

---

# **flywheel**

***Release 0.5.2***

**Mar 18, 2017**



---

## Contents

---

<b>1</b>	<b>User Guide</b>	<b>3</b>
1.1	Getting Started . . . . .	3
1.2	Models . . . . .	4
1.3	Table Queries . . . . .	10
1.4	CRUD . . . . .	13
1.5	Developing . . . . .	16
1.6	Changelog . . . . .	16
<b>2</b>	<b>API Reference</b>	<b>21</b>
2.1	flywheel package . . . . .	21
<b>3</b>	<b>Indices and tables</b>	<b>47</b>
	<b>Python Module Index</b>	<b>49</b>



Flywheel is a library for mapping python objects to DynamoDB tables. It uses a SQLAlchemy-like syntax for queries.  
Code lives here: <https://github.com/stevearc/flywheel>



## Getting Started

Flywheel can be installed with pip

```
pip install flywheel
```

Here are the steps to set up a simple example model with flywheel:

```
# Take care of some imports
from datetime import datetime
from flywheel import Model, Field, Engine

# Set up our data model
class Tweet(Model):
    userid = Field(hash_key=True)
    id = Field(range_key=True)
    ts = Field(type=datetime, index='ts-index')
    text = Field()

    def __init__(self, userid, id, ts, text):
        self.userid = userid
        self.id = id
        self.ts = ts
        self.text = text

# Create an engine and connect to an AWS region
engine = Engine()
engine.connect_to_region('us-east-1')

# Register our model with the engine so it can create the Dynamo table
engine.register(Tweet)

# Create the dynamo table for our registered model
engine.create_schema()
```

---

Now that you have your model, your engine, and the Dynamo table, you can begin adding tweets:

```
tweet = Tweet('myuser', '1234', datetime.utcnow(), text='@awscloud hey '
              'I found this cool new python library for AWS...')
engine.save(tweet)
```

To get data back out, query it using the engine:

```
# Get the 10 most recent tweets by 'myuser'
recent = engine.query(Tweet)\
    .filter(Tweet.ts <= datetime.utcnow(), userid='myuser')\
    .limit(10).all(desc=True)

# Get a specific tweet by a user
tweet = engine.query(Tweet).filter(userid='myuser', id='1234').first()
```

If you want to change a field, just make the change and sync it:

```
tweet.text = 'This tweet has been removed due to shameless promotion'
tweet.sync()
```

That's enough to give you a taste. The rest of the docs have more information on *creating models*, *writing queries*, or *how updates work*.

## Models

### Model Basics

This is what a model looks like:

```
class Tweet(Model):
    userid = Field(hash_key=True)
    id = Field(range_key=True)
    ts = Field(type=datetime, index='ts-index')
    text = Field()
```

The model declares the fields an object has, their *data types*, and the *schema* of the table.

Since models define the schema of a table, you can use them to create or delete tables. Every model has a `meta_` field attached to it which contains metadata about the model. This metadata object has the `create` and `delete` methods.

```
from dynamo3 import DynamoDBConnection

connection = DynamoDBConnection.connect_to_region('us-east-1')
Tweet.meta_.create_dynamo_schema(connection)
Tweet.meta_.delete_dynamo_schema(connection)
```

You can also register your models with the engine and create all the tables at once:

```
engine.register(User, Tweet, Message)
engine.create_schema()
```



## Data Types

DynamoDB supports three different data types: `STRING`, `NUMBER`, and `BINARY`. It also supports sets of these types: `STRING_SET`, `NUMBER_SET`, `BINARY_SET`.

You can use these values directly for the model declarations, though they require an import:

```
from flywheel import Model, Field, STRING, NUMBER

class Tweet(Model):
    userid = Field(type=STRING, hash_key=True)
    id = Field(type=STRING, range_key=True)
    ts = Field(type=NUMBER, index='ts-index')
    text = Field(type=STRING)
```

There are other settings for type that are represented by python primitives. Some of them (like `unicode`) are functionally equivalent to the DynamoDB option (`STRING`). Others, like `int`, enforce an additional application-level constraint on the data. Each option works transparently, so a `datetime` field would be set with `datetime` objects and you could query against it using other `datetime`'s.

Below is a table of python types, how they are stored in DynamoDB, and any special notes. For more information, the code for data types is located in [types](#).

PY2 Type	PY3 Type	Dynamo Type	Notes
unicode	str	STRING	Basic STRING type. This is the default for fields
str	bytes	BINARY	Binary data, (serialized objects, compressed data, etc)
int/long	int	NUMBER	Enforces integer constraint on data
float		NUMBER	
Decimal		NUMBER	
set		*_SET	This will use the appropriate type of DynamoDB set
bool		BOOL	
datetime		NUMBER	Stored with UTC timezone. See <a href="#">DateTimeType</a> for more.
date		NUMBER	
dict		MAP	
list		LIST	

If you attempt to set a field with a type that doesn't match, it will raise a `TypeError`. If a field was created with `coerce=True` it will first attempt to convert the value to the correct type. This means you could set an `int` field with the value `"123"` and it would perform the conversion for you.

**Note:** Certain fields will auto-coerce specific data types. For example, a `bytes` field will auto-encode a `unicode` to `utf-8` even if `coerce=False`. Similarly, a `unicode` field will auto-decode a `bytes` value to a `unicode` string.

**Warning:** If an `int` field is set to `coerce` values, it will still refuse to drop floating point data. This has the following effect:

```
>>> class Game(Model):
...     title = Field(hash_key=True)
...     points = Field(type=int, coerce=True)

>>> mygame = Game()
>>> mygame.points = 1.8
ValueError: Field 'points' refusing to convert 1.8 to int! Results in data loss!
```

## Set types

If you define a `set` field with no additional parameters `Field(type=set)`, flywheel will ensure that the field is a set, but will perform no type checking on the items within the set. This should work fine for basic uses when you are storing a number or string, but sets are able to contain any data type listed in the table above (and any *custom type* you declare). All you have to do is specify it in the `type` like so:

```
from flywheel import Model, Field, set_
from datetime import date

class Location(Model):
    name = Field(hash_key=True)
    events = Field(type=set_(date))
```

If you don't want to import `set_`, you can use an equivalent expression with the python `frozenset` builtin:

```
events = Field(type=frozenset([date]))
```

## Field Validation

You can apply one or more validators to a field. These are functions that enforce some constraint on the field value beyond the type. Unlike the type checking done above, the validation checks are only run when saving to the database. An example:

```
class Widget(Model):
    id = Field(type=int, check=lambda x: x > 0)
```

To apply multiple validation checks, pass them in as a list or tuple:

```
def is_odd(x):
    return x % 2 == 1

def is_natural(x):
    return x >= 0

class Widget(Model):
    odd_natural_num = Field(type=int, check=(is_odd, is_natural))
```

There is a special case for enforcing that a field is non-null, since it is a common case:

```
username = Field(nullable=False)
```

The `nullable=False` will generate an additional check to make sure the value is non-null.

## Custom Types

You can define your own custom data types and make them available across all of your models. All you need to do is create a subclass of `TypeDefinition`. Let's make a type that will store any python object in pickled format.

```
from flywheel.fields.types import TypeDefinition, BINARY, Binary
import cPickle as pickle

class PickleType(TypeDefinition):
    type = pickle # name you use to reference this type
    aliases = ['pickle'] # alternate names that reference this type
```

```

ddb_data_type = BINARY # data type of the field in dynamo

def coerce(self, value, force):
    # Perform no type checking because we can pickle ANYTHING
    return value

def ddb_dump(self, value):
    # Pickle and convert to a Binary object
    return Binary(pickle.dumps(value))

def ddb_load(self, value):
    # Convert from a Binary object and unpickle
    return pickle.loads(value.value)

```

Now that you have your type definition, you can either use it directly in your code:

```

class MyModel(Model):
    myobj = Field(type=PickleType())

```

Or you can register it globally and reference it by its type or any aliases that were defined.

```

from flywheel.fields.types import register_type

register_type(PickleType)

class MyModel(Model):
    myobj = Field(type='pickle')

```

## Schema

There are four main key concepts to understanding a DynamoDB table.

**Hash key:** This field will be sharded. Pick something with relatively random access (e.g. userid is good, timestamp is bad)

**Range key:** Optional. This field will be indexed, so you can query against it (within a specific hash key).

The hash key and range key together make the **Primary key**, which is the unique identifier for each object.

**Local Secondary Indexes:** Optional, up to 5. You may only use these if your table has a range key. These fields are indexed in a similar fashion as the range key. You may also query against them within a specific hash key. You can think of these as range keys with no uniqueness requirements.

**Global Secondary Indexes:** Optional, up to 5. These indexes have a hash key and optional range key, and can be put on any declared field. This allows you to shard your tables by more than one value.

For additional information on table design, read the [AWS docs on best practices](#)

Example declaration of hash and range key:

```

class Tweet(Model):
    userid = Field(hash_key=True)
    ts = Field(type=datetime, range_key=True)

```

For this version of a Tweet, each (userid, ts) pair is a unique value. The Dynamo table will be sharded across userids.

## Local Secondary Indexes

Indexes also have a Projection Type. Creating an index requires duplicating some amount of data in the storage, and the projection type allows you to optimize how much additional storage is used. The projection types are:

**All:** All fields are projected into the index

**Keys only:** Only the primary key and indexed keys are projected into the index

**Include:** Like the “keys only” projection, but allows you to specify additional fields to project into the index

This is how they it looks in the model declaration:

```
class Tweet(Model):
    userid = Field(hash_key=True)
    id = Field(range_key=True)
    ts = Field(type=datetime).all_index('ts-index')
    retweets = Field(type=int).keys_index('rt-index')
    likes = Field(type=int).include_index('like-index', ['text'])
    text = Field()
```

The default index projection is “All”, so you could replace the `ts` field above with:

```
ts = Field(type=datetime, index='ts-index')
```

## Global Secondary Indexes

Like their Local counterparts, Global Secondary Indexes can specify a projection type. Unlike their Local counterparts, Global Secondary Indexes are provisioned with a *separate* read/write throughput from the base table. This can be specified in the model declaration. Here are some examples below:

```
class Tweet(Model):
    __metadata__ = {
        'global_indexes': [
            GlobalIndex.all('ts-index', 'city', 'ts').throughput(read=10, write=2),
            GlobalIndex.keys('rt-index', 'city', 'retweets')\
                .throughput(read=10, write=2),
            GlobalIndex.include('like-index', 'city', 'likes',
                               includes=['text']).throughput(read=10, write=2),
        ],
    }
    userid = Field(hash_key=True)
    city = Field()
    id = Field(range_key=True)
    ts = Field(type=datetime)
    retweets = Field(type=int)
    likes = Field(type=int)
    text = Field()
```

If you want more on indexes, check out the [AWS docs on indexes](#).

## Composite Fields

Composite fields allow you to create fields that are combinations of multiple other fields. Suppose you’re creating a table where you plan to store a collection of social media items (tweets, facebook posts, instagram pics, etc). If you make the hash key the id of the item, there is the remote possibility that a tweet id will collide with a facebook id. Here is the solution:

```
class SocialMediaItem(Model):
    userid = Field(hash_key=True)
    type = Field()
    id = Field()
    uid = Composite('type', 'id', range_key=True)
```

This will automatically generate a `uid` field from the values of `type` and `id`. For example:

```
>>> item = SocialMediaItem(type='facebook', id='12345')
>>> print item.uid
facebook:12345
```

Note that setting a Composite field just doesn't work:

```
>>> item.uid = 'ILikeThisIDBetter'
>>> print item.uid
facebook:12345
```

By default, a Composite field simply joins its subfields with a `' : '`. You can change that behavior for fancier applications:

```
def score_merge(likes, replies, deleted):
    if deleted:
        return None
    return likes + 5 * replies

class Post(Model):
    userid = Field(hash_key=True)
    id = Field(range_key=True)
    likes = Field(type=int)
    replies = Field(type=int)
    deleted = Field(type=bool)
    score = Composite('likes', 'replies', 'deleted', type=int,
                      merge=score_merge, index='score-index')
```

So now you can update the `likes` or `replies` count, and the `score` will automatically change. Which will re-arrange it in the index that you created. Then, if you mark the post as “deleted”, it will remove the `score` field which removes it from the index.

*Whoooooaaahh...*

The last neat little thing about Composite fields is how you can query them. For numeric Composite fields you probably want to query directly on the `score` like any other field. But if you're merging strings like with `SocialMediaItem`, it can be cleaner to refer to the component fields themselves:

```
>>> fb_post = engine.query(SocialMediaItem).filter(userid='abc123',
...         type='facebook', id='12345').first()
```

The engine will automatically detect that you're trying to query on the range key, and construct the `uid` from the pieces you provided.

## Metadata

Part of the model declaration is the `__metadata__` attribute, which is a dict that configures the `Model.meta_` object. Models will inherit and merge the `__metadata__` fields from their ancestors. Keys that begin with an underscore will not be merged. For example:

```
class Vehicle(Model):
    __metadata__ = {
        '_name': 'all-vehicles',
        'throughput': {
            'read': 10,
            'write': 2,
        }
    }

class Car(Vehicle):
    pass
```

```
>>> print Car.__metadata__
{'throughput': {'read': 10, 'write': 2}}
```

Below is a list of all the values that may be set in the `__metadata__` attribute of a model.

Key	Type	Description
<code>_name</code>	str	The name of the DynamoDB table (defaults to class name)
<code>_abstract</code>	bool	If True, no DynamoDB table will be created for this model (useful if you just want a class to inherit from)
<code>throughput</code>	dict	The table read/write throughput (defaults to <code>{'read': 5, 'write': 5}</code> )
<code>global_indexes</code>	list	A list of GlobalIndex objects

## Table Queries

The query syntax is heavily inspired by [SQLAlchemy](#). In DynamoDB, queries must use one of the table's indexes. Queries are constrained to a single hash key value. This means that for a query there will always be at least one call to `filter` which will, at a minimum, set the hash key to search on.

```
# Fetch all tweets made by a user
engine.query(Tweet).filter(Tweet.userid == 'abc123').all()
```

You may also use inequality filters on range keys and secondary indexes

```
# Fetch all tweets made by a user in the past day
earlyts = datetime.utcnow() - timedelta(days=1)
engine.query(Tweet).filter(Tweet.userid == 'abc123',
                           Tweet.ts >= earlyts).all()
```

There are two finalizing statements that will return all results: `all()` and `gen()`. Calling `all()` will return a list of results. Calling `gen()` will return a generator. If your query will return a large number of results, using `gen()` can help you avoid storing them all in memory at the same time.

```
# Count how many retweets a user has in total
retweets = 0
all_tweets = engine.query(Tweet).filter(Tweet.userid == 'abc123').gen()
for tweet in all_tweets:
    retweets += tweet.retweets
```

There are two finalizing statements that retrieve a single item: `first()` and `one()`. Calling `first()` will return the first element of the results, or `None` if there are no results. Calling `one()` will return the first element of the results *only if there is exactly one result*. If there are no results or more results it will raise a `ValueError`.

```
# Get a single tweet by a user
tweet = engine.query(Tweet).filter(Tweet.userid == 'abc123').first()

# Get a specific tweet and fail if missing
tweet = engine.query(Tweet).filter(Tweet.userid == 'abc123',
                                    Tweet.id == '1234').one()
```

There is one more finalizing statement: `count()`. This will return the number of results that matched the query, instead of returning the results themselves.

```
# Get the number of tweets made by user abc123
num = engine.query(Tweet).filter(Tweet.userid == 'abc123').count()
```

You can set a `limit()` on a query to limit the number of results it returns:

```
# Get the first 10 tweets by a user after a timestamp
afterts = datetime.utcnow() - timedelta(hours=1)
tweets = engine.query(Tweet).filter(Tweet.userid == 'abc123',
                                    Tweet.ts >= afterts).limit(10).all()
```

One way to delete items from a table is with a query. Calling `delete()` will delete all items that match a query:

```
# Delete all of a user's tweets older than 1 year
oldts = datetime.utcnow() - timedelta(days=365)
engine.query(Tweet).filter(Tweet.userid == 'abc123',
                           Tweet.ts < oldts).delete()
```

Most of the time the query engine will be able to automatically detect which local or global secondary index you intend to use. If the index is ambiguous, you can manually specify the index. This can also be useful if you want the results to be sorted by a particular index when only querying the hash key.

```
# This is the schema for the following example
class Tweet(Model):
    userid = Field(hash_key=True)
    id = Field(range_key=True)
    ts = Field(type=datetime, index='ts-index')
    retweets = Field(type=int, index='rt-index')

# This returns 10 tweets in id order (more-or-less random)
ten_tweets = engine.query(Tweet).filter(userid='abc123').limit(10).all()

# Get the 10 most retweeted tweets for a user
top_ten = engine.query(Tweet).filter(userid='abc123').index('rt-index')\
    .limit(10).all(desc=True)

# Get The 10 most recent tweets for a user
ten_recent = engine.query(Tweet).filter(userid='abc123').index('ts-index')\
    .limit(10).all(desc=True)
```

### New in 0.2.1

Queries can filter on fields that are not the hash or range key. Filtering this way will strip out the results server-side, but it will not use an index. When filtering on these extra fields, you may use the additional filter operations that are listed under [Table Scans](#).

## Shorthand

If you want to avoid typing ‘query’ everywhere, you can simply call the engine:

```
# Long form query
engine.query(Tweet).filter(Tweet.userid == 'abc123').all()

# Abbreviated query
engine(Tweet).filter(Tweet.userid == 'abc123').all()
```

Filter constraints with == can be instead passed in as keyword arguments:

```
# Abbreviated filter
engine(Tweet).filter(userid='abc123').all()

engine(Tweet).filter(userid='abc123', id='1234').first()
```

You can still pass in other constraints as positional arguments to the same filter:

```
# Multiple filters in same statement
engine(Tweet).filter(Tweet.ts <= earlyts, userid='abc123').all()
```

## Table Scans

Table scans are similar to table queries, but they do not use an index. This means they have to read every item in the table. This is EXTREMELY SLOW. The benefit is that they do not have to filter based on the hash key, and they have a few additional filter arguments that may be used.

```
# Fetch all tweets ever
alltweets = engine.scan(Tweet).gen()

# Fetch all tweets that tag awscloud
tagged = engine.scan(Tweet).filter(Tweet.tags.contains_('awscloud')).all()

# Fetch all tweets with annoying, predictable text
annoying = set(['first post', 'hey guys', 'LOOK AT MY CAT'])
first = engine.scan(Tweets).filter(Tweet.text.in_(annoying)).all()

# Fetch all tweets with a link
linked = engine.scan(Tweet).filter(Tweet.link != None).all()
```

Since table scans don’t use indexes, you can filter on fields that are not declared in the model. Here are some examples:

```
# Fetch all tweets that link to wikipedia
educational = engine.scan(Tweet)\
    .filter(Tweet.field_('link').startswith('http://wikipedia')).all()

# You can also use the keyword arguments to filter
best_tweets = engine.scan(Tweet)\
    .filter(link='http://en.wikipedia.org/wiki/Morgan_freeman').all()
```



## CRUD

This section covers the operations you can do to save, read, update, and delete items from the database. All of these methods exist on the *Engine* object and can be called on one or many items. After being saved-to or loaded-from Dynamo, the items themselves will have these methods attached to them as well. For example, these are both valid:

```
>>> engine.sync(tweet)
>>> tweet.sync()
```

### Save

Save the item to Dynamo. This is intended for new items that were just created and need to be added to the database. If you `save()` an item that already exists in Dynamo, it will raise an exception. You may optionally use `save(overwrite=True)` to instead clobber existing data and write your version of the item to Dynamo.

```
>>> tweet = Tweet()
>>> engine.save(tweet)
>>> tweet.text = "Let's replace the whole item"
>>> tweet.save(overwrite=True)
```

### Refresh

Query dynamo to get the most up-to-date version of a model. Clobbers any existing data on the item. To force a consistent read use `refresh(consistent=True)`.

This call is very useful if you query indexes that use an incomplete projection type. The results won't have all of the item's fields, so you can call `refresh()` to get any attributes that weren't projected onto the index.

```
>>> tweet = engine.query(Tweet).filter(userid='abc123')\
...     .index('ts-index').first(desc=True)
>>> tweet.refresh()
```

### Get

Fetch an item from its primary key fields. This will be faster than a query, but requires you to know the primary keys of all items you want fetched.

```
>>> my_tweet = engine.get(Tweet, userid='abc123', id='1')
```

You can also fetch many at a time:

```
>>> key1 = {'userid': 'abc123', 'id': '1'}
>>> key2 = {'userid': 'abc123', 'id': '2'}
>>> key3 = {'userid': 'abc123', 'id': '3'}
>>> some_tweets = engine.get(Tweet, [key1, key2, key3])
```

### Delete

Deletes an item. You may pass in `delete(raise_on_conflict=True)`, which will only delete the item if none of the values have changed since it was read.

```
>>> tweet = engine.query(Tweet).filter(userid='abc123', id='123').first()
>>> tweet.delete()
```

You may also delete an item from a primary key specification:

```
>>> engine.delete_key(Tweet, userid='abc123', id='1')
```

And you may delete many at once:

```
>>> key1 = {'userid': 'abc123', 'id': '1'}
>>> key2 = {'userid': 'abc123', 'id': '2'}
>>> key3 = {'userid': 'abc123', 'id': '3'}
>>> engine.delete_key(Tweet, [key1, key2, key3])
```

## Sync

Save any fields that have been changed on an item. This will update changed fields in Dynamo and ensure that all fields exactly reflect the item in the database. This is usually used for updates, but it can be used to create new items as well.

```
>>> tweet = Tweet()
>>> engine.sync(tweet)
>>> tweet.text = "Update just this field"
>>> tweet.sync()
```

Models will automatically detect changes to mutable fields, such as `dict`, `list`, and `set`.

```
>>> tweet.tags.add('awscloud')
>>> tweet.sync()
```

Since `sync` does a partial update, it can tolerate concurrent writes of different fields.

```
>>> tweet = engine.query(Tweet).filter(userid='abc123', id='1234').first()
>>> tweet2 = engine.query(Tweet).filter(userid='abc123', id='1234').first()
>>> tweet.author = "The Pope"
>>> tweet.sync()
>>> tweet2.text = "Mo' money mo' problems"
>>> tweet2.sync() # it works!
>>> print tweet2.author
The Pope
```

This “merge” behavior is also what happens when you `sync()` items to create them. If the item to create already exists in Dynamo, that’s fine as long as there are no conflicting fields. Note that this behavior is distinctly different from `save()`, so make sure you pick the right call for your use case.

If you call `sync()` on an object that has not been changed, it is equivalent to calling `refresh()`.

## Safe Sync

If you use `sync(raise_on_conflict=True)`, the `sync` operation will check that the fields that you’re updating have not been changed since you last read them. This is very useful for preventing concurrent writes.

---

**Note:** If you change a key that is part of a *composite field*, flywheel will **force** the sync to raise on conflict. This avoids the risk of corrupting the value of the composite field.

---

## Atomic Increment

DynamoDB supports truly atomic increment/decrement of NUMBER fields. To use this functionality, there is a special call you need to make:

```
>>> # Increment the number of retweets by 1
>>> tweet.incr_(retweets=1)
>>> tweet.sync()
```

BOOM.

---

**Note:** Incrementing a field that is part of a composite field will also force the sync to raise on conflict.

---

## Atomic Add/Remove

DynamoDB also supports truly atomic add/remove to SET fields. To use this functionality, there is another special call:

```
>>> # Add two users to the set of tagged users
>>> tweet.add_(tags=set(['stevearc', 'dsa']))
>>> tweet.sync()
```

And to delete:

```
>>> tweet.remove_(tags='stevearc')
>>> tweet.sync()
```

Note that you can pass in a single value or a set of values to both `add_` and `remove_`.

## Sync-if-Constraints

### New in 0.2.1

You may pass in a list of constraints to check upon sync. If any of the constraints fail, then the sync will not complete. This should be used with `raise_on_conflict=True`. For example:

```
>>> account = engine.get(Account, username='dsa')
>>> account.incr_(moneys=-200)
>>> # atomically remove $200 from DSA's account, iff there is at least $200 to remove.
>>> account.sync(constraints=[Account.moneys >= 200])
```

## Default Conflict Behavior

You can configure the default behavior for each of these endpoints using `default_conflict`. The default setting will cause `sync()` to check for conflicts, `delete()` not to check for conflicts, and `save()` to overwrite existing values. Check the attribute docs for more options. You can, of course, pass in the argument to the calls directly to override this behavior on a case-by-case basis.

## Developing

To get started developing flywheel, run the following command:

```
wget https://raw.githubusercontent.com/stevearc/devbox/0.1.0/devbox/unbox.py && \
python unbox.py git@github.com:stevearc/flywheel
```

This will clone the repository and install the package into a virtualenv

## Running Tests

The command to run tests is `python setup.py nosetests`, or `tox`. Most of these tests require [DynamoDB Local](#). There is a nose plugin that will download and run the DynamoDB Local service during the tests. It requires the java 6/7 runtime, so make sure you have that installed.

## Changelog

### 0.5.2

- Bug fix: Change limit behavior to match docs. `query().limit()` will limit the number of results, `query().scan_limit()` will limit number of items scanned ([issue 57](#))

### 0.5.1

- Feature: Add `update_schema()` method to Engine ([:pr:53](#))

### 0.5.0

- **Breakage:** Removing support for overflow fields. The only fields flywheel will care about now are those that are explicitly set as a `Field()`
- Flywheel no longer forces `raise_on_conflict` to be `True` when you sync changes to fields that are part of a composite field. It is now up to the user to avoid putting their composite fields into an inconsistent state.
- Feature: `sync()` has a new argument, `no_read`, which changes the behavior for syncing models with no changes. Instead of performing a GET, it will leave them as-is. This should make it easier to perform batch syncs without worrying as much about wasted bandwidth on GETs.
- `Field` has renamed the `data_type` argument to `type` (`data_type` will still work)

### 0.4.11

- Bug fix: Boolean overflow fields no longer decoded as decimals ([:pr:46](#))

### 0.4.10

- Feature: Add `exists()` method to Engine ([issue 45](#))

### 0.4.9

- Feature: Add `save()` method to Models ([issue 40](#))
- Feature: Add `update_field()` method to Engine ([issue 43](#))

### 0.4.8

- Bug fix: Bad function call in `index_pk_dict_`

### 0.4.7

- New `index_pk_dict_` method for constructing *exclusive\_start\_key* for index queries ([issue 34](#))

### 0.4.6

- Pass `exclusive_start_key` through to `dynamo3` ([issue 34](#))

### 0.4.5

- Bug fix: Calling `refresh()` could sometimes crash from unordered results.

### 0.4.4

- Bug fix: Mutable field defaults are no longer shared among model instances

### 0.4.3

- Bug fix: Incorrect `ConditionalCheckFailedException` when syncing changes to a Composite field.
- Allow `DateTimeType` to be stored as a naive datetime.

### 0.4.2

- Make the `dict`, `list`, and `bool` types backwards-compatible with the old json-serialized format ([:pr:'24'](#))
- Allow queries to use `in`, `not null`, and a few other constraints that were missing ([commit 8b8854d](#))
- Models are smarter about marking fields as dirty for sync ([issue 26](#))
- Stopped using deprecated `expected` syntax for `dynamo3`

### 0.4.1

- **Warning:** Stored datetime objects will now be timezone-aware ([commit a7c253d](#))
- **Warning:** Stored datetime objects will now keep their microseconds ([commit fffe92c](#))

## 0.4.0

- **Breakage:** Dropping support for python 3.2 due to lack of botocore support
- **Breakage:** Changing the `list`, `dict`, and `bool` data types to use native DynamoDB types instead of JSON serializing
- **Breakage** and bug fix: Fixing serialization of `datetime` and `date` objects (for more info see the commit [\(commit df049af\)](#))
- Feature: Can now do ‘contains’ filters on lists
- Feature: Fields support multiple validation checks
- Feature: Fields have an easy way to enforce non-null values (`nullable=False`)

Data type changes are due to [an update in the DynamoDB API](#)

## 0.3.0

- **Breakage:** Engine namespace is slightly different. If you pass in a string it will be used as the table name prefix with no additional ‘-’ added.

## 0.2.1

- **Breakage:** Certain queries may now require you to specify an index where it was auto-detected before
- Feature: Queries can now filter on non-indexed fields
- Feature: More powerful “sync-if” constraints
- Feature: Can OR together filter constraints in queries

All changes are due to [an update in the DynamoDB API](#)

## 0.2.0

- **Breakage:** Engine no longer accepts boto connections (using `dynamo3` instead)
- **Breakage:** Removing `S3Type` (no longer have boto as dependency)
- Feature: Support Python 3.2 and 3.3
- Feature: `.count()` terminator for queries ([commit bf3261c](#))
- Feature: Can override throughputs in `Engine.create_schema()` ([commit 4d1abe0](#))
- Bug fix: Engine namespace is truly isolated ([commit 3b4fad7](#))

## 0.1.3

- Bug fix: Some queries fail when global index has no range key ([issue 9](#), [commit edce6e2](#))

## 0.1.2

- Bug fix: Field names can begin with an underscore ([commit 637f1ee](#), [issue 7](#))
- Feature: Models have a nice default `__init__` method ([commit 40068c2](#))

### 0.1.1

- Bug fix: Can call `incr_()` on models that have not been saved yet ([commit 0a1990f](#))
- Bug fix: Model comparison with `None` ([commit 374dda1](#))

### 0.1.0

- First public release





### flywheel package

#### Subpackages

##### flywheel.fields package

#### Submodules

##### flywheel.fields.conditions module

Query constraints

**class** flywheel.fields.conditions.**Condition**  
 Bases: `object`

A constraint that will be applied to a query or scan

#### Attributes

<code>eq_fields</code>	(dict) Mapping of field name to field value
<code>fields</code>	(dict) Mapping of field name to (operator, value) tuples
<code>limit</code>	(int or <code>dynamo3.Limit</code> ) Maximum number of results
<code>scan_limit</code>	(int) Maximum number of items to scan in DynamoDB
<code>index_name</code>	(str) Name of index to use for a query

**classmethod** **construct** (*field, op, other*)  
 Create a Condition on a field

**Parameters** **field** : str

Name of the field to constrain

**op** : str

Operator, such as 'eq', 'lt', or 'contains'

**other** : object

The value to constrain the field with

**Returns condition** : *Condition*

**classmethod construct\_index** (*name*)

Force the query to use a certain index

**Parameters name** : str

**Returns condition** : *Condition*

**classmethod construct\_limit** (*count*)

Create a condition that will limit the results to a count

**Parameters count** : int

**Returns condition** : *Condition*

**classmethod construct\_scan\_limit** (*count*)

Create a condition that will limit the number of items scanned

**Parameters count** : int

**Returns condition** : *Condition*

**query\_kwargs** (*model*)

Get the kwargs for doing a table query

**scan\_kwargs** ()

Get the kwargs for doing a table scan

## flywheel.fields.indexes module

Index definitions

**class** flywheel.fields.indexes.**GlobalIndex** (*name*, *hash\_key*, *range\_key=None*)

Bases: *object*

A global index for DynamoDB

**Parameters name** : str

The name of the index

**hash\_key** : str

The name of the field that is the hash key for the index

**range\_key** : str, optional

The name of the field that is the range key for the index

**throughput** : dict, optional

The read/write throughput of this global index. Used when creating a table. Dict has a 'read' and a 'write' key. (Default 5, 5)

**classmethod all** (*name*, *hash\_key*, *range\_key=None*)

Project all attributes into the index

**get\_ddb\_index** (*fields*)

Get the dynamo index class for this GlobalIndex

**classmethod include** (*name, hash\_key, range\_key=None, includes=None*)

Select which attributes to project into the index

**classmethod keys** (*name, hash\_key, range\_key=None*)

Project key attributes into the index

**throughput** (*read=5, write=5*)

Set the index throughput

**Parameters** **read** : int, optional

Amount of read throughput (default 5)

**write** : int, optional

Amount of write throughput (default 5)

## Notes

This is meant to be used as a chain:

```
class MyModel(Model):
    __metadata__ = {
        'global_indexes': [
            GlobalIndex('myindex', 'hkey', 'rkey').throughput(5, 2)
        ]
    }
```

## flywheel.fields.types module

Field type definitions

**class** flywheel.fields.types.**BinaryType**

Bases: *flywheel.fields.types.TypeDefinition*

Binary strings, stored as a str/bytes

**aliases** = ['B', <class 'dynamo3.types.Binary'>]

**coerce** (*value, force*)

**data\_type**

alias of *str*

**ddb\_data\_type** = 'B'

**ddb\_dump** (*value*)

**ddb\_load** (*value*)

**class** flywheel.fields.types.**BoolType**

Bases: *flywheel.fields.types.TypeDefinition*

Boolean type

**coerce** (*value, force*)

**data\_type**

alias of *bool*

**ddb\_data\_type** = 'BOOL'

**class** flywheel.fields.types.**DateTimeType** (*naive=False*)

Bases: *flywheel.fields.types.TypeDefinition*

Datetimes, stored as a unix timestamp

**Parameters** **naive** : bool, optional

If True, will load values from Dynamo with no timezone. If False, will add a UTC timezone. (Default False).

## Notes

If you want to use naive datetimes, you will need to reference the type class directly instead of going through an alias. For example:

```
from flywheel.fields.types import DateTimeType

field = Field(data_type=DateTimeType(naive=True))
```

**data\_type**

alias of datetime

**ddb\_data\_type** = 'N'

**ddb\_dump** (*value*)

**ddb\_load** (*value*)

**class** flywheel.fields.types.**DateType**

Bases: *flywheel.fields.types.TypeDefinition*

Dates, stored as timestamps

**data\_type**

alias of date

**ddb\_data\_type** = 'N'

**ddb\_dump** (*value*)

**ddb\_load** (*value*)

**class** flywheel.fields.types.**DecimalType**

Bases: *flywheel.fields.types.TypeDefinition*

Numerical values that use Decimal in the application layer.

This should be used if you want to work with floats but need the additional precision of the Decimal type.

**coerce** (*value, force*)

**data\_type**

alias of Decimal

**ddb\_data\_type** = 'N'

**class** flywheel.fields.types.**DictType**

Bases: *flywheel.fields.types.TypeDefinition*

Dict type, stored as a map

**coerce** (*value, force*)

```

    data_type
        alias of dict

    ddb_data_type = 'M'

    mutable = True

class flywheel.fields.types.FloatType
    Bases: flywheel.fields.types.TypeDefinition

    Float values

    coerce (value, force)

    data_type
        alias of float

    ddb_data_type = 'N'

    ddb_load (value)

class flywheel.fields.types.IntType
    Bases: flywheel.fields.types.TypeDefinition

    Integer values (includes longs)

    aliases = [<type 'int'>, <type 'long'>]

    coerce (value, force)

    data_type
        alias of int

    ddb_data_type = 'N'

    ddb_load (value)

class flywheel.fields.types.ListType
    Bases: flywheel.fields.types.TypeDefinition

    List type

    coerce (value, force)

    data_type
        alias of list

    ddb_data_type = 'L'

    mutable = True

class flywheel.fields.types.NumberType
    Bases: flywheel.fields.types.TypeDefinition

    Any kind of numerical value

    coerce (value, force)

    data_type = 'N'

    ddb_data_type = 'N'

    ddb_load (value)

class flywheel.fields.types.SetType (item_type=None, type_class=None)
    Bases: flywheel.fields.types.TypeDefinition

    Set types

```

**classmethod** `bind (item_type)`  
Create a set factory that will contain a specific data type

**coerce** (*value*, *force*)

**data\_type**  
alias of `set`

**ddb\_dump** (*value*)

**ddb\_dump\_inner** (*value*)  
We need to expose this for 'contains' and 'ncontains'

**ddb\_load** (*value*)

**mutable** = `True`

**class** `flywheel.fields.types.StringType`  
Bases: `flywheel.fields.types.TypeDefinition`

String values, stored as unicode

**aliases** = ['S']

**coerce** (*value*, *force*)

**data\_type**  
alias of `unicode`

**ddb\_data\_type** = 'S'

**class** `flywheel.fields.types.TypeDefinition`  
Bases: `flywheel.compat.UnicodeMixin`

Base class for all Field types

## Attributes

<code>data_type</code>	(object) The value you wish to pass in to Field as the <code>data_type</code> .
<code>aliases</code>	(list) Other values that will reference this type if passed to Field
<code>ddb_data_type</code>	{STRING, BINARY, NUMBER, STRING_SET, BINARY_SET, NUMBER_SET, BOOL, LIST, MAP}) The DynamoDB data type that backs this type
<code>mutable</code>	(bool) If True, flywheel will track updates to this field automatically when making calls to <code>sync()</code>
<code>allowed_filters</code>	(set) The set of filters that can be used on this field type

**aliases** = []

**coerce** (*value*, *force*)  
Check the type of a value and possible convert it

**Parameters** `value` : object

The value to check

**force** : bool

If True, always attempt to convert a bad type to the correct type

**Returns** `value` : object

A variable of the correct type

**Raises** `exc` : `TypeError` or `ValueError`

If the value is the incorrect type and could not be converted

**data\_type** = `None`

**ddb\_data\_type** = `None`

**ddb\_dump** (*value*)

Dump a value to a form that can be stored in DynamoDB

**ddb\_dump\_inner** (*value*)

If this is a set type, dump a value to the type contained in the set

**ddb\_load** (*value*)

Turn a value into this type from a DynamoDB value

**mutable** = `False`

**class** `flywheel.fields.types.UTCTimezone`

Bases: `datetime.tzinfo`

UTC

**dst** (*dt*)

**tzname** (*dt*)

**utcoffset** (*dt*)

`flywheel.fields.types.register_type` (*type\_class*, *allow\_in\_set=True*)

Register a type class for use with Fields

`flywheel.fields.types.set_` (*data\_type*)

Create an alias for a SetType that contains this data type

## Module contents

Field declarations for models

**class** `flywheel.fields.Composite` (*\*args*, *\*\*kwargs*)

Bases: `flywheel.fields.Field`

A field that is composed of multiple other fields

**Parameters** *\*fields* : list

List of names of fields that compose this composite field

**hash\_key** : bool, optional

This key is a DynamoDB hash key (default False)

**range\_key** : bool, optional

This key is a DynamoDB range key (default False)

**index** : str, optional

If present, create a local secondary index on this field with this as the name.

**data\_type** : str, optional

The dynamo data type. Valid values are (NUMBER, STRING, BINARY, NUMBER\_SET, STRING\_SET, BINARY\_SET, dict, list, bool, str, unicode, int, float, set, datetime, date, Decimal) (default unicode)

**coerce** : bool, optional

Attempt to coerce the value if it's the incorrect type (default False)

**check** : callable, optional

A function that takes the value and returns True if the value is valid (default None)

**merge** : callable, optional

The function that merges the subfields together. By default it simply joins them with a `'.'`.

**get\_cached\_value** (*obj*)

**resolve** (*obj=None, scope=None*)

Resolve a field value from an object or scope dict

**class** flywheel.fields.**Field**(*hash\_key=False, range\_key=False, index=None, data\_type=<object object>, type=<type 'unicode'>, coerce=False, check=None, nullable=True, default=<object object>*)

Bases: `object`

Declarative way to specify model fields

**Parameters** **hash\_key** : bool, optional

This key is a DynamoDB hash key (default False)

**range\_key** : bool, optional

This key is a DynamoDB range key (default False)

**index** : str, optional

If present, create a local secondary index on this field with this as the name.

**type** : object, optional

The field data type. You may use int, unicode, set, etc. or you may pass in an instance of `TypeDefinition` (default unicode)

**coerce** : bool, optional

Attempt to coerce the value if it's the incorrect type (default False)

**check** : callable or list, optional

A function that takes the value and returns True if the value is valid. May also be a list of such functions. (default None)

**nullable** : bool, optional

If false, will add a check (above) to ensure the value is not null (default True).

**default** : object, optional

The default value for this field that will be set when creating a model (default None, except for `set` data types which default to `set()`)

## Notes

`Field(index='my-index')`

Is shorthand for:



```
Field().all_index('my-index')
```

## Attributes

name	(str) The name of the attribute on the model
model	(class) The <i>Model</i> this field is attached to
composite	(bool) True if this is a composite field

**all\_index** (*name*)

Index this field and project all attributes

**Parameters** **name** : str

The name of the index

**beginswith\_** (*other*)

Create a query condition that this field must begin with a string

**between\_** (*low*, *high*)

Create a query condition that this field must be between two values (inclusive)

**betwixt\_** (*low*, *high*)

Poetic version of *between\_* ()

**can\_resolve** (*fields*)

Check if the provided fields are enough to fully resolve this field

**Parameters** **fields** : list or set

**Returns** **needed** : set

Set of the subfields needed to resolve this field. If empty, then it cannot be resolved.

**coerce** (*value*, *force\_coerce=None*)

Coerce the value to the field's data type

**contains\_** (*other*)

Create a query condition that this field must contain a value

**ddb\_data\_type**

Get the native DynamoDB data type

**ddb\_dump** (*value*)

Dump a value to its Dynamo format

**ddb\_dump\_for\_query** (*value*)

Dump a value to format for use in a Dynamo query

**ddb\_load** (*val*)

Decode a value retrieved from Dynamo

**default**

Get a shallow copy of the default value

**get\_cached\_value** (*obj*)

Get the cached value of a field before any local modifications

**get\_ddb\_index** ()

Construct a dynamo local index object

**in\_** (*other*)

Create a query condition that this field must be within a set of values

**include\_index** (*name*, *includes=None*)

Index this field and project selected attributes

**Parameters** **name** : str

The name of the index

**includes** : list, optional

List of non-key attributes to project into this index

**is\_mutable**

Return True if the data type is mutable

**is\_set**

Return True if data type is a set

**keys\_index** (*name*)

Index this field and project all key attributes

**Parameters** **name** : str

The name of the index

**ncontains\_** (*other*)

Create a query condition that this field cannot contain a value

**resolve** (*obj=None*, *scope=None*)

Resolve a field value from an object or scope dict

**validate** (*obj*)

Run the validation checks for this field on a model object.

**Parameters** **obj** : *Model*

**Raises** **err** : *ValueError*

Raised if any of the checks fail.

## Submodules

### flywheel.compat module

Utilities for Python 2/3 compatibility

**class** flywheel.compat.**UnicodeMixin**

Bases: *object*

Mixin that handles `__str__` and `__bytes__`. Just define `__unicode__`.

### flywheel.engine module

Query engine

**class** flywheel.engine.**Engine** (*dynamo=None*, *namespace=()*, *default\_conflict='update'*)

Bases: *object*

Query engine for models

**Parameters** **dynamo** : *dynamodb3.DynamoDBConnection*, optional

**namespace** : list or str, optional

String prefix or list of component parts of a prefix for models. All table names will be prefixed by this string or strings (joined by '-').

**default\_conflict** : { 'update', 'overwrite', 'raise' }, optional

Default setting for delete(), save(), and sync() (default 'update')

## Notes

The engine is used to save, sync, delete, and query DynamoDB. Here is a basic example of saving items:

```
item1 = MyModel()
engine.save(item1)
item1.foobar = 'baz'
item2 = MyModel()
engine.save([item1, item2], overwrite=True)
```

You can also use the engine to query tables:

```
user = engine.query(User).filter(User.id == 'abcdef').first()

# calling engine() is a shortcut for engine.query()
user = engine(User).filter(User.id == 'abcdef').first()

d_users = engine(User).filter(User.school == 'MIT',
                               User.name.startswith_('D')).all()

# You can pass in equality constraints as keyword args
user = engine(User).filter(id='abcdef').first()
```

Scans are like queries, except that they don't use an index. Scans iterate over the ENTIRE TABLE so they are REALLY SLOW. Scans have access to additional filter conditions such as "contains" and "in".

```
# This is suuuuuper slow!
user = engine.scan(User).filter(id='abcdef').first()

# If you're doing an extremely large scan, you should tell it to return
# a generator
all_users = engine.scan(User).gen()

# to filter a field not specified in the model declaration:
prince = engine.scan(User).filter(User.field_('bio').startswith(
    'Now this is a story all about how')).first()
```

**connect** (\*args, \*\*kwargs)

Connect to a specific host

**connect\_to\_host** (\*\*kwargs)

Connect to a specific host

**connect\_to\_region** (region, \*\*kwargs)

Connect to an AWS region

**create\_schema** (test=False, throughput=None)

Create the DynamoDB tables required by the registered models

**Parameters test** : bool, optional

If True, perform a dry run (default False)

**throughput** : dict, optional

If provided, override the throughputs of the Models when creating them. Details below.

**Returns** **names** : list

List of table names that were created

## Examples

The `throughput` argument is a mapping of table names to their throughputs. The throughput is a dict with a 'read' and 'write' value. It may also include the names of global indexes that map to their own dicts with a 'read' and 'write' value.

```
engine.create_schema(throughput={
    'table1': {
        'read': 4,
        'write': 10,
        'gindex-1': {
            'read': 6,
            'write': 3,
        }
    }
})
```

**default\_conflict**

Get the default\_conflict value

## Notes

The `default_conflict` setting configures the default behavior of `save()`, `sync()`, and `delete()`. Below is an explanation of the different values of `default_conflict`.

default_conflict	method	default
'update'		
	save	overwrite=True
	sync	raise_on_conflict=True
	delete	raise_on_conflict=False
'overwrite'		
	save	overwrite=True
	sync	raise_on_conflict=False
	delete	raise_on_conflict=False
'raise'		
	save	overwrite=False
	sync	raise_on_conflict=True
	delete	raise_on_conflict=True

**delete** (*items*, *raise\_on\_conflict*=None)

Delete items from dynamo

**Parameters** **items** : list or Model

List of *Model* objects to delete

**raise\_on\_conflict** : bool, optional

If True, raise exception if the object was changed concurrently in the database (default set by `default_conflict`)

**Raises** `exc`: `dynamo3.ConditionalCheckFailedException`

If `overwrite` is False and an item already exists in the database

## Notes

Due to the structure of the AWS API, deleting with `raise_on_conflict=False` is much faster because the requests can be batched.

**delete\_key** (*model*, *pkeys=None*, *\*\*kwargs*)

Delete one or more items from dynamo as specified by primary keys

**Parameters** *model*: *Model*

**pkeys**: list, optional

List of primary key dicts

**\*\*kwargs**: dict

If *pkeys* is None, delete only a single item and use *kwargs* as the primary key dict

**Returns** *count*: int

The number of deleted items

## Notes

If the model being deleted has no range key, you may use strings instead of primary key dicts. ex:

```
>>> class Item(Model):
...     id = Field(hash_key=True)
...
>>> items = engine.delete_key(Item, ['abc', 'def', '123', '456'])
```

**delete\_keys** (*model*, *pkeys=None*, *\*\*kwargs*)

Delete one or more items from dynamo as specified by primary keys

**Parameters** *model*: *Model*

**pkeys**: list, optional

List of primary key dicts

**\*\*kwargs**: dict

If *pkeys* is None, delete only a single item and use *kwargs* as the primary key dict

**Returns** *count*: int

The number of deleted items

## Notes

If the model being deleted has no range key, you may use strings instead of primary key dicts. ex:

```
>>> class Item(Model):  
...     id = Field(hash_key=True)  
...  
>>> items = engine.delete_key(Item, ['abc', 'def', '123', '456'])
```

**delete\_schema** (*test=False*)

Drop the DynamoDB tables for all registered models

**Parameters** **test** : bool, optional

If True, perform a dry run (default False)

**Returns** **names** : list

List of table names that were deleted

**exists** (*model, key\_or\_item, range\_key=None, consistent=False*)

Check if an item exists in the database

**Parameters** **model** : dynamodb3.Model

The model class of the item to check

**key\_or\_item** : dict or dynamodb3.Model or object

Either the value of the hash key, a model instance, or a dict that contains the primary key.

**range\_key** : object, optional

Value of the range key (if the previous argument is the hash key)

**consistent** : bool, optional

Perform a consistent read from dynamo (default False)

**get** (*model, pkeys=None, consistent=False, \*\*kwargs*)

Fetch one or more items from dynamo from the primary keys

**Parameters** **model** : *Model*

**pkeys** : list, optional

List of primary key dicts

**consistent** : bool, optional

Perform a consistent read from dynamo (default False)

**\*\*kwargs** : dict

If pkeys is None, fetch only a single item and use kwargs as the primary key dict.

**Returns** **items** : list or object

If pkeys is a list of key dicts, this will be a list of items. If pkeys is None and **\*\*kwargs** is used, this will be a single object.

## Notes

If the model being fetched has no range key, you may use strings instead of primary key dicts. ex:

```
>>> class Item(Model):
...     id = Field(hash_key=True)
...
>>> items = engine.get(Item, ['abc', 'def', '123', '456'])
```

**get\_schema()**

Get the schema for the registered models

**query(model)**

Create a table query for a specific model

**Returns** *query* : *Query*

**refresh(items, consistent=False)**

Overwrite model data with freshest from database

**Parameters** *items* : list or *Model*

Models to sync

**consistent** : bool, optional

If True, force a consistent read from the db. (default False)

**register(\*models)**

Register one or more models with the engine

Registering is required for schema creation or deletion

**save(items, overwrite=None)**

Save models to dynamo

**Parameters** *items* : list or *Model*

**overwrite** : bool, optional

If False, raise exception if item already exists (default set by *default\_conflict*)

**Raises** *exc* : *dynamo3.ConditionalCheckFailedException*

If overwrite is False and an item already exists in the database

## Notes

Overwrite will replace the *entire* item with the new one, not just different fields. After calling `save(overwrite=True)` you are guaranteed that the item in the database is exactly the item you saved.

Due to the structure of the AWS API, saving with `overwrite=True` is much faster because the requests can be batched.

**scan(model)**

Create a table scan for a specific model

**Returns** *scan* : *Scan*

**sync(items, raise\_on\_conflict=None, consistent=False, constraints=None, no\_read=False)**

Sync model changes back to database

This will push any updates to the database, and ensure that all of the synced items have the most up-to-date data.

**Parameters** *items* : list or *Model*

Models to sync

**raise\_on\_conflict** : bool, optional

If True, raise exception if any of the fields that are being updated were concurrently changed in the database (default set by *default\_conflict*)

**consistent** : bool, optional

If True, force a consistent read from the db. This will only take effect if the sync is only performing a read. (default False)

**constraints** : list, optional

List of more complex constraints that must pass for the update to complete. Must be used with *raise\_on\_conflict=True*. Format is the same as query filters (e.g. *Model.fieldname > 5*)

**no\_read** : bool, optional

If True, don't perform a GET on models with no changes. (default False)

**Raises** *exc* : *dynamo3.CheckFailed*

If *raise\_on\_conflict=True* and the data in dynamo fails the constraint checks.

**update\_field** (*item*, *name*, *value=<object object>*, *action='PUT'*, *constraints=None*)

Update the value of a single field

Note that this method bypasses field validators and will ignore any special behavior around Composite fields.

**Parameters** *item* : *Model*

The model to update

**name** : str

The name of the field to update

**value** : object, optional

The new value for the field. Default will use the value currently on the model.

**action** : str, optional

PUT, ADD, or DELETE. (default PUT)

**constraints** : list, optional

List of constraints that must pass for the update to complete. Format is the same as query filters (e.g. *Model.fieldname > 5*)

**update\_schema** (*test=False*, *throughput=None*)

Updates the DynamoDB table global indexes required by the registered models

**Parameters** *test* : bool, optional

If True, perform a dry run (default False)

**throughput** : dict, optional

If provided, override the throughputs of the Models when creating them. Details below.

**Returns** *names* : list

List of table names that were updated



## Examples

The `throughput` argument is a mapping of table names to their throughputs. The throughput is a dict with a 'read' and 'write' value. It may also include the names of global indexes that map to their own dicts with a 'read' and 'write' value.

```
engine.create_schema(throughput={
    'table1': {
        'gindex-1': {
            'read': 6,
            'write': 3,
        }
    }
})
```

## flywheel.model\_meta module

Model metadata and metaclass objects

**class** flywheel.model\_meta.**ModelMetaclass** (*name, bases, dct*)

Bases: `type`

Metaclass for Model objects

Merges model metadata, sets the `meta_` attribute, and performs validation checks.

**class** flywheel.model\_meta.**ModelMetadata** (*model*)

Bases: `object`

Container for model metadata

**Parameters** `model` : *Model*

## Attributes

<i>abstract</i>	Getter for abstract
-----------------	---------------------

<code>name</code>	(str) The unique name of the model. This is set by the <code>'_name'</code> field in <code>__metadata__</code> . Defaults to the name of the model class.
<code>global_indexes</code>	(dict) List of global indexes (hash_key, [range_key]) pairs.
<code>related_fields</code>	(dict) Mapping of field names to set of fields that change when that field changes (usually just that field name, but can be more if composite fields use it)
<code>orderings</code>	(list) List of <i>Ordering</i>
<code>throughput</code>	(dict) Mapping of 'read' and 'write' to the table throughput (default 5, 5)

## abstract

Getter for abstract

**create\_dynamo\_schema** (*connection, tablenames=None, test=False, wait=False, throughput=None, namespace=()*)

Create all Dynamo tables for this model

**Parameters** `connection` : *DynamoDBConnection*

`tablenames` : list, optional

List of tables that already exist. Will call 'describe' if not provided.

**test** : bool, optional

If True, don't actually create the table (default False)

**wait** : bool, optional

If True, block until table has been created (default False)

**throughput** : dict, optional

The throughput of the table and global indexes. Has the keys 'read' and 'write'. To specify throughput for global indexes, add the name of the index as a key and another 'read', 'write' dict as the value.

**namespace** : str or tuple, optional

The namespace of the table

**Returns table** : str

Table name that was created, or None if nothing created

**ddb\_tablename** (*namespace=()*)

The name of the DynamoDB table

**Parameters namespace** : list or str, optional

String prefix or list of component parts of a prefix for the table name. The prefix will be this string or strings (joined by '-').

**delete\_dynamo\_schema** (*connection, tablenames=None, test=False, wait=False, namespace=()*)

Drop all Dynamo tables for this model

**Parameters connection** : `DynamoDBConnection`

**tablenames** : list, optional

List of tables that already exist. Will call 'describe' if not provided.

**test** : bool, optional

If True, don't actually delete the table (default False)

**wait** : bool, optional

If True, block until table has been deleted (default False)

**namespace** : str or tuple, optional

The namespace of the table

**Returns table** : str

Table name that was deleted, or None if nothing deleted

**get\_ordering\_from\_fields** (*eq\_fields, fields*)

Get a unique ordering from constraint fields.

This does a best-effort guess of which index is being queried. It prioritizes indexes that have a constraint on the range key. It prioritizes the primary key over local and global indexes.

**Parameters eq\_fields** : list

List of field names that are constrained with '='.

**fields** : list

List of field names that are constrained with inequality operators ('>', '<', 'beginswith', etc)

**Returns** `ordering`: *Ordering*

**Raises** `exc`: *TypeError*

If more than one possible Ordering is found

**get\_ordering\_from\_index** (*index*)

Get the ordering with matching index name

**hk** (*obj=None, scope=None*)

Construct the primary key value

**index\_pk\_dict** (*index\_name, obj=None, scope=None, ddb\_dump=False*)

Get the primary key dict for an index (includes the table key)

**pk\_dict** (*obj=None, scope=None, ddb\_dump=False*)

Get the dynamo primary key dict for an item

**pk\_tuple** (*obj=None, scope=None, ddb\_dump=False, ddb\_load=False*)

Get a tuple that represents the primary key for an item

**post\_create** ()

Create the orderings

**post\_validate** ()

Build the dict of related fields

**rk** (*obj=None, scope=None*)

Construct the range key value

**update\_dynamo\_schema** (*connection, test=False, wait=False, throughput=None, namespace=()*)

Updates all Dynamo table global indexes for this model

**Parameters** `connection`: *DynamoDBConnection*

**test**: bool, optional

If True, don't actually create the table (default False)

**wait**: bool, optional

If True, block until table has been created (default False)

**throughput**: dict, optional

The throughput of the table and global indexes. Has the keys 'read' and 'write'. To specify throughput for global indexes, add the name of the index as a key and another 'read', 'write' dict as the value.

**namespace**: str or tuple, optional

The namespace of the table

**Returns** `table`: str

Table name that altered, or None if nothing altered

**validate\_model** ()

Perform validation checks on the model declaration

**class** `flywheel.model_meta.Ordering` (*meta, hash\_key, range\_key=None, index\_name=None*)

Bases: *object*

A way that the models are ordered

This will be a combination of a hash key and a range key. It may be the primary key, a local secondary index, or a global secondary index.

**pk\_dict** (*obj=None, scope=None, ddb\_dump=False*)

Get the dynamo primary key dict for this ordering

**query\_kwargs** (*eq\_fields, fields*)

Get the query and filter kwargs for querying against this index

**exception** `flywheel.model_meta.ValidationError`

Bases: `exceptions.Exception`

Model inconsistency

`flywheel.model_meta.merge_metadata(cls)`

Merge all the `__metadata__` dicts in a class's hierarchy

keys that do not begin with `'_'` will be inherited.

keys that begin with `'_'` will only apply to the object that defines them.

## flywheel.models module

Model code

**class** `flywheel.models.Model(*args, **kwargs)`

Bases: `object`

Base class for all tube models

For documentation on the metadata fields, check the attributes on the `ModelMetadata` class.

### Attributes

---

<code>__metadata_class__</code>	alias of <code>ModelMetadata</code>
---------------------------------	-------------------------------------

---

<code>__meta-data__</code>	(dict) For details see <i>Metadata</i>
<b>meta_</b>	( <i>ModelMetadata</i> ) The metadata for the model
<code>__engine__</code>	( <i>Engine</i> ) Cached copy of the Engine that was used to save/load the model. This will be set after saving or loading a model.
<code>__dirty__</code>	(set) The set of all immutable fields that have been changed since the last save operation.
<code>__cache__</code>	(dict) The last seen value that was stored in the database. This is used to construct the <code>expects</code> dict when making updates that raise on conflict.
<code>__in-crs__</code>	(dict) Mapping of fields to atomic add/delete operations for numbers and sets.

**add\_** (*\*\*kwargs*)

Atomically add to a set

**cached\_** (*name, default=None*)

Get the cached (server) value of a field

**construct\_ddb\_expects\_** (*fields=None*)

Construct a dynamo “expects” mapping based on the cached fields

**ddb\_dump\_** ()

Return a dict for inserting into DynamoDB

**ddb\_dump\_cached\_** (*name*)  
Dump a cached field to a Dynamo-friendly value

**ddb\_dump\_field\_** (*name*)  
Dump a field to a Dynamo-friendly value

**classmethod ddb\_load\_** (*engine, data*)  
Load a model from DynamoDB data

**delete** (*\*args, \*\*kwargs*)  
Delete the model from the database

**classmethod field\_** (*name*)  
Get Field or construct a placeholder for an undeclared field

This is used for creating scan filter constraints on fields that were not declared in the model

**hk\_**  
The value of the hash key

**incr\_** (*\*\*kwargs*)  
Atomically increment a number value

**index\_pk\_dict\_** (*index\_name*)  
The primary key dict for an index, encoded for dynamo

**keys\_** ()  
All declared fields

**loading\_** (*\*args, \*\*kws*)  
Context manager to speed up object load process

**mark\_dirty\_** (*name*)  
Mark that a field is dirty

**meta\_** = <flywheel.model\_meta.ModelMetadata object>

**mutate\_** (*action, \*\*kwargs*)  
Atomically mutate a set

**partial\_loading\_** (*\*args, \*\*kws*)  
For use when loading a partial object (i.e. from update\_field)

**persisted\_**  
True if the model exists in DynamoDB, False otherwise

**pk\_dict\_**  
The primary key dict, encoded for dynamo

**pk\_tuple\_**  
The primary key dict, encoded for dynamo

**post\_load\_** (*engine*)  
Called after model loaded from database

**post\_save\_** ()  
Called after item is saved to database

**post\_save\_fields\_** (*fields*)  
Called after update\_field or update\_fields

**pre\_save\_** (*engine*)  
Called before saving items

**refresh** (*consistent=False*)  
Overwrite model data with freshest from database

**remove\_** (*\*\*kwargs*)  
Atomically remove from a set

**rk\_**  
The value of the range key

**save** (*overwrite=None*)  
Save model data to database (see also: `sync`)

**set\_ddb\_val\_** (*key, val*)  
Decode and set a value retrieved from Dynamo

**sync** (*\*args, \*\*kwargs*)  
Sync model changes back to database

**class** `flywheel.models.SetDelta`  
Bases: `object`

Wrapper for an atomic change to a Dynamo set

Used to track the changes when using `add_()` and `remove_()`

**add** (*action, value*)  
Add another update to the delta

**Parameters** **action** : { 'ADD', 'DELETE' }

**value** : object  
The value to add or remove

**merge** (*other*)  
Merge the delta with a set

**Parameters** **other** : set  
The original set to merge the changes with

## flywheel.query module

Query and Scan builders

**exception** `flywheel.query.DuplicateEntityException`  
Bases: `exceptions.ValueError`

Raised when too many results are found.

**exception** `flywheel.query.EntityNotFoundException`  
Bases: `exceptions.ValueError`

Raised when results are expected and not found.

**class** `flywheel.query.Query` (*engine, model*)  
Bases: `object`

An object used to query dynamo tables

See the [Engine](#) for query examples

**Parameters** **engine** : `Engine`

**model** : class

Subclass of *Model*

**all** (*desc=False, consistent=False, attributes=None, filter\_or=False, exclusive\_start\_key=None*)

Return the query results as a list

**Parameters** **desc** : bool, optional

Return results in descending order (default False)

**consistent** : bool, optional

Force a consistent read of the data (default False)

**attributes** : list, optional

List of fields to retrieve from dynamo. If supplied, returns dicts instead of model objects.

**filter\_or** : bool, optional

If True, multiple filter() constraints will be joined with an OR (default AND).

**exclusive\_start\_key** : dict, optional

The ExclusiveStartKey to resume a previous query

**Returns** **results** : list

**count** (*filter\_or=False*)

Find the number of elements the match this query

**Parameters** **filter\_or** : bool, optional

If True, multiple filter() constraints will be joined with an OR (default AND).

**Returns** **count** : int

**delete** (*filter\_or=False*)

Delete all items that match the query

**Parameters** **filter\_or** : bool, optional

If True, multiple filter() constraints will be joined with an OR (default AND).

**dynamo**

Shortcut to access DynamoDBConnection

**filter** (*\*conditions, \*\*kwargs*)

Add a Condition to constrain the query

## Notes

The conditions may be passed in as positional arguments:

```
engine.query(User).filter(User.id == 12345)
```

Or they may be passed in as keyword arguments:

```
engine.query(User).filter(firstname='Monty', lastname='Python')
```

The limitations of the keyword method is that you may only create equality conditions. You may use both types in a single filter:

```
engine.query(User).filter(User.num_friends > 10, name='Monty')
```

**first** (*desc=False, consistent=False, attributes=None, filter\_or=False*)

Return the first result of the query, or None if no results

**Parameters** **desc** : bool, optional

Return results in descending order (default False)

**consistent** : bool, optional

Force a consistent read of the data (default False)

**attributes** : list, optional

List of fields to retrieve from dynamo. If supplied, returns dicts instead of model objects.

**filter\_or** : bool, optional

If True, multiple filter() constraints will be joined with an OR (default AND).

**Returns** **result** : *Model* or None

**gen** (*desc=False, consistent=False, attributes=None, filter\_or=False, exclusive\_start\_key=None*)

Return the query results as a generator

**Parameters** **desc** : bool, optional

Return results in descending order (default False)

**consistent** : bool, optional

Force a consistent read of the data (default False)

**attributes** : list, optional

List of fields to retrieve from dynamo. If supplied, gen() will iterate over dicts instead of model objects.

**filter\_or** : bool, optional

If True, multiple filter() constraints will be joined with an OR (default AND).

**exclusive\_start\_key** : dict, optional

The ExclusiveStartKey to resume a previous query

**Returns** **results** : generator

**index** (*name*)

Use a specific local or global index for the query

**limit** (*count*)

Limit the number of query results

**one** (*consistent=False, attributes=None, filter\_or=False*)

Return the result of the query. If there is not exactly one result, raises an exception (details below)

**Parameters** **consistent** : bool, optional

Force a consistent read of the data (default False)

**attributes** : list, optional

List of fields to retrieve from dynamo. If supplied, returns dicts instead of model objects.

**filter\_or** : bool, optional

If True, multiple filter() constraints will be joined with an OR (default AND).

**Returns** **result** : *Model*



**Raises** **e1**: *EntityNotFoundException*

If no entity is found. Subclasses *ValueError*.

**e2**: *DuplicateEntityException*

If more than one entity is found. Subclasses *ValueError*.

**scan\_limit** (*count*)

Limit the number of items scanned

**tablename**

Shortcut to access dynamo table name

**class** `flywheel.query.Scan` (*engine, model*)

Bases: *flywheel.query.Query*

An object used to scan dynamo tables

scans are like Queries except they don't use indexes. This means they iterate over all data in the table and are SLOW

**Parameters** **engine**: *Engine*

**model**: class

Subclass of *Model*

**count** (*filter\_or=False*)

**gen** (*attributes=None, desc=False, consistent=False, filter\_or=False, exclusive\_start\_key=None*)

**index** (*name*)

## flywheel.tests module

Unit and system tests for flywheel

**class** `flywheel.tests.DynamoSystemTest` (*methodName='runTest'*)

Bases: `unittest.case.TestCase`

Base class for tests that need an *Engine*

**dynamo** = `None`

**models** = `[]`

**classmethod** `setUpClass` ()

**tearDown** ()

**classmethod** `tearDownClass` ()

## Module contents

flywheel



## CHAPTER 3

---

### Indices and tables

---

- `genindex`
- `modindex`
- `search`



### f

- `flywheel`, [45](#)
- `flywheel.compat`, [30](#)
- `flywheel.engine`, [30](#)
- `flywheel.fields`, [27](#)
- `flywheel.fields.conditions`, [21](#)
- `flywheel.fields.indexes`, [22](#)
- `flywheel.fields.types`, [23](#)
- `flywheel.model_meta`, [37](#)
- `flywheel.models`, [40](#)
- `flywheel.query`, [42](#)
- `flywheel.tests`, [45](#)



## A

abstract (flywheel.model\_meta.ModelMetadata attribute), 37

add() (flywheel.models.SetDelta method), 42

add\_() (flywheel.models.Model method), 40

aliases (flywheel.fields.types.BinaryType attribute), 23

aliases (flywheel.fields.types.IntType attribute), 25

aliases (flywheel.fields.types.StringType attribute), 26

aliases (flywheel.fields.types.TypeDefinition attribute), 26

all() (flywheel.fields.indexes.GlobalIndex class method), 22

all() (flywheel.query.Query method), 43

all\_index() (flywheel.fields.Field method), 29

## B

startswith\_() (flywheel.fields.Field method), 29

between\_() (flywheel.fields.Field method), 29

betwixt\_() (flywheel.fields.Field method), 29

BinaryType (class in flywheel.fields.types), 23

bind() (flywheel.fields.types.SetType class method), 25

BoolType (class in flywheel.fields.types), 23

## C

cached\_() (flywheel.models.Model method), 40

can\_resolve() (flywheel.fields.Field method), 29

coerce() (flywheel.fields.Field method), 29

coerce() (flywheel.fields.types.BinaryType method), 23

coerce() (flywheel.fields.types.BoolType method), 23

coerce() (flywheel.fields.types.DecimalType method), 24

coerce() (flywheel.fields.types.DictType method), 24

coerce() (flywheel.fields.types.FloatType method), 25

coerce() (flywheel.fields.types.IntType method), 25

coerce() (flywheel.fields.types.ListType method), 25

coerce() (flywheel.fields.types.NumberType method), 25

coerce() (flywheel.fields.types.SetType method), 26

coerce() (flywheel.fields.types.StringType method), 26

coerce() (flywheel.fields.types.TypeDefinition method), 26

Composite (class in flywheel.fields), 27

Condition (class in flywheel.fields.conditions), 21

connect() (flywheel.engine.Engine method), 31

connect\_to\_host() (flywheel.engine.Engine method), 31

connect\_to\_region() (flywheel.engine.Engine method), 31

construct() (flywheel.fields.conditions.Condition class method), 21

construct\_ddb\_expects\_() (flywheel.models.Model method), 40

construct\_index() (flywheel.fields.conditions.Condition class method), 22

construct\_limit() (flywheel.fields.conditions.Condition class method), 22

construct\_scan\_limit() (flywheel.fields.conditions.Condition class method), 22

contains\_() (flywheel.fields.Field method), 29

count() (flywheel.query.Query method), 43

count() (flywheel.query.Scan method), 45

create\_dynamo\_schema() (flywheel.model\_meta.ModelMetadata method), 37

create\_schema() (flywheel.engine.Engine method), 31

## D

data\_type (flywheel.fields.types.BinaryType attribute), 23

data\_type (flywheel.fields.types.BoolType attribute), 23

data\_type (flywheel.fields.types.DateTimeType attribute), 24

data\_type (flywheel.fields.types.DateType attribute), 24

data\_type (flywheel.fields.types.DecimalType attribute), 24

data\_type (flywheel.fields.types.DictType attribute), 24

data\_type (flywheel.fields.types.FloatType attribute), 25

data\_type (flywheel.fields.types.IntType attribute), 25

data\_type (flywheel.fields.types.ListType attribute), 25

data\_type (flywheel.fields.types.NumberType attribute), 25

data\_type (flywheel.fields.types.SetType attribute), 26

data\_type (flywheel.fields.types.StringType attribute), 26

- `data_type` (flywheel.fields.types.TypeDefinition attribute), 27
  - `DateTimeType` (class in flywheel.fields.types), 24
  - `DateType` (class in flywheel.fields.types), 24
  - `ddb_data_type` (flywheel.fields.Field attribute), 29
  - `ddb_data_type` (flywheel.fields.types.BinaryType attribute), 23
  - `ddb_data_type` (flywheel.fields.types.BoolType attribute), 23
  - `ddb_data_type` (flywheel.fields.types.DateTimeType attribute), 24
  - `ddb_data_type` (flywheel.fields.types.DateType attribute), 24
  - `ddb_data_type` (flywheel.fields.types.DecimalType attribute), 24
  - `ddb_data_type` (flywheel.fields.types.DictType attribute), 25
  - `ddb_data_type` (flywheel.fields.types.FloatType attribute), 25
  - `ddb_data_type` (flywheel.fields.types.IntType attribute), 25
  - `ddb_data_type` (flywheel.fields.types.ListType attribute), 25
  - `ddb_data_type` (flywheel.fields.types.NumberType attribute), 25
  - `ddb_data_type` (flywheel.fields.types.StringType attribute), 26
  - `ddb_data_type` (flywheel.fields.types.TypeDefinition attribute), 27
  - `ddb_dump()` (flywheel.fields.Field method), 29
  - `ddb_dump()` (flywheel.fields.types.BinaryType method), 23
  - `ddb_dump()` (flywheel.fields.types.DateTimeType method), 24
  - `ddb_dump()` (flywheel.fields.types.DateType method), 24
  - `ddb_dump()` (flywheel.fields.types.SetType method), 26
  - `ddb_dump()` (flywheel.fields.types.TypeDefinition method), 27
  - `ddb_dump_()` (flywheel.models.Model method), 40
  - `ddb_dump_cached_()` (flywheel.models.Model method), 41
  - `ddb_dump_field_()` (flywheel.models.Model method), 41
  - `ddb_dump_for_query()` (flywheel.fields.Field method), 29
  - `ddb_dump_inner()` (flywheel.fields.types.SetType method), 26
  - `ddb_dump_inner()` (flywheel.fields.types.TypeDefinition method), 27
  - `ddb_load()` (flywheel.fields.Field method), 29
  - `ddb_load()` (flywheel.fields.types.BinaryType method), 23
  - `ddb_load()` (flywheel.fields.types.DateTimeType method), 24
  - `ddb_load()` (flywheel.fields.types.FloatType method), 25
  - `ddb_load()` (flywheel.fields.types.IntType method), 25
  - `ddb_load()` (flywheel.fields.types.NumberType method), 25
  - `ddb_load()` (flywheel.fields.types.SetType method), 26
  - `ddb_load()` (flywheel.fields.types.TypeDefinition method), 27
  - `ddb_load_()` (flywheel.models.Model class method), 41
  - `ddb_tablename()` (flywheel.model\_meta.ModelMetadata method), 38
  - `DecimalType` (class in flywheel.fields.types), 24
  - `default` (flywheel.fields.Field attribute), 29
  - `default_conflict` (flywheel.engine.Engine attribute), 32
  - `delete()` (flywheel.engine.Engine method), 32
  - `delete()` (flywheel.models.Model method), 41
  - `delete()` (flywheel.query.Query method), 43
  - `delete_dynamo_schema()` (flywheel.model\_meta.ModelMetadata method), 38
  - `delete_key()` (flywheel.engine.Engine method), 33
  - `delete_keys()` (flywheel.engine.Engine method), 33
  - `delete_schema()` (flywheel.engine.Engine method), 34
  - `DictType` (class in flywheel.fields.types), 24
  - `dst()` (flywheel.fields.types.UTCTimezone method), 27
  - `DuplicateEntityException`, 42
  - `dynamo` (flywheel.query.Query attribute), 43
  - `dynamo` (flywheel.tests.DynamoSystemTest attribute), 45
  - `DynamoSystemTest` (class in flywheel.tests), 45
- ## E
- `Engine` (class in flywheel.engine), 30
  - `EntityNotFoundException`, 42
  - `exists()` (flywheel.engine.Engine method), 34
- ## F
- `Field` (class in flywheel.fields), 28
  - `field_()` (flywheel.models.Model class method), 41
  - `filter()` (flywheel.query.Query method), 43
  - `first()` (flywheel.query.Query method), 43
  - `FloatType` (class in flywheel.fields.types), 25
  - `flywheel` (module), 45
  - `flywheel.compat` (module), 30
  - `flywheel.engine` (module), 30
  - `flywheel.fields` (module), 27
  - `flywheel.fields.conditions` (module), 21
  - `flywheel.fields.indexes` (module), 22
  - `flywheel.fields.types` (module), 23
  - `flywheel.model_meta` (module), 37
  - `flywheel.models` (module), 40
  - `flywheel.query` (module), 42
  - `flywheel.tests` (module), 45
- ## G
- `gen()` (flywheel.query.Query method), 44



[gen\(\)](#) (flywheel.query.Scan method), 45  
[get\(\)](#) (flywheel.engine.Engine method), 34  
[get\\_cached\\_value\(\)](#) (flywheel.fields.Composite method), 28  
[get\\_cached\\_value\(\)](#) (flywheel.fields.Field method), 29  
[get\\_ddb\\_index\(\)](#) (flywheel.fields.Field method), 29  
[get\\_ddb\\_index\(\)](#) (flywheel.fields.indexes.GlobalIndex method), 22  
[get\\_ordering\\_from\\_fields\(\)](#) (flywheel.model\_meta.ModelMetadata method), 38  
[get\\_ordering\\_from\\_index\(\)](#) (flywheel.model\_meta.ModelMetadata method), 39  
[get\\_schema\(\)](#) (flywheel.engine.Engine method), 35  
[GlobalIndex](#) (class in flywheel.fields.indexes), 22

## H

[hk\(\)](#) (flywheel.model\_meta.ModelMetadata method), 39  
[hk\\_](#) (flywheel.models.Model attribute), 41

## I

[in\\_\(\)](#) (flywheel.fields.Field method), 29  
[include\(\)](#) (flywheel.fields.indexes.GlobalIndex class method), 23  
[include\\_index\(\)](#) (flywheel.fields.Field method), 29  
[incr\\_\(\)](#) (flywheel.models.Model method), 41  
[index\(\)](#) (flywheel.query.Query method), 44  
[index\(\)](#) (flywheel.query.Scan method), 45  
[index\\_pk\\_dict\(\)](#) (flywheel.model\_meta.ModelMetadata method), 39  
[index\\_pk\\_dict\\_\(\)](#) (flywheel.models.Model method), 41  
[IntType](#) (class in flywheel.fields.types), 25  
[is\\_mutable](#) (flywheel.fields.Field attribute), 30  
[is\\_set](#) (flywheel.fields.Field attribute), 30

## K

[keys\(\)](#) (flywheel.fields.indexes.GlobalIndex class method), 23  
[keys\\_\(\)](#) (flywheel.models.Model method), 41  
[keys\\_index\(\)](#) (flywheel.fields.Field method), 30

## L

[limit\(\)](#) (flywheel.query.Query method), 44  
[ListType](#) (class in flywheel.fields.types), 25  
[loading\\_\(\)](#) (flywheel.models.Model method), 41

## M

[mark\\_dirty\\_\(\)](#) (flywheel.models.Model method), 41  
[merge\(\)](#) (flywheel.models.SetDelta method), 42  
[merge\\_metadata\(\)](#) (in module flywheel.model\_meta), 40  
[meta\\_](#) (flywheel.models.Model attribute), 41  
[Model](#) (class in flywheel.models), 40

[ModelMetaclass](#) (class in flywheel.model\_meta), 37  
[ModelMetadata](#) (class in flywheel.model\_meta), 37  
[models](#) (flywheel.tests.DynamoSystemTest attribute), 45  
[mutable](#) (flywheel.fields.types.DictType attribute), 25  
[mutable](#) (flywheel.fields.types.ListType attribute), 25  
[mutable](#) (flywheel.fields.types.SetType attribute), 26  
[mutable](#) (flywheel.fields.types.TypeDefinition attribute), 27  
[mutate\\_\(\)](#) (flywheel.models.Model method), 41

## N

[ncontains\\_\(\)](#) (flywheel.fields.Field method), 30  
[NumberType](#) (class in flywheel.fields.types), 25

## O

[one\(\)](#) (flywheel.query.Query method), 44  
[Ordering](#) (class in flywheel.model\_meta), 39

## P

[partial\\_loading\\_\(\)](#) (flywheel.models.Model method), 41  
[persisted\\_](#) (flywheel.models.Model attribute), 41  
[pk\\_dict\(\)](#) (flywheel.model\_meta.ModelMetadata method), 39  
[pk\\_dict\\_\(\)](#) (flywheel.model\_meta.Ordering method), 40  
[pk\\_dict\\_](#) (flywheel.models.Model attribute), 41  
[pk\\_tuple\(\)](#) (flywheel.model\_meta.ModelMetadata method), 39  
[pk\\_tuple\\_](#) (flywheel.models.Model attribute), 41  
[post\\_create\(\)](#) (flywheel.model\_meta.ModelMetadata method), 39  
[post\\_load\\_\(\)](#) (flywheel.models.Model method), 41  
[post\\_save\\_\(\)](#) (flywheel.models.Model method), 41  
[post\\_save\\_fields\\_\(\)](#) (flywheel.models.Model method), 41  
[post\\_validate\(\)](#) (flywheel.model\_meta.ModelMetadata method), 39  
[pre\\_save\\_\(\)](#) (flywheel.models.Model method), 41

## Q

[Query](#) (class in flywheel.query), 42  
[query\(\)](#) (flywheel.engine.Engine method), 35  
[query\\_kwargs\(\)](#) (flywheel.fields.conditions.Condition method), 22  
[query\\_kwargs\(\)](#) (flywheel.model\_meta.Ordering method), 40

## R

[refresh\(\)](#) (flywheel.engine.Engine method), 35  
[refresh\(\)](#) (flywheel.models.Model method), 41  
[register\(\)](#) (flywheel.engine.Engine method), 35  
[register\\_type\(\)](#) (in module flywheel.fields.types), 27  
[remove\\_\(\)](#) (flywheel.models.Model method), 42  
[resolve\(\)](#) (flywheel.fields.Composite method), 28  
[resolve\(\)](#) (flywheel.fields.Field method), 30

`rk()` (flywheel.model\_meta.ModelMetadata method), 39  
`rk_` (flywheel.models.Model attribute), 42

## S

`save()` (flywheel.engine.Engine method), 35  
`save()` (flywheel.models.Model method), 42  
`Scan` (class in flywheel.query), 45  
`scan()` (flywheel.engine.Engine method), 35  
`scan_kwargs()` (flywheel.fields.conditions.Condition method), 22  
`scan_limit()` (flywheel.query.Query method), 45  
`set_()` (in module flywheel.fields.types), 27  
`set_ddb_val_()` (flywheel.models.Model method), 42  
`SetDelta` (class in flywheel.models), 42  
`SetType` (class in flywheel.fields.types), 25  
`setUpClass()` (flywheel.tests.DynamoSystemTest class method), 45  
`StringType` (class in flywheel.fields.types), 26  
`sync()` (flywheel.engine.Engine method), 35  
`sync()` (flywheel.models.Model method), 42

## T

`tablename` (flywheel.query.Query attribute), 45  
`tearDown()` (flywheel.tests.DynamoSystemTest method), 45  
`tearDownClass()` (flywheel.tests.DynamoSystemTest class method), 45  
`throughput()` (flywheel.fields.indexes.GlobalIndex method), 23  
`TypeDefinition` (class in flywheel.fields.types), 26  
`tzname()` (flywheel.fields.types.UTCTimezone method), 27

## U

`UnicodeMixin` (class in flywheel.compat), 30  
`update_dynamo_schema()` (flywheel.model\_meta.ModelMetadata method), 39  
`update_field()` (flywheel.engine.Engine method), 36  
`update_schema()` (flywheel.engine.Engine method), 36  
`utcoffset()` (flywheel.fields.types.UTCTimezone method), 27  
`UTCTimezone` (class in flywheel.fields.types), 27

## V

`validate()` (flywheel.fields.Field method), 30  
`validate_model()` (flywheel.model\_meta.ModelMetadata method), 39  
`ValidationError`, 40