17</sup> Compare <https://docs.github.com/en/developers/webhooks-and-events/webhooks/securing-your-webhooks>

<sup>18</sup> [https://callback.example.com/callback/\\$RANDOM\\_STRING](https://callback.example.com/callback/$RANDOM_STRING)

```

1207   "secret":
1208   "MDEyMzQ1Njc4OWFiY2RlZjAxMjM0NTY3ODlhYmNkZWYwMTIzNDU2Nzg5YWJjZGVmMDEyMzQz
1209   NjU3ODlhYmNkZQ",
1210   ... additional filter parameters etc. ...
1211   }

```

### 1212 2.a Confirmation of the publisher if the `callbackURL` is valid

1213 Remark: As the subscription is not setup yet, not additional headers are provided.

```

1214 HEAD https://callback.example.com/callback/Ujh4kkQ9A

```

### 1215 2.b Response of the subscriber that the `callbackURL` is valid

```

1216
1217 HTTP/1.1 204 No Content

```

## 1218 3. Response from the publisher

```

1219 HTTP/1.1 201 Created
1220 Content-Type: application/json
1221 Content-Length: ...
1222 {
1223   "subscriptionId": "936DA01F-9ABD-4D9D-80C7-02AF85C822A8",
1224   "callbackURL": "https://callback.example.com/callback/Ujh4kkQ9A",
1225   ... additional optional content ...
1226 }

```

## 1227 **6.3.2 Performing a subscription call (informative)**

1228 A subscription call follows the following steps:

- 1229 1. The publisher SHALL perform a `POST` to the `callbackURL` of the subscriber.
  - 1230 • A `Subscription-ID` HTTP header containing the `subscriptionId` is added.
  - 1231 • A `Notification-Signature` HTTP header containing the computed
  - 1232 signature of the request body is added.
  - 1233 • The request-body is sent using the `application/json` format.
- 1234 2. The subscriber SHALL validate the `POST` request. It SHOULD be done in the
- 1235 following order. If any of the validation steps fail, the message SHALL be rejected.
  - 1236 • It is RECOMMENDED to start message parsing only if all of the validation
  - 1237 steps are performed without an error.
  - 1238 • The `Notification-Signature` HTTP header MUST be provided.
  - 1239 • The `Subscription-ID` HTTP header MUST be included. It MUST be a GUID.
  - 1240 • Additional provided custom information is RECOMMENDED to be
  - 1241 validated. (e.g. in the `callbackURL`)

- 1242                   • The subscriber uses the stored Shared Secret to compute the signature of the  
1243                   request body. The signature SHALL equal the provided signature.
- 1244                   • In case the callback was performed due to a a subscription of an event, the  
1245                   occurrence time of the event MUST be in the past. It MAY be a few seconds  
1246                   in the future to account for minor time synchronization issues.
- 1247           3. A successful callback is responded by the 204 No Content response code.

1248   Example for a subscription call using the secret from the example above

```
1249 POST https://callback.example.com/callback/Ujh4kkQ9A
1250 Subscription-ID: 936DA01F-9ABD-4D9D-80C7-02AF85C822A8
1251 Notification-Signature:
1252 sha256=66c2912069e6c9563d66fee4674cd23dd9dd00e6c08c985e964b11f92f477e48
1253 Content-Type: application/json
1254 Content-Length: ...
1255 {
1256   "id": "84db923d-2a19-4eb0-beb5-446c1ec57d34",
1257   "occurrenceDateTime": "2022-04-16T16:40:00+01:00",
1258   "typeCode": "ARRIVAL",
1259   "shipmentId": "123e4567-e89b-12d3-a456-426614174000"
1260 }
```

1261   Response

1262   HTTP/1.1 204 No Content

## 1263 **7 Appendix A: Examples**

1264 Printed JSON schema files of a realistic example can be very large, especially because of  
1265 the code lists used. Therefore, we have not included an example here.

1266 However, examples can be found on the web at the following address:

1267 <https://github.com/uncefact/spec-openAPI/examples>

## 8 Appendix B: Naming and Design Rules List (normative)

Rule #	Rule
[R 1 1]	<p>Conformance SHALL be determined through adherence to the content of the normative sections and rules. Furthermore, each rule is categorized to indicate the intended audience for the rule by the following:</p> <ol style="list-style-type: none"> <li>1. Rules, which must not be violated. Else, conformance and interoperability is lost.</li> <li>2. Rules, which may be modified, while still conformant to the NDR structure.</li> </ol> <p>Inf. Rules that are informative only. If a different implementation is chosen this does not have any impact on the conformance of the implementation towards this specification.</p>
[R 2 1]	All API specifications based on this OpenAPI Naming and Design Rules technical specification SHALL be compliant to the OpenAPI 3.1.x specification.
[R 3 1]	An API specification-claiming conformance to this specification SHALL define schema components as described in the JSON Schema Naming and Design Rules Technical Specification.
[R 4 1]	<p>Request body content and Response content used to transfer structured data information SHALL use the <b>application/json</b> media type for JavaScript Object Notation (JSON). This rule MAY only be deviated from, if the API implements a conversion service from or to JSON in another media type.</p> <p>Additional media types (e.g. <b>text/xml</b>) to transfer structured data information MAY be used. If non-structured information is transferred any valid media type MAY be used.</p>
[R 5 1]	Encoding SHALL be UTF-8.
[R 6 2]	The structure of the paths defined within APIs SHOULD be meaningful to the consumers. Paths SHOULD follow a predictable, hierarchical structure to enhance understandability and therefore usability.
[R 7 1]	<p>The API URLs SHOULD follow the standard naming convention as described below:</p> <pre>https://{env}.api.{dnsdomain}/v{m}/{service}/{resource}/{id}/{sub-resource}?{query}</pre> <p>The components are described as follows. If a rule is mandatory for a specific component of the URL is SHALL be applied to any conformant API specification, even if the basic URL structure is different from the one described above (e.g. if <b>api</b> is not used as a prefix to the <b>dnsdomain</b>).</p> <ul style="list-style-type: none"> <li>• <b>https://</b> SHALL be used as the web protocol.</li> <li>• <b>{env}</b> indicates the environment (e.g. <b>test</b>, <b>sandbox</b> or <b>dev</b>) and is usually omitted for production environment.</li> <li>• <b>{dnsdomain}</b> is the DNS domain of the API implementer (e.g. <b>unece.org</b>)</li> <li>• <b>{service}</b> is a logical grouping of API functions that represent a business service domain (e.g. <b>transport</b>). The <b>{service}</b> component is optional.</li> <li>• <b>v{m}</b> is the major version number of the API specification. This component SHALL be stated in the URL. It MAY be provided at a different place in the URL (e.g. as a prefix to the domain).</li> <li>• <b>{resource}</b> is the plural noun representing an API resource (e.g. <b>consignments</b>)</li> <li>• <b>{id}</b> is the unique identifier for the resource defined as a path parameter. Path parameters SHALL be used to identify a resource. This component is not part of the path if an operation is performed on a collection of the resource.</li> <li>• <b>{sub-resource}</b> is an optional sub-resource. Only used when there are contained collections or actions on a main resource (e.g. <b>consignmentItem</b>).</li> <li>• <b>{query}</b> is a list of additional parameters like filters that determine the results of a search (e.g. <b>consignments?loadingPort=AUSYD</b>).</li> </ul>

[R 8 1]	The total number of characters in the URL, including the path and the query, SHALL NOT exceed 2000 characters in length including any formatting codes such as commas, underscores, question marks, hyphens, plus or slashes.
[R 9 1]	Endpoints SHALL NOT be actions. Services and resources SHALL consist of nouns. HTTP verbs SHALL be used for actions.
[R 10 1]	Kebab-case SHALL be used in services.
[R 11 1]	Lower camelCase SHALL be used in resources, path parameters and query parameters.
[R 12 1]	Path parameters and query parameters with a relation to property names SHALL be consistent with property names.
[R 13 1]	Query parameters SHALL be URL safe.
[R 14 1]	Resource names SHALL be pluralised. Resource names SHOULD be consistent with schemas. If a schema is defined in singular, nevertheless the resource SHALL be pluralized. If the plural of a resource is non-standard, you MAY choose a more appropriate noun in its plural form.
[R 15 1]	Query parameters SHALL use ISO8601 compliant date and time representations that are defined in <b>UNTDID 2379 json</b> as defined in the JSON schema NDR technical specification. To represent a specific date, time or date-time the format SHALL comply with the JSON schema definition for date, time or date-time.
[R 16 1]	A prerequisite for an OpenAPI specification and its implementation to be fully compliant with this NDR TS is the use of UN/CEFACT semantics and UN/CEFACT syntax (e.g. UN/CEFACT XML, UN/CEFACT JSON Schema, and UN/CEFACT Vocabulary). An OpenAPI specification that does not use UN/CEFACT syntax or UN/CEFACT semantics may still be conformant to this NDR TS if it meets the criteria specified in [R 1 1].
[R 17 1]	Endpoints are RECOMMENDED to support CRUD operations. (Create, Read, Update, Delete). If an endpoint is not intended to support e.g. a delete operation, it SHALL return the HTTP response codes as defined in chapter 3.2.10.
[R 18 1]	APIs SHALL adhere to the idempotency of operations specified in Table 4.
[R 19 1]	APIs SHOULD implement the Idempotency-Key HTTP header field and the corresponding implementation advice in order to make non-idempotent operations like POST and PATCH fault-tolerant.
[R 20 1]	If pagination is used in an API, keyset-based pagination (cursor-pagination) SHALL be used. This means that the consumer cannot request a specific page, instead the consumer has to select a page-link provided by the server. The server SHALL provide links in the HTTP response header to the previous and next page and SHOULD provide links to the first and last page. More links MAY be provided. The cursor-value is a string, created by the server using whatever method it likes. It identifies a point in a list of results for a query containing filters and sorting parameters for a specific moment in time. Therefore, it divides the list into those that fall before the cursor and those that fall after the cursor. There may optionally be one result that falls "on" the cursor.
[R 21 1]	GET requests on collection results SHOULD implement pagination. The default and maximum page size SHOULD be 100, if not specified on the endpoint. If SHOULD be smaller, if the resulting page load is large. The default page size MAY be changed per endpoint. A consumer SHOULD be able to override the default page size. If the filter, sorting and/or page size used is changed when getting a result, the pagination SHALL BE reset to the first page. The query parameters described in the following table SHALL be used, rules SHALL be applied.
[R 22 1]	Sorting and filtering SHALL be done using query parameters. Using a path parameter is only allowed to identify a specific resource.
[R 23 1]	As a general guide, filtering SHOULD be done with case insensitivity. Whether you choose to filter with case insensitivity or not SHALL be clearly documented.

[R 24 1]	If an application needs to support a richer search and filter capability that includes logical operators, fuzzy search, grouping, and so on, API MAY apply a query string according to lucene query syntax . In that case, the filtering and query parameters normally are transmitted in the request body.
[R 25 1]	Sorting SHOULD be limited to specified fields. The sort direction MAY be omitted. The default sort direction is ascending. A colon <code>:</code> is used to separate the field name and the sort direction. Multiple sort fields are separated by comma <code>,</code> .
[R 26 1]	HTTP response codes SHALL be used. Table 6 defines HTTP response codes supported by conformant APIs. The column Response indicates whether an additional error response payload is RECOMMENDED to be returned as described in chapter 3.2.11.
[R 27 1]	Table 7 defines which HTTP response codes SHALL be supported for a specific HTTP request method by conformant APIs. Column Use indicates how a conformant API supports the specified http response code: <ul style="list-style-type: none"> <li>- <b>m</b> the code SHALL be supported</li> <li>- <b>ma</b> SHALL be supported for requests where the response is handled asynchronous, for instance due to forwarding or processing time. In that case, a <b>Location</b> HTTP response header SHALL be gives that points to the respective resource. In addition, a <b>Retry-After</b> HTTP response header is RECOMMENDED to be returned.</li> <li>- <b>r</b> the code is recommended to be supported.</li> </ul> The default response code for a positive response is marked in <b>bold</b> .
[R 28 1]	An API SHALL implement an error response schema to allow a standardised error handling. The response SHALL use the following JSON Schema. The JSON Schema MAY be extended.
[R 29 1]	The following rules are RECOMMENDED: <ul style="list-style-type: none"> <li>- The definitions in a conformant OpenAPI specification SHALL be considered as technical contracts between designers and developers and between consumers and providers.</li> <li>- Mock APIs SHOULD be created using the API description to allow early code integration for development.</li> <li>- The behaviour and intent of the API SHOULD be described with as much information as possible.</li> <li>- Operations SHOULD provide examples for request and response bodies.</li> <li>- Expected response codes and error messages SHOULD be provided in full.</li> <li>- Known issues or limitations SHOULD be clearly documented.</li> <li>- Expected performance, uptime and SLA/OLA SHOULD be clearly documented.</li> <li>- Although YAML is a supported file format of an OpenAPI specification, the JSON format SHOULD be used as the OpenAPI specification format.</li> </ul>
[R 30 1]	All APIs SHALL apply Semantic versioning 2.0.0 :  <b>MAJOR.MINOR.PATCH</b> The first version of an API SHALL start with a <b>MAJOR</b> version of 1. Pre-release version information and build metadata version information SHALL NOT be used in API versioning.
[R 31 1]	All APIs SHALL use URI versioning. They SHALL include the MAJOR version as part of the URI in the format of 'v{MAJOR}'. Example: <code>https://api.logistics.io/transport/v1/voyages</code> The minor and patch version SHALL NOT be used in the URI.
[R 32 1]	A custom header named API-Version SHALL be added to any response of the API. It SHALL be aligned with the URI version and SHALL state all three levels: <code>API-Version: 1.21.5</code>
[R 33 1]	An API-Version custom header MAY be added to a request. If added, it SHALL only contain the MAJOR version. <code>API-Version: 1</code>



[R 34 1]	An API SHALL implement a response to a GET request to the base URI of the API. The response SHALL use the JSON Schema defined in R 33.
[R 35 2]	APIs that are still in a <b>DRAFT</b> status SHOULD be placed in a sandbox environment. This could be done by changing the basis URL accordingly. Example for a productive base URL: <code>https://api.uncefact.unece.org/v1/</code> Examples for a development base URL: <code>https://sandbox.api.uncefact.unece.org/v1/</code> <code>https://staging.api.uncefact.unece.org/v1/</code>
[R 36 1]	Within a major release backward compatibility SHALL NOT be broken.
[R 37 1]	API clients and subscribers SHOULD be robust: <ul style="list-style-type: none"> <li>- Be conservative with API requests and data passed as input.</li> <li>- Be tolerant with unknown fields in the payload, but do not eliminate them from payload if needed for subsequent PUT requests.</li> </ul>
[R 38 1]	An API SHALL NOT be set to <b>DEPRECATED</b> until a replacement service is running with status <b>ACTIVE</b> . The root service of the API SHALL provide the <b>Deprecation Header Field</b> and the <b>Sunset HTTP Response Header Field</b> . A <b>Link header</b> SHALL be added in combination with the <b>Deprecation header</b> . It SHALL provide a link to the documentation. A second <b>Link header</b> SHALL be added linking to the replacement version of the API.
[R 39 1]	Deprecated endpoints SHALL be documented in the OpenAPI specification using the <b>DEPRECATED</b> property introduced since OpenAPI 3.0.0. Deprecated endpoints SHOULD provide the <b>Deprecation Header Field</b> and the <b>Sunset HTTP Response Header Field</b> . A <b>Link header</b> SHALL be added in combination with the <b>Deprecation header</b> . It SHALL provide a link to the documentation. Where possible, communication SHOULD be sent to consumers of deprecated endpoints.
[R 40 1]	The introduction of a major version SHOULD be avoided, whenever possible. This MAY be achieved as follows: <ul style="list-style-type: none"> <li>- Create a new service endpoint, if the process is changed.</li> <li>- Duplicate and Deprecate: add a <b>Deprecation Header</b> to the old service including a <b>Link Header</b> to documentation and to the new service. Eventually add a <b>Sunset Header</b>.</li> <li>- Create a new resource (a variant of the old) in addition to the old.</li> </ul>
[R 41 2]	New resources or service endpoints can be added during a minor release. In order to support the implementation of those new services a sandbox environment SHOULD be provided to the interested or affected consumers.
[R 42 1]	It is <b>RECOMMENDED</b> that no more than 3 parallel MAJOR versions are available. Implementers of the API SHALL NOT be more than 1 major version behind the latest version.
[R 43 1]	In order to provide a better experience for API consumers, APIs SHOULD provide a list of state transitions that are available for each resource. As possible values for link relation types the official IANA registry list SHALL be used. It MAY be extended. Any extension SHALL be documented in the API specification.
[R 44 1]	All API endpoints SHALL be secured. HTTPS SHALL be used. The OAUTH2 security scheme is <b>RECOMMENDED</b> . Other security schemes MAY be used. The receivers endpoints of subscription callbacks MAY be designed with different security measures like those described in chapter 6.3. The aspects described after rule 32 of API security are <b>RECOMMENDED</b> to be implemented.

1270 **9 Appendix C: Glossary**

Term	Definition
ABIE	Aggregate Business Information Entity – a term from CCTS that describes an information class such as “consignment”
API	Application Programming Interface – a term that references a machine-to-machine interface.
ASBIE	Association Business Information Entity – a term from CCTS that defines a directed relationship from source ABIE to target ABIE – e.g. “consignee” as a relationship between “consignment” and “party”
B2B	Business to Business
BBIE	Basic Business Information Entity – a term from CCTS that describes a property of a class such as party.name
BRS	Business requirement specification
CamelCase	CamelCase is a naming rule for a technical representation of identifiers consisting of several words. White spaces are removed and every new word begins with a capital letter. Example: <code>this identifier</code> is written as <code>thisIdentifier</code> in camelCase.
CCL	Core Component Library
CCTS	Core Component Technical Specification – a UN/CEFACT specification document that described the information management metamodel.
CDT	Core Data Type. A value domain for a BBIE that is a simple type such as “text” or “code”
HATEOS	Hypermedia as the Engine of Application State
IETF	Internet Engineering Task Force
JSON	JavaScript Object Notation – an IETF document syntax standard in common use by web developers for APIs.
JSON-LD	JSON-Linked Data – a JSON standard for linked data graphs / semantic vocabularies.
Kebab-case	Kebab-case is a naming rule for a technical representation of identifiers consisting of several words. Hyphens are used to connect words. Example: <code>this identifier</code> is written as <code>this-identifier</code> in kebab-case.
NDR	Naming & Design Rules – a set of rules for mapping one representation (e.g. RDM) to another (e.g. JSON-LD)
OpenAPI	An open source standard, language-agnostic interface to RESTful APIs.
OWL	Web Ontology Language
RDF	Resource Description Framework – a W3C semantic web standard
RDM	Reference Data Model- a UN/CEFACT semantic output.
RESTful API	See REST API
REST API	Representation State Transfer Application Programming Interface, a.k.a. RESTful API
RFC	Request for Comments
SDO	Standards Development Organisation
UN/CEFACT	United Nations Centre for Trade Facilitation and Electronic Business
UNECE	United Nations Economic Commission for Europe
URI	Uniform Resource Identifier – a namespace qualified string of characters that unambiguously identify a resource. AURL is one type of URI.

<b>Term</b>	<b>Definition</b>
URL	Uniform Resource Locator – the web address of a resource.
UNTDID	United Nations Trade Data Interchange Directory
XML	Extensible Markup Language
XMI	Xml Metadata Interchange - a well-established OMG standard for exchange of UML models between different tools.

1271

**Table 9 - Glossary**