

# GraphQL for MicroProfile

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

1.0-M5, November 21, 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	7
Scalars	7
Dates	7
Custom user defined scalars	8
Enumerable types	8
Complex objects	9
Types vs Input	9
Java interfaces as GraphQL entity types	10
GraphQL interfaces	10
Limitations	11
Generic types	12
Fields	12
Other annotations available on Complex Objects	16
Description	16
Default Values	16
Ignorable fields	16
Non-nullable fields	16
GraphQL Components	18
Component Definition	18
API Annotation	18
Basic Example	18
Queries	18
API Annotation	19
Basic POJO Example	19
Entity fields are also queries	19
Description	19
Void Queries	20
Query Names	21
Mutations	21
API Annotation	21

Basic POJO Examples .....	22
Void Mutations .....	23
Mutation Names .....	23
Generated Schema .....	23
Default Values .....	23
Lifecycle .....	24
Error Handling .....	25
Client Errors .....	25
Server Errors .....	25
Partial Results .....	25
JSON-B Integration .....	28
Schema Definition .....	28
Runtime Data Transformation .....	30
Release Notes for MicroProfile GraphQL 1.0 .....	31

Specification: GraphQL for MicroProfile

Version: 1.0-M5

Status: Draft

Release: November 21, 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.

There are 2 main requirements for all implementations of this specification, namely:

- Generate and make the GraphQL Schema available. This is done by looking at the annotations in the users code, and must include all GraphQL Queries and Mutations as well as all entities as

defined either explicitly by annotations or implicitly as the response type or argument(s) of Queries and Mutations.

- Execute GraphQL requests. This will be in the form of either a GraphQL Query or Mutation. As a minimum the specification must support executing these requests via HTTP.

# GraphQL Entities

Entities are the objects used in GraphQL. They can be:

- **Scalars**, or simple objects ("scalars" in GraphQL terminology),
- **Enumerable types** (similar to Java Enum),
- **Complex objects** that are composed of scalars or other objects or enums or a combination of these.

## Scalars

According to the [GraphQL documentation](#) a scalar has no sub-fields, and all GraphQL implementations are expected to handle, the following scalar types:

- **Short** - which maps to a Java `short/Short`
- **Int** - which maps to a Java `int/Integer`
- **Long** - which maps to a Java `long/Long`
- **Float** - which maps to a Java `float/Float` or `double/Double`.
- **Boolean** - which maps to a Java `boolean/Boolean`.
- **Char** - which maps to a Java `char, Character`.
- **String** - which maps to a Java `String`.
- **Byte** - which maps to a Java `byte/Byte`.
- **BigInteger** - which maps to a Java `BigInteger`.
- **BigDecimal** - which maps to a Java `BigDecimal`.
- **Date** - which maps to a Java `java.time.LocalDate`.
- **Time** - which maps to a Java `java.time.LocalTime`.
- **DateTime** - which maps to a Java `java.time.LocalDateTime`.
- **ID** - which is a specialized type serialized like a `String`. Usually, ID types are not intended to be human-readable.

Note that an ID scalar must map to a Java `String`, numerical primitive (`long`, `int`, `double`, `float`) or their object equivalents (`Long`, `Integer`, `Double`, `Float`), or a `java.util.UUID` - anything else is considered a deployment error.

The GraphQL specification also allows for users or implementers to define custom scalars:

## Dates

By default the date related scalars (`DateTime`, `Date`, and `Time`) will use a ISO format.

- `yyyy-MM-dd'T'HH:mm:ss'Z'` for `DateTime`
- `yyyy-MM-dd` for `Date`

- `HH:mm:ss` for Time

By adding the JSON-B annotation `@JsonDateFormat` a user can change the format.

## Custom user defined scalars

@TODO: Define how to create your own scalar.

## Enumerable types

GraphQL offers enumerable types similar to Java `enum` types. In order for an enum 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 (optionally annotated with `@Name`) of a query or mutation method,
- It must be annotated with `@Enum`

The implementation will produce the GraphQL `enum` type in the schema. For example:

```
@Type
public class SuperHero {
    private ShirtSize tshirtSize; // public getters/setters, ...

    @Enum("ClothingSize")
    public enum ShirtSize {
        S, M, L, XL
    }
}
```

The implementation would generate a schema that would include:

```

enum ClothingSize {
  L
  M
  S
  XL
}

type SuperHero {
  #...
  tshirtSize: ClothingSize
  #...
}

input SuperHeroInput {
  #...
  tshirtSize: ClothingSize
  #...
}
#...

```

When using an enumerated type, it is considered a validation error when the client enters a value that is not included in the enumerated type.

## Complex objects

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 (optionally annotated with `@Name`) of a query or mutation method,
- It must be annotated with `@Type`,
- It must be annotated with `@Input`

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 Input

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* the 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 `@Input` 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 for output types; for input types, the simple class name is used with "Input" appended. Thus, an entity class named `com.mypkg.Tree` would create a GraphQL schema type called "Tree" and an input type called "TreeInput".

## 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

GraphQL interfaces are very similar in concept to Java interfaces, in that other types may implement an interface. This allows the GraphQL schema to better align with the Java application's model and allows clients to retrieve the same data (fields) on multiple different entity types. GraphQL interfaces are created with a Java interface type is annotated with `@Interface`. The MP GraphQL implementation must then generate a schema where every class in the application that implements that Java interface must have a type in the schema that implements the GraphQL interface. For example:

```

@Interface
public interface Character {
    public String getName();
}

public class SuperHero implements Character {

    private String name;

    @Override
    @Description("Name of hero")
    public String getName() { return name; }

    // ...
}

public class Villain implements Character {

    private String name;

    @Override
    @Description("Name of villain")
    public String getName() { return name; }

    // ...
}

```

This should generate a schema like:

```

interface Character {
    name: String
}

type SuperHero implements Character {
    #Name of hero
    name: String
    #...
}

type Villain implements Character {
    #Name of villain
    name: String
    #...
}

```

## Limitations

## Generic types

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

## Fields

Fields in GraphQL are similar to fields in Java in that they are a child of a single entity. Thus, Java fields on entity classes are, by default, GraphQL fields of that entity. It is also possible for GraphQL fields that are not part of the Java entity object to be represented as a field of the GraphQL entity. This is because all GraphQL fields are also queries.

Consider the following example:

```
public class SuperHero {
    private String name;
    private String realName;
    private List<String> superPowers;
    // ...
}
```

The Java fields, `name`, `realName` and `superPowers` are all GraphQL fields of the `SuperHero` entity type. Now consider this example:

```
@GraphQLApi
public class MyQueries {

    @Query
    public Location currentLocation(@Source SuperHero hero) {
        return getLocationForHero(hero.getName());
    }
    // ...
}
```

The above query adds a new field to the `SuperHero` GraphQL entity type, called `currentLocation`. This field is not part of the `SuperHero` Java class, but *is* part of the GraphQL entity. This association is made by using the `@Source` annotation. Also note that the `currentLocation` method will only be invoked if the client requests the `currentLocation` field in the query. This is a useful way to prevent looking up data on the server that the client is not interested in.

Users can use the `@Name` annotation to specify a different field name for the field in the GraphQL schema. For example:

```
public class Widget {
    @Name("cost")
    private float price;
    // ... public getters/setters
}
```

This would result in a schema that looks something like:

```
type Widget {
  cost: Float!
}
input WidgetInput {
  cost: Float!
}
```

By putting the `@Name` annotation on the `getter` method, rather than the field, the name will only apply to the `Type`, eg:

```
public class Widget {
    private float price;

    @Name("cost")
    public float getPrice(){
        return this.price;
    }

    public void setPrice(float price){
        this.price = price;
    }
}
```

This would result in a schema that looks something like:

```
type Widget {
  cost: Float!
}
input WidgetInput {
  price: Float!
}
```

The input type keeps the default field name. Similarly, when the `@Name` annotation is only placed on the `setter` method, the name will only apply to the `Input`, eg:

```

public class Widget {

    private float price;

    public float getPrice(){
        return this.price;
    }

    @Name("cost")
    public void setPrice(float price){
        this.price = price;
    }
}

```

This would result in a schema that looks something like:

```

type Widget {
  price: Float!
}
input WidgetInput {
  cost: Float!
}

```

When the default name is used, i.e, there is no annotation specifying the name, the field name will always be used, and not the method name.

The same applies to **Query** and **Mutation** methods. If that method starts with **get**, **set** or **is**, that will be removed when determining the name. Eg:

```

@GraphQLApi
public class MyQueries {

    @Query
    public Location getCurrentLocation(@Source SuperHero hero) {
        // ...
    }
}

```

This would result in a schema that looks something like this:

```

#Query root
type Query {
  #...
  currentLocation(arg0: SuperHeroInput): String
  #...
}

```

Note that the `get` is removed from the name in the schema.

Even though `@Name` is not required on an input argument for a `@Query` or `@Mutation`, it is strongly recommended as it is the only guaranteed portable way to ensure the argument names.

If a user compiles with `-parameters` option, then the implementation should try to use the Java parameter names as the schema argument names, but this is not a requirement. Some implementations may still have trouble getting the parameter names even with the `-parameters` option.

Example recommended argument usage (with annotation):

```
@Query
public SuperHero superHero(@Name("name") String name) {
    return heroDB.getHero(name);
}
```

Above will result in:

```
#Query root
type Query {
  # ...
  superHero(name: String): SuperHero
  # ...
}
```

If the `@Name` annotation is not present, and the user did not compile with the `-parameters` option, or the implementation does not support the `-parameters` option, arguments will get generic names like `arg0`, `arg1` and so on.

Example argument usage (with no annotation):

```
@Query
public SuperHero superHero(String name) {
    return heroDB.getHero(name);
}
```

Above will result in:

```
#Query root
type Query {
  # ...
  superHero(arg0: String): SuperHero
  # ...
}
```

# Other annotations available on Complex Objects

## Description

The `@Description` annotation can be used to provide comments in the generated schema for entity types (both input and output types) and fields.

## 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 method parameters and will have no effect if specified on output types. The value specified in this annotation may be plain text for Java primitives and `String` types or JSON for complex types.

## 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.

If the `@Ignore` annotation is placed on the field itself, then the field will be excluded from both the input and output types in the generated schema. If the annotation is only placed on the "getter" method, then it will only be excluded from the input type. If the annotation is only placed on the "setter" method, then it will only be excluded from the output type.

## Non-nullable fields

The GraphQL specification states that fields may be marked as non-nullable - usually the field's type is marked with an exclamation point to indicate that null values are not allowed. Non-nullable fields may be present on types and input types, providing the client with the proper expectations for providing an input type and that they can expect a non-null value on the return type. If the client sends a null value for a required (non-nullable) field or sends an entity with the required (non-nullable) field unspecified, the implementation should respond with a validation error. Likewise, the implementation should return an error if a null is returned for a required (non-nullable) field from the application code.

By default all GraphQL fields generated from Java primitive properties (`boolean`, `int`, `double`, etc.) will automatically be marked as required. If a Java primitive property has a `@DefaultValue` annotation value, then null is allowed, but the implementation is expected to convert the value to be the default value specified in the annotation.

By default, all GraphQL fields generated from non-primitive properties will be considered nullable. A user may specify that a field is required/non-nullable by adding the `@NonNull` annotation. This annotation may be applied to an entity's getter method, setter method or field. The placement will determine whether it applies to the type, input type or both, respectively.

The implementation should ignore a `@NonNull` annotation when it is on the same field or setter method that also contains `@DefaultValue` annotation, as the "null" value would result in the default value being used.

One drawback to using non-nullable fields is that if there is an error loading a child field, that error could propagate itself up causing the field to be null - and since this is itself an error condition, the implementation must return the non-null field error, which means that the implementation would not be able to send partial results for other child fields.

# GraphQL Components

## Component Definition

### API Annotation

GraphQL endpoints must be annotated with the `@GraphQLApi` annotation.

The `@GraphQLApi` annotation is defined in the API:

*GraphQLApiAnnotation*

```
package org.eclipse.microprofile.graphql;

//...

@Target(ElementType.TYPE)
@Retention(RetentionPolicy.RUNTIME)
@Documented
public @interface GraphQLApi {
}
```

### Basic Example

*Example*

```
@GraphQLApi
@RequestScoped
public class MembershipGraphQLApi {

    @Inject
    private MembershipService membershipService;

    @Query("memberships")
    public List<Membership> getAllMemberships() {
        return getAllMemberships(Optional.empty());
    }

    // Other GraphQL queries and mutations
}
```

## Queries

Queries allows a user to ask for all or specific fields on an object.

## API Annotation

For classes that are annotated with `@GraphQLApi`, implementations must create a query in the schema for every method that is annotated with `@Query`.

The `@Query` annotation is defined in the API:

*QueryAnnotation*

```
package org.eclipse.microprofile.graphql;

//...

@Target(ElementType.METHOD)
@Retention(RetentionPolicy.RUNTIME)
@Documented
public @interface Query {
    String value() default "";
}
```

Table 1. Parameters

<b>value</b>	Overrides the field name in the resulting schema, default will be the method name
--------------	---

## Basic POJO Example

*Example*

```
@Query
public SuperHero superHero(@Name("name") String name) {
    return heroDB.getHero(name);
}
```

## Entity fields are also queries

In the previous example, an explicitly defined query named "superHero" returns a `SuperHero` entity. The fields on that entity class are also implicitly defined as queries. It is possible to define fields as queries explicitly by using the `@Source` annotation on a parameter to the query method. More information on this is available in the [Entity Fields](#) section.

## Description

Queries may be documented as comments in the schema by adding a `@Description` annotation with documentation text as the annotation value to the query method. For example:

### DescriptionExample

```
@Query
@Description("Returns the super hero with the specified name")
public SuperHero superHero(@Name("name") String name) {
    return heroDB.getHero(name);
}
```

This would generate a schema that would include:

### DescriptionSchemaExample

```
type Query {
  ...
  #Returns the super hero with the specified name
  superHero(name: String): [SuperHero]
  #...
  ----
```

The `@Description` annotation can also be placed on parameters of a query method to provide documentation for the arguments. For example:

### ArgumentDescriptionExample

```
@Query
@Description("Returns the super hero with the specified name")
public SuperHero superHero(@Name("name") @Description("Super hero name, not real name") String name) {
    return heroDB.getHero(name);
}
```

This would generate a schema that would include:

### ArgumentDescriptionSchemaExample

```
type Query {
  ...
  #Returns the super hero with the specified name
  superHero(
    #Super hero name, not real name
    name: String
  ): SuperHero
  #...
  ----
```

## Void Queries

By its very nature, query methods must return some value, thus it is considered a deployment

error for a method with a `void` return type to be annotated with `@Query`. If a void method is annotated with the `@Query` annotation, the implementation must prevent the application from starting and should provide a log message indicating that this is not allowed.

## Query Names

The name of a query in the schema is obtained using the following order:

- if the method is annotated with a `@Query` annotation containing a non-empty String for its value, that String value is used as the query name.
- if the method is annotated with a `@Name` annotation containing a non-empty String for its value, that String value is used as the query name.
- if the method is annotated with a `@JsonProperty` annotation containing a non-empty String for its value, that String value is used as the query name.
- if no other name can be determined, the Java method name is used as the query name.

Note that it is considered a deployment error if more than one query method has the same name with the same arguments.

## Mutations

While queries are intended for clients to read data, mutations are intended to modify data. Mutations may create new entities or update or delete existing entities.

### API Annotation

Like queries, mutation methods must be in a class annotated with the `@GraphQLApi` annotation.

The `@Mutation` annotation is defined in the API:

#### *Mutation Annotation*

```
package org.eclipse.microprofile.graphql;

//...

@Target(ElementType.METHOD)
@Retention(RetentionPolicy.RUNTIME)
@Documented
public @interface Mutation {
    String value() default "";
}
```

Table 2. Parameters

<b>value</b>	Overrides the mutation name in the resulting schema, default will be the method name
--------------	--

Mutations generally require arguments (parameters) in order to determine which entity to

modify/delete or the data necessary to create a new entity, etc. These arguments are the parameters of the mutation method and should be annotated with `@Name`. The `@Name` annotation's value is used to specify the name of the argument. The argument name will be used in the generated schema. Arguments can be GraphQL scalars or more complex input types (for more information on input types, see [Entities, Types vs Input](#)).

## Basic POJO Examples

### Create Example

```
@Mutation
public SuperHero createNewHero(@Name("hero") SuperHero newHero)
    throws DuplicateSuperHeroException {

    heroDB.addHero(newHero);
    return heroDB.getHero(newHero.getName());
}
```

### Delete Example

```
@Mutation
public SuperHero removeHero(@Name("hero") String heroName)
    throws UnknownHeroException {

    SuperHero removedHero = heroDB.removeHero(heroName);
    if (removedHero == null) {
        throw new UnknownHeroException(heroName);
    }
    return removedHero;
}
```

### Update Example

```
@Mutation
public SuperHero addNewPowerToHero(@Name("hero") SuperHero hero,
    @Name("newPower") String newPower)
    throws UnknownHeroException {

    SuperHero heroFromDB = heroDB.getHero(hero.getName());
    if (heroFromDB == null) {
        throw new UnknownHeroException(hero.getName());
    }
    heroFromDB.getSuperPowers().add(newPower);
    return heroFromDB;
}
```

## Void Mutations

Like query methods, mutation methods are expected to return some value, thus it is considered a deployment error for a method with a `void` return type to be annotated with `@Mutation`. If a void method is annotated with the `@Mutation` annotation, the implementation must prevent the application from starting and should provide a log message indicating that this is not allowed.

## Mutation Names

The name of a mutation in the schema is obtained using the following order:

- if the method is annotated with a `@Mutation` annotation containing a non-empty String for its value, that String value is used as the mutation name.
- if the method is annotated with a `@Name` annotation containing a non-empty String for its value, that String value is used as the mutation name.
- if the method is annotated with a `@JsonbProperty` annotation containing a non-empty String for its value, that String value is used as the mutation name.
- if no other name can be determined, the Java method name is used as the mutation name.

Note that it is considered a deployment error if more than one mutation method has the same name with the same arguments.

## Generated Schema

MicroProfile GraphQL uses a "code first" approach so that developers do not need to manually keep the code and schema in sync. Each MP GraphQL application will still have a schema but it is generated by the MP GraphQL implementation.

The schema must be available at the `/graphql/schema.graphql` location, relative to the context root.

For example, suppose your application was registered at host, "myhost" on TCP port "50080" with a context root of "MyApp" (usually the context root is the name of the WAR file but without the file extension), then the schema would be available at: <http://myhost:50080/MyApp/graphql/schema.graphql>

## Default Values

It is possible to specify a default value for query or mutation arguments so that the client does not need to specify a value. This is done via the `@DefaultValue` annotation and it is placed on the method parameter for the query/mutation method. The value of the annotation is used as the default value. For scalars, the value is taken as-is (i.e. a value of "foo" for a String would be "foo", a value of "25" for an Int would be 25, etc.). For complex types, the value of the annotation must represent the object as JSON.

```
@Query
public Collection<SuperHero> allHeroesIn(
    @DefaultValue("New York, NY") @Name("city") String city) {

    return allHeroesByFilter(hero -> {
        return city.equals(hero.getPrimaryLocation());});
    }
}

public final static String CAPE =
    "{" +
    "  \"id\": 1000," +
    "  \"name\": \"Cape\"," +
    "  \"powerLevel\": 3," +
    "  \"height\": 1.2," +
    "  \"weight\": 0.3," +
    "  \"supernatural\": false" +
    "}";

@Mutation
public SuperHero provisionHero(@Name("hero") String heroName,
    @DefaultValue(CAPE) @Name("item") Item item)
    throws UnknownHeroException {

    SuperHero hero = heroDB.getHero(heroName);
    if (hero == null) {
        throw new UnknownHeroException(heroName);
    }
    hero.getEquipment().add(item);
    return hero;
}
```

The `@DefaultValue` annotation may also be placed on fields and setters on entity classes to specify the default for GraphQL fields.

## Lifecycle

MicroProfile GraphQL components (POJOs annotated with `@GraphQLApi`) are CDI beans. As such, their lifecycle is managed by CDI. Request-scoped components should be constructed per-request, and application-scoped components should exist for the lifetime of the application. One exception to the normal scoping is `@Dependent` - this scope is treated as if it were a singleton.

# Error Handling

In GraphQL applications most errors will either be client, server or transport errors.

Client errors occur when the client has submitted an invalid request. Examples of client errors include specifying a query or mutation that does not exist, requesting a field on an entity that does not exist, specifying the wrong type of data (such as specifying an `Int` when the schema requires a `String`), etc.

Server errors occur when the request is valid and is properly transported to the server application but the response is unexpected or unable to be fulfilled. Examples of server errors include bugs in the application code, a back-end resource such as a database is down, etc.

Transport errors occur when the request cannot be delivered to the server or when the response cannot be delivered to the client. Examples of transport errors include network disruption, mis-configured firewalls, etc.

The MP GraphQL specification addresses the handling of client and server errors. Transport error handling is beyond the scope of this document.

## Client Errors

Client errors must be handled automatically by the implementation. Invalid requests must never result in user application code invocation. Instead, the implementation must provide the client with an error message that indicates why the client request was invalid.

## Server Errors

If the client request is valid, then the implementation must invoke the correct query or mutation method in the user application. The user application can indicate that an error has occurred by throwing an exception (checked or unchecked). When the user application throws an exception, the implementation must send back a response that includes an error message.

The user may determine the error message that is sent back to the client in two ways: - Throw an instance of `GraphQLException` or a subclass. The implementation must send the exception's message text to the client. Optionally, the user can specify an `ExceptionType` in the exception which must also be sent to the client if specified. - Specify the default error message to use when any other (non-`GraphQLException`) exception is thrown. This is set using the `MicroProfile Config` property, `mp.graphql.defaultErrorMessage`.

## Partial Results

It is possible in GraphQL to send back some results even though the overall request may have failed. This is possible by passing the partial results to the `GraphQLException` (or subclass of `GraphQLException`) that is thrown by the query or mutation method. For example:

```

@Query
public Collection<SuperHero> allHeroesFromCalifornia() throws GraphQLException {
    List<SuperHero> westCoastHeroes = new ArrayList<>();
    try {
        for (SuperHero hero : database.getAllHeroes()) {
            if (hero.getPrimaryLocation().contains("California")) {
                westCoastHeroes.add(hero);
            }
        }
    } catch (Exception ex) {
        throw new GraphQLException(ex, westCoastHeroes);
    }
    return westCoastHeroes;
}

```

If an exception is thrown while iterating over of the database collection of heroes or while checking a hero's location, all previously-processed heroes will still be in the list and will be displayed to the client along with the error data.

Note that the `partialResults` object passed to the `GraphQLException` must match the return type of the query/mutation method from which it is thrown. Otherwise the implementation must throw a `ClassCastException` internally resulting in a much less usable result returned to the client.

It is also possible to send partial results when using multiple methods and the `@Source` annotation. Here is an example:

```

@Query
public Collection<SuperHero> allHeroes() {
    return database.getAllHeroes();
}

@Query
public Location currentLocation(@Source SuperHero hero) throws GraphQLException {
    if (hero.hasLocationBlockingPower()) {
        throw new GraphQLException("Unable to determine location for " + hero.getName
());
    }
    return database.getLocationForHero(hero);
}

```

Suppose the client issued this query:

```
query allHeroes {  
  allHeroes {  
    name  
    currentLocation  
  }  
}
```

In this case, if there are any heroes that have a location blocking power, one or more errors will be returned to the client. However, the names of all of the heroes in the database will be returned as well as the location of all heroes to do not have a location blocking power.

# JSON-B Integration

The MicroProfile GraphQL project depends on JSON-B 1.0 APIs. Implementations of MP GraphQL are encouraged to use JSON-B implementations for consistency.

JSON-B annotations can be used to help determine schema information as well as data transformation at runtime. In all cases, annotations in the `org.eclipse.microprofile.graphql` package will trump JSON-B annotations if there is a conflict.

## Schema Definition

The `@JsonbProperty` annotation's `value` attribute can be used to modify the field name of a type, input type or both - depending on whether the annotation is placed on the getter method, setter method, or Java field, respectively.

For example, suppose you have the following Java code (with properties conforming to the JavaBean spec for `name`, `weight`, and `quantity`):

*JsonbProperty Java Code Example*

```
public class Widget {
    @JsonbProperty("widgetName")
    private String name;
    private double weight;
    private int quantity;
    //...

    @JsonbProperty("shippingWeight")
    public double getWeight() {
        return weight;
    }
    //...

    @JsonbProperty("qty")
    public void setQuantity(int quantity) {
        this.quantity = quantity;
    }
}
```

That would create a schema like this:

### JsonbProperty Schema Example

```
type Widget {
  widgetName: String
  quantity: Int!
  shippingWeight: Float!
}

input WidgetInput {
  widgetName: String
  qty: Int!
  weight: Float!
}
```

Similarly, the `@JsonbTransient` annotation can be used to ignore certain fields from the type or input type in the schema. The same rules apply: if the annotation is on the getter, then the field is ignored in the type; if the annotation is on the setter, then the field is ignored in the input type; if the annotation is on the Java field, it is ignored in both.

Here is an example, similar to the previous example:

### JsonbTransient Java Code Example

```
public class Widget {
  @JsonbTransient
  private String name;
  private double weight;
  private int quantity;
  //...

  @JsonbTransient
  public double getWeight() {
    return weight;
  }
  //...

  @JsonbTransient
  public void setQuantity(int quantity) {
    this.quantity = quantity;
  }
}
```

That would create a schema like this:

### *JsonbTransient Schema Example*

```
type Widget {
  quantity: Int!
}

input WidgetInput {
  weight: Float!
}
```

## Runtime Data Transformation

JSON-B annotations like `@JsonbDateFormat` and `@JsonbNumberFormat` can be used to transform data at runtime.

For example, suppose you have a Java entity like this:

### *JsonbDateFormat Java Code Example*

```
public class Widget {
  @JsonbDateFormat("yyyy-MM-dd")
  private LocalDate constructionDate;
}
```

A query that requests a Widget's construction date would get a result like this:

### *JsonbTransient Schema Example*

```
{
  "constructionDate": "2019-09-16"
}
```

The date format string specified in the `@JsonbDateFormat` annotation will be used as the field's description in the schema, if no `@Description` annotation is provided for that field. If the same field (or property) contains both `@JsonbDateFormat` and `@Description` annotations, it should be considered a best practice to document the date format in the description text so that the format is communicated to clients in the schema.

# 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.