

---

# **ndn-python-repo**

**Zhaoning Kong**

**Aug 02, 2021**



## CONTENTS:

<b>1</b>	<b>ndn-python-repo</b>	<b>1</b>
<b>2</b>	<b>Install and Run</b>	<b>3</b>
2.1	Install . . . . .	3
2.2	Migrate from repo-ng . . . . .	3
2.3	Instruction for developers . . . . .	3
<b>3</b>	<b>Configuration</b>	<b>5</b>
3.1	Repo namespace . . . . .	5
3.2	Repo prefix registration . . . . .	5
3.3	Choose the backend database . . . . .	6
3.4	TCP bulk insert . . . . .	6
3.5	Logging . . . . .	6
3.6	systemd . . . . .	7
<b>4</b>	<b>Specification</b>	<b>9</b>
4.1	Encoding . . . . .	9
4.2	Insert . . . . .	10
4.3	Delete . . . . .	11
4.4	Check . . . . .	12
4.5	TCP bulk insert . . . . .	12
<b>5</b>	<b>Miscellaneous packages</b>	<b>13</b>
5.1	Client-side packages . . . . .	13
5.2	ConcurrentFetcher package . . . . .	16
5.3	PubSub package . . . . .	16
5.4	Storage package . . . . .	19
<b>6</b>	<b>Examples</b>	<b>21</b>
<b>7</b>	<b>Authors</b>	<b>23</b>
<b>8</b>	<b>Future Plans</b>	<b>25</b>
<b>9</b>	<b>Indices and tables</b>	<b>27</b>
	<b>Python Module Index</b>	<b>29</b>
	<b>Index</b>	<b>31</b>



## **NDN-PYTHON-REPO**

A Named Data Networking (NDN) Repo implementation using [python-ndn](#).

Please see our [documentation](#) if you have any issues.



## INSTALL AND RUN

### 2.1 Install

Install the latest release with pip:

```
$ /usr/bin/pip3 install ndn-python-repo
```

Optionally, you can install the latest development version from source:

```
$ git clone https://github.com/JonnyKong/ndn-python-repo.git
$ cd ndn-python-repo && /usr/bin/pip3 install -e .
```

### 2.2 Migrate from repo-ng

ndn-python-repo provides a script to migrate existing data from repo-ng:

```
$ ndn-python-repo-port -d <path-to-repo-ng-dbfile> \
                       -a <ndn-python-repo-ipaddr> \
                       -p <ndn-python-repo-port>
```

It takes as input a repo-ng database file, reads the Data packets and pipe them through TCP bulk insert into the new repo.

### 2.3 Instruction for developers

Setup virtual environment with editable installation:

```
$ python3 -m venv venv
$ . venv/bin/activate
$ pip3 install -e .
```

Run all tests:

```
$ pip3 install pytest
$ pytest
```

Compile the documentation with Sphinx:

```
$ cd docs && pip3 install -r requirements.txt
$ make html
$ open _build/html/index.html
```



## CONFIGURATION

You can configure `ndn-python-repo` with a config file, by specifying the path to the file when starting a repo instance:

```
$ ndn-python-repo -c <config_file>
```

A sample config file is provided at `ndn_python_repo/ndn-python-repo.conf.sample`.

If no config file is given on the command line, this sample config file will be used by default.

### 3.1 Repo namespace

Specify the name of a repo in the config file. For example:

```
repo_config:  
  # the repo's routable prefix  
  repo_name: 'testrepo'
```

Another option is to specify the repo name when starting a repo on the command line. This overrides the repo name in the config file:

```
$ ndn-python-repo -r "/name_foo"
```

### 3.2 Repo prefix registration

By default, the repo registers the root prefix `/`.

Alternatively, you can configure repo such that it doesn't register the root prefix:

```
repo_config:  
  register_root: False
```

If `register_root` is set to `False`, the client is responsible of telling the repo which prefix to register or unregister every time in `RepoCommandParameter`. See *Insert* and *Delete* for details.

### 3.3 Choose the backend database

The ndn-python-repo uses one of the three backend databases:

- SQLite3 (default)
- leveldb
- MongoDB

To use non-default databases, perform the following steps:

1. Install ndn-python-repo with additional database support that you need:

```
$ /usr/bin/pip3 install ndn-python-repo[leveldb]
$ /usr/bin/pip3 install ndn-python-repo[mongodb]
```

2. Specify the database selection and database file in the config file. For example:

```
db_config:
  # choose one among sqlite3, leveldb, and mongodb
  db_type: 'mongodb'

  # only the chosen db's config will be read
  mongodb:
    'db': 'repo'
    'collection': 'data'
```

### 3.4 TCP bulk insert

By default, the repo listens on 0.0.0.0:7376 for TCP bulk insert. You can configure in the config file which address the repo listens on. For example:

```
tcp_bulk_insert:
  'addr': '127.0.0.1'
  'port': '7377'
```

### 3.5 Logging

Repo uses the python logging module, and by default logs all messages of and above level INFO to stdout. You can override the default options in the config file. For example:

```
logging_config:
  'level': 'WARNING'
  'file': '/var/log/ndn/ndn-python-repo/repo.log'
```

## 3.6 systemd

To run ndn-python-repo with systemd on Linux, perform the following steps:

1. Run the provided script to install the systemd script to /etc/systemd/system/:

```
$ sudo ndn-python-repo-install
```

2. Then, start, stop, and monitor a repo instance with systemd:

```
$ sudo systemctl start ndn-python-repo  
$ sudo systemctl stop ndn-python-repo  
$ sudo systemctl status ndn-python-repo
```

3. Examine logs:

```
$ sudo journalctl -u ndn-python-repo.service
```



## SPECIFICATION

### 4.1 Encoding

Most repo commands and status reports are Data packets whose Content contains RepoCommandParameter or RepoCommandResponse structure. These structures are defined as follows:

```
RepoCommandParameter =  
  [Name]  
  [ForwardingHint]  
  [StartBlockId]  
  [EndBlockId]  
  [ProcessId]  
  [RegisterPrefix]  
  [CheckPrefix]
```

```
RepoCommandResponse =  
  [Name]  
  [StartBlockId]  
  [EndBlockId]  
  [ProcessId]  
  [StatusCode]  
  [InsertNum]  
  [DeleteNum]
```

```
ForwardingHint = FORWARDING-HINT-TYPE TLV-LENGTH Name
```

```
StartBlockId = START-BLOCK-ID-TYPE TLV-LENGTH NonNegativeInteger
```

```
EndBlockId = END-BLOCK-ID-TYPE TLV-LENGTH NonNegativeInteger
```

```
ProcessId = PROCESS-ID-TYPE TLV-LENGTH NonNegativeInteger
```

```
RegisterPrefix = REGISTER-PREFIX-TYPE TLV-LENGTH Name
```

```
CheckPrefix = CHECK-PREFIX-TYPE TLV-LENGTH Name
```

```
StatusCode = STATUS-CODE-TYPE TLV-LENGTH NonNegativeInteger
```

```
InsertNum = INSERT-NUM-TYPE TLV-LENGTH NonNegativeInteger
```

```
DeleteNum = DELETE-NUM-TYPE TLV-LENGTH NonNegativeInteger
```

The type number assignments are as follows:

type	Assigned number (decimal)	Assigned number (hexadecimal)
START-BLOCK-ID-TYPE	204	0xCC
END-BLOCK-ID-TYPE	205	0xCD
PROCESS-ID-TYPE	206	0xCE
STATUS-CODE-TYPE	208	0xD0
INSERT-NUM-TYPE	209	0xD1
DELETE-NUM-TYPE	210	0xD2
FORWARDING-HINT-TYPE	211	0xD3
REGISTER-PREFIX-TYPE	212	0xD4
CHECK-PREFIX-TYPE	213	0xD5

## 4.2 Insert

Repo insertion process makes use of the *PubSub package*.

1. The repo subscribes to the topic `/<repo_name>/insert`.
2. The client publishes a message to the topic `/<repo_name>/insert`. The message payload is `RepoCommandParameter` with the following fields:
  - `name`: either a Data packet name, or a name prefix of Data packets.
  - `start_block_id` (Optional): inclusive start segment number.
  - `end_block_id` (Optional): inclusive end segment number.
  - `forwarding_hint` (Optional): forwarding hint for Data fetching. This is useful in two scenarios:
    - The producer choose not to announce its name prefix, but only allow the repo to reach it via forwarding hint.
    - The name prefix is already announced by repo node(s), but the producer in another node wants to insert to the repo.
  - `register_prefix` (Optional): if repo doesn't register the root prefix (*Configuration* `register_root` is disabled), client can tell repo to register this prefix.
  - `check_prefix`: a prefix of status check topic name. See *Check*.
  - `process_id`: a random byte string to identify this insertion process.
3. The repo fetches and inserts Data packets according to given parameters.
  - If both `start_block_id` and `end_block_id` are omitted, the repo fetches a single packet identified in `name` parameter. The insertion process succeeds when this packet is received.
  - If `start_block_id` is specified but `end_block_id` is omitted, the repo starts fetching segments starting from `/name/start_block_id`, and increments segment number after each packet. When an Interest receives timeout or nack 3 times, the insertion process stops and is considered successful.
  - Otherwise, the repo fetches all segments between `/name/start_block_id` and `/name/end_block_id`. If `start_block_id` is omitted, it defaults to 0. The insertion process succeeds when all packets are received.
  - Segment numbers are encoded in accordance with *NDN naming conventions rev2*.

### 4.2.1 Insert status check

The client can use the *Check* protocol to check the progress of an insertion process. The insertion check response message payload is *RepoCommandResponse* with the following fields:

- *status\_code*: status code, as defined on *Check*.
- *insert\_num*: number of Data packets received by the repo so far.

## 4.3 Delete

Repo deletion process makes use of the *PubSub package*.

1. The repo subscribes to the topic `/<repo_name>/delete`.
2. The client publishes a message to the topic `/<repo_name>/delete`. The message payload is *RepoCommandParameter* with the following fields:
  - *name*: either a Data packet name, or a name prefix of Data packets.
  - *start\_block\_id* (Optional): inclusive start segment number.
  - *end\_block\_id* (Optional): inclusive end segment number.
  - *register\_prefix* (Optional): if repo doesn't register the root prefix (*Configuration register\_root* is disabled), client can tell repo to unregister this prefix.
  - *check\_prefix*: a prefix of status check topic name. See *Check*.
  - *process\_id*: a random byte string to identify this deletion process.
3. The repo deletes Data packets according to given parameters.
  - If both *start\_block\_id* and *end\_block\_id* are omitted, the repo deletes a single packet identified in *name* parameter. The deletion process succeeds when this packet is deleted.
  - If *start\_block\_id* is specified but *end\_block\_id* is omitted, the repo starts deleting segments starting from `/name/start_block_id`, and increments segment number after each packet. When a name query does not find an existing segment, the deletion process stops and is considered successful.
  - Otherwise, the repo fetches all segments between `/name/start_block_id` and `/name/end_block_id`. If *start\_block\_id* is omitted, it defaults to 0. The deletion process succeeds when all packets are deleted.
  - Segment numbers are encoded in accordance with [NDN naming conventions rev2](#).

### 4.3.1 Delete status check

The client can use the *Check* protocol to check the progress of a deletion process. The deletion check response message payload is *RepoCommandResponse* with the following fields:

- *status\_code*: status code, as defined on *Check*.
- *delete\_num*: number of Data packets deleted by the repo so far.

## 4.4 Check

The check protocol is used by clients to check the progress of a insertion or deletion process.

1. Each insert/delete command has `check_prefix` and `process_id` parameters. Status check messages are published to the topic `/<check_prefix>/check/<process_id>`, derived from these parameters.
2. After receiving an insert/delete command, the repo periodically publishes the status of the insertion/deletion process to the topic. The message payload is `RepoCommandResponse`.
3. The client can subscribe to the topic to receive status updates.

### 4.4.1 Status Code Definition

The status code definitions are as follows:

StatusCode	
100	The command is OK
200	All the data has been inserted / deleted
300	The insertion / deletion is in progress

## 4.5 TCP bulk insert



## MISCELLANEOUS PACKAGES

### 5.1 Client-side packages

#### 5.1.1 Introduction

Application built with python-ndn can make use of the client packages provided.

There are four parts:

1. **PutfileClient**: insert files into the repo.
2. **GetfileClient**: get files from the repo.
3. **DeleteClient**: delete data packets from the repo.
4. **CommandChecker**: check process status from the repo.

The example programs in `examples/` illustrate how to use these packages.

Note that the type `Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]` in the documentation is equivalent to the `ndn.name.NonStrictName` type.

#### 5.1.2 Reference

**class** `ndn_python_repo.clients.putfile.PutfileClient`(*app, prefix, repo\_name*)

A client to insert files into the repo.

##### Parameters

- **app** (`NDNApp`) – `NDNApp`.
- **prefix** (`Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]`) – `NonStrictName`. The name of this client
- **repo\_name** (`Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]`) – `NonStrictName`. Routable name to remote repo.

**async insert\_file**(*file\_path, name\_at\_repo, segment\_size, freshness\_period, cpu\_count, forwarding\_hint=None, register\_prefix=None, check\_prefix=None*)

Insert a file to remote repo.

##### Parameters

- **file\_path** (`str`) – Local FS path to file to insert.
- **name\_at\_repo** (`Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]`) – `NonStrictName`. Name used to store file at repo.

- **segment\_size** (int) – Max size of data packets.
- **freshness\_period** (int) – Freshness of data packets.
- **cpu\_count** (int) – Cores used for converting file to TLV format.
- **forwarding\_hint** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview, None]) – NonStrictName. The forwarding hint the repo uses when fetching data.
- **register\_prefix** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview, None]) – NonStrictName. If repo is configured with `register_root=False`, it registers `register_prefix` after receiving the insertion command.
- **check\_prefix** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview, None]) – NonStrictName. The repo will publish process check messages under `<check_prefix>/check`. It is necessary to specify this value in the param, instead of using a predefined prefix, to make sure the subscriber can register this prefix under the NDN prefix registration security model. If not specified, default value is the client prefix.

**Return type** int

**Returns** Number of packets inserted.

**class** `ndn_python_repo.clients.getfile.GetfileClient`(*app, repo\_name*)  
 This client fetches a file from the repo, and save it to working directory.

A client to retrieve files from the remote repo.

**Parameters**

- **app** (NDNApp) – NDNApp.
- **repo\_name** – NonStrictName. Rutable name to remote repo.

**async** `fetch_file`(*name\_at\_repo, local\_filename=None, overwrite=False*)  
 Fetch a file from remote repo, and write to the current working directory.

**Parameters**

- **name\_at\_repo** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]) – NonStrictName. The name with which this file is stored in the repo.
- **local\_filename** (Optional[str]) – str. The filename of the retrieved file on the local file system.
- **overwrite** – If true, existing files are replaced.

**class** `ndn_python_repo.clients.delete.DeleteClient`(*app, prefix, repo\_name*)  
 This client deletes data packets from the remote repo.

**Parameters**

- **app** (NDNApp) – NDNApp.
- **repo\_name** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]) – NonStrictName. Rutable name to remote repo.

**async** `delete_file`(*prefix, start\_block\_id=None, end\_block\_id=None, register\_prefix=None, check\_prefix=None*)

Delete from repo packets between “<name\_at\_repo>/<start\_block\_id>” and “<name\_at\_repo>/<end\_block\_id>” inclusively.

**Parameters**

- **prefix** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]) – NonStrictName. The name of the file stored in the remote repo.
- **start\_block\_id** (Optional[int]) – int. Default value is 0.
- **end\_block\_id** (Optional[int]) – int. If not specified, repo will attempt to delete all data packets with segment number starting from *start\_block\_id* continuously.
- **register\_prefix** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview, None]) – If repo is configured with `register_root=False`, it unregisters `register_prefix` after receiving the deletion command.
- **check\_prefix** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview, None]) – NonStrictName. The repo will publish process check messages under `<check_prefix>/check`. It is necessary to specify this value in the param, instead of using a predefined prefix, to make sure the subscriber can register this prefix under the NDN prefix registration security model. If not specified, default value is the client prefix.

**Return type** int

**Returns** Number of deleted packets.

`class ndn_python_repo.clients.command_checker.CommandChecker(app)`

This client sends check interests to the repo.

**Parameters** `app` (NDNApp) – NDNApp.

`async check_delete(repo_name, process_id)`

Check the status of a delete process.

**Parameters**

- **repo\_name** – NonStrictName. The name of the remote repo.
- **process\_id** (int) – int. The process id of the process to check.

**Return type** RepoCommandResponse

**Returns** The response from the repo.

`async check_insert(repo_name, process_id)`

Check the status of an insert process.

**Parameters**

- **repo\_name** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]) – NonStrictName. The name of the remote repo.
- **process\_id** (int) – int. The process id of the process to check.

**Return type** RepoCommandResponse

**Returns** The response from the repo.

## 5.2 ConcurrentFetcher package

### 5.2.1 Introduction

Fetch data packets in parallel using a fixed window size.

Note that the type `Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]` in the documentation is equivalent to the `ndn.name.NonStrictName` type.

### 5.2.2 Reference

`ndn_python_repo.utils.concurrent_fetcher`(*app, name, start\_block\_id, end\_block\_id, semaphore, \*\*kwargs*)

An async-generator to fetch data packets between “*name/start\_block\_id*” and “*name/end\_block\_id*” concurrently.

#### Parameters

- **app** (NDNApp) – NDNApp.
- **name** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]) – NonStrictName. Name prefix of Data.
- **start\_block\_id** (int) – int. The start segment number.
- **end\_block\_id** (Optional[int]) – Optional[int]. The end segment number. If not specified, continue fetching until an interest receives timeout or nack or 3 times.

**Returns** Yield (FormalName, MetaInfo, Content, RawPacket) tuples in order.

## 5.3 PubSub package

### 5.3.1 Introduction

The PubSub package provides a pub-sub API with best-effort, at-most-once delivery guarantee.

If there are no subscribers reachable when a message is published, this message will not be re-transmitted.

If there are multiple subscribers reachable, the nearest subscriber will be notified of the published message in an any-cast style.

Note that the type `Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]` in the documentation is equivalent to the `ndn.name.NonStrictName` type.

### 5.3.2 Process

Under the hood the PubSub module transmits a series of Interest and Data packets:

1. The subscriber calls `subscribe(topic, cb)`. This makes the subscriber listen on “/`<topic>`/notify”.
2. The publisher invokes `publish(topic, msg)`. This method sends an Interest with name “/`<topic>`/notify”, which will be routed to a subscriber. The interest carries the following fields in its application parameters:
  - Publisher prefix: used by the subscriber to reach the publisher in the next step
  - NotifyNonce: a random bytes string, used by the publisher to de-multiplex among different publications

- Forwarding hint (optional): if publisher prefix is not announced in the routing system, publisher can provide a forwarding hint

Meanwhile, `msg` is wrapped into a Data packet named `"/<pub_prefix>/msg/<topic>/<notify_nonce>"`. Here, the data name contains `topic` to establish a binding between topic and nonce, to prevent man-in-the-middle attacks that changes the topic.

3. The subscriber receives the notification interest, constructs a new Interest `"/<pub_prefix>/msg/<topic>/<notify_nonce>"` and send it to the publisher.
4. The publisher receives the interest `"/<pub_prefix>/msg/<topic>/<notify_nonce>"`, and returns the corresponding data.
5. The subscriber receives the data, and invokes `cb(data.content)` to hand the message to the application.
6. The publisher receives the acknowledgement Data packet, and erases the soft state.

### 5.3.3 Encoding

The notify Interest's application parameter is encoded as follows:

```
NotifyAppParam = DATA-TYPE TLV-LENGTH
  [PublisherPrefix]
  [NotifyNonce]
  [PublisherFwdHint]

PublisherPrefix = Name

NotifyNonce = NOTIFY-NONCE-TYPE TLV-LENGTH Bytes

PublisherFwdHint = PUBLISHER-FWD-HINT-TYPE TLV-LENGTH Name
```

The type number assignments are as follows:

type	Assigned number (decimal)	Assigned number (hexadecimal)
NOTIFY-NONCE-TYPE	128	0x80
PUBLISHER-FWD-HINT-TYPE	211	0xD3

### 5.3.4 Reference

**class** `ndn_python_repo.utils.PubSub`(*app, prefix=None, forwarding\_hint=None*)

Initialize a PubSub instance with identity `prefix` and can be reached at `forwarding_hint`. TODO: support msg larger than MTU

**Parameters**

- **app** (NDNApp) – NDNApp.
- **prefix** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview, None]) – NonStrictName. The identity of this PubSub instance. The publisher needs a prefix under which can publish data. Note that you cannot initialize two PubSub instances with the same `prefix` on the same node, since it will cause double registration error.

- **forwarding\_hint** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview, None]) – NonStrictName. When working as publisher, if `prefix` is not reachable, the subscriber can use `forwarding_hint` to reach the publisher.

**async publish**(*topic, msg*)

Publish `msg` to `topic`. Make several attempts until the subscriber returns a response.

**Parameters**

- **topic** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]) – NonStrictName. The topic to publish `msg` to.
- **msg** (bytes) – bytes. The message to publish. The pub-sub API does not make any assumptions on the format of this message.

**Returns** Return true if received response from a subscriber.

**set\_base\_prefix**(*prefix*)

Avoid registering too many prefixes, by registering `prefix` with NFD. All other prefixes under `prefix` will be registered with interest filters, and will not have to be registered with NFD. Need to be called before `_wait_for_ready()`.

**Parameters** **prefix** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]) – NonStrictName. The base prefix to register.

**set\_publisher\_prefix**(*prefix*)

Set the identify of the publisher after initialization. Need to be called before `_wait_for_ready()`.

**Parameters** **prefix** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]) – NonStrictName. The identity of this PubSub instance.

**subscribe**(*topic, cb*)

Subscribe to `topic` with `cb`.

**Parameters**

- **topic** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]) – NonStrictName. The topic to subscribe to.
- **cb** (callable) – callable. A callback that will be called when a message under `topic` is received. This function takes one bytes argument.

**unsubscribe**(*topic*)

Unsubscribe from `topic`.

**Parameters** **topic** (Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]) – NonStrictName. The topic to unsubscribe from.

**async wait\_for\_ready**()

Need to be called to wait for pub-sub to be ready.

## 5.4 Storage package

ndn-python-repo supports 3 types of databases as backends. The Storage package provides a unified key-value storage API with the following features:

- Supports `MustBeFresh`
- Supports `CanBePrefix`
- Batched writes with periodic writebacks to improve performance

The Storage class provides an interface, and is implemented by:

- `SqliteStorage`
- `LevelDBStorage`
- `MongoDBStorage`

Note that the type `Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]` in the documentation is equivalent to the `ndn.name.NonStrictName` type.

### 5.4.1 Reference

**class** `ndn_python_repo.storage.Storage`

Interface for a unified key-value storage API.

**get\_data\_packet** (*name*, *can\_be\_prefix=False*, *must\_be\_fresh=False*)

Get a data packet named *name*.

**Parameters**

- **name** (`Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]`) – `NonStrictName`. The name of the data packet.
- **can\_be\_prefix** (`bool`) – `bool`. If true, use prefix match instead of exact match.
- **must\_be\_fresh** (`bool`) – `bool`. If true, ignore expired data.

**Return type** `Optional[bytes]`

**Returns** The value of the data packet.

**put\_data\_packet** (*name*, *data*)

Insert a data packet named *name* with value *data*. This method will parse *data* to get its `freshnessPeriod`, and compute its expiration time by adding the `freshnessPeriod` to the current time.

**Parameters**

- **name** (`Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]`) – `NonStrictName`. The name of the data packet.
- **data** (`bytes`) – `bytes`. The value of the data packet.

**remove\_data\_packet** (*name*)

Remove a data packet named *name*.

**Parameters** **name** (`Union[Iterable[Union[bytes, bytearray, memoryview, str]], str, bytes, bytearray, memoryview]`) – `NonStrictName`. The name of the data packet.

**Return type** `bool`

**Returns** True if a data packet is being removed.

**class** ndn\_python\_repo.storage.**SqliteStorage**(*db\_path*)  
Init table “data” with the attribute **key** being the primary key.

**Parameters** **db\_path** (str) – str. Path to database file.

**class** ndn\_python\_repo.storage.**LevelDBStorage**(*dir*)  
Creates a LevelDB storage instance at disk location **str**.

**Parameters** **dir** (str) – str. The disk location of the database directory.

**class** ndn\_python\_repo.storage.**MongoDBStorage**(*db, collection*)  
Init a MongoDB storage with unique index on **key**.

**Parameters**

- **db** (str) – str. Database name.
- **collection** (str) – str. Collection name.



**EXAMPLES**



## AUTHORS

- Zhaoning Kong <<https://jonnykong.com>>
- Xinyu Ma <<https://zjkmxy.github.io>>
- Yufeng Zhang <<https://yufengzh.io>>
- Zhiyi Zhang <<https://zhiyi-zhang.com>>
- Davide Pesavento
- Susmit Shannigrahi
- Saurab Dulal
- Junxiao Shi <<https://yoursunny.com>>



**FUTURE PLANS**



## INDICES AND TABLES

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





## PYTHON MODULE INDEX

### n

`ndn_python_repo.clients.command_checker`, 15

`ndn_python_repo.clients.delete`, 14

`ndn_python_repo.clients.getfile`, 14

`ndn_python_repo.clients.putfile`, 13



## INDEX

- C**
- `check_delete()` (*ndn\_python\_repo.clients.command\_checker.CommandChecker* method), 15
  - `check_insert()` (*ndn\_python\_repo.clients.command\_checker.CommandChecker* method), 15
  - `CommandChecker` (class in *ndn\_python\_repo.clients.command\_checker*), 15
  - `concurrent_fetcher()` (in *ndn\_python\_repo.utils* module), 16
- D**
- `delete_file()` (*ndn\_python\_repo.clients.delete.DeleteClient* method), 14
  - `DeleteClient` (class in *ndn\_python\_repo.clients.delete*), 14
- F**
- `fetch_file()` (*ndn\_python\_repo.clients.getfile.GetfileClient* method), 14
- G**
- `get_data_packet()` (*ndn\_python\_repo.storage.Storage* method), 19
  - `GetfileClient` (class in *ndn\_python\_repo.clients.getfile*), 14
- I**
- `insert_file()` (*ndn\_python\_repo.clients.putfile.PutfileClient* method), 13
- L**
- `LevelDBStorage` (class in *ndn\_python\_repo.storage*), 20
- M**
- module
    - ndn\_python\_repo.clients.command\_checker*, 15
    - ndn\_python\_repo.clients.delete*, 14
    - ndn\_python\_repo.clients.getfile*, 14
    - ndn\_python\_repo.clients.putfile*, 13
- MongoDBStorage** (class in *ndn\_python\_repo.storage*), 20
- N**
- ndn\_python\_repo.clients.command\_checker* module, 15
  - ndn\_python\_repo.clients.delete* module, 14
  - ndn\_python\_repo.clients.getfile* module, 14
  - ndn\_python\_repo.clients.putfile* module, 13
- P**
- `publish()` (*ndn\_python\_repo.utils.PubSub* method), 18
  - `PubSub` (class in *ndn\_python\_repo.utils*), 17
  - `put_data_packet()` (*ndn\_python\_repo.storage.Storage* method), 19
  - `PutfileClient` (class in *ndn\_python\_repo.clients.putfile*), 13
- R**
- `remove_data_packet()` (*ndn\_python\_repo.storage.Storage* method), 19
- S**
- `set_base_prefix()` (*ndn\_python\_repo.utils.PubSub* method), 18
  - `set_publisher_prefix()` (*ndn\_python\_repo.utils.PubSub* method), 18
  - `SqliteStorage` (class in *ndn\_python\_repo.storage*), 19
  - `Storage` (class in *ndn\_python\_repo.storage*), 19
  - `subscribe()` (*ndn\_python\_repo.utils.PubSub* method), 18
- U**
- `unsubscribe()` (*ndn\_python\_repo.utils.PubSub* method), 18
- W**
- `wait_for_ready()` (*ndn\_python\_repo.utils.PubSub* method), 18