

# Table of Contents

1. DNT RESTful API and Event Scheme Guidelines	1
Conventions Used in These Guidelines	1
Zalando specific information	1
2. Principles	1
API Design Principles	1
3. General Guidelines	2
<b>Must:</b> Provide API Reference Definition using OpenAPI	2
<b>Should:</b> Provide User Manual Documentation	3
<b>Must:</b> Write APIs in U.S. English	3
4. Security	3
<b>Must:</b> Secure Endpoints with DNT Connect	3
5. Compatibility	3
<b>Must:</b> Don't Break Backward Compatibility	3
<b>Should:</b> Prefer Compatible Extensions	4
<b>Must:</b> Prepare Clients To Not Crash On Compatible API Extensions	4
<b>Should:</b> Design APIs Conservatively	5
<b>Must:</b> Always Return JSON Objects As Top-Level Data Structures To Support Extensibility	6
<b>Must:</b> Treat Open API Definitions As Open For Extension By Default	6
<b>Should:</b> Used Open-Ended List of Values (x-extensible-enum) Instead of Enumerations	6
<b>Should:</b> Avoid Versioning	7
<b>Must:</b> Use Media Type Versioning	7
<b>Must:</b> Do Not Use URI Versioning	8
<b>Should:</b> Provide Version Information in OpenAPI Documentation	8
6. JSON Guidelines	9
<b>Must:</b> Use Consistent Property Names	9
<b>Must:</b> Property names must be snake_case (and never camelCase)	9
<b>Must:</b> Property names must match <code>^[a-z_][a-z_0-9]*\$</code>	9
<b>Should:</b> Array names should be pluralized	10
<b>Must:</b> Use Consistent Property Values	10
<b>Must:</b> Boolean property values must not be null	10
<b>Should:</b> Null values should have their fields removed	10
<b>Should:</b> Empty array values should not be null	10
<b>Should:</b> Enumerations should be represented as Strings	10
<b>Should:</b> Date property values should conform to RFC 3399	10
<b>May:</b> Time durations and intervals could conform to ISO 8601	11
<b>May:</b> Standards could be used for Language, Country and Currency	11
7. API Naming	11
<b>Must:</b> Use lowercase separate words with hyphens for Path Segments	11

<b>Must:</b> Use snake_case (never camelCase) for Query Parameters	12
<b>Must:</b> Use Hyphenated HTTP Headers	12
<b>Should:</b> Prefer Hyphenated-Pascal-Case for HTTP header Fields	12
<b>May:</b> Use Standardized Headers	12
<b>Must:</b> Pluralize Resource Names	12
<b>May:</b> Use /api as first Path Segment	12
<b>Must:</b> Avoid Trailing Slashes	13
<b>May:</b> Use Conventional Query Strings	13
8. Resources	13
<b>Must:</b> Avoid Actions — Think About Resources	13
<b>Should:</b> Model complete business processes	14
<b>Should:</b> Define <i>useful</i> resources	14
<b>Must:</b> Keep URLs Verb-Free	14
<b>Must:</b> Use Domain-Specific Resource Names	14
<b>Must:</b> Identify resources and Sub-Resources via Path Segments	14
<b>Should:</b> Only Use UUIDs If Necessary	15
<b>May:</b> Consider Using (Non-) Nested URLs	16
<b>Should:</b> Limit number of Resources	16
<b>Should:</b> Limit number of Sub-Resource Levels	16
9. HTTP	16
<b>Must:</b> Use HTTP Methods Correctly	16
<b>Must:</b> Fulfill Safeness and Idempotency Properties	20
<b>Must:</b> Use Specific HTTP Status Codes	20
<b>Must:</b> Provide Error Documentation	22
<b>Must:</b> Use 207 for Batch or Bulk Requests	23
<b>Must:</b> Use 429 with Headers for Rate Limits	24
<b>Should:</b> Explicitly define the Collection Format of Query Parameters	24
10. Performance	25
<b>Should:</b> Reduce Bandwidth Needs and Improve Responsiveness	25
<b>Should:</b> Use gzip Compression	25
<b>Should:</b> Support Filtering of Resource Fields	25
<b>Should:</b> Allow Optional Embedding of Sub-Resources	27
11. Pagination	28
<b>Must:</b> Support Pagination	28
<b>Should:</b> Prefer Cursor-Based Pagination, Avoid Offset-Based Pagination	28
<b>May:</b> Use Pagination Links Where Applicable	29
12. Hypermedia	30
<b>Must:</b> Use REST Maturity Level 2	30
<b>May:</b> Use REST Maturity Level 3 - HATEOAS	30
<b>Must:</b> Use Common Hypertext Controls	31
<b>Should:</b> Use Simple Hypertext Controls for Pagination and Self-References	32

<b>Must:</b> Not Use Link Headers with JSON entities .....	32
13. Data Formats .....	32
<b>Must:</b> Use JSON to Encode Structured Data .....	33
<b>May:</b> Use non JSON Media Types for Binary Data or Alternative Content Representations .....	33
<b>Must:</b> Use Standard Date and Time Formats .....	33
<b>May:</b> Use Standards for Country, Language and Currency Codes .....	33
<b>Must:</b> Define Format for Type Number and Integer .....	34
<b>Should:</b> Prefer standard Media type name application/json .....	34
14. Common Data Types .....	35
<b>Should:</b> Use a Common Money Object .....	35
<b>Must:</b> Use common field names and semantics .....	36
<b>Must:</b> Follow Hypertext Control Conventions .....	39
<b>Must:</b> Use Problem JSON .....	39
<b>Must:</b> Do not expose Stack Traces .....	40
15. Common Headers .....	40
<b>Must:</b> Use Content Headers Correctly .....	40
<b>May:</b> Use Content-Location Header .....	40
<b>Should:</b> Use Location Header instead of Content-Location Header .....	41
<b>May:</b> Use the Prefer header to indicate processing preferences .....	41
<b>May:</b> Consider using ETag together with If-(None-)Match header .....	42
16. Proprietary Headers .....	43
<b>Must:</b> Use Only the Specified Proprietary Zalando Headers .....	44
<b>Must:</b> Propagate Proprietary Headers .....	45
17. Deprecation .....	46
<b>Must:</b> Obtain Approval of Clients .....	46
<b>Must:</b> External Partners Must Agree on Deprecation Timespan .....	46
<b>Must:</b> Reflect Deprecation in API Definition .....	46
<b>Must:</b> Monitor Usage of Deprecated APIs .....	46
<b>Should:</b> Add a Warning Header to Responses .....	47
<b>Should:</b> Add Monitoring for Warning Header .....	47
<b>Must:</b> Not Start Using Deprecated APIs .....	47
18. API Operation .....	47
<b>Must:</b> Provide Online Access to OpenAPI Reference Definition .....	47
<b>Should:</b> Monitor API Usage .....	48
19. Events .....	48
<b>Must:</b> Treat Events as part of the service interface .....	49
<b>Must:</b> Make Events available for review .....	49
<b>Must:</b> Ensure Event Type schemas conform to Open API's Schema Object .....	49
<b>Must:</b> Ensure that Events are registered as Event Types .....	50
<b>Must:</b> Ensure Events conform to a well-known Event Category .....	53
<b>Must:</b> Ensure that Events define useful business resources .....	57

<b>Must:</b> Events must not provide sensitive customer personal data	57
<b>Must:</b> Use the General Event Category to signal steps and arrival points in business processes	58
<b>Must:</b> Use Data Change Events to signal mutations	58
<b>Should:</b> Provide a means for explicit event ordering	58
<b>Should:</b> Use the hash partition strategy for Data Change Events	59
<b>Should:</b> Ensure that Data Change Events match API representations	59
<b>Must:</b> Permissions on events must correspond to API permissions	60
<b>Must:</b> Indicate ownership of Event Types	60
<b>Must:</b> Define Event Payloads in accordance with the overall Guidelines	60
<b>Must:</b> Maintain backwards compatibility for Events	61
<b>Should:</b> Avoid additionalProperties in event type definitions	61
<b>Must:</b> Use unique Event identifiers	62
<b>Should:</b> Design for idempotent out-of-order processing	62
<b>Must:</b> Follow conventions for Event Type names	63
<b>Must:</b> Prepare for duplicate Events	63
Appendix A: References	63
OpenAPI Specification	63
Publications, specifications and standards	63
Dissertations	64
Books	64
Blogs	64
Appendix B: Tooling	64
API First Integrations	64
Support Libraries	65
Appendix C: Changelog	65
Rule Changes	65



# 1. DNT RESTful API and Event Scheme Guidelines

Den Norske Turistforening / The Norwegian Trekking Association / [dnt.no](http://dnt.no)

Other formats: [PDF](#), [EPUB3](#)

Based on the excellent work by [Zalando](#).

<div id="table-of-contents"></div>

## Conventions Used in These Guidelines

The requirement level keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" used in this document (case insensitive) are to be interpreted as described in [RFC 2119](#).

## Zalando specific information

The purpose of our “RESTful API guidelines” is to define standards to successfully establish “consistent API look and feel” quality.

These guidelines will, to some extent, remain work in progress as our work evolves.

In case guidelines are changing, following rules apply:

- existing APIs don't have to be changed, but we recommend it
- clients of existing APIs have to cope with these APIs based on outdated rules
- new APIs have to respect the current guidelines

Furthermore you should keep in mind that once an API becomes public externally available, it has to be re-reviewed and changed according to current guidelines - for sake of overall consistency.

## 2. Principles

### API Design Principles

Comparing SOA web service interfacing style of SOAP vs. REST, the former tend to be centered around operations that are usually use-case specific and specialized. In contrast, REST is centered around business (data) entities exposed as resources that are identified via URIs and can be manipulated via standardized CRUD-like methods using different representations, self-descriptive messages and hypermedia. RESTful APIs tend to be less use-case specific and comes with less rigid client / server coupling and are more suitable as a platform interface being open for diverse client applications.

- We prefer REST-based APIs with JSON payloads
- We prefer systems to be truly RESTful [1: Per definition of R.Fielding REST APIs have to support HATEOAS (maturity level 3). Our guidelines do not strongly advocate for full REST compliance, but limited hypermedia usage, e.g. for pagination (see [Hypermedia](#)). However, we still use the term "RESTful API", due to the absence of an alternative established term and to keep it like the very majority of web service industry that also use the term for their REST approximations — in fact, in today's industry full HATEOAS compliant APIs are a very rare exception.]

We apply the RESTful web service principles to all kind of application components, whether they provide functionality via the Internet or via the intranet as larger application elements. We strive to build interoperating distributed systems that different teams can evolve in parallel.

An important principle for (RESTful) API design and usage is Postel's Law, aka [the Robustness Principle](#) (RFC 1122):

- Be liberal in what you accept, be conservative in what you send

*Readings:* Read the following to gain additional insight on the RESTful service architecture paradigm and general RESTful API design style:

- Book: [Irresistable APIs: Designing web APIs that developers will love](#)
- Book: [REST in Practice: Hypermedia and Systems Architecture](#)
- Book: [Build APIs You Won't Hate](#)
- InfoQ eBook: [Web APIs: From Start to Finish](#)
- Lessons-learned blog: [Thoughts on RESTful API Design](#)
- Fielding Dissertation: [Architectural Styles and the Design of Network-Based Software Architectures](#)

## 3. General Guidelines

The titles are marked with the corresponding labels: **Must:**, **Should:**, **May:**.

### **Must: Provide API Reference Definition using OpenAPI**

We use the [OpenAPI specification](#) (aka Swagger spec) as standard for our REST API definitions. You **have to provide the API definition in YAML format** (instead of JSON) for the OpenAPI API definition files due to its improved readability.

Please stick to version 2.0 of the specification for now, until we succeed to update all our tooling to support the upcoming version 3.0.

We also call the OpenAPI API definition the "API Reference definition" (or "API definition"); as a reference manual it provides all information needed by an experienced API client developer to use this API.

The OpenAPI API specification file should be subject of version control together with source code

management. Services also have to support an [endpoint to access the API Reference definition](#) for their external API(s).

## Should: Provide User Manual Documentation

In addition to the API as OpenAPI Reference Definition, it's good practice to provide an API User Manual documentation to improve client developer experience, especially of engineers that are less experienced in using this API. A helpful API User Manual typically describes the following API aspects:

- API's scope, purpose and use cases
- concrete examples of API usage
- edge cases, error situation details and repair hints
- architecture context and major dependencies - including figures and sequence flows

The User Manual must be posted online, e.g. via GitHub Enterprise pages, on specific team web servers, or as a Google doc. And don't forget to include a link to this user manual into your OpenAPI definition using the "externalDocs" property.

## Must: Write APIs in U.S. English

### 4. Security

#### Must: Secure Endpoints with DNT Connect

If the API needs authentication, use DNT Connect in Sherpa as the OAuth2 provider. Sherpa handles OAuth2 clients, generates access tokens and the API service must verify the tokens towards Sherpa.

[OAuth2 flow] | [https://s3-eu-west-1.amazonaws.com/turistforeningen/API+security+flow+\(simple\).jpg](https://s3-eu-west-1.amazonaws.com/turistforeningen/API+security+flow+(simple).jpg)

### 5. Compatibility

#### Must: Don't Break Backward Compatibility

Change APIs, but keep all consumers running. Consumers usually have independent release lifecycles, focus on stability, and avoid changes that do not provide additional value. APIs are contracts between service providers and service consumers that cannot be broken via unilateral decisions.

There are two techniques to change APIs without breaking them:

- follow rules for compatible extensions
- introduce new API versions and still support older versions



We strongly encourage using compatible API extensions and discourage versioning. The below guideline rules for service providers and consumers enable us (having Postel's Law in mind) to make compatible changes without versioning.

**Hint:** Please note that the compatibility guarantees are for the "on the wire" format. Binary or source compatibility of code generated from an API definition is not covered by these rules. If client implementations update their generation process to a new version of the API definition, it has to be expected that code changes are necessary.

## Should: Prefer Compatible Extensions

API designers should apply the following rules to evolve RESTful APIs for services in a backward-compatible way:

- Add only optional, never mandatory fields.
- Never change the semantic of fields (e.g. changing the semantic from customer-number to customer-id, as both are different unique customer keys)
- Input fields may have (complex) constraints being validated via server-side business logic. Never change the validation logic to be more restrictive and make sure that all constraints are clearly defined in description.
- Enum ranges can be reduced when used as input parameters, only if the server is ready to accept and handle old range values too. Enum range can be reduced when used as output parameters.
- Enum ranges cannot not be extended when used for output parameters — clients may not be prepared to handle it. However, enum ranges can be extended when used for input parameters.
- Use `x-extensible-enum`, if range is used for output parameters and likely to be extended with growing functionality. It defines an open list of explicit values and clients must be agnostic to new values.
- Support redirection in case an URL has to change ([301 Moved Permanently](#)).

## Must: Prepare Clients To Not Crash On Compatible API Extensions

Service clients should apply the robustness principle:

- Be conservative with API requests and data passed as input, e.g. avoid to exploit definition deficits like passing megabytes for strings with unspecified maximum length.
- Be tolerant in processing and reading data of API responses, more specifically...

Service clients must be prepared for compatible API extensions of service providers:

- Be tolerant with unknown fields in the payload (see also Fowler's "[TolerantReader](#)" post), i.e. ignore new fields but do not eliminate them from payload if needed for subsequent PUT requests.
- Be prepared that `x-extensible-enum` return parameter may deliver new values; either be

agnostic or provide default behavior for unknown values.

- Be prepared to handle HTTP status codes not explicitly specified in endpoint definitions. Note also, that status codes are extensible. Default handling is how you would treat the corresponding x00 code (see [RFC7231 Section 6](#)).
- Follow the redirect when the server returns HTTP status [301 Moved Permanently](#).

## Should: Design APIs Conservatively

Designers of service provider APIs should be conservative and accurate in what they accept from clients:

- Unknown input fields in payload or URL should not be ignored; servers should provide error feedback to clients via an HTTP 400 response code.
- Be accurate in defining input data constraints (like formats, ranges, lengths etc.) — and check constraints and return dedicated error information in case of violations.
- Prefer being more specific and restrictive (if compliant to functional requirements), e.g. by defining length range of strings. It may simplify implementation while providing freedom for further evolution as compatible extensions.

Not ignoring unknown input fields is a specific deviation from Postel's Law (e.g. see also [The Robustness Principle Reconsidered](#)) and a strong recommendation. Servers might want to take different approach but should be aware of the following problems and be explicit in what is supported:

- Ignoring unknown input fields is actually not an option for PUT, since it becomes asymmetric with subsequent GET response and HTTP is clear about the PUT "replace" semantics and default roundtrip expectations (see [RFC7231 Section 4.3.4](#)). Note, accepting (i.e. not ignoring) unknown input fields and returning it in subsequent GET responses is a different situation and compliant to PUT semantics.
- Certain client errors cannot be recognized by servers, e.g. attribute name typing errors will be ignored without server error feedback. The server cannot differentiate between the client intentionally providing an additional field versus the client sending a mistakenly named field, when the client's actual intent was to provide an optional input field.
- Future extensions of the input data structure might be in conflict with already ignored fields and, hence, will not be compatible, i.e. break clients that already use this field but with different type.

In specific situations, where a (known) input field is not needed anymore, it either can stay in the API definition with "not used anymore" description or can be removed from the API definition as long as the server ignores this specific parameter.

## Must: Always Return JSON Objects As Top-Level Data Structures To Support Extensibility

In a response body, you must always return a JSON object (and not e.g. an array) as a top level data structure to support future extensibility. JSON objects support compatible extension by additional attributes. This allows you to easily extend your response and e.g. add pagination later, without breaking backwards compatibility.

## Must: Treat Open API Definitions As Open For Extension By Default

The Open API 2.0 specification is not very specific on default extensibility of objects, and redefines JSON-Schema keywords related to extensibility, like `additionalProperties`. Following our overall compatibility guidelines, Open API object definitions are considered open for extension by default as per [Section 5.18 "additionalProperties"](#) of JSON-Schema.

When it comes to Open API 2.0, this means an `additionalProperties` declaration is not required to make an object definition extensible:

- API clients consuming data must not assume that objects are closed for extension in the absence of an `additionalProperties` declaration and must ignore fields sent by the server they cannot process. This allows API servers to evolve their data formats.
- For API servers receiving unexpected data, the situation is slightly different. Instead of ignoring fields, servers *may* reject requests whose entities contain undefined fields in order to signal to clients that those fields would not be stored on behalf of the client. API designers must document clearly how unexpected fields are handled for PUT, POST and PATCH requests.

API formats must not declare `additionalProperties` to be false, as this prevents objects being extended in the future.

Note that this guideline concentrates on default extensibility and does not exclude the use of `additionalProperties` with a schema as a value, which might be appropriate in some circumstances.

## Should: Used Open-Ended List of Values (x-extensible-enum) Instead of Enumerations

Enumerations are per definition closed sets of values, that are assumed to be complete and not intended for extension. This closed principle of enumerations imposes compatibility issues when an enumeration must be extended. To avoid these issues, we strongly recommend to use an open-ended list of values instead of an enumeration unless:

1. the API has full control of the enumeration values, i.e. the list of values does not depend on any external tool or interface, and
2. the list of value is complete with respect to any thinkable and unthinkable future feature.

To specify an open-ended list of values use the marker `x-extensible-enum` as follows:

```
deliver_methods:
  type: string
  x-extensible-enum:
    - parcel
    - letter
    - email
```

**Note:** `x-extensible-enum` is not JSON Schema conform but will be ignored by most tools.

## Should: Avoid Versioning

When changing your RESTful APIs, do so in a compatible way and avoid generating additional API versions. Multiple versions can significantly complicate understanding, testing, maintaining, evolving, operating and releasing our systems ([supplementary reading](#)).

If changing an API can't be done in a compatible way, then proceed in one of these three ways:

- create a new resource (variant) in addition to the old resource variant
- create a new service endpoint — i.e. a new application with a new API (with a new domain name)
- create a new API version supported in parallel with the old API by the same microservice

As we discourage versioning by all means because of the manifold disadvantages, we strongly recommend to only use the first two approaches.

## Must: Use Media Type Versioning

However, when API versioning is unavoidable, you have to design your multi-version RESTful APIs using media type versioning (instead of URI versioning, see below). Media type versioning is less tightly coupled since it supports content negotiation and hence reduces complexity of release management.

Media type versioning: Here, version information and media type are provided together via the HTTP Content-Type header — e.g. `application/x.zalando.cart+json;version=2`. For incompatible changes, a new media type version for the resource is created. To generate the new representation version, consumer and producer can do content negotiation using the HTTP Content-Type and Accept headers. Note: This versioning only applies to the request and response content schema, not to URI or method semantics.

In this example, a client wants only the new version of the response:

```
Accept: application/x.zalando.cart+json;version=2
```

A server responding to this, as well as a client sending a request with content should use the Content-Type header, declaring that one is sending the new version:

```
Content-Type: application/x.zalando.cart+json;version=2
```

Using header versioning should:

- include versions in request and response headers to increase visibility
- include Content-Type in the Vary header to enable proxy caches to differ between versions

Hint: Until an incompatible change is necessary, it is recommended to stay with the standard `application/json` media type.

Hint: [OpenAPI currently doesn't support content negotiation](#), though [a comment in this issue](#) mentions a workaround (using a fragment identifier that gets stripped off). Another way would be to document just the new version, but let the server accept the old one (with the previous content-type).

Further reading: [API Versioning Has No "Right Way"](#) provides an overview on different versioning approaches to handle breaking changes without being opinionated.

## Must: Do Not Use URI Versioning

With URI versioning a (major) version number is included in the path, e.g. `/v1/customers`. The consumer has to wait until the provider has been released and deployed. If the consumer also supports hypermedia links — even in their APIs — to drive workflows (HATEOAS), this quickly becomes complex. So does coordinating version upgrades — especially with hyperlinked service dependencies — when using URL versioning. To avoid this tighter coupling and complexer release management we do not use URI versioning, and go instead with media type versioning and content negotiation (see above).

## Should: Provide Version Information in OpenAPI Documentation

Only the documentation, not the API itself, needs version information.

Example:

```
"swagger": "2.0",
"info": {
  "title": "Parcel service API",
  "description": "API for <...>",
  "version": "1.0.0",
  <...>
}
```

During a (possibly) long-running API review phase you need different versions of the API description. These versions may include changes that are incompatible with earlier draft versions. So we apply the following version schema MAJOR.MINOR.DRAFT that increments the...

- MAJOR version, when you make incompatible API changes
- MINOR version, when you add functionality in a backwards-compatible manner
- DRAFT version, when you make changes during the review phase that are not related to production releases

We recommend using the DRAFT version only for unreleased API definitions that are still under review; for example:

```
version 1.4.0 -- current version
version 1.4.1 -- first draft and call for review of API extensions compatible with
1.4.0
version 1.4.2 -- second draft and call for review of API extensions that are still
compatible with 1.4.0 but possibly incompatible with 1.4.1
version 1.5.0 -- approved version for implementation and release
version 1.5.1 -- first draft for next review and API change cycle; compatible with
1.4.0 and 1.5.0
```

Hint: This versioning scheme differs in the less strict DRAFT aspect from [semantic version information](#) used for released APIs and service applications.

## 6. JSON Guidelines

These guidelines provides recommendations for defining JSON data at Zalando. JSON here refers to [RFC 7159](#) (which updates [RFC 4627](#)), the “application/json” media type and custom JSON media types defined for APIs. The guidelines clarifies some specific cases to allow Zalando JSON data to have an idiomatic form across teams and services.

### Must: Use Consistent Property Names

#### Must: Property names must be snake\_case (and never camelCase)

No established industry standard exists, but many popular Internet companies prefer snake\_case: e.g. GitHub, Stack Exchange, Twitter. Others, like Google and Amazon, use both - but not only camelCase. It’s essential to establish a consistent look and feel such that JSON looks as if it came from the same hand.

#### Must: Property names must match $^[a-z_][a-z_0-9]*\$$

Property names are restricted to ASCII strings. The first character must be a letter, or an underscore, and subsequent characters can be a letter, an underscore, or a number.

(It is recommended to use `_` at the start of property names only for keywords like `_links`.)

## Should: Array names should be pluralized

To indicate they contain multiple values prefer to pluralize array names. This implies that object names should in turn be singular.

## Must: Use Consistent Property Values

### Must: Boolean property values must not be null

Schema based JSON properties that are by design booleans must not be presented as nulls. A boolean is essentially a closed enumeration of two values, true and false. If the content has a meaningful null value, strongly prefer to replace the boolean with enumeration of named values or statuses - for example `accepted_terms_and_conditions` with true or false can be replaced with `terms_and_conditions` with values yes, no and unknown.

### Should: Null values should have their fields removed

Swagger/OpenAPI, which is in common use, doesn't support null field values (it does allow omitting that field completely if it is not marked as required). However that doesn't prevent clients and servers sending and receiving those fields with null values. Also, in some cases null may be a meaningful value - for example, JSON Merge Patch [RFC 7382](#)) using null to indicate property deletion.

### Should: Empty array values should not be null

Empty array values can unambiguously be represented as the empty list, `[]`.

### Should: Enumerations should be represented as Strings

Strings are a reasonable target for values that are by design enumerations.

### Should: Date property values should conform to RFC 3399

Use the date and time formats defined by [RFC 3339](#):

- for "date" use strings matching `date-fullyear "-" date-month "-" date-mday`, for example: `2015-05-28`
- for "date-time" use strings matching `full-date "T" full-time`, for example `2015-05-28T14:07:17Z`

Note that the [OpenAPI format](#) "date-time" corresponds to "date-time" in the RFC) and `2015-05-28` for a date (note that the OpenAPI format "date" corresponds to "full-date" in the RFC). Both are specific profiles, a subset of the international standard [ISO 8601](#).

A zone offset may be used (both, in request and responses)—this is simply defined by the standards. However, we encourage restricting dates to UTC and without offsets. For example `2015-05-28T14:07:17Z` rather than `2015-05-28T14:07:17+00:00`. From experience we have learned that zone offsets are not easy to understand and often not correctly handled. Note also that zone offsets are different from local times that might be including daylight saving time. Localization of dates should be done by the services that provide user interfaces, if required.

When it comes to storage, all dates should be consistently stored in UTC without a zone offset. Localization should be done locally by the services that provide user interfaces, if required.

Sometimes it can seem data is naturally represented using numerical timestamps, but this can introduce interpretation issues with precision - for example whether to represent a timestamp as 1460062925, 1460062925000 or 1460062925.000. Date strings, though more verbose and requiring more effort to parse, avoid this ambiguity.

## May: Time durations and intervals could conform to ISO 8601

Schema based JSON properties that are by design durations and intervals could be strings formatted as recommended by ISO 8601 ([Appendix A of RFC 3399 contains a grammar for durations](#)).

## May: Standards could be used for Language, Country and Currency

- [ISO 3166-1-alpha2 country](#)
- (It's "GB", not "UK", even though "UK" has seen some use at Zalando)
- [ISO 639-1 language code](#)
- [BCP-47](#) (based on ISO 639-1) for language variants
- [ISO 4217 currency codes](#)

## 7. API Naming

### Must: Use lowercase separate words with hyphens for Path Segments

Example:

```
/shipment-orders/{shipment-order-id}
```

This applies to concrete path segments and not the names of path parameters. For example `{shipment_order_id}` would be ok as a path parameter.



## Must: Use snake\_case (never camelCase) for Query Parameters

Examples:

```
customer_number, order_id, billing_address
```

## Must: Use Hyphenated HTTP Headers

### Should: Prefer Hyphenated-Pascal-Case for HTTP header Fields

This is for consistency in your documentation (most other headers follow this convention). Avoid camelCase (without hyphens). Exceptions are common abbreviations like “ID.”

Examples:

```
Accept-Encoding  
Apply-To-Redirect-Ref  
Disposition-Notification-Options  
Original-Message-ID
```

See also: [HTTP Headers are case-insensitive \(RFC 7230\)](#).

## May: Use Standardized Headers

Use [this list](#) and mention its support in your OpenAPI definition.

## Must: Pluralize Resource Names

Usually, a collection of resource instances is provided (at least API should be ready here). The special case of a resource singleton is a collection with cardinality 1.

## May: Use /api as first Path Segment

In most cases, all resources provided by a service are part of the public API, and therefore should be made available under the root “/” base path. If the service should also support non-public, internal APIs — for specific operational support functions, for example — add “/api” as base path to clearly separate public and non-public API resources.

## Must: Avoid Trailing Slashes

The trailing slash must not have specific semantics. Resource paths must deliver the same results whether they have the trailing slash or not.

## May: Use Conventional Query Strings

If you provide query support for sorting, pagination, filtering functions or other actions, use the following standardized naming conventions:

- **q** — default query parameter (e.g. used by browser tab completion); should have an entity specific alias, like `sku`
- **limit** — to restrict the number of entries. See [Pagination](#) section below. Hint: You can use `size` as an alternate query string.
- **cursor** — key-based page start. See [Pagination](#) section below.
- **offset** — numeric offset page start. See [Pagination](#) section below. Hint: In combination with `limit`, you can use `page` as an alternative to `offset`.
- **sort** — comma-separated list of fields to sort. To indicate sorting direction, fields may be prefixed with `+` (ascending) or `-` (descending, default), e.g. `/sales-orders?sort=+id`
- **fields** — to retrieve a subset of fields. See [Support Filtering of Resource Fields](#) below.
- **embed** — to expand embedded entities (ie.: inside of an article entity, expand silhouette code into the silhouette object). Implementing “expand” correctly is difficult, so do it with care.

## 8. Resources

### Must: Avoid Actions — Think About Resources

REST is all about your resources, so consider the domain entities that take part in web service interaction, and aim to model your API around these using the standard HTTP methods as operation indicators. For instance, if an application has to lock articles explicitly so that only one user may edit them, create an article lock with PUT or POST instead of using a lock action.

Request:

```
PUT /article-locks/{article-id}
```

The added benefit is that you already have a service for browsing and filtering article locks.

## Should: Model complete business processes

An API should contain the complete business processes containing all resources representing the process. This enables clients to understand the business process, foster a consistent design of the business process, allow for synergies from description and implementation perspective, and eliminates implicit invisible dependencies between APIs.

In addition, it prevents services from being designed as thin wrappers around databases, which normally tends to shift business logic to the clients.

## Should: Define *useful* resources

As a rule of thumb resources should be defined to cover 90% of all its client's use cases. A *useful* resource should contain as much information as necessary, but as little as possible. A great way to support the last 10% is to allow clients to specify their needs for more/less information by supporting filtering and [embedding](#).

## Must: Keep URLs Verb-Free

The API describes resources, so the only place where actions should appear is in the HTTP methods. In URLs, use only nouns. Instead of thinking of actions (verbs), it's often helpful to think about putting a message in a letter box: e.g., instead of having the verb *cancel* in the url, think of sending a message to cancel an order to the *cancellations* letter box on the server side.

## Must: Use Domain-Specific Resource Names

API resources represent elements of the application's domain model. Using domain-specific nomenclature for resource names helps developers to understand the functionality and basic semantics of your resources. It also reduces the need for further documentation outside the API definition. For example, "sales-order-items" is superior to "order-items" in that it clearly indicates which business object it represents. Along these lines, "items" is too general.

## Must: Identify resources and Sub-Resources via Path Segments

Some API resources may contain or reference sub-resources. Embedded sub-resources, which are not top-level resources, are parts of a higher-level resource and cannot be used outside of its scope. Sub-resources should be referenced by their name and identifier in the path segments.

Composite identifiers must not contain / as a separator. In order to improve the consumer experience, you should aim for intuitively understandable URLs, where each sub-path is a valid reference to a resource or a set of resources. For example, if `/customers/12ev123bv12v/addresses/DE_100100101` is a valid path of your API, then `/customers/12ev123bv12v/addresses`, `/customers/12ev123bv12v` and `/customers` must be valid as well in principle.

Basic URL structure:

```
{resources}/{resource-id}/{sub-resources}/{sub-resource-id}
{resources}/{partial-id-1}[separator][partial-id-2]
```

Examples:

```
/carts/1681e6b88ec1/items
/carts/1681e6b88ec1/items/1
/customers/12ev123bv12v/addresses/DE_100100101
```

## Should: Only Use UUIDs If Necessary

Generating IDs can be a scaling problem in high frequency and near real time use cases. UUIDs solve this problem, as they can be generated without collisions in a distributed, non-coordinated way and without additional server roundtrips.

However, they also come with some disadvantages:

- pure technical key without meaning; not ready for naming or name scope conventions that might be helpful for pragmatic reasons, e.g. we learned to use names for product attributes, instead of UUIDs
- less usable, because...
- cannot be memorized and easily communicated by humans
- harder to use in debugging and logging analysis
- less convenient for consumer facing usage
- quite long: readable representation requires 36 characters and comes with higher memory and bandwidth consumption
- not ordered along their creation history and no indication of used id volume
- may be in conflict with additional backward compatibility support of legacy ids

UUIDs should be avoided were not needed for large scale id generation. Instead, for instance, server side support with id generation can be preferred (POST on id resource, followed by idempotent PUT on entity resource). Usage of UUIDs is especially discouraged as primary keys of master and configuration data, like brand-ids or attribute-ids which have low id volume but widespread steering functionality.

In any case, we should always use string rather than number type for identifiers. This gives us more flexibility to evolve the identifier naming scheme. Accordingly, if used as identifiers, UUIDs should not be qualified using a format property.

Hint: Usually, random UUID is used - see UUID version 4 in [RFC 4122](#). Though UUID version 1 also contains leading timestamps it is not reflected by its lexicographic sorting. This deficit is addressed by [ULID](#) (Universally Unique Lexicographically Sortable Identifier). You may favour ULID instead of

UUID, for instance, for pagination use cases ordered along creation time.

## May: Consider Using (Non-) Nested URLs

If a sub-resource is only accessible via its parent resource and may not exist without parent resource, consider using a nested URL structure, for instance:

```
/carts/1681e6b88ec1/cart-items/1
```

However, if the resource can be accessed directly via its unique id, then the API should expose it as a top-level resource. For example, customer has a collection for sales orders; however, sales orders have globally unique id and some services may choose to access the orders directly, for instance:

```
/customers/1681e6b88ec1  
/sales-orders/5273gh3k525a
```

## Should: Limit number of Resources

To keep maintenance and service evolution manageable, we should follow "functional segmentation" and "separation of concern" design principles and do not mix different business functionalities in same API definition. In this sense the number of resources exposed via API should be limited - our experience is that a typical range of resources for a well-designed API is between 4 and 8. There may be exceptions with more complex business domains that require more resources, but you should first check if you can split them into separate subdomains with distinct APIs.

Nevertheless one API should hold all necessary resources to model complete business processes helping clients to understand these flows.

## Should: Limit number of Sub-Resource Levels

There are main resources (with root url paths) and sub-resources (or "nested" resources with non-root url paths). Use sub-resources if their life cycle is (loosely) coupled to the main resource, i.e. the main resource works as collection resource of the subresource entities. You should use  $\leq 3$  sub-resource (nesting) levels— more levels increase API complexity and url path length. (Remember, some popular web browsers do not support URLs of more than 2000 characters)

# 9. HTTP

## Must: Use HTTP Methods Correctly

Be compliant with the standardized HTTP method semantics summarized as follows:

## GET

GET requests are used to read a single resource or query set of resources.

- GET requests for individual resources will usually generate a 404 if the resource does not exist
- GET requests for collection resources may return either 200 (if the listing is empty) or 404 (if the list is missing)
- GET requests must NOT have request body payload

**Note:** GET requests on collection resources should provide a sufficient filter mechanism as well as [Pagination](#).

## "GET with Body"

APIs sometimes face the problem, that they have to provide extensive structured request information with GET, that may even conflicts with the size limits of clients, load-balancers, and servers. As we require APIs to be standard conform (body in GET must be ignored on server side), API designers have to check the following two options:

1. GET with URL encoded query parameters: when it is possible to encode the request information in query parameters, respecting the usual size limits of clients, gateways, and servers, this should be the first choice. The request information can either be provided distributed to multiple query parameters or a single structured URL encoded string.
2. POST with body content: when a GET with URL encoded query parameters is not possible, a POST with body content must be used. In this case the endpoint must be documented with the hint **GET with body** to transport the GET semantic of this call.

**Note:** It is no option to encode the lengthy structured request information in header parameters. From a conceptual point of view, the semantic of an operation should always be expressed by resource name and query parameters, i.e. what goes into the URL. Request headers are reserved for general context information, e.g. FlowIDs. In addition, size limits on query parameters and headers are not reliable and depend on clients, gateways, server, and actual settings. Thus, switching to headers does not solve the original problem.

## PUT

PUT requests are used to create or update **entire** resources - single or collection resources. The semantic is best described as *»please put the enclosed representation at the resource mentioned by the URL, replacing any existing resource.«*.

- PUT requests are usually applied to single resources, and not to collection resources, as this would imply replacing the entire collection
- PUT requests are usually robust against non-existence of resources by implicitly creating before updating
- on successful PUT requests, the server will **replace the entire resource** addressed by the URL with the representation passed in the payload (subsequent reads will deliver the same payload)
- successful PUT requests will usually generate 200 or 204 (if the resource was updated - with or

without actual content returned), and 201 (if the resource was created)

**Note:** Resource IDs with respect to PUT requests are maintained by the client and passed as a URL path segment. Putting the same resource twice is required to be idempotent and to result in the same single resource instance. If PUT is applied for creating a resource, only URIs should be allowed as resource IDs. If URIs are not available POST should be preferred.

To prevent unnoticed concurrent updates when using PUT, the combination of **ETag** and **If-(None-)Match** headers should be considered to signal the server stricter demands to expose conflicts and prevent lost updates.

## POST

POST requests are idiomatically used to create single resources on a collection resource endpoint, but other semantics on single resources endpoint are equally possible. The semantic for collection endpoints is best described as *»please add the enclosed representation to the collection resource identified by the URL«*. The semantic for single resource endpoints is best described as *»please execute the given well specified request on the collection resource identified by the URL«*.

- POST request should only be applied to collection resources, and normally not on single resource, as this has an undefined semantic
- on successful POST requests, the server will create one or multiple new resources and provide their URI/URLs in the response
- successful POST requests will usually generate 200 (if resources have been updated), 201 (if resources have been created), and 202 (if the request was accepted but has not been finished yet)

**More generally:** POST should be used for scenarios that cannot be covered by the other methods sufficiently. For instance, GET with complex (e.g. SQL like structured) query that needs to be passed as request body payload because of the URL-length constraint. In such cases, make sure to document the fact that POST is used as a workaround.

**Note:** Resource IDs with respect to POST requests are created and maintained by server and returned with response payload. Posting the same resource twice is by itself **not** required to be idempotent and may result in multiple resource instances. Anyhow, if external URIs are present that can be used to identify duplicate requests, it is best practice to implement POST in an idempotent way.

## PATCH

PATCH request are only used for partial update of single resources, i.e. where only a specific subset of resource fields should be replaced. The semantic is best described as *»please change the resource identified by the URL according to my change request«*. The semantic of the change request is not defined in the HTTP standard and must be described in the API specification by using suitable media types.

- PATCH requests are usually applied to single resources, and not on collection resources, as this would imply patching on the entire collection

- PATCH requests are usually not robust against non-existence of resource instances
- on successful PATCH requests, the server will update parts of the resource addressed by the URL as defined by the change request in the payload
- successful PATCH requests will usually generate 200 or 204 (if resources have been updated)
- with or without updated content returned)

**Note:** since implementing PATCH correctly is a bit tricky, we strongly suggest to choose one and only one of the following patterns per endpoint, unless forced by a [backwards compatible change](#). In preference order:

1. use PUT with complete objects to update a resource as long as feasible (i.e. do not use PATCH at all).
2. use PATCH with partial objects to only update parts of a resource, when ever possible. (This is basically [JSON Merge Patch](#), a specialized media type `application/merge-patch+json` that is a partial resource representation.)
3. use PATCH with [JSON Patch](#), a specialized media type `application/json-patch+json` that includes instructions on how to change the resource.
4. use POST (with a proper description of what is happening) instead of PATCH if the request does not modify the resource in a way defined by the semantics of the media type.

In practice [JSON Merge Patch](#) quickly turns out to be too limited, especially when trying to update single objects in large collections (as part of the resource). In this cases [JSON Patch](#) can shown its full power while still showing readable patch requests (see also [JSON patch vs. merge](#)).

To prevent unnoticed concurrent updates when using PATCH, the combination of `ETag`and` `If-Match` headers should be considered to signal the server stricter demands to expose conflicts and prevent lost updates.

## DELETE

DELETE request are used to delete resources. The semantic is best described as *»please delete the resource identified by the URL«*.

- DELETE requests are usually applied to single resources, not on collection resources, as this would imply deleting the entire collection
- successful DELETE request will usually generate 200 (if the deleted resource is returned) or 204 (if no content is returned)
- failed DELETE request will usually generate 404 (if the resource cannot be found) or 410 (if the resource was already deleted before)

## HEAD

HEAD requests are used retrieve to header information of single resources and resource collections.

- HEAD has exactly the same semantics as GET, but returns headers only, no body.



## OPTIONS

OPTIONS are used to inspect the available operations (HTTP methods) of a given endpoint.

- OPTIONS requests usually either return a comma separated list of methods (provided by an **Allow**-Header) or as a structured list of link templates

**Note:** OPTIONS is rarely implemented, though it could be used to self-describe the full functionality of a resource.

## Must: Fulfill Safeness and Idempotency Properties

An operation can be...

- idempotent, i.e. operation will produce the same results if executed once or multiple times (note: this does not necessarily mean returning the same status code)
- safe, i.e. must not have side effects such as state changes

Method implementations must fulfill the following basic properties:

HTTP method	safe	idempotent
OPTIONS	Yes	Yes
HEAD	Yes	Yes
GET	Yes	Yes
PUT	No	Yes
POST	No	No
DELETE	No	Yes
PATCH	No	No

## Must: Use Specific HTTP Status Codes

This guideline groups the following rules for HTTP status codes usage:

- You must not invent new HTTP status codes; only use standardized HTTP status codes and consistent with its intended semantics.
- You should use the most specific HTTP status code for your concrete resource request processing status or error situation.
- When using HTTP status codes that are less commonly used and not listed below, you must provide good documentation in the API definition.

There are 60 different HTTP status codes with specific semantics defined in the HTTP standards (mainly [RFC7231](#) and [RFC-6585](#)) - and there are upcoming new ones, e.g. [draft legally-restricted-status](#) (see overview on all error codes on [Wikipedia](#) or via <https://httpstatuses.com/>). And there are unofficial ones, e.g. used by specific web servers like Nginx.

List of most commonly used and best understood HTTP status codes:

## Success Codes

Code	Meaning	Methods
200	OK - this is the standard success response	All
201	Created - Returned on successful entity creation. You are free to return either an empty response or the created resource in conjunction with the Location header. (More details found in the <a href="#">Common Headers</a> .) <i>Always</i> set the Location header.	POST, PUT
202	Accepted - The request was successful and will be processed asynchronously.	POST, PUT, DELETE, PATCH
204	No content - There is no response body	PUT, DELETE
207	Multi-Status - The response body contains multiple status informations for different parts of a batch/bulk request. See <a href="#">Must: Provide Error Documentation</a> .	POST

## Redirection Codes

Code	Meaning	Methods
301	Moved Permanently - This and all future requests should be directed to the given URI.	All
303	See Other - The response to the request can be found under another URI using a GET method.	PATCH, POST, PUT, DELETE
304	Not Modified - resource has not been modified since the date or version passed via request headers If-Modified-Since or If-None-Match.	GET

## Client Side Error Codes

Code	Meaning	Methods
400	Bad request - generic / unknown error	All
401	Unauthorized - the users must log in (this often means “Unauthenticated”)	All
403	Forbidden - the user is not authorized to use this resource	All
404	Not found - the resource is not found	All
405	Method Not Allowed - the method is not supported, see OPTIONS	All
406	Not Acceptable - resource can only generate content not acceptable according to the Accept headers sent in the request	All
408	Request timeout - the server times out waiting for the resource	All
409	Conflict - request cannot be completed due to conflict, e.g. when two clients try to create the same resource or if there are concurrent, conflicting updates	PUT, DELETE, PATCH
410	Gone - resource does not exist any longer, e.g. when accessing a resource that has intentionally been deleted	All

Code	Meaning	Methods
412	Precondition Failed - returned for conditional requests, e.g. If-Match if the condition failed. Used for optimistic locking.	PUT, DELETE, PATCH
415	Unsupported Media Type - e.g. clients sends request body without content type	PUT, DELETE, PATCH
423	Locked - Pessimistic locking, e.g. processing states	PUT, DELETE, PATCH
428	Precondition Required - server requires the request to be conditional (e.g. to make sure that the “lost update problem” is avoided).	All
429	Too many requests - the client does not consider rate limiting and sent too many requests. See <a href="#">Must: Use 207 for Batch or Bulk Requests</a> .	All

### Server Side Error Codes:

Code	Meaning	Methods
500	Internal Server Error - a generic error indication for an unexpected server execution problem (here, client retry may be senseful)	All
501	Not Implemented - server cannot fulfill the request (usually implies future availability, e.g. new feature).	All
503	Service Unavailable - server is (temporarily) not available (e.g. due to overload) — client retry may be senseful.	All

## Must: Provide Error Documentation

APIs should define the functional, business view and abstract from implementation aspects. Errors become a key element providing context and visibility into how to use an API. The error object should be extended by an application-specific error identifier if and only if the HTTP status code is not specific enough to convey the domain-specific error semantic. For this reason, we use a standardized error return object definition — see [Must: Follow Hypertext Control Conventions](#).

The OpenAPI specification shall include definitions for error descriptions that will be returned; they are part of the interface definition and provide important information for service clients to handle exceptional situations and support troubleshooting.

Service providers should differentiate between technical and functional errors. In most cases it's not useful to document technical errors that are not in control of the service provider unless the status code convey application-specific semantics. The list of status code that can be omitted from API specifications includes but is not limited to:

- **401 Unauthorized**
- **403 Forbidden**
- **404 Not Found** unless it has some additional semantics
- **405 Method Not Allowed**

- 406 Not Acceptable
- 408 Request Timeout
- 413 Payload Too Large
- 414 URI Too Long
- 415 Unsupported Media Type
- 500 Internal Server Error
- 502 Bad Gateway
- 503 Service Unavailable
- 504 Gateway Timeout

Even though they might not be documented - they may very much occur in production, so clients should be prepared for unexpected response codes, and in case of doubt handle them like they would handle the corresponding x00 code. Adding new response codes (specially error responses) should be considered a compatible API evolution.

Functional errors on the other hand, that convey domain-specific semantics, must be documented and are strongly encouraged to be expressed with [Must: Follow Hypertext Control Conventions](#).

## Must: Use 207 for Batch or Bulk Requests

Some APIs are required to provide either *batch* or *bulk* requests using POST for performance reasons, i.e. for communication and processing efficiency. In this case services may be in need to signal multiple response codes for each part of an batch or bulk request. As HTTP does not provide proper guidance for handling batch/bulk requests and responses, we herewith define the following approach:

- A batch or bulk request **always** has to respond with HTTP status code **207**, unless it encounters a generic or unexpected failure before looking at individual parts.
- A batch or bulk response with status code 207 **always** returns a multi-status object containing sufficient status and/or monitoring information for each part of the batch or bulk request.
- A batch or bulk request may result in a status code 400/500, only if the service encounters a failure before looking at individual parts or, if an unanticipated failure occurs.

The before rules apply *even in the case* that processing of all individual part *fail* or each part is executed *asynchronously*! They are intended to allow clients to act on batch and bulk responses by inspecting the individual results in a consistent way.

**Note:** while a *batch* defines a collection of requests triggering independent processes, a *bulk* defines a collection of independent resources created or updated together in one request. With respect to response processing this distinction normally does not matter.

# Must: Use 429 with Headers for Rate Limits

APIs that wish to manage the request rate of clients must use the '429 Too Many Requests' response code if the client exceeded the request rate and therefore the request can't be fulfilled. Such responses must also contain header information providing further details to the client. There are two approaches a service can take for header information:

- Return a 'Retry-After' header indicating how long the client ought to wait before making a follow-up request. The Retry-After header can contain a HTTP date value to retry after or the number of seconds to delay. Either is acceptable but APIs should prefer to use a delay in seconds.
- Return a trio of 'X-RateLimit' headers. These headers (described below) allow a server to express a service level in the form of a number of allowing requests within a given window of time and when the window is reset.

The 'X-RateLimit' headers are:

- **X-RateLimit-Limit**: The maximum number of requests that the client is allowed to make in this window.
- **X-RateLimit-Remaining**: The number of requests allowed in the current window.
- **X-RateLimit-Reset**: The relative time in seconds when the rate limit window will be reset.

The reason to allow both approaches is that APIs can have different needs. Retry-After is often sufficient for general load handling and request throttling scenarios and notably, does not strictly require the concept of a calling entity such as a tenant or named account. In turn this allows resource owners to minimise the amount of state they have to carry with respect to client requests. The 'X-RateLimit' headers are suitable for scenarios where clients are associated with pre-existing account or tenancy structures. 'X-RateLimit' headers are generally returned on every request and not just on a 429, which implies the service implementing the API is carrying sufficient state to track the number of requests made within a given window for each named entity.

# Should: Explicitly define the Collection Format of Query Parameters

There are different ways of supplying a set of values as a query parameter. One particular type should be selected and stated explicitly in the API definition. The OpenAPI property `collectionFormat` is used to specify the format of the query parameter.

Only the `csv` or `multi` formats should be used for multi-value query parameters as described below.

Collection Format	Description	Example
<code>csv</code>	Comma separated values	<code>?parameter=value1,value2,value3</code>
<code>multi</code>	Multiple parameter instances	<code>?parameter=value1&amp;parameter=value2&amp;parameter=value3</code>

When choosing the collection format, take into account the tool support, the escaping of special characters and the maximal URL length.

## 10. Performance

### Should: Reduce Bandwidth Needs and Improve Responsiveness

APIs should support techniques for reducing bandwidth based on client needs. This holds for APIs that (might) have high payloads and/or are used in high-traffic scenarios like the public Internet and telecommunication networks. Typical examples are APIs used by mobile web app clients with (often) less bandwidth connectivity. (Zalando is a 'Mobile First' company, so be mindful of this point.)

Common techniques include:

- [gzip compression](#)
- querying field filters to retrieve a subset of resource attributes (see [Should: Support Filtering of Resource Fields](#) below)
- paginate lists of data items (see [Pagination](#) below)
- [ETag](#) and [If-\(None-\)Match](#) headers to avoid re-fetching of unchanged resources (see [May: Use the Prefer header to indicate processing preferences](#))
- pagination for incremental access of larger (result) lists

Each of these items is described in greater detail below.

### Should: Use gzip Compression

Compress the payload of your API's responses with gzip, unless there's a good reason not to — for example, you are serving so many requests that the time to compress becomes a bottleneck. This helps to transport data faster over the network (fewer bytes) and makes frontends respond faster.

Though gzip compression might be the default choice for server payload, the server should also support payload without compression and its client control via [Accept-Encoding](#) request header — see also [RFC 7231 Section 5.3.4](#). The server should indicate used gzip compression via the [Content-Encoding](#) header.

### Should: Support Filtering of Resource Fields

Depending on your use case and payload size, you can significantly reduce network bandwidth need by supporting filtering of returned entity fields. Here, the client can determine the subset of fields he wants to receive via the [fields query parameter](#) — example see [Google AppEngine API's partial response](#):

## Unfiltered

```
GET http://api.example.org/resources/123 HTTP/1.1
```

```
HTTP/1.1 200 OK
```

```
Content-Type: application/json
```

```
{
  "id": "cddd5e44-dae0-11e5-8c01-63ed66ab2da5",
  "name": "John Doe",
  "address": "1600 Pennsylvania Avenue Northwest, Washington, DC, United States",
  "birthday": "1984-09-13",
  "partner": {
    "id": "1fb43648-dae1-11e5-aa01-1fbc3abb1cd0",
    "name": "Jane Doe",
    "address": "1600 Pennsylvania Avenue Northwest, Washington, DC, United States",
    "birthday": "1988-04-07"
  }
}
```

## Filtered

```
GET http://api.example.org/resources/123?fields=(name,partner(name)) HTTP/1.1
```

```
HTTP/1.1 200 OK
```

```
Content-Type: application/json
```

```
{
  "name": "John Doe",
  "partner": {
    "name": "Jane Doe"
  }
}
```

As illustrated by this example, field filtering should be done via request parameter "fields" with value range defined by the following [BNF](#) grammar.

```

<fields> ::= <negation> <fields_expression> | <fields_expression>

<negation> ::= "!"

<fields_expression> ::= "(" <field_set> ")"

<field_set> ::= <qualified_field> | <qualified_field> "," <field_set>

<qualified_field> ::= <field> | <field> <fields_expression>

<field> ::= <DASH_LETTER_DIGIT> | <DASH_LETTER_DIGIT> <field>

<DASH_LETTER_DIGIT> ::= <DASH> | <LETTER> | <DIGIT>

<DASH> ::= "-" | "_"

<LETTER> ::= "A" | "B" | "C" | "D" | "E" | "F" | "G" | "H" | "I" | "J" | "K" | "L" |
"M" | "N" | "O" | "P" | "Q" | "R" | "S" | "T" | "U" | "V" | "W" | "X" | "Y" | "Z" |
"a" | "b" | "c" | "d" | "e" | "f" | "g" | "h" | "i" | "j" | "k" | "l" | "m" | "n" |
"o" | "p" | "q" | "r" | "s" | "t" | "u" | "v" | "w" | "x" | "y" | "z"

<DIGIT> ::= "0" | "1" | "2" | "3" | "4" | "5" | "6" | "7" | "8" | "9"

```

A `fields_expression` as defined by the grammar describes the properties of an object, i.e. `(name)` returns only the `name` property of the root object. `(name,partner(name))` returns the `name` and `partner` properties where `partner` itself is also an object and only its `name` property is returned.

Hint: OpenAPI doesn't allow you to formally specify whether depending on a given parameter will return different parts of the specified result schema. Explain this in English in the parameter description.

## Should: Allow Optional Embedding of Sub-Resources

Embedding related resources (also know as *Resource expansion*) is a great way to reduce the number of requests. In cases where clients know upfront that they need some related resources they can instruct the server to prefetch that data eagerly. Whether this is optimized on the server, e.g. a database join, or done in a generic way, e.g. an HTTP proxy that transparently embeds resources, is up to the implementation.

See [Must: Avoid Trailing Slashes](#) for naming, e.g. "embed" for steering of embedded resource expansion. Please use the [BNF](#) grammar, as already defined above for filtering, when it comes to an embedding query syntax.

Embedding a sub-resource can possibly look like this where an order resource has its order items as sub-resource (`/order/{orderId}/items`):



```
GET /order/123?embed=(items) HTTP/1.1
```

```
{
  "id": "123",
  "_embedded": {
    "items": [
      {
        "position": 1,
        "sku": "1234-ABCD-7890",
        "price": {
          "amount": 71.99,
          "currency": "EUR"
        }
      }
    ]
  }
}
```

## 11. Pagination

### Must: Support Pagination

Access to lists of data items must support pagination for best client side batch processing and iteration experience. This holds true for all lists that are (potentially) larger than just a few hundred entries.

There are two page iteration techniques:

- [Offset/Limit-based pagination](#): numeric offset identifies the first page entry
- [Cursor-based](#) — aka key-based — pagination: a unique key element identifies the first page entry (see also [Facebook's guide](#))

The technical conception of pagination should also consider user experience related issues. As mentioned in this [article](#), jumping to a specific page is far less used than navigation via next/previous page links. This favours cursor-based over offset-based pagination.

### Should: Prefer Cursor-Based Pagination, Avoid Offset-Based Pagination

Cursor-based pagination is usually better and more efficient when compared to offset-based pagination. Especially when it comes to high-data volumes and / or storage in NoSQL databases.

Before choosing cursor-based pagination, consider the following trade-offs:

- Usability/framework support:
  - Offset / limit based pagination is more known than cursor-based pagination, so it has more

framework support and is easier to use for API clients

- Use case: Jump to a certain page
  - If jumping to a particular page in a range (e.g., 51 of 100) is really a required use case, cursor-based navigation is not feasible
- Variability of data may lead to anomalies in result pages
  - Offset-based pagination may create duplicates or lead to missing entries if rows are inserted or deleted between two subsequent paging requests.
  - When using cursor-based pagination, paging cannot continue when the cursor entry has been deleted while fetching two pages
- Performance considerations - efficient server-side processing using offset-based pagination is hardly feasible for:
  - Higher data list volumes, especially if they do not reside in the database's main memory
  - Sharded or NoSQL databases
- Cursor-based navigation may not work if you need the total count of results and / or backward iteration support

Further reading:

- [Twitter](#)
- [Use the Index, Luke](#)
- [Paging in PostgreSQL](#)

## May: Use Pagination Links Where Applicable

- API implementing [HATEOS](#) may use [simplified hypertext controls](#) for pagination within collections.

Those collections should then have an `items` attribute holding the items of the current page. The collection may contain additional metadata about the collection or the current page (e.g. `index`, `page_size`) when necessary.

You should avoid providing a total count in your API unless there's a clear need to do so. Very often, there are systems and performance implications to supporting full counts, especially as datasets grow and requests become complex queries or filters that drive full scans (e.g., your database might need to look at all candidate items to count them). While this is an implementation detail relative to the API, it's important to consider your ability to support serving counts over the life of a service.

If the collection consists of links to other resources, the collection name should use [IANA registered link relations](#) as names whenever appropriate, but use plural form.

E.g. a service for articles could represent the collection of hyperlinks to an article's `authors` like that:

```
{
  "self": "https://.../articles/xyz/authors/",
  "index": 0,
  "page_size": 5,
  "items": [
    {
      "href": "https://...",
      "id": "123e4567-e89b-12d3-a456-426655440000",
      "name": "Kent Beck"
    },
    {
      "href": "https://...",
      "id": "987e2343-e89b-12d3-a456-426655440000",
      "name": "Mike Beedle"
    },
    ...
  ],
  "first": "https://...",
  "next": "https://...",
  "prev": "https://...",
  "last": "https://..."
}
```

## 12. Hypermedia

### Must: Use REST Maturity Level 2

We strive for a good implementation of [REST Maturity Level 2](#) as it enables us to build resource-oriented APIs that make full use of HTTP verbs and status codes. You can see this expressed by many rules throughout these guidelines, e.g.:

- **May:** Use Conventional Query Strings
- **Should:** Define *useful* resources
- **Should:** Limit number of Sub-Resource Levels
- **Must:** Fulfill Safeness and Idempotency Properties

Although this is not HATEOAS, it should not prevent you from designing proper link relationships in your APIs as stated in rules below.

### May: Use REST Maturity Level 3 - HATEOAS

We do not generally recommend to implement [REST Maturity Level 3](#). HATEOAS comes with additional API complexity without real value in our SOA context where client and server interact via REST APIs and provide complex business functions as part of our e-commerce SaaS platform.

Our major concerns regarding the promised advantages of HATEOAS (see also [RESTistential Crisis](#)

over [Hypermedia APIs](#), [Why I Hate HATEOAS](#) and others): \* We follow the [API First principle](#) with APIs explicitly defined outside the code with standard specification language. HATEOAS does not really add value for SOA client engineers in terms of API self-descriptiveness: a client engineer finds necessary links and usage description (depending on resource state) in the API reference definition anyway. \* Generic HATEOAS clients which need no prior knowledge about APIs and explore API capabilities based on hypermedia information provided, is a theoretical concept that we haven't seen working in practise and does not fit to our SOA set-up. The OpenAPI description format (and tooling based on OpenAPI) doesn't provide sufficient support for HATEOAS either. \* In practice relevant HATEOAS approximations (e.g. following specifications like HAL or JSON API) support API navigation by abstracting from URL endpoint and HTTP method aspects via link types. So, Hypermedia does not prevent clients from required manual changes when domain model changes over time. \* Hypermedia make sense for humans, less for SOA machine clients. We would expect use cases where it may provide value more likely in the frontend and human facing service domain. \* Hypermedia does not prevent API clients to implement shortcuts and directly target resources without 'discovering' them.

However, we do not forbid HATEOAS; you could use it, if you checked its limitations and still see clear value for your usage scenario that justifies its additional complexity. If you use HATEOAS please share experience and present your findings in the [API Guild \[internal link\]](#).

## Must: Use Common Hypertext Controls

When embedding links to other resources into representations you must use the common hypertext control object. It contains at least one attribute:

- **href**: The URI of the resource the hypertext control is linking to. All our API are using HTTP(s) as URI scheme.

In API that contain any hypertext controls, the attribute name **href** is reserved for usage within hypertext controls.

The schema for hypertext controls can be derived from this model:

```
HttpLink:
  description: A base type of objects representing links to resources.
  type: object
  properties:
    href:
      description: Any URI that is using http or https protocol
      type: string
      format: uri
  required: [ "href" ]
```

The name of an attribute holding such a **HttpLink** object specifies the relation between the object that contains the link and the linked resource. Implementations should use names from the [IANA Link Relation Registry](#) whenever appropriate. As IANA link relation names use hyphen-case notation, while this guide enforces snake\_case notation for attribute names, hyphens in IANA names have to be replaced with underscores (e.g. the IANA link relation type **version-history** would

become the attribute `version_history`)

Specific link objects may extend the basic link type with additional attributes, to give additional information related to the linked resource or the relationship between the source resource and the linked one.

E.g. a service providing "Person" resources could model a person who is married with some other person with a hypertext control that contains attributes which describe the other person (`id`, `name`) but also the relationship "spouse" between between the two persons (`since`):

```
{
  "id": "446f9876-e89b-12d3-a456-426655440000",
  "name": "Peter Mustermann",
  "spouse": {
    "href": "https://...",
    "since": "1996-12-19",
    "id": "123e4567-e89b-12d3-a456-426655440000",
    "name": "Linda Mustermann"
  }
}
```

Hypertext controls are allowed anywhere within a JSON model. While this specification would allow [HAL](#), we actually don't recommend/enforce the usage of HAL anymore as the structural separation of meta-data and data creates more harm than value to the understandability and usability of an API.

## Should: Use Simple Hypertext Controls for Pagination and Self-References

Hypertext controls for pagination inside collections and self-references should use a simple URI value in combination with their corresponding [link relations](#) (`next`, `prev`, `first`, `last`, `self`) instead of the extensible common hypertext control

See [Pagination](#) for information how to best represent paginateable collections.

## Must: Not Use Link Headers with JSON entities

We don't allow the use of the [Link Header defined by RFC 5988](#) in conjunction with JSON media types. We prefer links directly embedded in JSON payloads to the uncommon link header syntax.

# 13. Data Formats

## Must: Use JSON to Encode Structured Data

Use JSON-encoded body payload for transferring structured data. The JSON payload must follow [RFC-7159](#) by having (if possible) a serialized object as the top-level structure, since it would allow for future extension. This also applies for collection resources where one naturally would assume an array. See [Should: Prefer Cursor-Based Pagination, Avoid Offset-Based Pagination](#) for an example.

## May: Use non JSON Media Types for Binary Data or Alternative Content Representations

Other media types may be used in following cases:

- Transferring binary data or data whose structure is not relevant. This is the case if payload structure is not interpreted and consumed by clients as is. Example of such use case is downloading images in formats JPG, PNG, GIF.
- In addition to JSON version alternative data representations (e.g. in formats PDF, DOC, XML) may be made available through content negotiation.

## Must: Use Standard Date and Time Formats

### JSON Payload

Read more about date and time format in [Should: Enumerations should be represented as Strings](#).

### HTTP headers

Http headers including the proprietary headers use the [HTTP date format defined in RFC 7231](#).

## May: Use Standards for Country, Language and Currency Codes

Use the following standard formats for country, language and currency codes:

- [ISO 3166-1-alpha2 country codes](#)
  - (It is “GB”, not “UK”, even though “UK” has seen some use at Zalando)
- [ISO 639-1 language code](#)
  - [BCP-47](#) (based on ISO 639-1) for language variants
- [ISO 4217 currency codes](#)

## Must: Define Format for Type Number and Integer

Whenever an API defines a property of type `number` or `integer`, the precision must be defined by the format as follows to prevent clients from guessing the precision incorrectly, and thereby changing the value unintentionally:

type	format	specified value range
integer	int32	integer between $-2^{31}$ and $2^{31}-1$
integer	int64	integer between $-2^{63}$ and $2^{63}-1$
integer	bigint	arbitrarily large signed integer number
number	float	IEEE 754-2008/ISO 60559:2011 binary64 decimal number
number	double	IEEE 754-2008/ISO 60559:2011 binary128 decimal number
number	decimal	arbitrarily precise signed decimal number

The precision must be translated by clients and servers into the most specific language types. E.g. for the following definitions the most specific language types in Java will translate to `BigDecimal` for `Money.amount` and `int` or `Integer` for the `OrderList.page_size`:

```
Money:
  type: object
  properties:
    amount:
      type: number
      description: Amount expressed as a decimal number of major currency units
      format: decimal
      example: 99.95
    ...

OrderList:
  type: object
  properties:
    page_size:
      type: integer
      description: Number of orders in list
      format: int32
      example: 42
```

## Should: Prefer standard Media type name `application/json`

Previously, this guideline allowed the use of custom media types like `application/x.zalando.article+json`. This usage is not recommended anymore and should be avoided, except where it is necessary for cases of [media type versioning](#). Instead, the standard media type name `application/json` (or [Must: Follow Hypertext Control Conventions](#)).

Custom media types with subtypes beginning with `x` bring no advantage compared to the standard

media type for JSON, and make automated processing more difficult. They are also [discouraged by RFC 6838](#).

## 14. Common Data Types

Definitions of data objects that are good candidates for wider usage:

### Should: Use a Common Money Object

Use the following common money structure:

```
Money:
  type: object
  properties:
    amount:
      type: number
      description: Amount expressed as a decimal number of major currency units
      format: decimal
      example: 99.95
    currency:
      type: string
      description: 3 letter currency code as defined by ISO-4217
      format: iso-4217
      example: EUR
  required:
    - amount
    - currency
```

The decimal values for "amount" describe unit and subunit of the currency in a single value, where the digits before the decimal point are for the major unit and the digits after the decimal point are for the minor unit. Note that some business cases (e.g. transactions in Bitcoin) call for a higher precision, so applications must be prepared to accept values with unlimited precision, unless explicitly stated otherwise in the API specification. Examples for correct representations (in EUR):

- `42.20` or `42.2` = 42 Euros, 20 Cent
- `0.23` = 23 Cent
- `42.0` or `42` = 42 Euros
- `1024.42` = 1024 Euros, 42 Cent
- `1024.4225` = 1024 Euros, 42.25 Cent

Make sure that you don't convert the "amount" field to `float` / `double` types when implementing this interface in a specific language or when doing calculations. Otherwise, you might lose precision. Instead, use exact formats like Java's `BigDecimal`. See [Stack Overflow](#) for more info.

Some JSON parsers (NodeJS's, for example) convert numbers to floats by default. After discussing the [pros and cons \(internal link\)](#), we've decided on "decimal" as our amount format. It is not a



standard OpenAPI format, but should help us to avoid parsing numbers as float / doubles.

## Must: Use common field names and semantics

There exist a variety of field types that are required in multiple places. To achieve consistency across all API implementations, you must use common field names and semantics whenever applicable.

### Generic Fields

There are some data fields that come up again and again in API data:

- **id**: the identity of the object. If used, IDs must opaque strings and not numbers. IDs are unique within some documented context, are stable and don't change for a given object once assigned, and are never recycled cross entities.
- **xyz\_id**: an attribute within one object holding the identifier of another object must use a name that corresponds to the type of the referenced object or the relationship to the referenced object followed by **\_id** (e.g. **customer\_id** not **customer\_number**; **parent\_node\_id** for the reference to a parent node from a child node, even if both have the type **Node**)
- **created**: when the object was created. If used, this must be a **date-time** construct.
- **modified**: when the object was updated. If used, this must be a **date-time** construct.
- **type**: the kind of thing this object is. If used, the type of this field should be a string. Types allow runtime information on the entity provided that otherwise requires examining the Open API file.

Example JSON schema:

```

tree_node:
  type: object
  properties:
    id:
      description: the identifier of this node
      type: string
    created:
      description: when got this node created
      type: string
      format: 'date-time'
    modified:
      description: when got this node last updated
      type: string
      format: 'date-time'
    type:
      type: string
      enum: [ 'LEAF', 'NODE' ]
    parent_node_id:
      description: the identifier of the parent node of this node
      type: string
  example:
    id: '123435'
    created: '2017-04-12T23:20:50.52Z'
    modified: '2017-04-12T23:20:50.52Z'
    type: 'LEAF'
    parent_node_id: '534321'

```

These properties are not always strictly necessary, but making them idiomatic allows API client developers to build up a common understanding of Zalando's resources. There is very little utility for API consumers in having different names or value types for these fields across APIs.

## Address Fields

Address structures play a role in different functional and use-case contexts, including country variances. All attributes that relate to address information should follow the naming and semantics defined below.

```

addressee:
  description: a (natural or legal) person that gets addressed
  type: object
  required:
    - first_name
    - last_name
    - street
    - city
    - zip
    - country_code
  properties:
    salutation:

```

```
description: |
  a salutation and/or title used for personal contacts to some
  addressee; not to be confused with the gender information!
type: string
example: Mr
first_name:
description: |
  given name(s) or first name(s) of a person; may also include the
  middle names.
type: string
example: Hans Dieter
last_name:
description: |
  family name(s) or surname(s) of a person
type: string
example: Mustermann
business_name:
description: |
  company name of the business organization. Used when a business is
  the actual addressee; for personal shipments to office addresses, use
  `care_of` instead.
type: string
example: Consulting Services GmbH
required:
- first_name
- last_name

address:
description:
  an address of a location/destination
type: object
properties:
  care_of:
description: |
  (aka c/o) the person that resides at the address, if different from
  addressee. E.g. used when sending a personal parcel to the
  office /someone else's home where the addressee resides temporarily
type: string
example: Consulting Services GmbH
  street:
description: |
  the full street address including house number and street name
type: string
example: Schönhauser Allee 103
  additional:
description: |
  further details like building name, suite, apartment number, etc.
type: string
example: 2. Hinterhof rechts
  city:
description: |
```

```
    name of the city / locality
  type: string
  example: Berlin
zip:
  description: |
    zip code or postal code
  type: string
  example: 14265
country_code:
  description: |
    the country code according to
    [iso-3166-1-alpha-2](https://en.wikipedia.org/wiki/ISO_3166-1_alpha-2)
  type: string
  example: DE
required:
- street
- city
- zip
- country_code
```

Grouping and cardinality of fields in specific data types may vary based on the specific use case (e.g. combining addressee and address fields into a single type when modeling an address label vs distinct addressee and address types when modeling users and their addresses).

## Must: Follow Hypertext Control Conventions

APIs that provide hypertext controls (links) to interconnect API resources must follow the conventions for naming and modeling of hypertext controls as defined in section [Hypermedia](#).

## Must: Use Problem JSON

[RFC 7807](#) defines the media type `application/problem+json`. Operations should return that (together with a suitable status code) when any problem occurred during processing and you can give more details than the status code itself can supply, whether it be caused by the client or the server (i.e. both for 4xx or 5xx errors).

A previous version of this guideline (before the publication of that RFC and the registration of the media type) told to return `application/x.problem+json` in these cases (with the same contents). Servers for APIs defined before this change should pay attention to the `Accept` header sent by the client and set the `Content-Type` header of the problem response correspondingly. Clients of such APIs should accept both media types.

APIs may define custom problems types with extension properties, according to their specific needs.

The Open API schema definition can be found [on github](#). You can reference it by using:

```
responses:
  503:
    description: Service Unavailable
    schema:
      $ref: 'https://zalando.github.io/problem/schema.yaml#/Problem'
```

## Must: Do not expose Stack Traces

Stack traces contain implementation details that are not part of an API, and on which clients should never rely. Moreover, stack traces can leak sensitive information that partners and third parties are not allowed to receive and may disclose insights about vulnerabilities to attackers.

## 15. Common Headers

This section describes a handful of headers, which we found raised the most questions in our daily usage, or which are useful in particular circumstances but not widely known.

### Must: Use Content Headers Correctly

Content or entity headers are headers with a **Content-** prefix. They describe the content of the body of the message and they can be used in both, HTTP requests and responses. Commonly used content headers include but are not limited to:

- **Content-Disposition** can indicate that the representation is supposed to be saved as a file, and the proposed file name.
- **Content-Encoding** indicates compression or encryption algorithms applied to the content.
- **Content-Length** indicates the length of the content (in bytes).
- **Content-Language** indicates that the body is meant for people literate in some human language(s).
- **Content-Location** indicates where the body can be found otherwise ([Must: Use Content Headers Correctly](#) for more details).
- **Content-Range** is used in responses to range requests to indicate which part of the requested resource representation is delivered with the body.
- **Content-Type** indicates the media type of the body content.

### May: Use Content-Location Header

The Content-Location header is *optional* and can be used in successful write operations (PUT, POST or PATCH) or read operations (GET, HEAD) to guide caching and signal a receiver the actual location of the resource transmitted in the response body. This allows clients to identify the resource and to update their local copy when receiving a response with this header.

The Content-Location header can be used to support the following use cases:

- For reading operations GET and HEAD, a different location than the requested URI can be used to indicate that the returned resource is subject to content negotiations, and that the value provides a more specific identifier of the resource.
- For writing operations PUT and PATCH, an identical location to the requested URI, can be used to explicitly indicate that the returned resource is the current representation of the newly created or updated resource.
- For writing operations POST and DELETE, a content location can be used to indicate that the body contains a status report resource in response to the requested action, which is available at provided location.

**Note:** When using the Content-Location header, the Content-Type header has to be set as well. For example:

```
GET /products/123/images HTTP/1.1

HTTP/1.1 200 OK
Content-Type: image/png
Content-Location: /products/123/images?format=raw
```

## Should: Use Location Header instead of Content-Location Header

As the correct usage of Content-Location with respect to semantics and caching is difficult, we *discourage* the use of Content-Location. In most cases it is sufficient to direct clients to the resource location by using the Location header instead without hitting the Content-Location specific ambiguities and complexities.

More details in RFC 7231 [7.1.2 Location](#), [3.1.4.2 Content-Location](#)

## May: Use the Prefer header to indicate processing preferences

The **Prefer** header defined in [RFC7240](#) allows clients to request processing behaviors from servers. [RFC7240](#) pre-defines a number of preferences and is extensible, to allow others to be defined. Support for the Prefer header is entirely optional and at the discretion of API designers, but as an existing Internet Standard, is recommended over defining proprietary "X-" headers for processing directives.

The **Prefer** header can be defined like this in an API definition:

#### Prefer:

name: Prefer

description: |

The RFC7240 Prefer header indicates that particular server behaviors are preferred by the client but are not required for successful completion of the request.

# (indicate the preferences supported by the API)

in: header

type: string

required: false

Supporting APIs may return the **Preference-Applied** header also defined in [RFC7240](#) to indicate whether the preference was applied.

## May: Consider using ETag together with If-(None-)Match header

When creating or updating resources it may be necessary to expose conflicts and to prevent the lost update problem. This can be best accomplished by using the **ETag** header together with the **If-Match** and **If-None-Match**. The contents of an **ETag: <entity-tag>** header is either (a) a hash of the response body, (b) a hash of the last modified field of the entity, or (c) a version number or identifier of the entity version.

To expose conflicts between concurrent update operations via PUT, POST, or PATCH, the **If-Match: <entity-tag>** header can be used to force the server to check whether the version of the updated entity is conforming to the requested **<entity-tag>**. If no matching entity is found, the operation is supposed a to respond with status code 412 - precondition failed.

Beside other use cases, the **If-None-Match:** header with parameter **\*** can be used in a similar way to expose conflicts in resource creation. If any matching entity is found, the operation is supposed a to respond with status code 412 - precondition failed.

The **ETag**, **If-Match**, and **If-None-Match** headers can be defined as follows in the API definition:

### Etag:

name: Etag

description: |

The RFC7232 ETag header field in a response provides the current entity-tag for the selected resource. An entity-tag is an opaque identifier for different versions of a resource over time, regardless whether multiple versions are valid at the same time. An entity-tag consists of an opaque quoted string, possibly prefixed by a weakness indicator.

in: header

type: string

required: false

example: W/"xy", "5", "7da7a728-f910-11e6-942a-68f728c1ba70"

### IfMatch:

name: If-Match

description: |

The RFC7232 If-Match header field in a request requires the server to only operate on the resource that matches at least one of the provided entity-tags. This allows clients express a precondition that prevent the method from being applied, if there have been any changes to the resource.

in: header

type: string

required: false

example: "5", "7da7a728-f910-11e6-942a-68f728c1ba70"

### IfNoneMatch:

name: If-None-Match

description: |

The RFC7232 If-None-Match header field in a request requires the server to only operate on the resource if it does not match any of the provided entity-tags. If the provided entity-tag is `\*`, it is required that the resource does not exist at all.

in: header

type: string

required: false

example: "7da7a728-f910-11e6-942a-68f728c1ba70", \*

## 16. Proprietary Headers

This section shares definitions of proprietary headers that should be named consistently because they address overarching service-related concerns. Whether services support these concerns or not is optional; therefore, the OpenAPI API specification is the right place to make this explicitly visible. Use the parameter definitions of the resource HTTP methods.



# Must: Use Only the Specified Proprietary Zalando Headers

As a general rule, proprietary HTTP headers should be avoided. Still they can be useful in cases where context needs to be passed through multiple services in an end-to-end fashion. As such, a valid use-case for a proprietary header is providing context information, which is not a part of the actual API, but is needed by subsequent communication.

From a conceptual point of view, the semantics and intent of an operation should always be expressed by URLs path and query parameters, the method, and the content. Headers are more often used to implement functions close to the protocol considerations, such as flow control, content negotiation, and authentication. Thus, headers are reserved for general context information ([RFC-7231](#)).

X- headers were initially reserved for unstandardized parameters, but the usage of X- headers is deprecated ([RFC-6648](#)). This complicates the contract definition between consumer and producer of an API following these guidelines, since there is no aligned way of using those headers. Because of this, the guidelines restrict which X- headers can be used and how they are used.

The Internet Engineering Task Force's states in [RFC-6648](#) that company specific header' names should incorporate the organization's name. We aim for backward compatibility, and therefore keep the X- prefix.

The following proprietary headers have been specified by this guideline for usage so far. Remember that HTTP header field names are not case-sensitive.

Header field name	Type	Description	Header field value example
X-Flow-ID	String	The flow id of the request, which is written into the logs and passed to called services. Helpful for operational troubleshooting and log analysis. It supports traceability of requests and identifying request flows through system of many services. It should be a string consisting of just printable ASCII characters (i.e. without whitespace). Verify in a received request that it fits to a specific format, has a sensible maximum length and possibly throw out or escape characters/bytes which could crash your log parsing (line breaks, tabs, spaces, NULL). If a legacy subsystem can only work with flow IDs of a specific format, it needs to define this in its API, or make its own ones.	GKY7oDhpSi KY_gAAAABZ _A
X-Tenant-ID	String	The tenant id for future platform multitenancy support. <i>Should not be used unless new platform multitenancy is truly supported. But should be used by New Platform Prototyping services.</i> Must be validated for external retailer, supplier, etc. tenant users via OAuth2; details in clarification. Currently only used by New Platform Prototyping services.	9f8b3ca3- 4be5-436c- a847- 9cd55460c495

Header field name	Type	Description	Header field value example
X-Sales-Channel	String	Sales channels are owned by retailers and represent a specific consumer segment being addressed with a specific product assortment that is offered via CFA retailer catalogs to consumers (see <a href="#">platform glossary [internal link]</a> )	101
X-Frontend-Type	String	Consumer facing applications (CFAs) provide business experience to their customers via different frontend application types, for instance, mobile app or browser. Info should be passed-through as generic aspect — there are diverse concerns, e.g. pushing mobiles with specific coupons, that make use of it. Current range is mobile-app, browser, facebook-app, chat-app	mobile-app
X-Device-Type	String	There are also use cases for steering customer experience (incl. features and content) depending on device type. Via this header info should be passed-through as generic aspect. Current range is smartphone, tablet, desktop, other	tablet
X-Device-OS	String	On top of device type above, we even want to differ between device platform, e.g. smartphone Android vs. iOS. Via this header info should be passed-through as generic aspect. Current range is iOS, Android, Windows, Linux, MacOS	Android
X-App-Domain	Integer	The app domain (i.e. shop channel context) of the request. Note, app-domain is a legacy concept that will be replaced in new platform by combinations of main CFA concerns like retailer, sales channel, country	16

**Exception:** The only exception to this guideline are the conventional hop-by-hop **X-RateLimit-**headers which can be used as defined in **Must: Use 207 for Batch or Bulk Requests**.

## Must: Propagate Proprietary Headers

All Zalando's proprietary headers are end-to-end headers. [2: HTTP/1.1 standard ([RFC-7230](#)) defines two types of headers: end-to-end and hop-by-hop headers. End-to-end headers must be transmitted to the ultimate recipient of a request or response. Hop-by-hop headers, on the contrary, are meaningful for a single connection only.]

All headers specified above must be propagated to the services down the call chain. The header names and values must remain unchanged.

For example, the values of the custom headers like **X-Device-Type** can affect the results of queries by using device type information to influence recommendation results. Besides, the values of the custom headers can influence the results of the queries (e.g. the device type information influences the recommendation results).

Sometimes the value of a proprietary header will be used as part of the entity in a subsequent request. In such cases, the proprietary headers must still be propagated as headers with the subsequent request, despite the duplication of information.

# 17. Deprecation

Sometimes it is necessary to phase out an API endpoint (or version), for instance, if a field is no longer supported in the result or a whole business functionality behind an endpoint has to be shut down. There are many other reasons as well. As long as these endpoints are still used by consumers these are breaking changes and not allowed. Deprecation rules have to be applied to make sure that necessary consumer changes are aligned and deprecated endpoints are not used before API changes are deployed.

## Must: Obtain Approval of Clients

Before shutting down an API (or version of an API) the producer must make sure, that all clients have given their consent to shut down the endpoint. Producers should help consumers to migrate to a potential new endpoint (i.e. by providing a migration manual). After all clients are migrated, the producer may shut down the deprecated API.

## Must: External Partners Must Agree on Deprecation Timespan

If the API is consumed by any external partner, the producer must define a reasonable timespan that the API will be maintained after the producer has announced deprecation. The external partner (client) must agree to this minimum after-deprecation-lifespan before he starts using the API.

## Must: Reflect Deprecation in API Definition

API deprecation must be part of the OpenAPI definition. If a method on a path, a whole path or even a whole API endpoint (multiple paths) should be deprecated, the producers must set `deprecated=true` on each method / path element that will be deprecated (OpenAPI 2.0 only allows you to define deprecation on this level). If deprecation should happen on a more fine grained level (i.e. query parameter, payload etc.), the producer should set `deprecated=true` on the affected method / path element and add further explanation to the `description` section.

If `deprecated` is set to `true`, the producer must describe what clients should use instead and when the API will be shut down in the `description` section of the API definition.

## Must: Monitor Usage of Deprecated APIs

Owners of APIs used in production must monitor usage of deprecated APIs until the API can be shut down in order to align deprecation and avoid uncontrolled breaking effects. See also the [Must: Provide Online Access to OpenAPI Reference Definition](#).

## Should: Add a Warning Header to Responses

During deprecation phase, the producer should add a **Warning** header (see [RFC 7234 - Warning header](#)) field. When adding the **Warning** header, the **warn-code** must be **299** and the **warn-text** should be in form of *"The path/operation/parameter/... {name} is deprecated and will be removed by {date}. Please see {link} for details."* with a link to a documentation describing why the API is no longer supported in the current form and what clients should do about it. Adding the **Warning** header is not sufficient to gain client consent to shut down an API.

## Should: Add Monitoring for Warning Header

Clients should monitor the **Warning** header in HTTP responses to see if an API will be deprecated in future.

## Must: Not Start Using Deprecated APIs

Clients must not start using deprecated parts of an API.

# 18. API Operation

## Must: Provide Online Access to OpenAPI Reference Definition

All service applications must support access to the OpenAPI Reference Definitions of their external APIs — it is optional for internal APIs — via the following two API endpoints:

- endpoint(s) for GET access on its OpenAPI definition(s), for instance <https://example.com/swagger.json> or <https://example.com/swagger.yaml>.
- API Discovery endpoint <https://example.com/.well-known/schema-discovery> that delivers a JSON discovery document of the following form (adapt the **schema\_url** to point to the endpoint mentioned above):

```
{
  "schema_url": "/swagger.json",
  "schema_type": "swagger-2.0",
  "ui_url": "/ui/"
}
```

- whether these endpoints have to be secured by OAuth depends on context of your API and product management
- if you secure these endpoint, only use **uid** scope

Hint: Though discovery endpoints have to be supported, they should not be specified in the OpenAPI definition as they are generic and provide no API specific information.

We distinguish between internal and external APIs of an application which is owned by a specific team and often implemented via small set of services. An external API is used by clients outside the team - usually another application owned by a different team or even an external business partner user of our platform. An internal API is only used within the application and only by the owning team, for instance, for operational or implementation internal purposes.

Background: In our dynamic and complex service infrastructure, it is important to provide API client developers a central place with online access to the OpenAPI reference definitions of all running applications. As a part of the Zalando infrastructure, the `.well-known/schema-discovery` endpoint is used by the API Discovery to detect all API definitions. It checks all running applications via the endpoint above and stores the discovered API definitions. API Discovery itself provides a RESTful API as well as an API Viewer (Swagger-UI) for central access to all discovered API definitions.

Editorial: For the time being, this document is an appropriate place to mention this rule, even though it is not a RESTful API definition rule but related to service implementation obligations to support client developer API discovery.

Further reading:

- [Library to make your Spring Boot service crawlable via API Discovery](#)

## Should: Monitor API Usage

Owners of APIs used in production should monitor API service to get information about its using clients. This information, for instance, is useful to identify potential review partner for API changes.

Hint: A preferred way of client detection implementation is by logging of the client-id retrieved from the OAuth token.

# 19. Events

Zalando's architecture centers around decoupled microservices and in that context we favour asynchronous event driven approaches. The guidelines in this section focus on how to design and publish events intended to be shared for others to consume.

### Events, Event Types and Categories.

Events are defined using an item called an *Event Type*. The Event Type allows events to have their structure declared with a schema by producers and understood by consumers. An Event Type declares standard information, such as a name, an owning application (and by implication, an owning team), a schema defining the event's custom data, and a compatibility mode declaring how the schema will be evolved. Event Types also allow the declaration of validation and enrichment strategies for events, along with supplemental information such as how events can be partitioned in an event stream.

Event Types belong to a well known *Event Category* (such as a data change category), which provides extra information that is common to that kind of event.

Event Types can be published and made available as API resources for teams to use, typically in an *Event Type Registry*. Each event published can then be validated against the overall structure of its event type and the schema for its custom data.

The basic model described above was originally developed in the [Nakadi project](#), which acts as a reference implementation of the event type registry, and as a validating publish/subscribe broker for event producers and consumers.

## Must: Treat Events as part of the service interface

Events are part of a service's interface to the outside world equivalent in standing to a service's REST API. Services publishing data for integration must treat their events as a first class design concern, just as they would an API. For example this means approaching events with the "API first" principle in mind as described in the [\[introduction\]](#).

## Must: Make Events available for review

Services publishing event data for use by others must make the event schema available for review.

## Must: Ensure Event Type schemas conform to Open API's Schema Object

Event type schema are defined in accordance with the Open API Schema Object specification, which uses a subset of [JSON Schema Draft 4](#), and also adds other features. This allows events to properly align with API resource representations (it's particularly useful for events that represent data changes about resources). The guideline exists since declaring event type schema using JSON-Schema syntax is common practice (because Open API doesn't yet allow standalone object definitions).

In the rest of this section we'll call out some of the more notable differences between Open API and JSON-Schema. Please note this is not a complete list - in general it's recommended to familiarise yourself with Open API's Schema Object (the details are available from the "[Schema Object](#)" section of that specification).

Open API *removes* some JSON-Schema keywords. These must not be used in event type schemas. The list of Open API object keywords can be seen [here](#), but for convenience the list of non-available keywords relative to JSON-Schema are:

- `additionalItems`
- `contains`
- `patternProperties`
- `dependencies`
- `propertyNames`
- `const`
- `not`
- `oneOf`

Open API *redefines* some keywords:

- **additionalProperties**: For event types that declare compatibility guarantees, there are recommended constraints around the use of this field. See the guideline "Avoid **additionalProperties** in event type definitions" for details.

Open API *extends* JSON-Schema with some keywords:

- **readOnly**: events are logically immutable, so **readOnly** can be considered redundant, but harmless.
- **discriminator**: discriminators exist to support polymorphism and act as an alternative to **oneOf**.
- **^x-**: patterned objects in the form of **vendor extensions** can be used in event type schema, but it might be the case that general purpose validators do not understand them to enforce a validation check, and fall back to must-ignore processing. A future version of the guidelines may define well known vendor extensions for events.

## Must: Ensure that Events are registered as Event Types

In Zalando's architecture, events are registered using a structure called an *Event Type*. The Event Type declares standard information as follows:

- A well known event category, such as a general or data change category.
- The name of the event type.
- An owning application, and by implication, an owning team.
- A schema defining the event payload.
- The compatibility mode for the type.

Event Types allow easier discovery of event information and ensure that information is well-structured, consistent, and can be validated.

Event type owners must pay attention to the choice of compatibility mode. The mode provides a means to evolve the schema. The range of modes are designed to be flexible enough so that producers can evolve schemas while not inadvertently breaking existing consumers:

- **none**: Any schema modification is accepted, even if it might break existing producers or consumers. When validating events, undefined properties are accepted unless declared in the schema.
- **forward**: A schema **S1** is forward compatible if the previously registered schema, **S0** can read events defined by **S1** - that is, consumers can read events tagged with the latest schema version using the previous version as long as consumers follow the robustness principle described in the guideline's [API Design Principles](#).
- **compatible**: This means changes are fully compatible. A new schema, **S1**, is fully compatible when every event published since the first schema version will validate against the latest schema. In compatible mode, only the addition of new optional properties and definitions to an

existing schema is allowed. Other changes are forbidden.

The compatibility mode interact with revision numbers in the schema `version` field, which follows semantic versioning (MAJOR.MINOR.PATCH):

- Changing an event type with compatibility mode `compatible` can lead to a PATCH or MINOR version revision. MAJOR breaking changes are not allowed.
- Changing an event type with compatibility mode `forward` can lead to a PATCH or MINOR version revision. MAJOR breaking changes are not allowed.
- Changing an event type with compatibility mode `none` can lead to PATCH, MINOR or MAJOR level changes.

The following examples illustrate this relations:

- Changes to the event type's `title` or `description` are considered PATCH level.
- Adding new optional fields to an event type's schema is considered a MINOR level change.
- All other changes are considered MAJOR level, such as renaming or removing fields, or adding new required fields.

The core Event Type structure is shown below as an Open API object definition:

```
EventType:
  description: |
    An event type defines the schema and its runtime properties. The required
    fields are the minimum set the creator of an event type is expected to
    supply.
  required:
    - name
    - category
    - owning_application
    - schema
  properties:
    name:
      description: |
        Name of this EventType. Note: the name can encode the
        owner/responsible for this EventType and ideally should follow a
        naming pattern that makes it easy to read and understand.
      type: string
      pattern: '[a-zA-Z][-0-9a-zA-Z_]*(\.[a-zA-Z][-0-9a-zA-Z_]*)*'
      example: order.order_cancelled, business_partner.contract
    owning_application:
      description: |
        Name of the application (eg, as would be used in infrastructure
        application or service registry) owning this `EventType`.
      type: string
      example: price-service
    category:
      description: Defines the category of this EventType.
      type: string
```



```

x-extensible-enum:
  - data
  - general
compatibility_mode:
  description: |
    The compatibility mode to evolve the schema.
  type: string
x-extensible-enum:
  - compatible
  - forward
  - none
default: forward
schema:
  description: The most recent payload schema for this EventType.
  type: object
  properties:
    version:
      description: Values are based on semantic versioning (eg "1.2.1").
      type: string
      default: '1.0.0'
    created_at:
      description: Creation timestamp of the schema.
      type: string
      readOnly: true
      format: date-time
      example: '1996-12-19T16:39:57-08:00'
    type:
      description: |
        The schema language of schema definition. Currently only
        json_schema (JSON Schema v04) syntax is defined, but in the
        future there could be others.
      type: string
      x-extensible-enum:
        - json_schema
    schema:
      description: |
        The schema as string in the syntax defined in the field type.
      type: string
  required:
    - type
    - schema
created_at:
  description: When this event type was created.
  type: string
  pattern: date-time
updated_at:
  description: When this event type was last updated.
  type: string
  pattern: date-time

```

APIs such as registries supporting event types, may extend the model, including the set of supported categories and schema formats. For example the Nakadi API's event category registration also allows the declaration of validation and enrichment strategies for events, along with supplemental information, such as how events are partitioned in the stream.

## Must: Ensure Events conform to a well-known Event Category

An *event category* describes a generic class of event types. The guidelines define two such categories:

- General Event: a general purpose category.
- Data Change Event: a category used for describing changes to data entities used for data replication based data integration.

The set of categories is expected to evolve in the future.

A category describes a predefined structure that event publishers must conform to along with standard information about that kind of event (such as the operation for a data change event).

### The General Event Category.

The structure of the *General Event Category* is shown below as an Open API Schema Object definition:

```
GeneralEvent:
  description: |
    A general kind of event. Event kinds based on this event define their
    custom schema payload as the top level of the document, with the
    "metadata" field being required and reserved for standard metadata. An
    instance of an event based on the event type thus conforms to both the
    EventMetadata definition and the custom schema definition. Previously
    this category was called the Business Category.
  required:
    - metadata
  properties:
    metadata:
      $ref: '#/definitions/EventMetadata'
```

Event types based on the General Event Category define their custom schema payload at the top-level of the document, with the `metadata` field being reserved for standard information (the contents of `metadata` are described further down in this section).

In the example fragment below, the reserved `metadata` field is shown with fields "a" and "b" being defined as part of the custom schema:

Note:

- The General Event in a previous version of the guidelines was called a *Business Event*. Implementation experience has shown that the category's structure gets used for other kinds of events, hence the name has been generalized to reflect how teams are using it.
- The General Event is still useful and recommended for the purpose of defining events that drive a business process.
- The Nakadi broker still refers to the General Category as the Business Category and uses the keyword "business" for event type registration. Other than that, the JSON structures are identical.

See [Must: Events must not provide sensitive customer personal data](#) for more guidance on how to use the category.

### **The Data Change Event Category.**

The *Data Change Event Category* structure is shown below as an Open API Schema Object:

## DataChangeEvent:

### description: |

Represents a change to an entity. The required fields are those expected to be sent by the producer, other fields may be added by intermediaries such as a publish/subscribe broker. An instance of an event based on the event type conforms to both the DataChangeEvent's definition and the custom schema definition.

### required:

- metadata
- data\_op
- data\_type
- data

### properties:

#### metadata:

description: The metadata for this event.

\$ref: '#/definitions/EventMetadata'

#### data:

##### description: |

Contains custom payload for the event type. The payload must conform to a schema associated with the event type declared in the metadata object's `event\_type` field.

type: object

#### data\_type:

description: name of the (business) data entity that has been mutated

type: string

example: 'sales\_order.order'

#### data\_op:

type: string

enum: ['C', 'U', 'D', 'S']

##### description: |

The type of operation executed on the entity:

- C: Creation of an entity
- U: An update to an entity.
- D: Deletion of an entity.
- S: A snapshot of an entity at a point in time.

The Data Change Event Category is structurally different to the General Event Category. It defines a field called `data` for placing the custom payload information, as well as specific information related to data changes in the `data_type`. In the example fragment below, the fields `a` and `b` are part of the custom payload housed inside the `data` field:

See the following guidelines for more guidance on how to use the Data Change Event Category:

- **Should:** Use the hash partition strategy for Data Change Events
- **Must:** Use the General Event Category to signal steps and arrival points in business processes
- **Should:** Provide a means for explicit event ordering

## Event Metadata.

The General and Data Change event categories share a common structure for *metadata*. The metadata structure is shown below as an Open API Schema Object:

```
EventMetadata:
  type: object
  description: |
    Carries metadata for an Event along with common fields. The required
    fields are those expected to be sent by the producer, other fields may be
    added by intermediaries such as publish/subscribe broker.
  required:
    - eid
    - occurred_at
  properties:
    eid:
      description: Identifier of this event.
      type: string
      format: uuid
      example: '105a76d8-db49-4144-ace7-e683e8f4ba46'
    event_type:
      description: The name of the EventType of this Event.
      type: string
      example: 'example.important-business-event'
    occurred_at:
      description: When the event was created according to the producer.
      type: string
      format: date-time
      example: '1996-12-19T16:39:57-08:00'
    received_at:
      description: |
        When the event was seen by an intermediary such as a broker.
      type: string
      readOnly: true
      format: date-time
      example: '1996-12-19T16:39:57-08:00'
    version:
      description: |
        Version of the schema used for validating this event. This may be
        enriched upon reception by intermediaries. This string uses semantic
        versioning.
      type: string
      readOnly: true
    parent_eids:
      description: |
        Event identifiers of the Event that caused the generation of
        this Event. Set by the producer.
      type: array
      items:
        type: string
        format: uuid
      example: '105a76d8-db49-4144-ace7-e683e8f4ba46'
```

```
flow_id:
  description: |
    A flow-id for this event (corresponds to the X-Flow-Id HTTP header).
  type: string
  example: 'JAh6xH40QhCJ9PutIV_RYw'
partition:
  description: |
    Indicates the partition assigned to this Event. Used for systems
    where an event type's events can be sub-divided into partitions.
  type: string
  example: '0'
```

Please note that intermediaries acting between the producer of an event and its ultimate consumers, may perform operations like validation of events and enrichment of an event's **metadata**. For example brokers such as Nakadi, can validate and enrich events with arbitrary additional fields that are not specified here and may set default or other values, if some of the specified fields are not supplied. How such systems work is outside the scope of these guidelines but producers and consumers working with such systems should look into their documentation for additional information.

## Must: Ensure that Events define useful business resources

Events are intended to be used by other services including business process/data analytics and monitoring. They should be based around the resources and business processes you have defined for your service domain and adhere to its natural lifecycle (see also [Should: Model complete business processes](#)).

As there is a cost in creating an explosion of event types and topics, prefer to define event types that are abstract/generic enough to be valuable for multiple use cases, and avoid publishing event types without a clear need.

## Must: Events must not provide sensitive customer personal data

Similar to API permission scopes, there will be Event Type permissions passed via an OAuth token supported in near future. In the meantime, teams are asked to note the following:

- Sensitive data, such as (e-mail addresses, phone numbers, etc) are subject to strict access and data protection controls.
- Event type owners **must not** publish sensitive information unless it's mandatory or necessary to do so. For example, events sometimes need to provide personal data, such as delivery addresses in shipment orders (as do other APIs), and this is fine.

## Must: Use the General Event Category to signal steps and arrival points in business processes

When publishing events that represent steps in a business process, event types must be based on the General Event category.

All your events of a single business process will conform to the following rules:

- Business events must contain a specific identifier field (a business process id or "bp-id") similar to flow-id to allow for efficient aggregation of all events in a business process execution.
- Business events must contain a means to correctly order events in a business process execution. In distributed settings where monotonically increasing values (such as a high precision timestamp that is assured to move forwards) cannot be obtained, the `parent_eids` data structure allows causal relationships to be declared between events.
- Business events should only contain information that is new to the business process execution at the specific step/arrival point.
- Each business process sequence should be started by a business event containing all relevant context information.
- Business events must be published reliably by the service.

At the moment we cannot state whether it's best practice to publish all the events for a business process using a single event type and represent the specific steps with a state field, or whether to use multiple event types to represent each step. For now we suggest assessing each option and sticking to one for a given business process.

## Must: Use Data Change Events to signal mutations

When publishing events that represents created, updated, or deleted data, change event types must be based on the Data Change Event category.

- Change events must identify the changed entity to allow aggregation of all related events for the entity.
- Change events [Must: Use Data Change Events to signal mutations](#).
- Change events must be published reliably by the service.

## Should: Provide a means for explicit event ordering

Some common error cases may require event consumers to reconstruct event streams or replay events from a position within the stream. Events *should* therefore contain a way to restore their partial order of occurrence.

This can be done - among other ways - by adding - a strictly monotonically increasing entity version (e.g. as created by a database) to allow for partial ordering of all events for an entity - a strictly monotonically increasing message counter

System timestamps are not necessarily a good choice, since exact synchronization of clocks in distributed systems is difficult, two events may occur in the same microsecond and system clocks may jump backward or forward to compensate drifts or leap-seconds. If you use system timestamps to indicate event ordering, you must carefully ensure that your designated event order is not messed up by these effects.

**Note** that basing events on data structures that can be converged upon in a distributed setting (such as [CRDTs](#), [logical clocks](#) and [vector clocks](#)) is outside the scope of this guidance.

## **Should: Use the hash partition strategy for Data Change Events**

The **hash** partition strategy allows a producer to define which fields in an event are used as input to compute a logical partition the event should be added to. Partitions are useful as they allow supporting systems to scale their throughput while provide local ordering for event entities.

The **hash** option is particularly useful for data changes as it allows all related events for an entity to be consistently assigned to a partition, providing a relative ordered stream of events for that entity. This is because while each partition has a total ordering, ordering across partitions is not assured by a supporting system, thus it is possible for events sent across partitions to appear in a different order to consumers than the order they arrived at the server.

When using the **hash** strategy the partition key in almost all cases should represent the entity being changed and not a per event or change identifier such as the **eid** field or a timestamp. This ensures data changes arrive at the same partition for a given entity and can be consumed effectively by clients.

There may be exceptional cases where data change events could have their partition strategy set to be the producer defined or random options, but generally **hash** is the right option - that is while the guidelines here are a "should", they can be read as "must, unless you have a very good reason".

## **Should: Ensure that Data Change Events match API representations**

A data change event's representation of an entity should correspond to the REST API representation.

There's value in having the fewest number of published structures for a service. Consumers of the service will be working with fewer representations, and the service owners will have less API surface to maintain. In particular, you should only publish events that are interesting in the domain and abstract away from implementation or local details - there's no need to reflect every change that happens within your system.

There are cases where it could make sense to define data change events that don't directly correspond to your API resource representations. Some examples are -

- Where the API resource representations are very different from the datastore representation, but the physical data are easier to reliably process for data integration.



- Publishing aggregated data. For example a data change to an individual entity might cause an event to be published that contains a coarser representation than that defined for an API
- Events that are the result of a computation, such as a matching algorithm, or the generation of enriched data, and which might not be stored as entity by the service.

## **Must: Permissions on events must correspond to API permissions**

If a resource can be read synchronously via a REST API and read asynchronously via an event, the same read-permission must apply: We want to protect access to data, not the way data is accessed.

## **Must: Indicate ownership of Event Types**

Event definitions must have clear ownership - this can be indicated via the `owning_application` field of the EventType.

Typically there is one producer application, which owns the EventType and is responsible for its definition, akin to how RESTful API definitions are managed. However, the owner may also be a particular service from a set of multiple services that are producing the same kind of event.

## **Must: Define Event Payloads in accordance with the overall Guidelines**

Events must be consistent with other API data and the API Guidelines in general.

Everything expressed in the [\[introduction\]](#) to these Guidelines is applicable to event data interchange between services. This is because our events, just like our APIs, represent a commitment to express what our systems do and designing high-quality, useful events allows us to develop new and interesting products and services.

What distinguishes events from other kinds of data is the delivery style used, asynchronous publish-subscribe messaging. But there is no reason why they could not be made available using a REST API, for example via a search request or as a paginated feed, and it will be common to base events on the models created for the service's REST API.

The following existing guideline sections are applicable to events:

- [General Guidelines](#)
- [API Naming](#)
- [Data Formats](#)
- [Common Data Types](#)
- [Hypermedia](#)

## Must: Maintain backwards compatibility for Events

Changes to events must be based around making additive and backward compatible changes. This follows the guideline, "Must: Don't Break Backward Compatibility" from the [Compatibility guidelines](#).

In the context of events, compatibility issues are complicated by the fact that producers and consumers of events are highly asynchronous and can't use content-negotiation techniques that are available to REST style clients and servers. This places a higher bar on producers to maintain compatibility as they will not be in a position to serve versioned media types on demand.

For event schema, these are considered backward compatible changes, as seen by consumers -

- Adding new optional fields to JSON objects.
- Changing the order of fields (field order in objects is arbitrary).
- Changing the order of values with same type in an array.
- Removing optional fields.
- Removing an individual value from an enumeration.

These are considered backwards-incompatible changes, as seen by consumers -

- Removing required fields from JSON objects.
- Changing the default value of a field.
- Changing the type of a field, object, enum or array.
- Changing the order of values with different type in an array (also known as a tuple).
- Adding a new optional field to redefine the meaning of an existing field (also known as a co-occurrence constraint).
- Adding a value to an enumeration (note that `x-extensible-enum` is not available in JSON Schema)

## Should: Avoid `additionalProperties` in event type definitions

Event type schema should avoid using `additionalProperties` declarations, in order to support schema evolution.

Events are often intermediated by publish/subscribe systems and are commonly captured in logs or long term storage to be read later. In particular, the schemas used by publishers and consumers can drift over time. As a result, compatibility and extensibility issues that happen less frequently with client-server style APIs become important and regular considerations for event design. The guidelines recommend the following to enable event schema evolution:

- Publishers who intend to provide compatibility and allow their schemas to evolve safely over time **must not** declare an `additionalProperties` field with a value of `true` (i.e., a wildcard extension point). Instead they must define new optional fields and update their schemas in advance of publishing those fields.

- Consumers **must** ignore fields they cannot process and not raise errors. This can happen if they are processing events with an older copy of the event schema than the one containing the new definitions specified by the publishers.

The above constraint does not mean fields can never be added in future revisions of an event type schema - additive compatible changes are allowed, only that the new schema for an event type must define the field first before it is published within an event. By the same turn the consumer must ignore fields it does not know about from its copy of the schema, just as they would as an API client - that is, they cannot treat the absence of an `additionalProperties` field as though the event type schema was closed for extension.

Requiring event publishers to define their fields ahead of publishing avoids the problem of *field redefinition*. This is when a publisher defines a field to be of a different type that was already being emitted, or, is changing the type of an undefined field. Both of these are prevented by not using `additionalProperties`.

See also "Treat Open API Definitions As Open For Extension By Default" in the [Compatibility](#) section for further guidelines on the use of `additionalProperties`.

## Must: Use unique Event identifiers

The `eid` (event identifier) value of an event must be unique.

The `eid` property is part of the standard metadata for an event and gives the event an identifier. Producing clients must generate this value when sending an event and it must be guaranteed to be unique from the perspective of the owning application. In particular events within a given event type's stream must have unique identifiers. This allows consumers to process the `eid` to assert the event is unique and use it as an idempotency check.

Note that uniqueness checking of the `eid` might be not enforced by systems consuming events and it is the responsibility of the producer to ensure event identifiers do in fact distinctly identify events. A straightforward way to create a unique identifier for an event is to generate a UUID value.

## Should: Design for idempotent out-of-order processing

Events that are designed for [idempotent](#) out-of-order processing allow for extremely resilient systems: If processing an event fails, consumers and producers can skip/delay/retry it without stopping the world or corrupting the processing result.

To enable this freedom of processing, you must explicitly design for idempotent out-of-order processing: Either your events must contain enough information to infer their original order during consumption or your domain must be designed in a way that order becomes irrelevant.

As common example similar to data change events, idempotent out-of-order processing can be supported by sending the following information:

- the process/resource/entity identifier,
- a [monotonically increasing ordering key](#) and

- the process/resource state after the change.

A receiver that is interested in the current state can then ignore events that are older than the last processed event of each resource. A receiver interested in the history of a resource can use the ordering key to recreate a (partially) ordered sequence of events.

## Must: Follow conventions for Event Type names

Event types can follow these naming conventions (each convention has its own should, must or could conformance level) -

- Event type names must be url-safe. This is because the event type names may appear in URLs published by other systems and APIs.
- Event type names should be lowercase words and numbers, using hyphens, underscores or periods as separators.

## Must: Prepare for duplicate Events

Event consumers must be able to process duplicate events.

Most message brokers and data streaming systems offer “at-least-once” delivery. That is, one particular event is delivered to the consumers one or more times. Other circumstances can also cause duplicate events.

For example, these situations occur if the publisher sends an event and doesn't receive the acknowledgment (e.g. due to a network issue). In this case, the publisher will try to send the same event again. This leads to two identical events in the event bus which have to be processed by the consumers. Similar conditions can appear on consumer side: an event has been processed successfully, but the consumer fails to confirm the processing.

# Appendix A: References

This section collects links to documents to which we refer, and base our guidelines on.

## OpenAPI Specification

- [OpenAPI Specification](#)

## Publications, specifications and standards

- [RFC 3339: Date and Time on the Internet: Timestamps](#)
- [RFC 5988: Web Linking](#)
- [RFC 7159: The JavaScript Object Notation \(JSON\) Data Interchange Format](#)
- [RFC 7230: Hypertext Transfer Protocol \(HTTP/1.1\): Message Syntax and Routing](#)
- [RFC 7231: Hypertext Transfer Protocol \(HTTP/1.1\): Semantics and Content](#)

- [RFC 7807: Problem Details for HTTP APIs](#)
- [ISO 8601: Date and time format](#)
- [ISO 3166-1 alpha-2: Two letter country codes](#)
- [ISO 639-1: Two letter language codes](#)
- [ISO 4217: Currency codes](#)
- [BCP 47: Tags for Identifying Languages](#)

## Dissertations

- [Roy Thomas Fielding - Architectural Styles and the Design of Network-Based Software Architectures](#). This is the text which defines what REST is.

## Books

- [REST in Practice: Hypermedia and Systems Architecture](#)
- [Build APIs You Won't Hate](#)
- [InfoQ eBook - Web APIs: From Start to Finish](#)

## Blogs

- [Lessons-learned blog: Thoughts on RESTful API Design](#)

## Appendix B: Tooling

This is not a part of the actual guidelines, but might be helpful for following them. Using a tool mentioned here doesn't automatically ensure you follow the guidelines.

## API First Integrations

The following frameworks were specifically designed to support the API First workflow with OpenAPI YAML files (sorted alphabetically):

- [Connexion: OpenAPI First framework for Python on top of Flask](#)
- [Friboo: utility library to write microservices in Clojure with support for Swagger and OAuth](#)
- [Api-First-Hand: API-First Play Bootstrapping Tool for Swagger/OpenAPI specs](#)
- [Swagger Codegen: template-driven engine to generate client code in different languages by parsing Swagger Resource Declaration](#)
- [Swagger Codegen Tooling: plugin for Maven that generates pieces of code from OpenAPI specification](#)
- [Swagger Plugin for IntelliJ IDEA: plugin to help you easily edit Swagger specification files inside IntelliJ IDEA](#)

The Swagger/OpenAPI homepage lists more [Community-Driven Language Integrations](#), but most of them do not fit our API First approach.

## Support Libraries

These utility libraries support you in implementing various parts of our RESTful API guidelines (sorted alphabetically):

- [Problem: Java library that implements application/problem+json](#)
- [Problems for Spring Web MVC: library for handling Problems in Spring Web MVC](#)
- [Jackson Datatype Money: extension module to properly support datatypes of javax.money](#)
- [JSON fields: framework for limiting fields of JSON objects exposed by Rest APIs](#)
- [Tracer: call tracing and log correlation in distributed systems](#)
- [TWINTIP: API definition crawler for the STUPS ecosystem](#)
- [TWINTIP Spring Integration: API discovery endpoint for Spring Web MVC](#)

## Appendix C: Changelog

This change log only contains major changes made after October 2016.

Non-major changes are editorial-only changes or minor changes of existing guidelines, e.g. adding new error code. Major changes are changes that come with additional obligations, or even change an existing guideline obligation. The latter changes are additionally labeled with "Rule Change" here.

To see a list of all changes, please have a look at the [commit list in Github](#).

### Rule Changes

- **2017-08-22:** Migration to AsciiDoc
- **2017-07-20:** Be more precise on client vs. server obligations for compatible API extensions.
- **2017-06-06:** Made money object guideline clearer.
- **2017-05-17:** Added guideline on query parameter collection format.
- **2017-05-10:** Added the convention of using RFC2119 to describe guideline levels, and replaced `book.could` with `book.may`.
- **2017-03-30:** Added rule that permissions on resources in events must correspond to permissions on API resources
- **2017-03-30:** Added rule that APIs should be modelled around business processes
- **2017-02-28:** Extended information about how to reference sub-resources and the usage of composite identifiers in the [Must: Use Domain-Specific Resource Names](#) part.
- **2017-02-22:** Added guidance for conditional requests with If-Match/If-None-Match

- 2017-02-02: Added guideline for batch and bulk request
- 2017-02-01: **May:** Use Content-Location Header
- 2017-01-18: Removed "Avoid Javascript Keywords" rule
- 2017-01-05: Clarification on the usage of the term "REST/RESTful"
- 2016-12-07: Introduced "API as a Product" principle
- 2016-12-06: New guideline: "Should Only Use UUIDs If Necessary"
- 2016-12-04: Changed OAuth flow example from implicit to password in [Security](#).
- 2016-10-13: **Must:** Define Format for Type Number and Integer
- 2016-10-10: Introduced the changelog. From now on all rule changes on API guidelines will be recorded here.