

---

# **hdfs3 Documentation**

*Release 0.3.0*

**Continuum Analytics**

**Sep 17, 2018**



---

## Contents

---

<b>1</b>	<b>Motivation</b>	<b>3</b>
<b>2</b>	<b>Related Work</b>	<b>5</b>



Use HDFS natively from Python.

The [Hadoop File System](#) (HDFS) is a widely deployed, distributed, data-local file system written in Java. This file system backs most clusters running Hadoop and Spark.

Pivotal produced `libhdfs3`, an alternative native C/C++ HDFS client that interacts with HDFS without the JVM, exposing first class support to non-JVM languages like Python.

This library, `hdfs3`, is a lightweight Python wrapper around the C/C++ `libhdfs3` library. It provides both direct access to `libhdfs3` from Python as well as a typical Pythonic interface.

```
>>> from hdfs3 import HDFFileSystem
>>> hdfs = HDFFileSystem(host='localhost', port=8020)
>>> hdfs.ls('/user/data')
>>> hdfs.put('local-file.txt', '/user/data/remote-file.txt')
>>> hdfs.cp('/user/data/file.txt', '/user2/data')
```

HDFS3 files comply with the Python File interface. This enables interactions with the broader ecosystem of PyData projects.

```
>>> with hdfs.open('/user/data/file.txt') as f:
...     data = f.read(1000000)

>>> with hdfs.open('/user/data/file.csv.gz') as f:
...     df = pandas.read_csv(f, compression='gzip', nrows=1000)
```



---

## Motivation

---

We choose to use an alternative *C/C++/Python* HDFS client rather than the default JVM client for the following reasons:

- **Convenience:** Interactions between Java libraries and Native (*C/C++/Python*) libraries can be cumbersome. Using a native library from Python smoothes over the experience in development, maintenance, and debugging.
- **Performance:** Native libraries like `libhdfs3` do not suffer the long JVM startup times, improving interaction.





- `libhdfs3`: The underlying C++ library from Apache HAWQ
- `snakebite`: Another Python HDFS library using Protobufs
- `Dask`: Parent project, a parallel computing library in Python
- `Dask.distributed`: Distributed computing in Python

## 2.1 Installation

### 2.1.1 Conda

Both the `hdfs3` Python library and the compiled `libhdfs3` library (and its dependencies) are available from the `conda-forge` repository using `conda`:

```
$ conda install hdfs3 -c conda-forge
```

Note that `conda` packages are only available for the `linux-64` platform.

### 2.1.2 PyPI and apt-get

Alternatively you can install the `libhdfs3.so` library using a system installer like `apt-get`:

```
echo "deb https://dl.bintray.com/wangzw/deb trusty contrib" | sudo tee /etc/apt/  
↪sources.list.d/bintray-wangzw-deb.list  
sudo apt-get install -y apt-transport-https  
sudo apt-get update  
sudo apt-get install libhdfs3 libhdfs3-dev
```

And install the Python wrapper library, `hdfs3`, with `pip`:

```
pip install hdfs3
```

### 2.1.3 Build from Source

See the [libhdfs3 installation instructions](#) to install the compiled library.

You can download the hdfs3 Python wrapper library from github and install normally:

```
git clone git@github.com:dask/hdfs3
cd hdfs3
python setup.py install
```

## 2.2 Quickstart

Import hdfs3 and connect to an HDFS cluster:

```
>>> from hdfs3 import HDFFileSystem
>>> hdfs = HDFFileSystem(host='localhost', port=8020)
```

Write data to file:

```
>>> with hdfs.open('/tmp/myfile.txt', 'wb') as f:
...     f.write(b'Hello, world!')
```

Read data back from file:

```
>>> with hdfs.open('/tmp/myfile.txt') as f:
...     print(f.read())
```

Interact with files on HDFS:

```
>>> hdfs.ls('/tmp')

>>> hdfs.put('local-file.txt', '/tmp/remote-file.txt')

>>> hdfs.cp('/tmp/remote-file.txt', '/tmp/copied-file.txt')
```

## 2.3 Examples

### 2.3.1 Word count

#### Setup

In this example, we'll use the hdfs3 library to count the number of words in text files (Enron email dataset, 6.4 GB) stored in HDFS.

Copy the text data from Amazon S3 into HDFS on the cluster:

```
$ hadoop distcp s3n://AWS_SECRET_ID:AWS_SECRET_KEY@blaze-data/enron-email hdfs:///tmp/
↪enron
```

where `AWS_SECRET_ID` and `AWS_SECRET_KEY` are valid AWS credentials.

### Code example

Import `hdfs3` and other standard libraries used in this example:

```
>>> import hdfs3
>>> from collections import defaultdict, Counter
```

Initialize a connection to HDFS, replacing `NAMENODE_HOSTNAME` and `NAMENODE_PORT` with the hostname and port (default: 8020) of the HDFS namenode.

```
>>> hdfs = hdfs3.HDFFileSystem('NAMENODE_HOSTNAME', port=NAMENODE_PORT)
```

Generate a list of filenames from the text data in HDFS:

```
>>> filenames = hdfs.glob('/tmp/enron/**/*.txt')
>>> print(filenames[:5])

['/tmp/enron/edrm-enron-v2_nemec-g_xml.zip/merged.txt',
'/tmp/enron/edrm-enron-v2_ring-r_xml.zip/merged.txt',
'/tmp/enron/edrm-enron-v2_bailey-s_xml.zip/merged.txt',
'/tmp/enron/edrm-enron-v2_fischer-m_xml.zip/merged.txt',
'/tmp/enron/edrm-enron-v2_geaccone-t_xml.zip/merged.txt']
```

Print the first 1024 bytes of the first text file:

```
>>> print(hdfs.head(filenames[0]))

b'Date: Wed, 29 Nov 2000 09:33:00 -0800 (PST)\r\nFrom: Xochitl-Alexis Velasc
o\r\nTo: Mark Knippa, Mike D Smith, Gerald Nemec, Dave S Laipple, Bo Barnwel
l\r\nCc: Melissa Jones, Iris Waser, Pat Radford, Bonnie Shumaker\r\nSubject:
Finalize ECS/EES Master Agreement\r\nX-SDOC: 161476\r\nX-ZLID: zl-edrm-enro
n-v2-nemec-g-2802.eml\r\n\r\nPlease plan to attend a meeting to finalize the
ECS/EES Master Agreement \r\ntomorrow 11/30/00 at 1:30 pm CST.\r\n\r\nI wi
ll email everyone tomorrow with location.\r\n\r\nDave-I will also email you
the call in number tomorrow.\r\n\r\nThanks\r\nXochitl\r\n\r\n*****\r\n
EDRM Enron Email Data Set has been produced in EML, PST and NSF format by ZL
Technologies, Inc. This Data Set is licensed under a Creative Commons Attri
bution 3.0 United States License <http://creativecommons.org/licenses/by/3.0
/us/> . To provide attribution, please cite to "ZL Technologies, Inc. (http:
//www.zlti.com)." \r\n*****\r\nDate: Wed, 29 Nov 2000 09:40:00 -0800 (P
ST)\r\nFrom: Jill T Zivley\r\nTo: Robert Cook, Robert Crockett, John Handley
, Shawna'
```

Create a function to count words in each file:

```
>>> def count_words(file):
...     word_counts = defaultdict(int)
...     for line in file:
...         for word in line.split():
...             word_counts[word] += 1
...     return word_counts

>>> print(count_words(['apple banana apple', 'apple orange']))

defaultdict(int, {'apple': 3, 'banana': 1, 'orange': 1})
```

Count the number of words in the first text file:

```
>>> with hdfs.open(filenamees[0]) as f:
...     counts = count_words(f)

>>> print(sorted(counts.items(), key=lambda k_v: k_v[1], reverse=True)[:10])

[(b'the', 1065320),
 (b'of', 657220),
 (b'to', 569076),
 (b'and', 545821),
 (b'or', 375132),
 (b'in', 306271),
 (b'shall', 255680),
 (b'be', 210976),
 (b'any', 206962),
 (b'by', 194780)]
```

Count the number of words in all of the text files. This operation required about 10 minutes to run on a single machine with 4 cores and 16 GB RAM:

```
>>> all_counts = Counter()
>>> for fn in filenamees:
...     with hdfs.open(fn) as f:
...         counts = count_words(f)
...         all_counts.update(counts)
```

Print the total number of words and the words with the highest frequency from all of the text files:

```
>>> print(len(all_counts))

8797842

>>> print(sorted(all_counts.items(), key=lambda k_v: k_v[1], reverse=True)[:10])

[(b'0', 67218380),
 (b'the', 19586868),
 (b'-' , 14123768),
 (b'to', 11893464),
 (b'N/A', 11814665),
 (b'of', 11724827),
 (b'and', 10253753),
 (b'in', 6684937),
 (b'a', 5470371),
 (b'or', 5227805)]
```

The complete Python script for this example is shown below:

```
# word-count.py

import hdfs3
from collections import defaultdict, Counter

hdfs = hdfs3.HDFFileSystem('NAMENODE_HOSTNAME', port=NAMENODE_PORT)

filenamees = hdfs.glob('/tmp/enron/*/*')
print(filenamees[:5])
print(hdfs.head(filenamees[0]))
```

(continues on next page)

(continued from previous page)

```

def count_words(file):
    word_counts = defaultdict(int)
    for line in file:
        for word in line.split():
            word_counts[word] += 1
    return word_counts

print(count_words(['apple banana apple', 'apple orange']))

with hdfs.open(filenamees[0]) as f:
    counts = count_words(f)

print(sorted(counts.items(), key=lambda k_v: k_v[1], reverse=True)[:10])

all_counts = Counter()

for fn in filenamees:
    with hdfs.open(fn) as f:
        counts = count_words(f)
        all_counts.update(counts)

print(len(all_counts))
print(sorted(all_counts.items(), key=lambda k_v: k_v[1], reverse=True)[:10])

```

## 2.4 HDFS Configuration

### 2.4.1 Defaults

Several methods are available for configuring HDFS3.

The simplest is to load values from `core-site.xml` and `hdfs-site.xml` files. HDFS3 will search typical locations and reads default configuration parameters from there. The file locations may also be specified with the environment variables `HADOOP_CONF_DIR`, which is the directory containing the XML files, `HADOOP_INSTALL`, in which case the files are expected in subdirectory `hadoop/conf/` or `LIBHDFS3_CONF`, which should explicitly point to the `hdfs-site.xml` file you wish to use.

It is also possible to pass parameters to HDFS3 when instantiating the file system. You can either provide individual common overrides (e.g., `host='myhost'`) or provide a whole configuration as a dictionary (`pars={}`) with the same key names as typically contained in the XML config files. These parameters will take precedence over any loaded from files, or you can disable using the default configuration at all with `autoconf=False`.

The special environment variable `LIBHDFS3_CONF` will be automatically set when parsing the config files, if possible. Since the library is only loaded upon the first instantiation of a `HDFFileSystem`, you still have the option to change its value in `os.environ`.

### 2.4.2 Short-circuit reads in HDFS

Typically in HDFS, all data reads go through the datanode. Alternatively, a process that runs on the same node as the data can bypass or *short-circuit* the communication path through the datanode and instead read directly from a file.

HDFS and `hdfs3` can be configured for short-circuit reads. The easiest method is to edit the `hdfs-site.xml` file whose location you specify as above.

- Configure the appropriate settings in `hdfs-site.xml` on all of the HDFS nodes:

```
<configuration>
  <property>
    <name>dfs.client.read.shortcircuit</name>
    <value>true</value>
  </property>

  <property>
    <name>dfs.domain.socket.path</name>
    <value>/var/lib/hadoop-hdfs/dn_socket</value>
  </property>
</configuration>
```

The above configuration changes should allow for short-circuit reads. If you continue to receive warnings to retry the same node but disable read shortcircuit feature, check the above settings. Note that the HDFS reads should still function despite the warning, but performance might be impacted.

For more information about configuring short-circuit reads, refer to the [HDFS Short-Circuit Local Reads](#) documentation.

## 2.4.3 High-availability mode

Although HDFS is resilient to failure of data-nodes, the name-node is a single repository of metadata for the system, and so a single point of failure. High-availability (HA) involves configuring fall-back name-nodes which can take over in the event of failure. A good description can be found [‘here’](#).

In the case of [‘libhdfs3’](#), the library used by `hdfs3`, the configuration required for HA can be passed to the client directly in python code, or included in configuration files, as with any other configuration options.

In python code, this could look like the following:

```
host = "nameservice1"
conf = {"dfs.nameservices": "nameservice1",
        "dfs.ha.namenodes.nameservice1": "namenode113,namenode188",
        "dfs.namenode.rpc-address.nameservice1.namenode113": "hostname_of_server1:8020",
        ↵,
        "dfs.namenode.rpc-address.nameservice1.namenode188": "hostname_of_server2:8020",
        ↵,
        "dfs.namenode.http-address.nameservice1.namenode188": "hostname_of_
        ↵server1:50070",
        "dfs.namenode.http-address.nameservice1.namenode188": "hostname_of_
        ↵server2:50070",
        "hadoop.security.authentication": "kerberos"
}
fs = HDFSFileSystem(host=host, pars=conf)
```

Note that no port is specified (requires `hdfs` version 0.1.3), it's value should be `None`.

## 2.5 API

<code>HDFSFileSystem([host, port, connect, ...])</code>	Connection to an HDFS namenode
<code>HDFSFileSystem.cat(path)</code>	Return contents of file
<code>HDFSFileSystem.chmod(path, mode)</code>	Change access control of given path

Continued on next page

Table 1 – continued from previous page

<code>HDFFileSystem.chown(path, owner, group)</code>	Change owner/group
<code>HDFFileSystem.df()</code>	Used/free disc space on the HDFS system
<code>HDFFileSystem.du(path[, total, deep])</code>	Returns file sizes on a path.
<code>HDFFileSystem.exists(path)</code>	Is there an entry at path?
<code>HDFFileSystem.get(hdfs_path, local_path[, ...])</code>	Copy HDFS file to local
<code>HDFFileSystem.getmerge(path, filename[, ...])</code>	Concat all files in path (a directory) to local output file
<code>HDFFileSystem.get_block_locations(path[, ...])</code>	Fetch physical locations of blocks
<code>HDFFileSystem.glob(path)</code>	Get list of paths mathing glob-like pattern (i.e., with “*”s).
<code>HDFFileSystem.info(path)</code>	File information (as a dict)
<code>HDFFileSystem.ls(path[, detail])</code>	List files at path
<code>HDFFileSystem.mkdir(path)</code>	Make directory at path
<code>HDFFileSystem.mv(path1, path2)</code>	Move file at path1 to path2
<code>HDFFileSystem.open(path[, mode, replication, ...])</code>	Open a file for reading or writing
<code>HDFFileSystem.put(filename, path[, chunk, ...])</code>	Copy local file to path in HDFS
<code>HDFFileSystem.read_block(fn, offset, length)</code>	Read a block of bytes from an HDFS file
<code>HDFFileSystem.rm(path[, recursive])</code>	Use recursive for <code>rm -r</code> , i.e., delete directory and contents
<code>HDFFileSystem.set_replication(path, replication)</code>	Instruct HDFS to set the replication for the given file.
<code>HDFFileSystem.tail(path[, size])</code>	Return last bytes of file
<code>HDFFileSystem.touch(path)</code>	Create zero-length file
<code>HDFFile(fs, path, mode[, replication, buff, ...])</code>	File on HDFS
<code>HDFFile.close()</code>	Flush and close file, ensuring the data is readable
<code>HDFFile.flush()</code>	Send buffer to the data-node; actual write may happen later
<code>HDFFile.info()</code>	Filesystem metadata about this file
<code>HDFFile.read([length])</code>	Read bytes from open file
<code>HDFFile.readlines()</code>	Return all lines in a file as a list
<code>HDFFile.seek(offset[, from_what])</code>	Set file read position.
<code>HDFFile.tell()</code>	Get current byte location in a file
<code>HDFFile.write(data)</code>	Write bytes to open file (which must be in w or a mode)
<code>HDFSMap(hdfs, root[, check])</code>	Wrap a HDFFileSystem as a mutable mapping.

```
class hdfs3.core.HDFFileSystem (host=<class 'hdfs3.utils.MyNone'>, port=<class 'hdfs3.utils.MyNone'>, connect=True, autoconf=True, pars=None, **kwargs)
```

Connection to an HDFS namenode

```
>>> hdfs = HDFFileSystem(host='127.0.0.1', port=8020) # doctest: +SKIP
```

```
cancel_token (token=None)
```

Revoke delegation token

#### Parameters

**token:** `str` or `None` If `None`, uses the instance’s token. It is an error to do that if there is no

token.

**cat** (*path*)

Return contents of file

**chmod** (*path, mode*)

Change access control of given path

Exactly what permissions the file will get depends on HDFS configurations.

#### Parameters

**path** [string] file/directory to change

**mode** [integer] As with the POSIX standard, each octal digit refers to user-group-all, in that order, with read-write-execute as the bits of each group.

#### Examples

Make read/writeable to all >>> `hdfs.chmod('/path/to/file', 0o777) # doctest: +SKIP`

Make read/writeable only to user >>> `hdfs.chmod('/path/to/file', 0o700) # doctest: +SKIP`

Make read-only to user >>> `hdfs.chmod('/path/to/file', 0o100) # doctest: +SKIP`

**chown** (*path, owner, group*)

Change owner/group

**concat** (*destination, paths*)

Concatenate inputs to destination

Source files *should* all have the same block size and replication. The destination file must be in the same directory as the source files. If the target exists, it will be appended to.

Some HDFSs impose that the target file must exist and be an exact number of blocks long, and that each concated file except the last is also a whole number of blocks.

The source files are deleted on successful completion.

**connect** ()

Connect to the name node

This happens automatically at startup

**delegate\_token** (*user=None*)

Generate delegate auth token.

#### Parameters

**user: bytes/str** User to pass to delegation (defaults to user supplied to instance); this user is the only one that can renew the token.

**df** ()

Used/free disc space on the HDFS system

**disconnect** ()

Disconnect from name node

**du** (*path, total=False, deep=False*)

Returns file sizes on a path.

#### Parameters

**path** [string] where to look



**total** [bool (False)] to add up the sizes to a grand total

**deep** [bool (False)] whether to recurse into subdirectories

**exists** (*path*)

Is there an entry at path?

**get** (*hdfs\_path*, *local\_path*, *blocksize=65536*)

Copy HDFS file to local

**get\_block\_locations** (*path*, *start=0*, *length=0*)

Fetch physical locations of blocks

**getmerge** (*path*, *filename*, *blocksize=65536*)

Concat all files in path (a directory) to local output file

**glob** (*path*)

Get list of paths mathing glob-like pattern (i.e., with “\*”s).

If passed a directory, gets all contained files; if passed path to a file, without any “\*”, returns one-element list containing that filename. Does not support python3.5’s “\*\*” notation.

**head** (*path*, *size=1024*)

Return first bytes of file

**info** (*path*)

File information (as a dict)

**isdir** (*path*)

Return True if path refers to an existing directory.

**isfile** (*path*)

Return True if path refers to an existing file.

**list\_encryption\_zones** ()

Get list of all the encryption zones

**ls** (*path*, *detail=False*)

List files at path

#### Parameters

**path** [string/bytes] location at which to list files

**detail** [bool (=True)] if True, each list item is a dict of file properties; otherwise, returns list of filenames

**makedirs** (*path*, *mode=457*)

Create directory together with any necessary intermediates

**mkdir** (*path*)

Make directory at path

**mv** (*path1*, *path2*)

Move file at path1 to path2

**open** (*path*, *mode='rb'*, *replication=0*, *buff=0*, *block\_size=0*)

Open a file for reading or writing

#### Parameters

**path:** string Path of file on HDFS

**mode:** string One of ‘rb’, ‘wb’, or ‘ab’

**replication:** int Replication factor; if zero, use system default (only on write)

**buf: int (=0)** Client buffer size (bytes); if 0, use default.

**block\_size: int** Size of data-node blocks if writing

**put** (*filename, path, chunk=65536, replication=0, block\_size=0*)  
Copy local file to path in HDFS

**read\_block** (*fn, offset, length, delimiter=None*)  
Read a block of bytes from an HDFS file

Starting at *offset* of the file, read *length* bytes. If *delimiter* is set then we ensure that the read starts and stops at delimiter boundaries that follow the locations *offset* and *offset + length*. If *offset* is zero then we start at zero. The bytestring returned will not include the surrounding delimiter strings.

If *offset+length* is beyond the eof, reads to eof.

#### Parameters

**fn: string** Path to filename on HDFS

**offset: int** Byte offset to start read

**length: int** Number of bytes to read

**delimiter: bytes (optional)** Ensure reading starts and stops at delimiter bytestring

#### See also:

`hdfs3.utils.read_block`

#### Examples

```
>>> hdfs.read_block('/data/file.csv', 0, 13) # doctest: +SKIP
b'Alice, 100\nBo'
>>> hdfs.read_block('/data/file.csv', 0, 13, delimiter=b'\n') # doctest: _
↪+SKIP
b'Alice, 100\nBob, 200'
```

**renew\_token** (*token=None*)  
Renew delegation token

#### Parameters

**token: str or None** If None, uses the instance's token. It is an error to do that if there is no token.

#### Returns

**New expiration time for the token**

**rm** (*path, recursive=True*)  
Use recursive for *rm -r*, i.e., delete directory and contents

**set\_replication** (*path, replication*)  
Instruct HDFS to set the replication for the given file.

If successful, the head-node's table is updated immediately, but actual copying will be queued for later. It is acceptable to set a replication that cannot be supported (e.g., higher than the number of data-nodes).

**tail** (*path, size=1024*)  
Return last bytes of file

**touch** (*path*)  
Create zero-length file

**walk** (*path*)  
Directory tree generator, see `os.walk`

**class** `hdfs3.core.HDFFile` (*fs, path, mode, replication=0, buff=0, block\_size=0*)  
File on HDFS

Matches the standard Python file interface.

## Examples

```
>>> with hdfs.open('/path/to/hdfs/file.txt') as f: # doctest: +SKIP
...     bytes = f.read(1000) # doctest: +SKIP
>>> with hdfs.open('/path/to/hdfs/file.csv') as f: # doctest: +SKIP
...     df = pd.read_csv(f, nrows=1000) # doctest: +SKIP
```

**close** ()  
Flush and close file, ensuring the data is readable

**flush** ()  
Send buffer to the data-node; actual write may happen later

**info** ()  
Filesystem metadata about this file

**next** ()  
Enables reading a file as a buffer in pandas

**read** (*length=None*)  
Read bytes from open file

**readline** (*chunksize=256, lineterminator='\n'*)  
Return a line using buffered reading.

A line is a sequence of bytes between ““

““ **markers (or given** line-terminator).

Line iteration uses this method internally.

Note: this function requires many calls to HDFS and is slow; it is in general better to wrap an `HDFFile` with an `io.TextIOWrapper` for buffering, text decoding and newline support.

**readlines** ()  
Return all lines in a file as a list

**seek** (*offset, from\_what=0*)  
Set file read position. Read mode only.

Attempt to move out of file bounds raises an exception. Note that, by the convention in python file seek, offset should be  $\leq 0$  if `from_what` is 2.

### Parameters

**offset** [int] byte location in the file.

**from\_what** [int 0, 1, 2] if 0 (default), relative to file start; if 1, relative to current location; if 2, relative to file end.

### Returns

**new position****tell** ()

Get current byte location in a file

**write** (*data*)

Write bytes to open file (which must be in w or a mode)

**class** hdfs3.mapping.**HDFSMap** (*hdfs, root, check=False*)

Wrap a HDFSFileSystem as a mutable mapping.

The keys of the mapping become files under the given root, and the values (which must be bytes) the contents of those files.

**Parameters****hdfs** [HDFSFileSystem]**root** [string] path to contain the stored files (directory will be created if it doesn't exist)**check** [bool (=True)] performs a touch at the location, to check writeability.**Examples**

```
>>> hdfs = hdfs3.HDFSFileSystem() # doctest: +SKIP
>>> mw = HDFSMap(hdfs, '/writable/path/') # doctest: +SKIP
>>> mw['loc1'] = b'Hello World' # doctest: +SKIP
>>> list(mw.keys()) # doctest: +SKIP
['loc1']
>>> mw['loc1'] # doctest: +SKIP
b'Hello World'
```

## 2.6 Known Limitations

### 2.6.1 Forked processes

The `libhdfs3` library is not fork-safe. If you start using `hdfs3` in a parent process and then fork a child process, using the library from the child process may produce random errors because of system resources that will not be available (such as background threads). Common solutions include the following:

- Use threads instead of processes
- Use Python 3 and a multiprocessing context using either the “spawn” or “forkserver” method (see [multiprocessing docs](#))
- Only instantiate `HDFSFileSystem` within the forked processes, do not ever start an `HDFSFileSystem` within the parent processes

**C**

cancel\_token() (hdfs3.core.HDFFileSystem method), 11  
cat() (hdfs3.core.HDFFileSystem method), 12  
chmod() (hdfs3.core.HDFFileSystem method), 12  
chown() (hdfs3.core.HDFFileSystem method), 12  
close() (hdfs3.core.HDFFile method), 15  
concat() (hdfs3.core.HDFFileSystem method), 12  
connect() (hdfs3.core.HDFFileSystem method), 12

**D**

delegate\_token() (hdfs3.core.HDFFileSystem method), 12  
df() (hdfs3.core.HDFFileSystem method), 12  
disconnect() (hdfs3.core.HDFFileSystem method), 12  
du() (hdfs3.core.HDFFileSystem method), 12

**E**

exists() (hdfs3.core.HDFFileSystem method), 13

**F**

flush() (hdfs3.core.HDFFile method), 15

**G**

get() (hdfs3.core.HDFFileSystem method), 13  
get\_block\_locations() (hdfs3.core.HDFFileSystem method), 13  
getmerge() (hdfs3.core.HDFFileSystem method), 13  
glob() (hdfs3.core.HDFFileSystem method), 13

**H**

HDFFile (class in hdfs3.core), 15  
HDFFileSystem (class in hdfs3.core), 11  
HDFSMap (class in hdfs3.mapping), 16  
head() (hdfs3.core.HDFFileSystem method), 13

**I**

info() (hdfs3.core.HDFFile method), 15  
info() (hdfs3.core.HDFFileSystem method), 13  
isdir() (hdfs3.core.HDFFileSystem method), 13  
isfile() (hdfs3.core.HDFFileSystem method), 13

**L**

list\_encryption\_zones() (hdfs3.core.HDFFileSystem method), 13  
ls() (hdfs3.core.HDFFileSystem method), 13

**M**

makedirs() (hdfs3.core.HDFFileSystem method), 13  
mkdir() (hdfs3.core.HDFFileSystem method), 13  
mv() (hdfs3.core.HDFFileSystem method), 13

**N**

next() (hdfs3.core.HDFFile method), 15

**O**

open() (hdfs3.core.HDFFileSystem method), 13

**P**

put() (hdfs3.core.HDFFileSystem method), 14

**R**

read() (hdfs3.core.HDFFile method), 15  
read\_block() (hdfs3.core.HDFFileSystem method), 14  
readline() (hdfs3.core.HDFFile method), 15  
readlines() (hdfs3.core.HDFFile method), 15  
renew\_token() (hdfs3.core.HDFFileSystem method), 14  
rm() (hdfs3.core.HDFFileSystem method), 14

**S**

seek() (hdfs3.core.HDFFile method), 15  
set\_replication() (hdfs3.core.HDFFileSystem method), 14

**T**

tail() (hdfs3.core.HDFFileSystem method), 14  
tell() (hdfs3.core.HDFFile method), 16  
touch() (hdfs3.core.HDFFileSystem method), 14

**W**

walk() (hdfs3.core.HDFFileSystem method), 15  
write() (hdfs3.core.HDFFile method), 16