Watson - DB

Release 2.0.0

September 30, 2014

Contents

SqlAlchemy integration for Watson-Framework.

CHAPTER 1

Build Status

CHAPTER 2

Dependencies

- watson-framework
- sqlalchemy
- alembic

CHAPTER 3

Installation

```
pip install watson-db
```

CHAPTER 4

Testing

Watson can be tested with py.test. Simply activate your virtualenv and run `python setup.py test`.

CHAPTER 5

Contributing

If you would like to contribute to Watson, please feel free to issue a pull request via Github with the associated tests for your code. Your name will be added to the AUTHORS file under contributors.

CHAPTER 6

Table of Contents

6.1 Usage

6.1.1 Configuration

Before being able to integrate SqlAlchemy with Watson, there are a few things that must be implemented first within your applications config.

1. Add the init event to your applications configuration.

```
'events': {
     events.INIT: [
         ('watson.db.listeners.Init', 1, True)
     ],
 }
```

2. Create a default configuration for a database session.

```
db = {
    'connections': {
      'default': {
          'connection_string': 'sqlite:///:memory:',
          'engine_options': {},
          'session_options': {}
        }
     }
 }
```

`engine_options` and `session_options` are optional values and can contain any kwarg values that `create_session` and `sessionmaker` from SqlAlchemy take.

A full example configuration might look like this:

```
db = {
   'connections': {
        'default': {
            'connection_string': 'sqlite:///../data/db/default.db',
            'metadata': 'app.models.Base',
            'engine_options': {
                'encoding': 'utf-8',
                'echo': False,
                'pool_recycle': 3600
            }
```

```
        },
    },
    'migrations': {
        'path': '../data/db/migrations',
        'use_twophase': False
    },
    'fixtures': {
        'path': '../data/db/fixtures',
        'data': (
            # Fixtures will be located in ../data/db/fixtures/model.json
            # and will be inserted into the 'default' database.
            ('model', None),
        )
    }
}
```

6.1.2 Fixtures

Fixtures are a way of inserting some initial data into a database to populate it. They are stored in basic JSON format, and can be defined as follows.

```
[
    {
        "class": "app.models.Model",
        "fields": {
            "id": 1,
            "column": "Value"
        }
    }
    // .. more records
]
```

Each fixture that is to be loaded via the *populate* command should be included in the `data` value of the `fixtures` in the format `(FIXTURE_NAME, DATABASE_CONNECTION_NAME)`. If DATABASE_CONNECTION_NAME is set to None, then the default connection will be used.

6.1.3 Migrations

Watson DB utilizes Alembic to handle migrations, which can be run via the command line. See the commands section of this document for more information on the individual commands.

6.1.4 Commands

The commands available to you are split into two namespaces, db, and db:migrate. These can be accessed via `./console.py db` and `./console.py db:migrate` respectively.

db

create

Creates the databases against the associated model metadata and connections.

dump

Prints out the SQL statements used to create the database.

populate

Inserts the data from the fixtures into the databases.

db:migrate

These commands are essentially wrappers to the Alembic command line. Additional arguments that can be specified can be found by appending –help to the command.

branches

current

downgrade

history

init

revision

stamp

upgrade

6.1.5 Services

Services provide a straightforward way to interact with the models in your application without having to directly call against the SqlAlchemy session itself. Each service should be defined within the configuration to use the relevant SqlAlchemy session in it's constructor.

```
dependencies = {
    'definitions': {
        'myservice': {
            'item': 'myapp.services.MyService',
            'init': ['sqlalchemy_session_default']
        },
        'mycontroller': {
            'item': 'myapp.controllers.MyController',
            'property': {
                'service': 'myservice'
            }
        }
    }
}
```

6.1.6 Example

Once configured, the session can be retrieved from the container via `container.get('sqlalchemy_session_[SESSION_NAM`

watson.db also provides a paginator class for paginating a set of results back from SQLAlchemy. Basic usage includes:

```
# within myapp.models
from watson.db import models

class MyModel(models.Model):
    # .. columns
```

```
# within myapp.services
from watson.db import services
from myapp import models

class MyService(services.Base):
    __model__ = models.MyModel

# within myapp.controllers, assuming the MyService object has
# been injected into the controller as the 'service' attribute.
from watson.db import utils
from watson.framework import controllers

class MyController(controllers.Rest):
    def GET(self):
        return {
            'paginator': utils.Pagination(self.service.query, limit=50)
        }

# within view
{% for item in paginator %}
{% endfor %}
<div class="pagination">
{% for page in paginator.iter_pages() %}
    {% if page == paginator.page %}
    <a href="#" class="current">{{ page }}</a>
    {% else %}
    <a href="#">{{ page }}</a>
    {% endif %}
{% endfor %}
</div>
```

6.2 Reference Library

6.2.1 watson.db.commands

class `watson.db.commands.`**`Database`**(*config*)

Database commands.

`_session_or_engine`(*type_*)

Retrieves all the sessions or engines from the container.

`create`(*drop*)

Create the relevant databases.

`dump`()

Print the Schema of the database.

`populate`()

Add data from fixtures to the database(s).

class `watson.db.commands.`**`Migrate`**(*config*)

Alembic integration with Watson.

`branches`()

Show current un-spliced branch points.

current()
Display the current revision for each database.

downgrade(*sql=False*, *tag=None*, *revision='-1'*)
Revert to a previous version.

history(*rev_range*)
List changeset scripts in chronological order.

Parameters **rev_range** – Revision range in format [start]:[end]

init()
Initializes Alembic migrations for the project.

revision(*sql=False*, *autogenerate=False*, *message=None*)
Create a new revision file.

Parameters

- **sql** (*bool*) – Don't emit SQL to database - dump to standard output instead
- **autogenerate** (*bool*) – Populate revision script with andidate migration operatons, based on comparison of database to model
- **message** (*string*) – Message string to use with 'revision'

stamp(*sql=False*, *tag=None*, *revision='head'*)
'stamp' the revision table with the given revision; don't run any migrations.

Parameters

- **sql** (*bool*) – Don't emit SQL to database - dump to standard output instead
- **tag** (*string*) – Arbitrary 'tag' name - can be used by custom env.py scripts
- **revision** (*string*) – Revision identifier

upgrade(*sql=False*, *tag=None*, *revision='head'*)
Upgrade to a later version.

Parameters

- **sql** (*bool*) – Don't emit SQL to database - dump to standard output instead
- **tag** (*string*) – Arbitrary 'tag' name - can be used by custom env.py scripts
- **revision** (*string*) – Revision identifier

6.2.2 watson.db.contextmanagers

watson.db.contextmanagers.**transaction_scope**(*session*, *should_close=False*)
Provides a transactional scope for session calls.

See:

- http://docs.sqlalchemy.org/en/latest/orm/session.html

Example:

```
class MyController(controllers.Rest):

    def GET(self):
        with transaction_scope(self.db):
            session.add(Model())
```

6.2.3 watson.db.engine

`watson.db.engine.`**`create_db`**(*engine*, *model*, *drop=False*)

Creates a new database on the given engine based on the models metadata.

Parameters

- **engine** (*Engine*) – A SQLAlchemy engine object
- **model** (*object*) – The model base containing the associated metadata.

`watson.db.engine.`**`make_engine`**(***kwargs*)

Create a new engine for SqlAlchemy.

Remove the container argument that is sent through from the DI container.

6.2.4 watson.db.fixtures

6.2.5 watson.db.listeners

class `watson.db.listeners.`**`Complete`**

Cleanups the db session at the end of each request.

class `watson.db.listeners.`**`Init`**

Bootstraps watson.db into the event system of watson.

Each session and engine can be retrieved from the container by using sqlalchemy_engine_[name of engine] and sqlalchemy_session_[name of session] respectively.

`_load_default_commands`(*config*)

Load some existing

`_validate_config`(*config*)

Validates the config and sets some standard defaults.

6.2.6 watson.db.meta

class `watson.db.meta.`**`_DeclarativeMeta`**(*classname*, *bases*, *dict_*)

Responsible for automatically assigning a tablename to a model.

Tablenames will be pluralized.

6.2.7 watson.db.models

6.2.8 watson.db.panels

6.2.9 watson.db.session

6.2.10 watson.db.services

class `watson.db.services.`**`Base`**(*session*)

Provides common interactions with the SQLAlchemy session.

Example:

```
class MyService(Base):
    __model__ = models.MyModel

# sqlalchemy_session is a reference to a Session object
service = MyService(sqlalchemy_session)
mymodel = service.new(attr='Value')
print(mymodel.attr)  # 'Value'
service.save(mymodel)
```

session
Session

The SqlAlchemy session

__model__
mixed

The model object the server interacts with

__init__(*session*)

Parameters **session** (*Session*) – The SqlAlchemy session

all()

Returns A list of all model objects

Return type list

count()
Returns the total number of model objects

Returns int

delete(*model*)
Deletes a model from the database.

Parameters **model** (*mixed*) – The model to delete

delete_all(**models*)
Delete a list of models.

If deleting more than a single model of the same type, then a Service.find(**kwargs).delete() should be called (and wrapped in a transaction_scope) instead.

Parameters **models** (*list*) – The models to delete

find(***kwargs*)
Shorthand for the filter_by method.

Should be used when performing query specific operations (such as bulk deletion)

Parameters **kwargs** – The fields to search for

first(***kwargs*)
Return the first matching result for the query.

Parameters **kwargs** – The fields to search for

Returns The object found, or None if nothing was returned

Return type mixed

get(*id*, *error_on_not_found=False*)
Retrieve a single object based on it's ID.

Parameters

- **id** (*int*) – The primary key of the record
- **error_on_not_found** (*bool*) – Raise an exception if not found

Returns The matching model object

Return type mixed

Raises **Exception when no matching results are found.** –

new(***kwargs*)

Creates a new instance of the model object

Parameters **kwargs** (*mixed*) – The initial values for the model

Returns The newly created model

Return type mixed

save(*model*)

Add the model to the session and save it to the database.

Parameters **model** (*mixed*) – The object to save

Returns The saved model

Return type mixed

save_all(**models*)

Save a list of models.

Parameters **models** (*list, tuple*) – The models to save

6.2.11 watson.db.utils

class watson.db.utils.**Pagination**(*query*, *page=1*, *limit=20*)

Provides simple pagination for query results.

query

Query

The SQLAlchemy query to be paginated

page

int

The page to be displayed

limit

int

The maximum number of results to be displayed on a page

total

int

The total number of results

items

list

The items returned from the query

Example:

```
# within controller
query = session.query(Model)
paginator = Pagination(query, limit=50)

# within view
{% for item in paginator %}
{% endfor %}
<div class="pagination">
{% for page in paginator.iter_pages() %}
    {% if page == paginator.page %}
    <a href="#" class="current">{{ page }}</a>
    {% else %}
    <a href="#">{{ page }}</a>
    {% endif %}
{% endfor %}
</div>
```

__init__(*query*, *page=1*, *limit=20*)

Initialize the paginator and set some default values.

has_next

Return whether or not there are more pages from the currently displayed page.

Returns boolean

has_previous

Return whether or not there are previous pages from the currently displayed page.

Returns boolean

iter_pages()

An iterable containing the number of pages to be displayed.

Example:

```
{% for page in paginator.iter_pages() %}{% endfor %}
```

pages

The total amount of pages to be displayed based on the number of results and the limit being displayed.

Returns int

Python Module Index

W