
DynamoDB-mock Documentation

Release 0.3.1

Jean-Tiare Le Bigot

October 11, 2012

CONTENTS

OVERVIEW

DynamoDB is a minimalistic NoSQL engine provided by Amazon as a part of their AWS product.

DynamoDB is great in production environment but sucks when testing your application. Tables need roughly 1 min to be created, deleted or updated. Items operation rates depend on how much you pay and tests will conflict if 2 developers run them at the same time.

ddbmock brings a tiny in-memory(tm) implementation of DynamoDB API. It can either be run as a stand alone server or as a regular library helping you to build lightning fast unit and functional tests :)

ddbmock does *not* intend to be production ready. It *will* **lose** you data. you've been warned! I currently recommend the “boto extension” mode for unit-tests and the “server” mode for functional tests.

DOCUMENTATION

2.1 User guide

2.1.1 Getting started with DynamoDB-mock

DynamoDB is a minimalistic NoSQL engine provided by Amazon as a part of their AWS product.

DynamoDB is great in production environment but sucks when testing your application. Tables needs roughly 1 min to be created, deleted or updated. Items operation rates depends on how much you pay and tests will conflict if 2 developers run them at the same time.

ddbmock brings a tiny in-memory(tm) implementation of DynamoDB API. It can either be run as a stand alone server or as a regular library helping you to build lightning fast unit and functional tests :)

ddbmock does *not* intend to be production ready. It *will loose* you data. you've been warned! I currently recommend the "boto extension" mode for unit-tests and the "server" mode for functional tests.

Installation

```
$ pip install ddbmock
```

Example usage

Run as Regular client-server

Ideal for test environment. For stage and production I highly recommend using DynamoDB servers. ddbmock comes with no warranty and *will loose your data(tm)*.

```
$ pserve development.ini # launch the server on 0.0.0.0:6543
```

```
import boto
from ddbmock import connect_ddbmock

# Use the provided helper to connect your *own* endpoint
db = connect_ddbmock()

# Done ! just use it wherever in your project as usual.
db.list_tables() # get list of tables (empty at this stage)
```

Note: if you do not want to import ddbmock only for the helper, here is a reference implementation:

```
def connect_ddbmock(host='localhost', port=6543):
    import boto
    from boto.regioninfo import RegionInfo
    endpoint = '{}:{}'.format(host, port)
    region = RegionInfo(name='ddbmock', endpoint=endpoint)
    return boto.connect_dynamodb(region=region, port=port, is_secure=False)
```

Run as a standalone library

Ideal for unit testing or small scale automated functional tests. Nice to play around with boto DynamoDB API too :)

```
import boto
from ddbmock import connect_boto

# Wire-up boto and ddbmock together
db = connect_boto()

# Done ! just use it wherever in your project as usual.
db.list_tables() # get list of tables (empty at this stage)
```

2.1.2 Current Status

This documents reflects ddbmock status as of 3/10/12. It may be outdated.

Methods support

- CreateTable DONE
- DeleteTable DONE
- UpdateTable DONE
- DescribeTable DONE
- GetItem DONE
- PutItem DONE
- DeleteItem DONE
- UpdateItem ALMOST
- BatchGetItem WIP
- BatchWriteItem TODO
- Query WIP
- Scan WIP

There is basically no support for Limit, ExclusiveStartKey, ScanIndexForward and their associated features at all in ddbmock. This affects all “WIP” functions.

UpdateItem has a different behavior when the target item did not exist prior the update operation. In particular, the ADD operator will always behave as though the item existed before.

Comparison Operators

Some comparison might not work as expected on binary data as it is performed on the base64 representation instead of the binary one. Please report a bug if this is a problem for you, or, even better, open a pull request :)

All operators exists as lower case functions in `ddbmock.database.comparison`. This list can easily be extended to add new/custom operators.

Common to Query and Scan

- `EQ DONE`
- `LE DONE`
- `LT DONE`
- `GE DONE`
- `GT DONE`
- `BEGINS_WITH DONE`
- `BETWEEN DONE`

Specific to Scan

- `NULL DONE`
- `NOT_NULL DONE`
- `CONTAINS DONE`
- `NOT_CONTAINS DONE`
- `IN DONE`

`IN` operator is the only that can not be imported directly as it overlaps builtin `in` keyword. If you need it, either import it with `getattr` on the module or as `in_test` which, anyway, is its internal name.

Return value specifications

- `NONE DONE`
- `ALL_OLD DONE`
- `ALL_NEW DONE`
- `UPDATED_OLD DONE`
- `UPDATED_NEW DONE`

Please note that only `UpdateItem` supports the 5. Other compatible methods understands only the 2 first.

Rates and size limitations

basically, none are supported yet

Request rate

- Throttle read operations when provisioned throughput exceeded. TODO
- Throttle write operations when provisioned throughput exceeded. TODO
- Maximum throughput is 10,000. DONE
- Minimum throughput is 1. DONE
- Report accurate throughput. WIP (low level foundations are in place)

ddbmock currently reports the consumed throughput based on item count. Their size is ignored from the computation. Actually, this is even trickier as the throughput is normally spreaded among partitions which ddbmock does not handle.

Request size

- Limit response size to 1MB. TODO
- Limit request size to 1MB. TODO
- Limit `BatchGetItem` to 100 per request. TODO
- Limit `BatchWriteItem` to 25 per request. TODO

Table managment

- No more than 256 tables. DONE
- No more than 10 `CREATING` tables. TODO
- No more than 10 `DELETING` tables. TODO
- No more than 1 `UPDATING` table. TODO
- No more than 1 Throughput decrease/calendar day. DONE
- No more than *2 Throughput increase/update. DONE
- At least 10% change per update. DONE

Types and items Limitations

- Table names can only be between 3 and 255 bytes long. DONE
- Table names can only contains a-z, A-Z, 0-9, '_', '-', and '.'. DONE
- No more than 64kB/Item including fieldname but not indexing overhead. DONE
- Primary key names can only be between 1 and 255 bytes long. DONE
- Attribute value can *not* be Null. DONE
- `hash_key` value maximu 2048 bytes. DONE
- `range_key` value maximu 1024 bytes. DONE
- Numbers max 38 digits precision; between 10^{-128} and 10^{+126} . DONE

Table description

- item count. DONE
- data size. DONE
- date: creation. DONE
- date: last throughput increase. DONE
- date: last throughput decrease. DONE

Dates are represented as float timestamps using scientific notation by DynamoDB but we only send them as plain number, not caring about the representation. Most parsers won't do any difference anyway.

2.1.3 Change log - Migration guide.

ddbmock 0.3.1

This section documents all user visible changes included between ddbmock versions 0.3.0 and versions 0.3.1.dev

Additions

- 100% tests coverage
- add basic tests for pyramid entry-point (#1)
- add plenty of unit and functional tests. Coverage is 100%
- add support for all ALL_OLD ALL_NEW UPDATED_OLD UPDATED_NEW in UpdateItem
- add accurate field size calculation
- add accurate item size calculation
- add accurate table size calculation
- add MAX_TABLES check at table creation

Changes

- accurate table statuses
- fix pyramid entry-point
- fix list validations. Len limitation was not working
- attempt to store empty field/set raise ValidationError (#4)
- accurate exception detection and reporting in UpdateTable
- accurate hash_key and range_key size validation
- accurate number limitations (max 38 digits precision; between 10^{-128} and 10^{+126})
- rename connect_boto to connect_boto_patch + compat layer
- rename connect_ddbmock to connect_boto_network + compat layer
- block PutItem/UpdateItem when bigger than MAX_ITEM_SIZE

Upgrade

Nothing mandatory as this is a minor release but, I recommend that you:

- rename `connect_boto` to `connect_boto_patch`
- rename `connect_ddbmock` to `connect_boto_network`

ddbmock 0.3.0

Initial ddbmock release. This is *alpha quality* software. Some import features such as “Exclusive Start Key”, “Reverse” and “Limit” as well as `BatchWriteItem` have not been written (yet).

Additions

- entry-point WEB (network mode)
- entry-point Boto (standalone mode)
- support for `CreateTable` method
- support for `DeleteTable` method
- support for `UpdateTable` method
- support for `DescribeTable` method
- support for “`GetItem`” method
- support for `PutItem` method
- support for `DeleteItem` method
- support for `UpdateItem` method (small approximations)
- support for `BatchGetItem` method (initial)
- support for `Query` method (initial)
- support for `Scan` method (initial)
- all comparison operators
- aggressive input validation

Known bugs - limitations

- no support for `BatchWriteItem`
- no support for “Exclusive Start Key”, “Reverse” and “Limit” in

`Query` and `Scan` - no support for “UnprocessedKeys” in `BatchGetItem` - Web entry-point is untested, fill bugs if necessary :)

2.2 Indices and tables

- *genindex*
- *modindex*

- *search*

CONTRIBUTE

Want to contribute, report a but of request a feature ? The development goes on BitBucket:

- **Download:** <http://pypi.python.org/pypi/ddbmock>
- **Report bugs:** <https://bitbucket.org/Ludia/dynamodb-mock/issues>
- **Fork the code:** <https://bitbucket.org/Ludia/dynamodb-mock/overview>