

GraphQL for MicroProfile

Jean-Francois James, Phillip Krüger, Andy McCright, Jean-Baptiste Roux, Bojan Tomic, Adam Anderson

1.0-M2, June 17, 2019

Table of Contents

MicroProfile GraphQL	2
Introduction to MicroProfile GraphQL	3
About GraphQL	3
Why GraphQL	3
GraphQL and REST	3
What make GraphQL different?	4
GraphQL and Databases	5
MicroProfile GraphQL	5
GraphQL Entities	6
POJOs	6
Types vs InputTypes	6
Java interfaces as GraphQL entity types	6
GraphQL interfaces	7
Limitations	7
Generic types	7
Fields	7
Scalars	7
Deprecation	7
Default Values	7
FieldsOrder / InputFieldsOrder	7
Ignorable fields	7
GraphQL Components	8
Component Definition	8
Queries	8
Mutations	8
Generated Schema	8
Deprecation	8
Default Values	8
Lifecycle	8
Error Handling	9
Release Notes for MicroProfile GraphQL 1.0	10

Specification: GraphQL for MicroProfile

Version: 1.0-M2

Status: Draft

Release: June 17, 2019

Copyright (c) 2019 Contributors to the Eclipse Foundation

Licensed under the Apache License, Version 2.0 (the "License");
you may not use this file except in compliance with the License.
You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software
distributed under the License is distributed on an "AS IS" BASIS,
WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
See the License for the specific language governing permissions and
limitations under the License.

MicroProfile GraphQL

Introduction to MicroProfile GraphQL

About GraphQL

GraphQL is an open-source data query and manipulation language for APIs, and a runtime for fulfilling queries with existing data. GraphQL interprets strings from the client, and returns data in an understandable, predictable, pre-defined manner. GraphQL is an alternative, though not necessarily a replacement for REST.

GraphQL was developed internally by Facebook in 2012 before being publicly released in 2015. Facebook delivered both a [specification](#) and a [reference implementation](#) in JavaScript.

On 7 November 2018, Facebook moved the GraphQL project to the newly-established [GraphQL foundation](#), hosted by the non-profit Linux Foundation. This is a significant milestone in terms of industry and community adoption. GraphQL is widely used by [many customers](#).

- More info: <https://en.wikipedia.org/wiki/GraphQL>
- Home page: <https://graphql.org/>
- Specification: <https://facebook.github.io/graphql/draft/>

Why GraphQL

The main reasons for using GraphQL are:

- Avoiding over-fetching or under-fetching data. Clients can retrieve several types of data in a single request or can limit the response data based on specific criteria.
- Enabling data models to evolve. Changes to the schema can be made so as to not require changes on existing clients, and vice versa - this can be done without a need for a new version of the application.
- Advanced developer experience:
 - The schema defines how the data can be accessed and serves as the contract between the client and the server. Development teams on both sides can work without further communication.
 - Native schema introspection enables users to discover APIs and to refine the queries on the client-side. This advantage is increased with graphical tools such as [GraphiQL](#) and [GraphQL Voyager](#) enabling smooth and easy API discovery.
 - On the client-side, the query language provides flexibility and efficiency enabling developers to adapt to the constraints of their technical environments while reducing server round-trips.

GraphQL and REST

GraphQL and REST have many similarities and are both widely used in modern microservice applications. The two technologies also have some differences.

REST stands for "Representational State Transfer". It is an architectural style for network-based software specified by Roy Fielding in 2000 in a [dissertation](#) defining 6 theoretical constraints:

1. uniform interface
2. stateless
3. client-server
4. cacheable
5. layered system
6. code on demand (optional).

REST is often implemented as JSON over HTTP, but REST is fundamentally technically agnostic to data type and transport; it is an architectural style. In particular, it doesn't require to use HTTP. However, it recommends using the maximum capacity of the underlying network protocol to apply the 6 basic principles. For instance, REST implementations can utilize HTTP semantics with a proper use of verbs (POST, GET, PUT, PATCH, DELETE) and response codes (2xx, 4xx, 5xx).

GraphQL takes its roots from a Facebook [specification](#) published in 2015. As of this date, GraphQL has been subject to 5 releases:

- June 2018
- October 2016
- April 2016
- October 2015
- July 2015

According to its definition: "GraphQL is a query language for describing the capabilities and requirements of data models for client-server applications."

Like REST, GraphQL is independent from particular transport protocols or data models:

- it does not endorse the use of HTTP though in practice, and like REST, it is clearly the most widely used protocol,
- it is not tied to any specific database technology or storage engine and is instead backed by existing code and data.

What make GraphQL different?

In practice, here are the main differentiating features of GraphQL compared to REST:

- **schema-driven:** a GraphQL API natively exposes a schema describing the structure of the data and operations (queries and mutations) exposed. This schema acts as a contract between the server and its clients. In a way GraphQL provides an explicit answer to the API discovery problem where REST relies on the ability of developers to properly use other mechanisms such as HATEOS and/or OpenAPI,
- **single HTTP endpoint:** a typical GraphQL API is made of a single endpoint and access to data

and operations is achieved through the query language. In a HTTP context, the endpoint is defined as a URL and the query can be transported as a query string (GET request) or in the request body (POST request),

- **flexible data retrieval:** by construction the query language enables the client to select the expected data in the response with a fine level of granularity, thus avoiding over- or under-fetching data,
- **reduction of server requests:** the language allows the client to aggregate the expected data into a single request,
- **easier version management:** thanks to the native capabilities to create new data while deprecating old ones,
- **partial results:** partial results are delivered by the GraphQL server in case of errors. A GraphQL result is made of data and errors. Clients are responsible for processing the partial results,
- **low coupling with HTTP:** GraphQL does not try to make the most of HTTP semantics. Queries can be made using GET or POST requests. The HTTP result code does not reflect the GraphQL response,
- **challenging authorization handling:** an appropriate data access authorization policy must be defined and implemented to counter the extreme flexibility of the query language. For example, one client may be authorized to access some data that others are not,
- **challenging API management:** most API management solutions are based on REST capabilities and allow for endpoint (URL-based) policies to be established. GraphQL API has a single entry point. It may be necessary to analyze the client request data to ensure it conforms to established policies. For example, it may be necessary to validate mutations or to prevent the client from executing an overly complex request that would crash the server.

GraphQL and Databases

GraphQL is about data query and manipulation but it is not a database technology:

- It is a query language for APIs,
- It is database and storage agnostic,
- It can be used in front of any kind of backend, with or without a database.

One of GraphQL's strength is its multi-datasource capability enabling a single endpoint to aggregate data from various sources with a single API.

MicroProfile GraphQL

The intent of the MicroProfile GraphQL specification is provide a "code-first" set of APIs that will enable users to quickly develop portable GraphQL-based applications in Java.

GraphQL Entities

Entities are the objects used in GraphQL. They can be simple objects ("scalars" in GraphQL terminology) or more complex objects that are composed of scalars. According to the [GraphQL documentation](#) a scalar has no sub-fields, and all GraphQL implementations are expected to handle, the following scalar types: - `Int` - which maps to a Java `int/Integer` or `long/Long`. - `Float` - which maps to a Java `float/Float` or `double/Double`. - `String` - which maps to a Java `char`, `char[]`, `String`, etc. - `Boolean` - which maps to a Java `boolean/Boolean`. - `ID` - which is a specialized type serialized like a `String`.

In order for an entity class to be defined in the GraphQL schema, it must meet at least one of the following criteria: - It must be the return type or parameter (annotated with `@Argument`) of a query or mutation method, - It must be annotated with `@Type`, - It must be annotated with `@InputType`

POJOs

Any Plain Old Java Object (POJO) can be an entity. No special annotations are required. Implementations of MicroProfile GraphQL must use JSON-B to serialize and deserialize entities to JSON, so it is possible to further define entities using JSON-B annotations.

If the entity cannot be serialized by JSON-B, the implementation must return in an internal server error to the client.

Types vs InputTypes

GraphQL differentiates types from input types. Input types are entities that are sent by the client as arguments to queries or mutations. Types are entities that are sent from the server to the client as return types from queries or mutations.

In many cases the same Java type can be used for input (sent *from* client) and output (sent *to* the client), but there are cases where an application may need two different Java types to handle input and output.

The `@Type` annotation is used for output entities while the `@InputType` annotation is used for input entities.

Normally these annotations are unnecessary if the type can be serialized and/or deserialized by JSON-B, and if the type is specified in a query or mutation method. These annotations can be used to specify the name of the type in the GraphQL schema; by default, the entity name in the schema will be the same as the simple class name of the entity type. These annotations can also be used to add a description to the entity in the schema.

Java interfaces as GraphQL entity types

It is possible for entities (types and input types) to be defined as a Java interfaces. In order for JSON-B to deserialize an interface, the interface may need a `JsonbDeserializer` in order to instantiate a concrete type.

GraphQL interfaces

Limitations

Generic types

TODO: info on limitations to generic types (collections only?)

Fields

Scalars

TODO: info on scalars/primitives

Deprecation

Application developers can mark entity fields as deprecated in the GraphQL schema with the `@Deprecated` annotation - either the annotation that comes from the JDK in the `java.lang` package or the MicroProfile GraphQL annotation in the `org.eclipse.microprofile.graphql` package. The latter annotation allows the developer to specify additional text that might provide more information to the consumers of the schema.

Default Values

The `@DefaultValue` annotation may be used to specify a value in an input type to be used if the client did not specify a value. Default values may only be specified on input types (and also as `@Argument` parameters) and will have no effect if specified on output types. The value specified in this annotation may be

FieldsOrder / InputFieldsOrder

TODO: info on fields ordering

Ignorable fields

There may be cases where a developer wants to use a class as a GraphQL type or input type, but use fields that should not be part of the exposed schema. The `@Ignore` annotation can be placed on the field to prevent it from being part of the schema.

GraphQL Components

Component Definition

TODO: info on annotating components with `@GraphQLApi`

Queries

TODO: info on how to create queries

Mutations

TODO: info on how to create mutations

Generated Schema

TODO: info on how the schema should look based on annotations in component classes

Deprecation

TODO: info on how to use the `@Deprecated` annotation appropriately

Default Values

TODO: info on how to use the `@DefaultValue` annotation appropriately

Lifecycle

TODO: info on the components' lifecycle - based on CDI scoping

Error Handling

TODO: info on error handling - including partial results

Release Notes for MicroProfile GraphQL 1.0

[MicroProfile GraphQL Spec PDF](#) [MicroProfile GraphQL Spec HTML](#) [MicroProfile GraphQL Spec Javadocs](#)

Key features:

- Code-first approach to GraphQL schema generation.