

---

# **unqlite-python Documentation**

*Release 0.9.3*

**charles leifer**

**Nov 25, 2021**



---

# Contents

---

<b>1</b>	<b>Installation</b>	<b>3</b>
<b>2</b>	<b>Quick-start</b>	<b>5</b>
2.1	Key/value features . . . . .	5
2.2	Cursors . . . . .	6
2.3	Document store features . . . . .	6
2.4	Collections . . . . .	7
2.5	Transactions . . . . .	8
<b>3</b>	<b>API Documentation</b>	<b>9</b>
<b>4</b>	<b>Indices and tables</b>	<b>25</b>
	<b>Index</b>	<b>27</b>



# unqlite-python

Fast Python bindings for [UnQLite](#), a lightweight, embedded NoSQL database and JSON document store.

**Warning:** Read the issue tracker for [this database](#) before considering using it. UnQLite has not seen any meaningful development since 2014. It is **strongly** recommended that you use [Sqlite](#). Sqlite has robust support for [json](#) and is actively developed and maintained.

UnQLite features:

- Embedded, zero-conf database
- Transactional (ACID)
- Single file or in-memory database
- Key/value store
- Cursor support and linear record traversal
- JSON document store
- Thread-safe
- Terabyte-sized databases

UnQLite-Python features:

- Compiled library, extremely fast with minimal overhead.
- Supports key/value operations, cursors, and transactions using Pythonic APIs.
- Support for Jx9 scripting.
- APIs for working with Jx9 JSON document collections.

The source code for unqlite-python is [hosted on GitHub](#).

---

**Note:** If you encounter any bugs in the library, please [open an issue](#), including a description of the bug and any related traceback.

---

Contents:



# CHAPTER 1

---

## Installation

---

You can use `pip` to install `unqlite`:

```
pip install cython unqlite
```

The project is hosted at <https://github.com/coleifer/unqlite-python> and can be installed from source:

```
git clone https://github.com/coleifer/unqlite-python
cd unqlite-python
python setup.py build
python setup.py install
```

---

**Note:** `unqlite-python` depends on `Cython` to generate the Python extension. As of 0.5.0, `unqlite-python` no longer ships with a generated C source file, so it is necessary to install `Cython` in order to compile `unqlite-python`.

---

After installing `unqlite-python`, you can run the unit tests by executing the `tests` module:

```
python tests.py
```





Below is a sample interactive console session designed to show some of the basic features and functionality of the unqlite-python library. Also check out the [full API documentation](#).

To begin, instantiate an *UnQLite* object. You can specify either the path to a database file, or use UnQLite as an in-memory database.

```
>>> from unqlite import UnQLite
>>> db = UnQLite() # Create an in-memory database.
```

## 2.1 Key/value features

UnQLite can be used as a key/value store.

```
>>> db['foo'] = 'bar' # Use as a key/value store.
>>> db['foo'] # The key/value deals in byte-strings.
b'bar'

>>> for i in range(4):
...     db['k%s' % i] = str(i)
...

>>> 'k3' in db
True
>>> 'k4' in db
False
>>> del db['k3']

>>> db.append('k2', 'XXXX')
>>> db['k2']
b'2XXXX'
```

The database can also be iterated through directly. Note that keys are decoded while values are left as bytestrings.

```
>>> [item for item in db]
[('foo', b'bar'), ('k0', b'0'), ('k1', b'1'), ('k2', b'2XXXX')]
```

UnQLite databases support common dict APIs, such as `keys()`, `values()`, `items()`, and `update()`.

## 2.2 Cursors

For finer-grained record traversal, you can use cursors.

```
>>> with db.cursor() as cursor:
...     for key, value in cursor:
...         print(key, '=>', value.decode('utf8'))
...
k0 => 0
k1 => 1
k2 => 2XXXX

>>> with db.cursor() as cursor:
...     cursor.seek('k2')
...     print(cursor.value())
...
b'2XXXX'

>>> with db.cursor() as cursor:
...     cursor.seek('k0')
...     print(list(cursor.fetch_until('k2', include_stop_key=False)))
...
[('k0', b'0'), ('k1', b'1')]
```

For more information, see the [Cursor](#) API documentation.

## 2.3 Document store features

In my opinion the most interesting feature of UnQLite is its JSON document store. The [Jx9 scripting language](#) is used to interact with the document store, and it is a wacky mix of C, JavaScript and PHP.

Interacting with the document store basically consists of creating a Jx9 script (you might think of it as an imperative SQL query), compiling it, and then executing it.

---

**Note:** As of v0.8.0 the document store and collections APIs treat all strings as unicode.

---

```
>>> script = """
...     db_create('users');
...     db_store('users', $list_of_users);
...     $users_from_db = db_fetch_all('users');
...     """

>>> list_of_users = [
...     {'username': 'Huey', 'age': 3},
...     {'username': 'Mickey', 'age': 5}
... ]
```

(continues on next page)

(continued from previous page)

```

>>> with db.vm(script) as vm:
...     vm['list_of_users'] = list_of_users
...     vm.execute()
...     users_from_db = vm['users_from_db']
...
True

>>> users_from_db # UnQLite assigns items in a collection an ID.
[{'username': 'Huey', 'age': 3, '__id': 0},
 {'username': 'Mickey', 'age': 5, '__id': 1}]

```

This is just a taste of what is possible with Jx9. More information can be found in the [VM](#) documentation.

## 2.4 Collections

To simplify working with JSON document collections, unqlite-python provides a light API for executing Jx9 queries on collections. A collection is an ordered list of JSON objects (records). Records can be appended, updated or deleted. be support for updates as well.

To begin working with *Collection*, you can use the *UnQLite.collection()* factory method:

```

>>> users = db.collection('users')
>>> users.create() # Create the collection if it does not exist.
>>> users.exists()
True

```

You can use the *Collection.store()* method to add one or many records. To add a single record just pass in a python dict. To add multiple records, pass in a list of dicts. Records can be fetched and deleted by ID using *fetch()* and *delete()*.

```

>>> users.store([
...     {'name': 'Charlie', 'color': 'green'},
...     {'name': 'Huey', 'color': 'white'},
...     {'name': 'Mickey', 'color': 'black'}])
True
>>> users.store({'name': 'Leslie', 'color': 'also green'})
True

>>> users.fetch(0) # Fetch the first record, user with "__id" = 0.
{'__id': 0, 'color': 'green', 'name': 'Charlie'}

>>> users.delete(0) # Delete the first record (user "__id" = 0).
True
>>> users.delete(users.last_record_id()) # Delete the last record.
True

```

You can retrieve all records in the collection, or specify a filtering function. The filtering function will be registered as a foreign function with the Jx9 VM and called *from* the VM.

```

>>> users.all()
[{'__id': 1, 'color': 'white', 'name': 'Huey'},
 {'__id': 2, 'color': 'black', 'name': 'Mickey'}]

```

(continues on next page)

(continued from previous page)

```
>>> users.filter(lambda obj: obj['name'].startswith('H'))
[{'_id': 1, 'color': 'white', 'name': 'Huey'}]
```

More information can be found in the *Collection* documentation.

## 2.5 Transactions

UnQLite supports transactions for file-backed databases (since transactions occur at the filesystem level, they have no effect on in-memory databases).

The easiest way to create a transaction is with the context manager:

```
>>> db = UnQLite('/tmp/test.db')
>>> with db.transaction():
...     db['k1'] = 'v1'
...     db['k2'] = 'v2'
...
>>> db['k1']
b'v1'
```

You can also use the transaction decorator which will wrap a function call in a transaction and commit upon successful execution (rolling back if an exception occurs).

```
>>> @db.commit_on_success
... def save_value(key, value, exc=False):
...     db[key] = value
...     if exc:
...         raise Exception('uh-oh')
...
>>> save_value('k3', 'v3')
>>> save_value('k3', 'vx', True)
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "unqlite/core.py", line 312, in wrapper
    return fn(*args, **kwargs)
  File "<stdin>", line 5, in save_value
Exception: uh-oh
>>> db['k3']
b'v3'
```

For finer-grained control you can call *begin()*, *rollback()* and *commit()* manually.

```
>>> db.begin()
>>> db['k3'] = 'v3-xx'
>>> db.commit()
True
>>> db['k3']
b'v3-xx'
```

```
class UnQLite ([filename=':mem:', flags=UNQLITE_OPEN_CREATE[, open_database=True]])
```

The *UnQLite* object provides a pythonic interface for interacting with *UnQLite* databases. *UnQLite* is a lightweight, embedded NoSQL database and JSON document store.

#### Parameters

- **filename** (*str*) – The path to the database file.
- **flags** (*int*) – How the database file should be opened.
- **open\_database** (*bool*) – When set to `True`, the database will be opened automatically when the class is instantiated. If set to `False` you will need to manually call `open()`.

**Note:** *UnQLite* supports in-memory databases, which can be created by passing in `:mem:` as the database file. This is the default behavior if no database file is specified.

Example usage:

```
>>> db = UnQLite() # Create an in-memory database.
>>> db['foo'] = 'bar' # Use as a key/value store.
>>> print db['foo']
bar

>>> for i in range(4):
...     db['k%s' % i] = str(i)
...

>>> 'k3' in db
True
>>> 'k4' in db
False
>>> del db['k3']

>>> db.append('k2', 'XXXX')
```

(continues on next page)

(continued from previous page)

```

>>> db['k2']
'2XXXX'

>>> with db.cursor() as cursor:
...     for key, value in cursor:
...         print key, '=>', value
...
foo => bar
k0 => 0
k1 => 1
k2 => 2XXXX

>>> script = """
...     db_create('users');
...     db_store('users', $list_of_users);
...     $users_from_db = db_fetch_all('users');
... """

>>> list_of_users = [
...     {'username': 'Huey', 'age': 3},
...     {'username': 'Mickey', 'age': 5}
... ]

>>> with db.vm(script) as vm:
...     vm['list_of_users'] = list_of_users
...     vm.execute()
...     users_from_db = vm['users_from_db']
...
True

>>> users_from_db # UnQLite assigns items in a collection an ID.
[{'username': 'Huey', 'age': 3, '__id': 0},
 {'username': 'Mickey', 'age': 5, '__id': 1}]

```

**open()**

Open the database. This method should only be called if the database was manually closed, or if the database was instantiated with `open_database=False`.

Valid flags:

- UNQLITE\_OPEN\_CREATE
- UNQLITE\_OPEN\_READONLY
- UNQLITE\_OPEN\_READWRITE
- UNQLITE\_OPEN\_CREATE
- UNQLITE\_OPEN\_EXCLUSIVE
- UNQLITE\_OPEN\_TEMP\_DB
- UNQLITE\_OPEN\_NOMUTEX
- UNQLITE\_OPEN\_OMIT\_JOURNALING
- UNQLITE\_OPEN\_IN\_MEMORY
- UNQLITE\_OPEN\_MMAP

Detailed descriptions of these flags can be found in the [unqlite\\_open docs](#).

**close()**

Close the database.

**Warning:** If you are using a file-based database, by default any uncommitted changes will be committed when the database is closed. If you wish to discard uncommitted changes, you can use `disable_autocommit()`.

**\_\_enter\_\_()**

Use the database as a context manager, opening the connection and closing it at the end of the wrapped block:

```
with UnQLite('my_db.ldb') as db:
    db['foo'] = 'bar'

# When the context manager exits, the database is closed.
```

**disable\_autocommit()**

When the database is closed, prevent any uncommitted writes from being saved.

---

**Note:** This method only affects file-based databases.

---

**store(key, value)**

Store a value in the given key.

**Parameters**

- **key** (*str*) – Identifier used for storing data.
- **value** (*str*) – A value to store in UnQLite.

Example:

```
db = UnQLite()
db.store('some key', 'some value')
db.store('another key', 'another value')
```

You can also use the dictionary-style `db[key] = value` to store a value:

```
db['some key'] = 'some value'
```

**fetch(key)**

Retrieve the value stored at the given key. If no value exists in the given key, a `KeyError` will be raised.

**Parameters** **key** (*str*) – Identifier to retrieve

**Returns** The data stored at the given key

**Raises** `KeyError` if the given key does not exist.

Example:

```
db = UnQLite()
db.store('some key', 'some value')
value = db.fetch('some key')
```

You can also use the dictionary-style `value = db[key]` lookup to retrieve a value:

```
value = db['some key']
```

**delete** (*key*)

Remove the key and its associated value from the database.

**Parameters** **key** (*str*) – The key to remove from the database.

**Raises** `KeyError` if the given key does not exist.

Example:

```
def clear_cache():
    db.delete('cached-data')
```

You can also use the python `del` keyword combined with a dictionary lookup:

```
def clear_cache():
    del db['cached-data']
```

**append** (*key, value*)

Append the given value to the data stored in the key. If no data exists, the operation is equivalent to `store()`.

**Parameters**

- **key** (*str*) – The identifier of the value to append to.
- **value** – The value to append.

**exists** (*key*)

Return whether the given key exists in the database.

**Parameters** **key** (*str*) –

**Returns** A boolean value indicating whether the given key exists in the database.

Example:

```
def get_expensive_data():
    if not db.exists('cached-data'):
        db.set('cached-data', calculate_expensive_data())
    return db.get('cached-data')
```

You can also use the python `in` keyword to determine whether a key exists:

```
def get_expensive_data():
    if 'cached-data' not in db:
        db['cached-data'] = calculate_expensive_data()
    return db['cached-data']
```

**begin** ()

Begin a transaction.

**rollback** ()

Roll back the current transaction.

**commit** ()

Commit the current transaction.

**transaction** ()

Create a context manager for performing multiple operations in a transaction.



**Warning:** Transactions occur at the disk-level and have no effect on in-memory databases.

Example:

```
# Transfer $100 in a transaction.
with db.transaction():
    db['from_acct'] = db['from_account'] - 100
    db['to_acct'] = db['to_acct'] + 100

# Make changes and then roll them back.
with db.transaction():
    db['foo'] = 'bar'
db.rollback() # Whoops, do not commit these changes.
```

#### `commit_on_success (fn)`

Function decorator that will cause the wrapped function to have all statements wrapped in a transaction. If the function returns without an exception, the transaction is committed. If an exception occurs in the function, the transaction is rolled back.

Example:

```
>>> @db.commit_on_success
... def save_value(key, value, exc=False):
...     db[key] = value
...     if exc:
...         raise Exception('uh-oh')
...
>>> save_value('k3', 'v3')
>>> save_value('k3', 'vx', True)
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "unqlite/core.py", line 312, in wrapper
    return fn()
  File "<stdin>", line 5, in save_value
Exception: uh-oh
>>> db['k3']
'v3'
```

#### `cursor()`

**Returns** a *Cursor* instance.

Create a cursor for traversing database records.

#### `vm(code)`

**Parameters** `code` (*str*) – a Jx9 script.

**Returns** a *VM* instance with the compiled script.

Compile the given Jx9 script and return an initialized *VM* instance.

Usage:

```
script = "$users = db_fetch_all('users');"
with db.vm(script) as vm:
    vm.execute()
    users = vm['users']
```

#### `collection(name)`

**Parameters** `name` (*str*) – The name of the collection.

Factory method for instantiating a *Collection* for working with a collection of JSON objects.

Usage:

```
Users = db.collection('users')

# Fetch all records in the collection.
all_users = Users.all()

# Create a new record.
Users.store({'name': 'Charlie', 'activities': ['reading', 'programming']})
```

See the *Collection* docs for more examples.

**keys** ()

**Returns** A generator that successively yields the keys in the database.

**values** ()

**Returns** A generator that successively yields the values in the database.

**items** ()

**Returns** A generator that successively yields tuples containing the keys and values in the database.

**update** (*data*)

**Parameters** `data` (*dict*) – Dictionary of data to store in the database. If any keys in `data` already exist, the values will be overwritten.

**\_\_iter\_\_** ()

UnQLite databases can be iterated over. The iterator is a *Cursor*, and will yield 2-tuples of keys and values:

```
db = UnQLite('my_db.udb')
for (key, value) in db:
    print key, '=>', value
```

**range** (*start\_key*, *end\_key* [, *include\_end\_key=True* ])

Iterate over a range of key/value pairs in the database.

```
for key, value in db.range('d.20140101', 'd.20140201', False):
    calculate_daily_aggregate(key, value)
```

**\_\_len\_\_** ()

Return the number of records in the database.

**Warning:** This method calculates the length by iterating and counting every record. At the time of writing, there is no C API for calculating the size of the database.

**flush** ()

Delete all records in the database.

**Warning:** This method works by iterating through all the records and deleting them one-by-one. At the time of writing there is no API for bulk deletes. If you are worried about speed, simply delete the database file and re-open it.

**random\_string** (*nbytes*)

**Parameters** *nbytes* (*int*) – number of bytes to generate

**Returns** a string consisting of random lower-case letters (a-z).

**random\_number** ()

**Returns** a random positive integer

**lib\_version** ()

**Returns** The UnQLite library version.

**class Transaction** (*unqlite*)

**Parameters** *unqlite* (*UnQLite*) – An *UnQLite* instance.

Context-manager for executing wrapped blocks in a transaction. Rather than instantiating this object directly, it is recommended that you use *UnQLite.transaction()*.

Example:

```
with db.transaction():
    db['from_acct'] = db['from_acct'] + 100
    db['to_acct'] = db['to_acct'] - 100
```

To roll back changes inside a transaction, call *UnQLite.rollback()*:

```
with db.transaction():
    db['from_acct'] = db['from_acct'] + 100
    db['to_acct'] = db['to_acct'] - 100
    if int(db['to_acct']) < 0:
        db.rollback() # Not enough funds!
```

**class Cursor** (*unqlite*)

**Parameters** *unqlite* (*UnQLite*) – An *UnQLite* instance.

Create a cursor. Cursors should generally be used as context managers.

Rather than instantiating this class directly, it is preferable to call the factory method *UnQLite.cursor()*.

```
for i in range(4):
    db['k%d' % i] = str(i)

# Cursor support iteration, which returns key/value pairs.
with db.cursor() as cursor:
    all_items = [(key, value) for key, value in cursor]

    # You can seek to a record, then iterate to retrieve a portion
    # of results.
    cursor.seek('k2')
    k2, k3 = [key for key, _ in cursor]

# Previous cursor was closed automatically, open a new one.
with db.cursor() as cursor:
```

(continues on next page)

(continued from previous page)

```

cursor.seek('k1') # Jump to the 2nd record, k1
assert cursor.key() == 'k1' # Use the key()/value() methods.
assert cursor.value() == '1'

cursor.delete() # Delete k1/v1
cursor.first() # Cursor now points to k0/0
cursor.next() # Cursor jumps to k2/2 since k1/1 is deleted.
assert cursor.key() == 'k2'

keys = [key for key, value in cursor] # Cursor iterates from k2->k3
assert keys == ['k2', 'k3']

```

**reset()**

Reset the cursor, which also resets the pointer to the first record.

**seek**(key[, flags=UNQLITE\_CURSOR\_MATCH\_EXACT])

Advance the cursor to the given key using the comparison method described in the flags.

A detailed description of alternate flags and their usage can be found in the `unqlite_kv_cursor` docs.

Usage:

```

with db.cursor() as cursor:
    cursor.seek('item.20140101')
    while cursor.is_valid():
        data_for_day = cursor.value()
        # do something with data for day
        handle_data(data_for_day)
        if cursor.key() == 'item.20140201':
            break
        else:
            cursor.next()

```

**first()**

Place cursor at the first record.

**last()**

Place cursor at the last record.

**next\_entry()**

Move the cursor to the next record.

**Raises** `StopIteration` if you have gone past the last record.

**previous\_entry()**

Move the cursor to the previous record.

**Raises** `StopIteration` if you have gone past the first record.

**is\_valid()**

**Return type** `bool`

Indicate whether this cursor is pointing to a valid record.

**\_\_iter\_\_()**

Iterate over the keys in the database, returning 2-tuples of key/value.

---

**Note:** Iteration will begin wherever the cursor is currently pointing, rather than starting at the first record.

---

**key()**

Return the key of the current record.

**value()**

Return the value of the current record.

**delete()**

Delete the record currently pointed to by the cursor.

**Warning:** The `delete()` method is a little weird in that it only seems to work if you explicitly call `seek()` beforehand.

**fetch\_until**(*stop\_key*[, *include\_stop\_key=True*])**Parameters**

- **stop\_key** (*str*) – The key at which the cursor should stop iterating.
- **include\_stop\_key** (*bool*) – Whether the stop key/value pair should be returned.

Yield successive key/value pairs until the `stop_key` is reached. By default the `stop_key` and associated value will be returned, but this behavior can be controlled using the `include_stop_key` flag.

**class VM**(*unqlite*, *code*)**Parameters**

- **unqlite** (`UnQLite`) – An `UnQLite` instance.
- **code** (*str*) – A Jx9 script.

Python wrapper around an UnQLite virtual machine. The VM is the primary means of executing Jx9 scripts and interacting with the JSON document store.

VM instances should not be instantiated directly, but created by calling `UnQLite.vm()`.

---

**Note:** For information on Jx9 scripting, see the [Jx9 docs](#).

---

Example of passing values into a Jx9 script prior to execution, then extracting values afterwards:

```
script = """
    $collection = 'users';
    db_create($collection);
    db_store($collection, $values);
    $users = db_fetch_all($collection);
    """

# We can pass all sorts of interesting data in to our script.
values = [
    {'username': 'huey', 'color': 'white'},
    {'username': 'mickey', 'color': 'black'},
]

with db.vm(script) as vm:
    # Set the value of the `values` variable in the Jx9 script:
    vm['values'] = values

    # Execute the script, which creates the collection and stores
    # the two records.
```

(continues on next page)

(continued from previous page)

```

vm.execute()

# After execution, we can extract the value of the `users` variable.
users = vm['users']

# Jx9 document store assigns a unique 0-based id to each record
# in a collection. The extracted variable `users` will now equal:
print users == [
    {'username': 'huey', 'color': 'white', '__id': 0},
    {'username': 'mickey', 'color': 'black', '__id': 1},
] # prints `True`

```

When using the VM outside of a context-manager, the following steps should be followed:

1. instantiate *VM* with a Jx9 script.
2. call *VM.compile()* to compile the script.
3. optional: set one or more values to pass to the Jx9 script using *VM.set\_value()* or *VM.set\_values()*.
4. call *VM.execute()* to execute the script.
5. optional: read one or more values back from the VM context, for example a return value for a function call, using *VM.get\_value()*.
6. call *VM.reset()* and return to step 4 if you intend to re-execute the script, or call *VM.close()* to free the VM and associated resources.

**execute()**

Execute the compiled Jx9 script.

**close()**

Release the VM, deallocating associated memory.

---

**Note:** When using the VM as a context manager, this is handled automatically.

---

**\_\_enter\_\_()**

Typically the VM should be used as a context manager. The context manager API handles compiling the Jx9 code and releasing the data-structures afterwards.

```

with db.vm(jx9_script) as vm:
    vm.execute()

```

**set\_value** (*name*, *value*)**Parameters**

- **name** (*str*) – A variable name
- **value** – Value to pass in to the scope of the Jx9 script, which should be either a string, int, float, bool, list, dict, or None (basically a valid JSON type).

Set the value of a Jx9 variable. You can also use dictionary-style assignment to set the value.

**set\_values** (*mapping*)

**Parameters** **mapping** (*dict*) – Dictionary of name to value to pass in to the Jx9 script. This method is short-hand for making multiple calls to *set\_value()*.

Set multiple Jx9 variables.

**get\_value** (*name*)

**Parameters** **name** (*str*) – A variable name

Retrieve the value of a variable after the execution of a Jx9 script. You can also use dictionary-style lookup to retrieve the value.

**compile** (*code*)

**Parameters** **code** (*str*) – A Jx9 script.

Compile the Jx9 script and initialize the VM.

**Warning:** It is not necessary to call this method yourself, as it is called automatically when the VM is used as a context manager.

---

**Note:** This does not execute the code. To execute the code, you must also call `VM.execute()`.

---

**class Collection** (*unqlite, name*)

**Parameters**

- **unqlite** – a *UnQLite* instance
- **name** (*str*) – the name of the collection

Perform common operations on a JSON document collection.

---

**Note:** Rather than instantiating this class directly, use the factory method *UnQLite.collection()*.

---

Basic operations:

```
>>> users = db.collection('users')
>>> users.create() # Create the collection if it does not exist.
>>> users.exists()
True

>>> users.store([
...     {'name': 'Charlie', 'color': 'green'},
...     {'name': 'Huey', 'color': 'white'},
...     {'name': 'Mickey', 'color': 'black'}])
True
>>> users.store({'name': 'Leslie', 'color': 'also green'})
True

>>> users.fetch(0) # Fetch the first record (user "__id" = 0).
{'__id': 0, 'color': 'green', 'name': 'Charlie'}

>>> users.delete(0) # Delete the first record (user "__id" = 0).
True
>>> users.delete(users.last_record_id()) # Delete the last record.
True

>>> users.update(1, {'color': 'white', 'name': 'Baby Huey'})
```

(continues on next page)

(continued from previous page)

```
True

>>> users.all()
[{'__id': 1, 'color': 'white', 'name': 'Baby Huey'},
 {'__id': 2, 'color': 'black', 'name': 'Mickey'}]

>>> for user in users:
...     print(user)
{'__id': 1, 'color': 'white', 'name': 'Baby Huey'}
{'__id': 2, 'color': 'black', 'name': 'Mickey'}

>>> users.filter(lambda obj: obj['name'].startswith('B'))
[{'__id': 1, 'color': 'white', 'name': 'Baby Huey'}]
```

**all()****Returns** list containing all records in the collection.As of 0.9.0, it is also possible to iterate the collection using a Python iterable. See *iterator()*.**iterator()****Returns** *CollectionIterator* for iterating over the records in the collection.

```
>>> reg = db.collection('register')
>>> reg.create()
>>> reg.store([{'key': 'k0'}, {'key': 'k1'}, {'key': 'k2'}])

>>> it = reg.iterator()
>>> for row in it:
...     print(row)
{'__id': 0, 'key': 'k0'}
{'__id': 1, 'key': 'k1'}
{'__id': 2, 'key': 'k2'}

>>> list(it) # We can re-use the iterator.
[{'__id': 0, 'key': 'k0'},
 {'__id': 1, 'key': 'k1'},
 {'__id': 2, 'key': 'k2'}]

>>> for row in reg: # Iterating over collection is fine, too.
...     print(row)
{'__id': 0, 'key': 'k0'}
{'__id': 1, 'key': 'k1'}
{'__id': 2, 'key': 'k2'}
```

**filter** (*filter\_fn*)

Filter the list of records using the provided function (or lambda). Your filter function should accept a single parameter, which will be the record, and return a boolean value indicating whether the record should be returned.

**Parameters** *filter\_fn* – callable that accepts record and returns boolean.**Returns** list of matching records.

Example:

```
>>> users.filter(lambda user: user['is_admin'] == True)
[{'__id': 0, 'username': 'Huey', 'is_admin': True},
```

(continues on next page)



(continued from previous page)

```
{'__id': 3, 'username': 'Zaizee', 'is_admin': True},
{'__id': 4, 'username': 'Charlie', 'is_admin': True}]
```

**create()**

Create the collection if it does not exist.

**Returns** true on success, false if collection already exists.**drop()**

Drop the collection, deleting all records.

**Returns** true on success, false if collection does not exist.**exists()****Returns** boolean value indicating whether the collection exists.**creation\_date()****Returns** the timestamp the collection was created (if exists) or None.**set\_schema** ([\_schema=None[, \*\*kwargs]])Set the schema metadata associated with the collection. The schema is **not enforced by the database engine**, and is for metadata purposes.**Parameters**

- **\_schema** (*dict*) – a mapping of field-name to data-type, or
- **kwargs** – key/value mapping of field to data-type.

**Returns** true on success, false on failure.**get\_schema()**

Get the schema metadata associated with the collection.

**Returns** mapping of field-name to data-type on success, or None.**last\_record\_id()****Returns** The integer ID of the last record stored in the collection.**current\_record\_id()****Returns** The integer ID of the record pointed to by the active cursor.**reset\_cursor()**

Reset the collection cursor to point to the first record in the collection.

**\_\_len\_\_()**

Return the number of records in the collection.

**\_\_iter\_\_()**Return a *CollectionIterator* for iterating over the records in the collection.**fetch(record\_id)**

Return the record with the given id.

```
>>> users = db.collection('users')
>>> users.fetch(0) # Fetch the first record in collection (id=0).
{'name': 'Charlie', 'color': 'green', '__id': 0}

>>> users[1] # You can also use dictionary-style lookup.
{'name': 'Huey', 'color': 'white', '__id': 1}
```

You can also use the dictionary API:

```
>>> users[0]
{'name': 'Charlie', 'color': 'green', '__id': 0}
```

**store** (*record* [, *return\_id=True* ])

**Parameters**

- **record** – Either a dictionary (single-record), or a list of dictionaries.
- **return\_id** (*bool*) – Return the ID of the newly-created object.

**Returns** New object's ID, or a boolean indicating if the record was stored successfully.

Store the record(s) in the collection.

```
>>> users = db.collection('users')
>>> users.store({'name': 'Charlie', 'color': 'green'})
True
>>> users.store([
...     {'name': 'Huey', 'color': 'white'},
...     {'name': 'Mickey', 'color': 'black'}])
True
```

**update** (*record\_id*, *record*)

**Parameters**

- **record\_id** – The ID of the record to update.
- **record** – A dictionary of data to update the given record ID.

**Returns** Boolean value indicating if the update was successful.

Update the data stored for the given *record\_id*. The data is completely replaced, rather than being appended to.

```
>>> users = db.collection('users')
>>> users.store({'name': 'Charlie'})
True
>>> users.update(users.last_record_id(), {'name': 'Chuck'})
True
>>> users.fetch(users.last_record_id())
{'__id': 0, 'name': 'Chuck'}
```

You can also use dictionary-style assignment using the record ID:

```
>>> users[0] = {'name': 'Charles'} # Can also use item assignment by id.
>>> users[0]
{'__id': 0, 'name': 'Charles'}
```

**delete** (*record\_id*)

**Parameters** **record\_id** – The database-provided ID of the record to delete.

**Returns** Boolean indicating if the record was deleted successfully.

Delete the record with the given id.

```
>>> data = db.collection('data')
>>> data.create()
```

(continues on next page)

(continued from previous page)

```
>>> data.store({'foo': 'bar'})
True
>>> data.delete(data.last_record_id())
True
>>> data.all()
[]
```

You can also use the dictionary API:

```
>>> del users[1] # Delete user object with `__id=1`.
```

**fetch\_current()**

Fetch the record pointed to by the collection cursor.

**..warning::** This method does not work as intended as the VM is reset for each script execution.

**class CollectionIterator(Collection)**

Python iterator that returns rows from a collection. This class should not be instantiated directly, but via `Collection.iterator()` or implicitly by iterating directly over a `Collection`.



## CHAPTER 4

---

### Indices and tables

---

- genindex
- modindex
- search



## Symbols

\_\_enter\_\_() (*UnQLite method*), 11  
 \_\_enter\_\_() (*VM method*), 18  
 \_\_iter\_\_() (*Collection method*), 21  
 \_\_iter\_\_() (*Cursor method*), 16  
 \_\_iter\_\_() (*UnQLite method*), 14  
 \_\_len\_\_() (*Collection method*), 21  
 \_\_len\_\_() (*UnQLite method*), 14

## A

all() (*Collection method*), 20  
 append() (*UnQLite method*), 12

## B

begin() (*UnQLite method*), 12

## C

close() (*UnQLite method*), 10  
 close() (*VM method*), 18  
 Collection (*built-in class*), 19  
 collection() (*UnQLite method*), 13  
 CollectionIterator (*built-in class*), 23  
 commit() (*UnQLite method*), 12  
 commit\_on\_success() (*UnQLite method*), 13  
 compile() (*VM method*), 19  
 create() (*Collection method*), 21  
 creation\_date() (*Collection method*), 21  
 current\_record\_id() (*Collection method*), 21  
 Cursor (*built-in class*), 15  
 cursor() (*UnQLite method*), 13

## D

delete() (*Collection method*), 22  
 delete() (*Cursor method*), 17  
 delete() (*UnQLite method*), 12  
 disable\_autocommit() (*UnQLite method*), 11  
 drop() (*Collection method*), 21

## E

execute() (*VM method*), 18

exists() (*Collection method*), 21  
 exists() (*UnQLite method*), 12

## F

fetch() (*Collection method*), 21  
 fetch() (*UnQLite method*), 11  
 fetch\_current() (*Collection method*), 23  
 fetch\_until() (*Cursor method*), 17  
 filter() (*Collection method*), 20  
 first() (*Cursor method*), 16  
 flush() (*UnQLite method*), 14

## G

get\_schema() (*Collection method*), 21  
 get\_value() (*VM method*), 19

## I

is\_valid() (*Cursor method*), 16  
 items() (*UnQLite method*), 14  
 iterator() (*Collection method*), 20

## K

key() (*Cursor method*), 16  
 keys() (*UnQLite method*), 14

## L

last() (*Cursor method*), 16  
 last\_record\_id() (*Collection method*), 21  
 lib\_version() (*UnQLite method*), 15

## N

next\_entry() (*Cursor method*), 16

## O

open() (*UnQLite method*), 10

## P

previous\_entry() (*Cursor method*), 16

## R

random\_number() (*UnQLite method*), 15  
random\_string() (*UnQLite method*), 15  
range() (*UnQLite method*), 14  
reset() (*Cursor method*), 16  
reset\_cursor() (*Collection method*), 21  
rollback() (*UnQLite method*), 12

## S

seek() (*Cursor method*), 16  
set\_schema() (*Collection method*), 21  
set\_value() (*VM method*), 18  
set\_values() (*VM method*), 18  
store() (*Collection method*), 22  
store() (*UnQLite method*), 11

## T

Transaction (*built-in class*), 15  
transaction() (*UnQLite method*), 12

## U

UnQLite (*built-in class*), 9  
update() (*Collection method*), 22  
update() (*UnQLite method*), 14

## V

value() (*Cursor method*), 17  
values() (*UnQLite method*), 14  
VM (*built-in class*), 17  
vm() (*UnQLite method*), 13