Graphene Documentation

Release 1.0

Syrus Akbary

Jul 25, 2017

Contents

Contents:

Getting started

What is GraphQL?

For an introduction to GraphQL and an overview of its concepts, please refer to the official introduction.

Let's build a basic GraphQL schema from scratch.

Requirements

- Python (2.7, 3.2, 3.3, 3.4, 3.5, pypy)
- Graphene (1.0)

Project setup

```
pip install "graphene>=1.0"
```

Creating a basic Schema

A GraphQL schema describes your data model, and provides a GraphQL server with an associated set of resolve methods that know how to fetch data.

We are going to create a very simple schema, with a `Query` with only one field: `hello` and an input name. And when we query it, it should return `"Hello {name}"`.

```python
import graphene

class Query(graphene.ObjectType):
```

```
    hello = graphene.String(name=graphene.Argument(graphene.String, default_value=
↪"stranger"))

    def resolve_hello(self, args, context, info):
        return 'Hello ' + args['name']

schema = graphene.Schema(query=Query)
```

Querying

Then we can start querying our schema:

```
result = schema.execute('{ hello }')
print result.data['hello'] # "Hello stranger"
```

Congrats! You got your first graphene schema working!

Types Reference

Enums

A `Enum` is a special `GraphQL` type that represents a set of symbolic names (members) bound to unique, constant values.

Definition

You can create an `Enum` using classes:

```python
import graphene

class Episode(graphene.Enum):
    NEWHOPE = 4
    EMPIRE = 5
    JEDI = 6
```

But also using instances of Enum:

```python
Episode = graphene.Enum('Episode', [('NEWHOPE', 4), ('EMPIRE', 5), ('JEDI', 6)])
```

Value descriptions

It's possible to add a description to a enum value, for that the the enum value needs to have the `description` property on it.

```python
class Episode(graphene.Enum):
    NEWHOPE = 4
    EMPIRE = 5
    JEDI = 6
```

```
    @property
    def description(self):
        if self == Episode.NEWHOPE:
            return 'New Hope Episode'
        return 'Other episode'
```

Usage with Python Enums

In case the Enums are already defined it's possible to reuse them using the `Enum.from_enum` function.

```
graphene.Enum.from_enum(AlreadyExistingPyEnum)
```

Notes

`graphene.Enum` uses `enum.Enum` internally (or a backport if that's not available) and can be used in a similar way, with the exception of member getters.

In the Python `Enum` implementation you can access a member by initing the Enum.

```
from enum import Enum
class Color(Enum):
    RED = 1
    GREEN = 2
    BLUE = 3

assert Color(1) == Color.RED
```

However, in Graphene `Enum` you need to call get to have the same effect:

```
from graphene import Enum
class Color(Enum):
    RED = 1
    GREEN = 2
    BLUE = 3

assert Color.get(1) == Color.RED
```

Scalars

Graphene defines the following base Scalar Types:

- `graphene.String`

- `graphene.Int`

- `graphene.Float`

- `graphene.Boolean`

- `graphene.ID`

Graphene also provides custom scalars for Dates, Times, and JSON:

- `graphene.types.datetime.DateTime`

- `graphene.types.datetime.Time`

- `graphene.types.json.JSONString`

Custom scalars

You can create custom scalars for your schema. The following is an example for creating a DateTime scalar:

```python
import datetime
from graphene.types import Scalar
from graphql.language import ast


class DateTime(Scalar):
    '''DateTime Scalar Description'''

    @staticmethod
    def serialize(dt):
        return dt.isoformat()

    @staticmethod
    def parse_literal(node):
        if isinstance(node, ast.StringValue):
            return datetime.datetime.strptime(
                node.value, "%Y-%m-%dT%H:%M:%S.%f")

    @staticmethod
    def parse_value(value):
        return datetime.datetime.strptime(value, "%Y-%m-%dT%H:%M:%S.%f")
```

Mounting Scalars

Scalars mounted in a `ObjectType`, `Interface` or `Mutation` act as `Field`s.

```python
class Person(graphene.ObjectType):
    name = graphene.String()

# Is equivalent to:
class Person(graphene.ObjectType):
    name = graphene.Field(graphene.String)
```

Note: when using the `Field` constructor directly, pass the type and not an instance.

Types mounted in a `Field` act as `Argument`s.

```python
graphene.Field(graphene.String, to=graphene.String())

# Is equivalent to:
graphene.Field(graphene.String, to=graphene.Argument(graphene.String))
```

Lists and Non-Null

Object types, scalars, and enums are the only kinds of types you can define in Graphene. But when you use the types in other parts of the schema, or in your query variable declarations, you can apply additional type modifiers that affect validation of those values.

NonNull

```python
import graphene

class Character(graphene.ObjectType):
    name = graphene.NonNull(graphene.String)
```

Here, we're using a `String` type and marking it as Non-Null by wrapping it using the `NonNull` class. This means that our server always expects to return a non-null value for this field, and if it ends up getting a null value that will actually trigger a GraphQL execution error, letting the client know that something has gone wrong.

The previous `NonNull` code snippet is also equivalent to:

```python
import graphene

class Character(graphene.ObjectType):
    name = graphene.String(required=True)
```

List

```python
import graphene

class Character(graphene.ObjectType):
    appears_in = graphene.List(graphene.String)
```

Lists work in a similar way: We can use a type modifier to mark a type as a `List`, which indicates that this field will return a list of that type. It works the same for arguments, where the validation step will expect a list for that value.

Interfaces

An Interface contains the essential fields that will be implemented by multiple ObjectTypes.

The basics:

- Each Interface is a Python class that inherits from `graphene.Interface`.
- Each attribute of the Interface represents a GraphQL field.

Quick example

This example model defines a `Character` interface with a name. `Human` and `Droid` are two implementations of that interface.

```python
import graphene

class Character(graphene.Interface):
    name = graphene.String()

# Human is a Character implementation
class Human(graphene.ObjectType):
    class Meta:
        interfaces = (Character, )
```

```
    born_in = graphene.String()

# Droid is a Character implementation
class Droid(graphene.ObjectType):
    class Meta:
        interfaces = (Character, )

    function = graphene.String()
```

`name` is a field on the `Character` interface that will also exist on both the `Human` and `Droid` ObjectTypes (as those implement the `Character` interface). Each ObjectType may define additional fields.

The above types have the following representation in a schema:

```
interface Character {
  name: String
}

type Droid implements Character {
  name: String
  function: String
}

type Human implements Character {
  name: String
  bornIn: String
}
```

AbstractTypes

An AbstractType contains fields that can be shared among `graphene.ObjectType`, `graphene.Interface`, `graphene.InputObjectType` or other `graphene.AbstractType`.

The basics:

- Each AbstractType is a Python class that inherits from `graphene.AbstractType`.

- Each attribute of the AbstractType represents a field (a `graphene.Field` or `graphene.InputField` depending on where it is mounted)

Quick example

In this example UserFields is an `AbstractType` with a name. `User` and `UserInput` are two types that have their own fields plus the ones defined in `UserFields`.

```
import graphene

class UserFields(graphene.AbstractType):
    name = graphene.String()

class User(graphene.ObjectType, UserFields):
    pass

class UserInput(graphene.InputObjectType, UserFields):
    pass
```

```
type User {
  name: String
}

inputtype UserInput {
  name: String
}
```

ObjectTypes

An ObjectType is the single, definitive source of information about your data. It contains the essential fields and behaviors of the data you're querying.

The basics:

- Each ObjectType is a Python class that inherits from `graphene.ObjectType`.
- Each attribute of the ObjectType represents a `Field`.

Quick example

This example model defines a Person, with a first and a last name:

```python
import graphene


class Person(graphene.ObjectType):
    first_name = graphene.String()
    last_name = graphene.String()
    full_name = graphene.String()

    def resolve_full_name(self, args, context, info):
        return '{} {}'.format(self.first_name, self.last_name)
```

first_name and **last_name** are fields of the ObjectType. Each field is specified as a class attribute, and each attribute maps to a Field.

The above `Person` ObjectType has the following schema representation:

```
type Person {
  firstName: String
  lastName: String
  fullName: String
}
```

Resolvers

A resolver is a method that resolves certain fields within a `ObjectType`. If not specifed otherwise, the resolver of a field is the `resolve_{field_name}` method on the `ObjectType`.

By default resolvers take the arguments `args`, `context` and `info`.

NOTE: The resolvers on a `ObjectType` are always treated as `staticmethod``s, so the first argument to the resolver method ``self (or root)` need not be an actual instance of the `ObjectType`.

Quick example

This example model defines a `Query` type, which has a reverse field that reverses the given `word` argument using the `resolve_reverse` method in the class.

```python
import graphene

class Query(graphene.ObjectType):
    reverse = graphene.String(word=graphene.String())

    def resolve_reverse(self, args, context, info):
        word = args.get('word')
        return word[::-1]
```

Resolvers outside the class

A field can use a custom resolver from outside the class:

```python
import graphene

def reverse(root, args, context, info):
    word = args.get('word')
    return word[::-1]

class Query(graphene.ObjectType):
    reverse = graphene.String(word=graphene.String(), resolver=reverse)
```

Instances as data containers

Graphene `ObjectTypes` can act as containers too. So with the previous example you could do:

```python
peter = Person(first_name='Peter', last_name='Griffin')

peter.first_name # prints "Peter"
peter.last_name # prints "Griffin"
```

Schema

A Schema is created by supplying the root types of each type of operation, query and mutation (optional). A schema definition is then supplied to the validator and executor.

```python
my_schema = Schema(
    query=MyRootQuery,
    mutation=MyRootMutation,
)
```

Types

There are some cases where the schema cannot access all of the types that we plan to have. For example, when a field returns an `Interface`, the schema doesn't know about any of the implementations.

In this case, we need to use the `types` argument when creating the Schema.

```
my_schema = Schema(
    query=MyRootQuery,
    types=[SomeExtraObjectType, ]
)
```

Querying

To query a schema, call the `execute` method on it.

```
my_schema.execute('{ lastName }')
```

Auto CamelCase field names

By default all field and argument names (that are not explicitly set with the `name` arg) will be converted from `snake_case` to `camelCase` (as the API is usually being consumed by a js/mobile client)

For example with the ObjectType

```
class Person(graphene.ObjectType):
    last_name = graphene.String()
    other_name = graphene.String(name='_other_Name')
```

the `last_name` field name is converted to `lastName`.

In case you don't want to apply this transformation, provide a `name` argument to the field constructor. `other_name` converts to `_other_Name` (without further transformations).

Your query should look like

```
{
    lastName
    _other_Name
}
```

To disable this behavior, set the `auto_camelcase` to `False` upon schema instantiation.

```
my_schema = Schema(
    query=MyRootQuery,
    auto_camelcase=False,
)
```

Mutations

A Mutation is a special ObjectType that also defines an Input.

Quick example

This example defines a Mutation:

```python
import graphene


class CreatePerson(graphene.Mutation):
    class Input:
        name = graphene.String()

    ok = graphene.Boolean()
    person = graphene.Field(lambda: Person)

    @staticmethod
    def mutate(root, args, context, info):
        person = Person(name=args.get('name'))
        ok = True
        return CreatePerson(person=person, ok=ok)
```

person and **ok** are the output fields of the Mutation when is resolved.

Input attributes are the arguments that the Mutation `CreatePerson` needs for resolving, in this case **name** will be the only argument for the mutation.

mutate is the function that will be applied once the mutation is called.

So, we can finish our schema like this:

```python
# ... the Mutation Class


class Person(graphene.ObjectType):
    name = graphene.String()
    age = graphene.Int()


class MyMutations(graphene.ObjectType):
    create_person = CreatePerson.Field()

# We must define a query for our schema
class Query(graphene.ObjectType):
    person = graphene.Field(Person)

schema = graphene.Schema(query=Query, mutation=MyMutations)
```

Executing the Mutation

Then, if we query (`schema.execute(query_str)`) the following:

```
mutation myFirstMutation {
    createPerson(name:"Peter") {
        person {
            name
        }
        ok
    }
}
```

We should receive:

```
{
    "createPerson": {
        "person" : {
```

```
            name: "Peter"
        },
        "ok": true
    }
}
```

InputFields and InputObjectTypes

InputFields are used in mutations to allow nested input data for mutations

To use an InputField you define an InputObjectType that specifies the structure of your input data

```python
import graphene

class PersonInput(graphene.InputObjectType):
    name = graphene.String()
    age = graphene.Int()

class CreatePerson(graphene.Mutation):
    class Input:
        person_data = graphene.Argument(PersonInput)

    person = graphene.Field(lambda: Person)

    @staticmethod
    def mutate(root, args, context, info):
        p_data = args.get('person_data')

        name = p_data.get('name')
        age = p_data.get('age')

        person = Person(name=name, age=age)
        return CreatePerson(person=person)
```

Note that **name** and **age** are part of **person_data** now

Using the above mutation your new query would look like this:

```
mutation myFirstMutation {
    createPerson(personData: {name:"Peter", age: 24}) {
        person {
            name,
            age
        }
    }
}
```

InputObjectTypes can also be fields of InputObjectTypes allowing you to have as complex of input data as you need

```python
import graphene

class LatLngInput(graphene.InputObjectType):
    lat = graphene.Float()
    lng = graphene.Float()

#A location has a latlng associated to it
class LocationInput(graphene.InputObjectType):
```

```
    name = graphene.String()
    latlng = graphene.InputField(LatLngInput)
```

Execution

Executing a query

For executing a query a schema, you can directly call the `execute` method on it.

```
schema = graphene.Schema(...)
result = schema.execute('{ name }')
```

`result` represents the result of execution. `result.data` is the result of executing the query, `result.errors` is `None` if no errors occurred, and is a non-empty list if an error occurred.

Context

You can pass context to a query via `context_value`.

```python
class Query(graphene.ObjectType):
    name = graphene.String()

    def resolve_name(self, args, context, info):
        return context.get('name')

schema = graphene.Schema(Query)
result = schema.execute('{ name }', context_value={'name': 'Syrus'})
```

Middleware

You can use `middleware` to affect the evaluation of fields in your schema.

A middleware is any object that responds to `resolve(*args, next_middleware)`.

Inside that method, it should either:

- Send `resolve` to the next middleware to continue the evaluation; or

- Return a value to end the evaluation early.

Resolve arguments

Middlewares `resolve` is invoked with several arguments:

- `next` represents the execution chain. Call `next` to continue evalution.

- `root` is the root value object passed throughout the query.

- `args` is the hash of arguments passed to the field.

- `context` is the context object passed throughout the query.

- `info` is the resolver info.

Example

This middleware only continues evaluation if the `field_name` is not `'user'`

```python
class AuthorizationMiddleware(object):
    def resolve(self, next, root, args, context, info):
        if info.field_name == 'user':
            return None
        return next(root, args, context, info)
```

And then execute it with:

```python
result = schema.execute('THE QUERY', middleware=[AuthorizationMiddleware()])
```

Dataloader

DataLoader is a generic utility to be used as part of your application's data fetching layer to provide a simplified and consistent API over various remote data sources such as databases or web services via batching and caching.

Batching

Batching is not an advanced feature, it's DataLoader's primary feature. Create loaders by providing a batch loading function.

```python
from promise import Promise
from promise.dataloader import DataLoader

class UserLoader(DataLoader):
    def batch_load_fn(self, keys):
        # Here we return a promise that will result on the
        # corresponding user for each key in keys
        return Promise.resolve([get_user(id=key) for key in keys])
```

A batch loading function accepts an list of keys, and returns a `Promise` which resolves to an list of `values`.

Then load individual values from the loader. `DataLoader` will coalesce all individual loads which occur within a single frame of execution (executed once the wrapping promise is resolved) and then call your batch function with all requested keys.

```
user_loader = UserLoader()

user_loader.load(1).then(lambda user: user_loader.load(user.best_friend_id))

user_loader.load(2).then(lambda user: user_loader.load(user.best_friend_id))
```

A naive application may have issued *four* round-trips to a backend for the required information, but with `DataLoader` this application will make at most *two*.

`DataLoader` allows you to decouple unrelated parts of your application without sacrificing the performance of batch data-loading. While the loader presents an API that loads individual values, all concurrent requests will be coalesced and presented to your batch loading function. This allows your application to safely distribute data fetching requirements throughout your application and maintain minimal outgoing data requests.

Using with Graphene

DataLoader pairs nicely well with Graphene/GraphQL. GraphQL fields are designed to be stand-alone functions. Without a caching or batching mechanism, it's easy for a naive GraphQL server to issue new database requests each time a field is resolved.

Consider the following GraphQL request:

```
{
  me {
    name
    bestFriend {
      name
    }
    friends(first: 5) {
      name
      bestFriend {
        name
      }
    }
  }
}
```

Naively, if `me`, `bestFriend` and `friends` each need to request the backend, there could be at most 13 database requests!

When using DataLoader, we could define the User type using our previous example with leaner code and at most 4 database requests, and possibly fewer if there are cache hits.

```
class User(graphene.ObjectType):
    name = graphene.String()
    best_friend = graphene.Field(lambda: User)
    friends = graphene.List(lambda: User)

    def resolve_best_friend(self, args, context, info):
        return user_loader.load(self.best_friend_id)
```

```
    def resolve_friends(self, args, context, info):
        return user_loader.load_many(self.friend_ids)
```

Relay

Graphene has complete support for Relay and offers some utils to make integration from Python easy.

Nodes

A `Node` is an Interface provided by `graphene.relay` that contains a single field `id` (which is a `ID!`). Any object that inherits from it has to implement a `get_node` method for retrieving a `Node` by an *id*.

Quick example

Example usage (taken from the Starwars Relay example):

```python
class Ship(graphene.ObjectType):
    '''A ship in the Star Wars saga'''
    class Meta:
        interfaces = (relay.Node, )

    name = graphene.String(description='The name of the ship.')

    @classmethod
    def get_node(cls, id, context, info):
        return get_ship(id)
```

The `id` returned by the `Ship` type when you query it will be a scalar which contains enough info for the server to know its type and its id.

For example, the instance `Ship(id=1)` will return `U2hpcDox` as the id when you query it (which is the base64 encoding of `Ship:1`), and which could be useful later if we want to query a node by its id.

Custom Nodes

You can use the predefined `relay.Node` or you can subclass it, defining custom ways of how a node id is encoded (using the `to_global_id` method in the class) or how we can retrieve a Node given a encoded id (with the `get_node_from_global_id` method).

Example of a custom node:

```python
class CustomNode(Node):

    class Meta:
        name = 'Node'

    @staticmethod
    def to_global_id(type, id):
        return '{}:{}'.format(type, id)

    @staticmethod
    def get_node_from_global_id(global_id, context, info, only_type=None):
        type, id = global_id.split(':')
        if only_node:
            # We assure that the node type that we want to retrieve
            # is the same that was indicated in the field type
            assert type == only_node._meta.name, 'Received not compatible node.'

        if type == 'User':
            return get_user(id)
        elif type == 'Photo':
            return get_photo(id)
```

The `get_node_from_global_id` method will be called when `CustomNode.Field` is resolved.

Accessing node types

If we want to retrieve node instances from a `global_id` (scalar that identifies an instance by it's type name and id), we can simply do `Node.get_node_from_global_id(global_id, context, info)`.

In the case we want to restrict the instance retrieval to a specific type, we can do: `Node.get_node_from_global_id(global_id, context, info, only_type=Ship)`. This will raise an error if the `global_id` doesn't correspond to a Ship type.

Node Root field

As is required in the Relay specification, the server must implement a root field called `node` that returns a `Node` Interface.

For this reason, `graphene` provides the field `relay.Node.Field`, which links to any type in the Schema which implements `Node`. Example usage:

```python
class Query(graphene.ObjectType):
    # Should be CustomNode.Field() if we want to use our custom Node
    node = relay.Node.Field()
```

Connection

A connection is a vitaminized version of a List that provides ways of slicing and paginating through it. The way you create Connection types in `graphene` is using `relay.Connection` and `relay.ConnectionField`.

Quick example

If we want to create a custom Connection on a given node, we have to subclass the `Connection` class.

In the following example, `extra` will be an extra field in the connection, and `other` an extra field in the Connection Edge.

```python
class ShipConnection(Connection):
    extra = String()

    class Meta:
        node = Ship

    class Edge:
        other = String()
```

The `ShipConnection` connection class, will have automatically a `pageInfo` field, and a `edges` field (which is a list of `ShipConnection.Edge`). This `Edge` will have a `node` field linking to the specified node (in `ShipConnection.Meta`) and the field `other` that we defined in the class.

Connection Field

You can create connection fields in any Connection, in case any ObjectType that implements `Node` will have a default Connection.

```python
class Faction(graphene.ObjectType):
    name = graphene.String()
    ships = relay.ConnectionField(ShipConnection)

    def resolve_ships(self, args, context, info):
        return []
```

Mutations

Most APIs don't just allow you to read data, they also allow you to write.

In GraphQL, this is done using mutations. Just like queries, Relay puts some additional requirements on mutations, but Graphene nicely manages that for you. All you need to do is make your mutation a subclass of `relay.ClientIDMutation`.

```python
class IntroduceShip(relay.ClientIDMutation):

    class Input:
        ship_name = graphene.String(required=True)
        faction_id = graphene.String(required=True)

    ship = graphene.Field(Ship)
    faction = graphene.Field(Faction)
```

```
    @classmethod
    def mutate_and_get_payload(cls, input, context, info):
        ship_name = input.get('ship_name')
        faction_id = input.get('faction_id')
        ship = create_ship(ship_name, faction_id)
        faction = get_faction(faction_id)
        return IntroduceShip(ship=ship, faction=faction)
```

Accepting Files

Mutations can also accept files, that's how it will work with different integrations:

```python
class UploadFile(graphene.ClientIDMutation):
    class Input:
        pass
        # nothing needed for uploading file

    # your return fields
    success = graphene.String()

    @classmethod
    def mutate_and_get_payload(cls, input, context, info):
        # When using it in Django, context will be the request
        files = context.FILES
        # Or, if used in Flask, context will be the flask global request
        # files = context.files

        # do something with files

        return UploadFile(success=True)
```

Useful links

- Getting started with Relay
- Relay Global Identification Specification
- Relay Cursor Connection Specification
- Relay input Object Mutation

Testing in Graphene

Automated testing is an extremely useful bug-killing tool for the modern developer. You can use a collection of tests – a test suite – to solve, or avoid, a number of problems:

- When you're writing new code, you can use tests to validate your code works as expected.

- When you're refactoring or modifying old code, you can use tests to ensure your changes haven't affected your application's behavior unexpectedly.

Testing a GraphQL application is a complex task, because a GraphQL application is made of several layers of logic – schema definition, schema validation, permissions and field resolution.

With Graphene test-execution framework and assorted utilities, you can simulate GraphQL requests, execute mutations, inspect your application's output and generally verify your code is doing what it should be doing.

Testing tools

Graphene provides a small set of tools that come in handy when writing tests.

Test Client

The test client is a Python class that acts as a dummy GraphQL client, allowing you to test your views and interact with your Graphene-powered application programmatically.

Some of the things you can do with the test client are:

- Simulate Queries and Mutations and observe the response.

- Test that a given query request is rendered by a given Django template, with a template context that contains certain values.

Overview and a quick example

To use the test client, instantiate `graphene.test.Client` and retrieve GraphQL responses:

```python
from graphene.test import Client

def test_hey():
    client = Client(my_schema)
    executed = client.execute('''{ hey }''')
    assert executed == {
        'data': {
            'hey': 'hello!'
        }
    }
```

Execute parameters

You can also add extra keyword arguments to the `execute` method, such as `context_value`, `root_value`, `variable_values`, ...:

```python
from graphene.test import Client

def test_hey():
    client = Client(my_schema)
    executed = client.execute('''{ hey }''', context_value={'user': 'Peter'})
    assert executed == {
        'data': {
            'hey': 'hello Peter!'
        }
    }
```

Snapshot testing

As our APIs evolve, we need to know when our changes introduce any breaking changes that might break some of the clients of our GraphQL app.

However, writing tests and replicate the same response we expect from our GraphQL application can be tedious and repetitive task, and sometimes it's easier to skip this process.

Because of that, we recommend the usage of SnapshotTest.

SnapshotTest let us write all this tests in a breeze, as creates automatically the `snapshots` for us the first time the test is executed.

Here is a simple example on how our tests will look if we use `pytest`:

```python
def test_hey(snapshot):
    client = Client(my_schema)
    # This will create a snapshot dir and a snapshot file
    # the first time the test is executed, with the response
    # of the execution.
    snapshot.assert_match(client.execute('''{ hey }'''))
```

If we are using `unittest`:

```
from snapshottest import TestCase

class APITestCase(TestCase):
    def test_api_me(self):
        """Testing the API for /me"""
        client = Client(my_schema)
        self.assertMatchSnapshot(client.execute('''{ hey }'''))
```

Integrations

- Graphene-Django (source)
- Graphene-SQLAlchemy (source)
- Graphene-GAE (source)