
S3QL Documentation

Release 3.2

Nikolaus Rath

Jul 17, 2019

CONTENTS

1	S3QL	1
1.1	Features	1
1.2	Development Status	2
1.3	Supported Platforms	2
1.4	Contributing	2
2	Installation	3
2.1	Dependencies	3
2.2	Installing S3QL	4
2.3	Development Version	4
2.4	Running tests requiring remote servers	4
3	Storage Backends	7
3.1	Google Storage	7
3.2	Amazon S3	8
3.3	OpenStack/Swift	9
3.4	Rackspace CloudFiles	10
3.5	S3 compatible	10
3.6	Local	11
4	Important Rules to Avoid Losing Data	13
4.1	Rules in a Nutshell	13
4.2	Consistency Window List	14
4.3	Data Consistency	14
4.4	Data Durability	15
5	File System Creation	17
6	Managing File Systems	19
6.1	Changing the Passphrase	19
6.2	Upgrading the file system	20
6.3	Deleting a file system	20
6.4	Restoring Metadata Backups	20
7	Mounting	21
7.1	Permission Checking	22
7.2	Compression Algorithms	22
7.3	Notes about Caching	23
7.4	NFS Support	23
7.5	Failure Modes	23
7.6	Automatic Mounting	24

8	Advanced S3QL Features	25
8.1	Snapshotting and Copy-on-Write	25
8.2	Getting Statistics	26
8.3	Immutable Trees	26
8.4	Fast Recursive Removal	27
8.5	Runtime Configuration	27
9	Unmounting	29
10	Checking for Errors	31
10.1	Checking and repairing internal file system errors	31
10.2	Detecting and handling backend data corruption	32
11	Storing Backend Information and Credentials	35
12	Contributed Programs	37
12.1	benchmark.py	37
12.2	clone_fs.py	37
12.3	pcp.py	37
12.4	s3ql_backup.sh	37
12.5	expire_backups.py	38
12.6	remove_objects.py	39
13	Tips & Tricks	41
13.1	SSH Backend	41
13.2	Permanently mounted backup file system	41
13.3	Improving copy performance	41
14	Frequently Asked Questions	43
14.1	What does “python-apsw must be linked dynamically to sqlite3” mean?	43
14.2	How can I improve the file system throughput?	43
14.3	Why does df show 1 TB of free space?	43
14.4	Which operating systems are supported?	44
14.5	Is there a file size limit?	44
14.6	Suppose I want to make a small change in a very large file. Will S3QL download and re-upload the entire file?	44
14.7	I don’t quite understand this de-duplication feature.	44
14.8	What does the “Transport endpoint not connected” error mean?	44
14.9	What does “Backend reports that fs is still mounted elsewhere, aborting” mean?	44
14.10	Can I access an S3QL file system on multiple computers simultaneously?	45
14.11	What maximum object size should I use?	45
14.12	Is there a way to make the cache persistent / access the file system offline?	46
14.13	I would like to use S3QL with Hubic, but.	46
14.14	What’s a reasonable metadata upload interval?	46
15	Known Issues	49
16	Manpages	51
16.1	The <code>mkfs.s3ql</code> command	51
16.2	The <code>s3qladm</code> command	52
16.3	The <code>mount.s3ql</code> command	54
16.4	The <code>s3qlstat</code> command	56
16.5	The <code>s3qlctrl</code> command	57
16.6	The <code>s3qlcp</code> command	59
16.7	The <code>s3qlrm</code> command	60

16.8	The <code>s3q1lock</code> command	61
16.9	The <code>umount.s3q1</code> command	63
16.10	The <code>fsck.s3q1</code> command	64
16.11	The <code>s3q1_oauth_client</code> command	66
16.12	The <code>s3q1_verify</code> command	67
16.13	The <code>pcp</code> command	68
16.14	The <code>expire_backups</code> command	69
17	Further Resources / Getting Help	73
18	Implementation Details	75
18.1	Metadata Storage	75
18.2	Data Storage	75
18.3	Data De-Duplication	76
18.4	Caching	76
18.5	Eventual Consistency Handling	76
18.6	Encryption	76

S3QL

S3QL is a file system that stores all its data online using storage services like [Google Storage](#), [Amazon S3](#), or [Open-Stack](#). S3QL effectively provides a hard disk of dynamic, infinite capacity that can be accessed from any computer with internet access.

S3QL is a standard conforming, full featured UNIX file system that is conceptually indistinguishable from any local file system. Furthermore, S3QL has additional features like compression, encryption, data de-duplication, immutable trees and snapshotting which make it especially suitable for online backup and archival.

S3QL is designed to favor simplicity and elegance over performance and feature-creep. Care has been taken to make the source code as readable and serviceable as possible. Solid error detection and error handling have been included from the very first line, and S3QL comes with extensive automated test cases for all its components.

1.1 Features

- **Transparency.** Conceptually, S3QL is indistinguishable from a local file system. For example, it supports hardlinks, symlinks, standard unix permissions, extended attributes and file sizes up to 2 TB.
- **Dynamic Size.** The size of an S3QL file system grows and shrinks dynamically as required.
- **Compression.** Before storage, all data may be compressed with the LZMA, bzip2 or deflate (gzip) algorithm.
- **Encryption.** After compression (but before upload), all data can be AES encrypted with a 256 bit key. An additional SHA256 HMAC checksum is used to protect the data against manipulation.
- **Data De-duplication.** If several files have identical contents, the redundant data will be stored only once. This works across all files stored in the file system, and also if only some parts of the files are identical while other parts differ.
- **Immutable Trees.** Directory trees can be made immutable, so that their contents can no longer be changed in any way whatsoever. This can be used to ensure that backups can not be modified after they have been made.
- **Copy-on-Write/Snapshotting.** S3QL can replicate entire directory trees without using any additional storage space. Only if one of the copies is modified, the part of the data that has been modified will take up additional storage space. This can be used to create intelligent snapshots that preserve the state of a directory at different points in time using a minimum amount of space.
- **High Performance independent of network latency.** All operations that do not write or read file contents (like creating directories or moving, renaming, and changing permissions of files and directories) are very fast because they are carried out without any network transactions.

S3QL achieves this by saving the entire file and directory structure in a database. This database is locally cached and the remote copy updated asynchronously.

- **Support for low bandwidth connections.** S3QL splits file contents into smaller blocks and caches blocks locally. This minimizes both the number of network transactions required for reading and writing data, and the amount of data that has to be transferred when only parts of a file are read or written.

1.2 Development Status

S3QL is considered stable and suitable for production use. Starting with version 2.17.1, S3QL uses semantic versioning. This means that backwards-incompatible versions (e.g., versions that require an upgrade of the file system revision) will be reflected in an increase of the major version number.

1.3 Supported Platforms

S3QL is developed and tested under Linux. Users have also reported running S3QL successfully on OS-X, FreeBSD and NetBSD. We try to maintain compatibility with these systems, but (due to lack of pre-release testers) we cannot guarantee that every release will run on all non-Linux systems. Please report any bugs you find, and we will try to fix them.

1.4 Contributing

The S3QL source code is available on [GitHub](#).

INSTALLATION

S3QL depends on several other programs and libraries that have to be installed first. The best method to satisfy these dependencies depends on your distribution.

2.1 Dependencies

The following is a list of the programs and libraries required for running S3QL. Generally, you should first check if your distribution already provides a suitable packages and only install from source if that is not the case.

- Kernel: Linux 2.6.9 or newer or FreeBSD with **FUSE4BSD**. Starting with kernel 2.6.26 you will get significantly better write performance, so under Linux you should actually use *2.6.26 or newer whenever possible*.
- The `psmisc` utilities.
- **SQLite** version 3.7.0 or newer. SQLite has to be installed as a *shared library* with development headers.
- **Python** 3.5.0 or newer. Make sure to also install the development headers.
- The following Python modules:
 - `setuptools`, version 1.0 or newer.
 - `cryptography`
 - `defusedxml`
 - `apsw`, version 3.7.0 or newer.
 - `llfuse`, any version between 1.0 (inclusive) and 2.0 (exclusive)
 - `dugong`, any version between 3.4 (inclusive) and 4.0 (exclusive)
 - `pytest`, version 3.7 or newer (optional, to run unit tests)
 - `systemd` (optional, for enabling systemd support). Do *not* install the module from PyPi, this is from a third-party developer and incompatible with the official module from the systemd developers.
 - `requests` (optional, required for OAuth2 authentication with Google Storage)
 - `google-auth` (optional, required for ADC authentication with Google Storage)
 - `google-auth-oauthlib` (optional, required for browser-based authentication with Google Storage)

To check if a specific module `<module>` is installed, execute `python3 -c 'import <module>; print (<module>.__version__)'`. This will result in an `ImportError` if the module is not installed, and will print the installed version if the module is installed.

2.2 Installing S3QL

To build and install S3QL itself, proceed as follows:

1. Download S3QL from <https://github.com/s3ql/s3ql/releases>
2. Unpack it into a folder of your choice
3. Run `python3 setup.py build_ext --inplace` to build S3QL.
4. Run `python3 -m pytest tests/` to run a self-test. If this fails, ask for help on the [mailing list](#) or report a bug in the [issue tracker](#).

Now you have three options:

- You can run the S3QL commands from the `bin/` directory.
- You can install S3QL system-wide for all users. To do that, you have to run `sudo python3 setup.py install`.
- You can install S3QL into `~/.local` by executing `python3 setup.py install --user`. In this case you should make sure that `~/.local/bin` is in your `$PATH` environment variable.

2.3 Development Version

If you have checked out the unstable development version from the Mercurial repository, a bit more effort is required. You'll also need:

- Version 0.24 or newer of the [Cython](#) compiler.
- Version 1.2b1 or newer of the [Sphinx](#) document processor.

With these additional dependencies installed, S3QL can be build and tested with

```
python3 setup.py build_cython
python3 setup.py build_ext --inplace
python3 -m pytest test
```

Note that when building from the Mercurial or Git repository, building and testing is done with several additional checks. This may cause compilation and/or tests to fail even though there are no problems with functionality. For example, any use of functions that are scheduled for deprecation in future Python version will cause tests to fail. If you would rather just check for functionality, you can delete the `MANIFEST.in` file. In that case, the build system will behave as it does for a regular release.

The HTML and PDF documentation can be generated with

```
python3 setup.py build_sphinx
```

and S3QL can be installed as usual with

```
python3 setup.py install [--user]
```

2.4 Running tests requiring remote servers

By default, tests requiring a connection to a remote storage backend are skipped. If you would like to run these tests too (which is always a good idea), you have to create additional entries in your `~/.s3ql/authinfo2` file that tell S3QL what server and credentials to use for these tests. These entries have the following form:

```
[<BACKEND>-test]
backend-login: <user>
backend-password: <password>
test-fs: <storage-url>
```

Here *<BACKEND>* specifies the backend that you want to test (e.g. *s3*, *s3c*, *gs*, or *swift*), *<user>* and *<password>* are the backend authentication credentials, and *<storage-url>* specifies the full storage URL that will be used for testing. **Any existing S3QL file system in this storage URL will be destroyed during testing.**

For example, to run tests that need connection to a Google Storage server, you would add something like

```
[gs-test]
backend-login: GOOGIGWLONT238MD7HZ4
backend-password: rmEbstjscoeunt1249oes1298gauidbs3h1
test-fs: gs://joes-gs-bucket/s3ql_tests/
```

On the next run of `runtest.py` (or `py.test` when using the development version), the additional tests will be run. If the tests are still skipped, you can get more information about why tests are being skipped by passing the `-rs` argument to `runtest.py/py.test`.

STORAGE BACKENDS

S3QL supports different *backends* to store data at different service providers and using different protocols. A *storage url* specifies a backend together with some backend-specific information and uniquely identifies an S3QL file system. The form of the storage url depends on the backend and is described for every backend below.

Furthermore, every S3QL commands that accepts a storage url also accepts a `--backend-options` parameter that can be used to pass backend-specific options to the backend module. The available options are documented with the respective backends below.

All storage backends respect the `http_proxy` (for plain HTTP connections) and `https_proxy` (for SSL connections) environment variables.

Note: Storage backends are not necessarily compatible. Don't expect that you can e.g. copy the data stored by the local backend into Amazon S3 using some non-S3QL tool and then access it with S3QL's S3 backend. If you want to copy file systems from one backend to another, you need to use the `clone_fs.py` script (from the `contrib` directory in the S3QL tarball).

3.1 Google Storage

Google Storage is an online storage service offered by Google. In order to use it with S3QL, make sure that you enable the JSON API in the [GCP Console API Library](#)

The Google Storage backend uses OAuth2 authentication or ADC (Application Default Credentials).

To use OAuth2 authentication, specify `oauth2` as the backend login and a valid OAuth2 refresh token as the backend password. To obtain a refresh token, you can use the `s3ql_oauth_client` program. It will instruct you to open a specific URL in your browser, enter a code and authenticate with your Google account. Once this procedure is complete, `s3ql_oauth_client` will print out the refresh token. Note that you need to do this procedure only once, the refresh token will remain valid until you explicitly revoke it.

To use ADC, specify `adc` as the backend login and use an arbitrary value for the backend password.

To create a Google Storage bucket, you can use e.g. the [Google Storage Manager](#). The storage URL for accessing the bucket in S3QL is then

```
gs://<bucketname>/<prefix>
```

Here *bucketname* is the name of the bucket, and *prefix* can be an arbitrary prefix that will be prepended to all object names used by S3QL. This allows you to store several S3QL file systems in the same Google Storage bucket.

The Google Storage backend accepts the following backend options:

ssl-ca-path=<path>

Instead of using the system's default certificate store, validate the server certificate against the specified CA certificates. <path> may be either a file containing multiple certificates, or a directory containing one certificate per file.

tcp-timeout

Specifies the timeout used for TCP connections. If no data can be exchanged with the remote server for longer than this period, the TCP connection is closed and re-established (default: 20 seconds).

3.2 Amazon S3

Amazon S3 is the online storage service offered by Amazon Web Services (AWS). To use the S3 backend, you first need to sign up for an AWS account. The account is free, you will pay only for the amount of storage and traffic that you actually use. After that, you need to create a bucket that will hold the S3QL file system, e.g. using the [AWS Management Console](#). For best performance, it is recommend to create the bucket in the geographically closest storage region, but not the US Standard region (see *Important Rules to Avoid Losing Data* for the reason).

The storage URL for accessing S3 buckets in S3QL has the form

```
s3://<region>/<bucket>/<prefix>
```

prefix can be an arbitrary prefix that will be prepended to all object names used by S3QL. This allows you to store several S3QL file systems in the same S3 bucket. For example, the storage URL

```
s3://ap-south-1/foomart.net/data/s3ql_backup/
```

refers to the *foomart.net* bucket in the *ap-south-1* region. All storage objects that S3QL stores in this bucket will be prefixed with *data/s3ql_backup/*.

Note that the backend login and password for accessing S3 are not the user id and password that you use to log into the Amazon Webpage, but the *AWS access key id* and *AWS secret access key* shown under [My Account/Access Identifiers](#).

The Amazon S3 backend accepts the following backend options:

no-ssl

Disable encrypted (https) connections and use plain HTTP instead.

ssl-ca-path=<path>

Instead of using the system's default certificate store, validate the server certificate against the specified CA certificates. <path> may be either a file containing multiple certificates, or a directory containing one certificate per file.

tcp-timeout

Specifies the timeout used for TCP connections. If no data can be exchanged with the remote server for longer than this period, the TCP connection is closed and re-established (default: 20 seconds).

sse

Enable server side encryption. Both costs & benefits of S3 server side encryption are probably rather small, and this option does *not* affect any client side encryption performed by S3QL itself.

it

Use INTELLIGENT_TIERING storage class for new objects. See [AWS S3 Storage classes](#)

ia

Use STANDARD_IA (infrequent access) storage class for new objects. See [AWS S3 Storage classes](#)

oia

Use ONEZONE_IA (infrequent access) storage class for new objects. See [AWS S3 Storage classes](#)

rrs

Enable reduced redundancy storage for newly created objects (overwrites the *ia* option).

When enabling this option, it is strongly recommended to periodically run *s3ql_verify*, because objects that are lost by the storage backend may cause subsequent data loss even later in time due to the data de-duplication feature of S3QL (see *Data Durability* for details).

3.3 OpenStack/Swift

OpenStack is an open-source cloud server application suite. Swift is the cloud storage module of OpenStack. Swift/OpenStack storage is offered by many different companies.

There are two different storage URL for the OpenStack backend that make use of different authentication APIs. For legacy (v1) authentication, the storage URL is

```
swift://<hostname>[:<port>]/<container>[/<prefix>]
```

for Keystone (v2) authentication, the storage URL is

```
swiftks://<hostname>[:<port>]/<region>:<container>[/<prefix>]
```

Note that when using Keystone authentication, you can (and have to) specify the storage region of the container as well.

In both cases, *hostname* name should be the name of the authentication server. The storage container must already exist (most OpenStack providers offer either a web frontend or a command line tool for creating containers). *prefix* can be an arbitrary prefix that will be prepended to all object names used by S3QL, which can be used to store multiple S3QL file systems in the same container.

When using legacy authentication, the backend login and password correspond to the OpenStack username and API Access Key. When using Keystone authentication, the backend password is your regular OpenStack password and the backend login combines you OpenStack username and tenant name in the form `<tenant>:<user>`. If no tenant is required, the OpenStack username alone may be used as backend login.

The OpenStack backend accepts the following backend options:

no-ssl

Use plain HTTP to connect to the authentication server. This option does not directly affect the connection to the storage server. Whether HTTPS or plain HTTP is used to connect to the storage server is determined by the authentication server.

ssl-ca-path=<path>

Instead of using the system's default certificate store, validate the server certificate against the specified CA certificates. `<path>` may be either a file containing multiple certificates, or a directory containing one certificate per file.

tcp-timeout

Specifies the timeout used for TCP connections. If no data can be exchanged with the remote server for longer than this period, the TCP connection is closed and re-established (default: 20 seconds).

disable-expect100

If this option is specified, S3QL does not use the `Expect: continue` header (cf. RFC2616, section 8.2.3) when uploading data to the server. This can be used to work around broken storage servers that don't fully support HTTP 1.1, but may decrease performance as object data will be transmitted to the server more than once in some circumstances.

no-feature-detection

If this option is specified, S3QL does not try to dynamically detect advanced features of the Swift backend. In this case S3QL can only use the least common denominator of supported Swift versions and configurations.

Note: The Swift API unfortunately lacks a number of features that S3QL normally makes use of. S3QL works around these deficiencies as much as possible. However, this means that storing data using the Swift backend generally requires more network round-trips and transfer volume than the other backends. Also, S3QL requires Swift storage servers to provide immediate consistency for newly created objects.

3.4 Rackspace CloudFiles

Rackspace CloudFiles uses [OpenStack](#) internally, so it is possible to just use the OpenStack/Swift backend (see above) with `auth.api.rackspacecloud.com` as the host name. For convenience, there is also a special `rackspace` backend that uses a storage URL of the form

```
rackspace://<region>/<container>[/<prefix>]
```

The storage container must already exist in the selected region. *prefix* can be an arbitrary prefix that will be prepended to all object names used by S3QL and can be used to store several S3QL file systems in the same container.

You can create a storage container for S3QL using the [Cloud Control Panel](#) (click on *Files* in the topmost menu bar).

The Rackspace backend accepts the same backend options as the *OpenStack backend*.

3.5 S3 compatible

The S3 compatible backend allows S3QL to access any storage service that uses the same protocol as Amazon S3. The storage URL has the form

```
s3c://<hostname>:<port>/<bucketname>/<prefix>
```

Here *bucketname* is the name of an (existing) bucket, and *prefix* can be an arbitrary prefix that will be prepended to all object names used by S3QL. This allows you to store several S3QL file systems in the same bucket.

The S3 compatible backend accepts the following backend options:

no-ssl

Disable encrypted (https) connections and use plain HTTP instead.

ssl-ca-path=<path>

Instead of using the system's default certificate store, validate the server certificate against the specified CA certificates. `<path>` may be either a file containing multiple certificates, or a directory containing one certificate per file.

tcp-timeout

Specifies the timeout used for TCP connections. If no data can be exchanged with the remote server for longer than this period, the TCP connection is closed and re-established (default: 20 seconds).

disable-expect100

If this option is specified, S3QL does not use the `Expect: continue` header (cf. [RFC2616, section 8.2.3](#)) when uploading data to the server. This can be used to work around broken storage servers that don't fully support HTTP 1.1, but may decrease performance as object data will be transmitted to the server more than once in some circumstances.

dumb-copy

If this option is specified, S3QL assumes that a COPY request to the storage server has succeeded as soon as the server returns a 200 OK status. The [S3 COPY API](#) specifies that the storage server may still return an error in the request body (see the [copy proposal](#) for the rationale), so this option should only be used if you are certain that your storage server only returns 200 OK when the copy operation has been completely and successfully carried out. Using this option may be necessary if your storage server does not return a valid response body for a successful copy operation.

3.6 Local

S3QL is also able to store its data on the local file system. This can be used to backup data on external media, or to access external services that S3QL can not talk to directly (e.g., it is possible to store data over SSH by first mounting the remote system using [sshfs](#) and then using the local backend to store the data in the sshfs mountpoint).

The storage URL for local storage is

```
local://<path>
```

Note that you have to write three consecutive slashes to specify an absolute path, e.g. `local:///var/archive`. Also, relative paths will automatically be converted to absolute paths before the authentication file (see [Storing Backend Information and Credentials](#)) is read, i.e. if you are in the `/home/john` directory and try to mount `local:/s3ql`, the corresponding section in the authentication file must match the storage url `local:///home/john/s3ql`.

The local backend does not accept any backend options.

IMPORTANT RULES TO AVOID LOSING DATA

Most S3QL backends store data in distributed storage systems. These systems differ from a traditional, local hard disk in several important ways. In order to avoid losing data, this section should be read very carefully.

4.1 Rules in a Nutshell

To avoid losing your data, obey the following rules:

1. Know what durability you can expect from your chosen storage provider. The durability describes how likely it is that a stored object becomes damaged over time. Such data corruption can never be prevented completely, techniques like geographic replication and RAID storage just reduce the likelihood of it to happen (i.e., increase the durability).
2. When choosing a backend and storage provider, keep in mind that when using S3QL, the effective durability of the file system data will be reduced because of S3QL's data de-duplication feature.
3. Determine your storage service's consistency window. The consistency window that is important for S3QL is the smaller of the times for which:
 - a newly created object may not yet be included in the list of stored objects
 - an attempt to read a newly created object may fail with the storage service reporting that the object does not exist

If *one* of the above times is zero, we say that as far as S3QL is concerned the storage service has *immediate* consistency.

If your storage provider claims that *neither* of the above can ever happen, while at the same time promising high durability, you should choose a respectable provider instead.

4. When mounting the same file system on different computers (or on the same computer but with different `--cachedir` directories), the time that passes between the first and second of invocation of `mount.s3ql` must be at least as long as your storage service's consistency window. If your storage service offers immediate consistency, you do not need to wait at all.
5. Before running `fsck.s3ql` or `s3qladm`, the file system must have been left untouched for the length of the consistency window. If your storage service offers immediate consistency, you do not need to wait at all.

The rest of this section explains the above rules and the reasons for them in more detail. It also contains a list of the consistency windows for a number of larger storage providers.

4.2 Consistency Window List

The following is a list of the consistency windows (as far as S3QL is concerned) for a number of storage providers. This list doesn't come with any guarantees and may be outdated. If your storage provider is not included, or if you need more reliable information, check with your storage provider.

Storage Provider	Consistency Window
Amazon S3 in the US standard region	No guarantees
Amazon S3 in other regions	Immediate
Google Storage	Immediate

4.3 Data Consistency

In contrast to the typical hard disk, most storage providers do not guarantee *immediate consistency* of written data. This means that:

- after an object has been stored, requests to read this object may still fail or return the prior contents for a little while.
- after an object has been deleted, attempts to read it may still return the (old) data for some time, and it may still remain in the list of stored objects for some time.
- after a new object has been created, it may still not be included when retrieving the list of stored objects for some time.

Of course, none of this is acceptable for a file system, and S3QL generally handles any of the above situations internally so that it always provides a fully consistent file system to the user. However, there are some situations where an S3QL user nevertheless needs to be aware of the peculiarities of his chosen storage service.

Suppose that you mount the file system, store some new data, delete some old data and unmount it. If you then mount the file system again right away on another computer, there is no guarantee that S3QL will see any of the changes that the first S3QL process has made. At least in theory it is therefore possible that on the second mount, S3QL does not see any of the changes that you have done and presents you an “old version” of the file system without them. Even worse, if you notice the problem and unmount the file system, S3QL will upload the old status (which S3QL necessarily has to consider as current) and thereby permanently override the newer version (even though this change may not become immediately visible either). S3QL uses several techniques to reduce the likelihood of this to happen (see [Implementation Details](#) for more information on this), but without support from the storage service, the possibility cannot be eliminated completely.

The same problem of course also applies when checking the file system. If the storage service provides S3QL with only partially updated data, S3QL has no way to find out if this a real consistency problem that needs to be fixed or if it is only a temporary problem that will resolve itself automatically (because there are still changes that have not become visible yet).

This is where the so called *consistency window* comes in. The consistency window is the maximum time (after writing or deleting the object) for which any of the above “outdated responses” may be received. If the consistency window is zero, i.e. all changes are immediately effective, the storage service is said to have *immediate consistency*. If the window is infinite, i.e. there is no upper bound on the time it may take for changes to become effect, the storage service is said to be *eventually consistent*. Note that often there are different consistency windows for the different operations. For example, Google Storage offers immediate consistency when reading data, but only eventual consistency for the list of stored objects.

To prevent the problem of S3QL working with an outdated copy of the file system data, it is therefore sufficient to simply wait for the consistency window to pass before mounting the file system again (or running a file system check). The length of the consistency window changes from storage service to storage service, and if your service is not

included in the list below, you should check the web page or ask the technical support of your storage provider. The window that is important for S3QL is the smaller of the times for which

- a newly created object may not yet be included in the list of stored objects
- an attempt to read a newly created object may fail with the storage service reporting that the object does not exist

Unfortunately, many storage providers are hesitant to guarantee anything but eventual consistency, i.e. the length of the consistency window is potentially infinite. In that case you simply have to pick a length that you consider “safe enough”. For example, even though Amazon is only guaranteeing eventual consistency, the ordinary consistency window for data stored in S3 is just a few seconds, and only in exceptional circumstances (i.e., core network outages) it may rise up to hours ([source](#)).

4.4 Data Durability

The durability of a storage service a measure of the average probability of a storage object to become corrupted over time. The lower the chance of data loss, the higher the durability. Storage services like Amazon S3 claim to achieve a durability of up to 99.999999999% over a year, i.e. if you store 100000000 objects for 100 years, you can expect that at the end of that time one object will be corrupted or lost.

S3QL is designed to reduce redundancy and store data in the smallest possible form. Therefore, S3QL is generally not able to compensate for any such losses, and when choosing a storage service you should carefully review if the offered durability matches your requirements. When doing this, there are two factors that should be kept in mind.

Firstly, even though S3QL is not able to compensate for storage service failures, it is able to detect them: when trying to access data that has been lost or corrupted by the storage service, an IO error will be returned and the mount point will become inaccessible to ensure that the problem is noticed.

Secondly, the consequences of a data loss by the storage service can be significantly more severe than you may expect because of S3QL’s data de-duplication feature: a data loss in the storage service at time x may cause data that is written *after* time x to be lost as well. Consider the following scenario:

1. You store an important file in the S3QL file system.
2. The storage service loses the data blocks of this file. As long as you do not access the file or run `fsck.s3ql`, S3QL is not aware that the data has been lost by the storage service.
3. You save an additional copy of the important file in a different location on the same S3QL file system.
4. S3QL detects that the contents of the new file are identical to the data blocks that have been stored earlier. Since at this point S3QL is not aware that these blocks have been lost by the storage service, it does not save another copy of the file contents in the storage service but relies on the (presumably) existing blocks instead.
5. Therefore, even though you saved another copy, you still do not have a backup of the important file (since both copies refer to the same data blocks that have been lost by the storage service).

For some storage services, `fsck.s3ql` can mitigate this effect. When `fsck.s3ql` runs, it asks the storage service for a list of all stored objects. If objects are missing, it can then mark the damaged files and prevent the problem from spreading forwards in time. Figuratively speaking, this establishes a “checkpoint”: data loss that occurred before running `fsck.s3ql` can not affect any file system operations that are performed after the check. Unfortunately, many storage services only “discover” that objects are missing or broken when the object actually needs to be retrieved. In this case, `fsck.s3ql` will not learn anything by just querying the list of objects.

This effect can be mitigated to some degree by using the `s3ql_verify` command in addition to `fsck.s3ql`. `s3ql_verify` asks the storage service to look up every stored object and may therefore take much longer than running `fsck.s3ql`, but can also offer a much stronger assurance that no data has been lost by the storage service. To “recover” from damaged storage objects in the backend, the damaged objects found by `s3ql_verify` have to be explicitly deleted (so that a successive `fsck.s3ql` is able detect them as missing, correct the file system metadata,

and move any affected files to `lost+found`). This procedure is currently not automated, so it is generally a good idea to choose a storage service where the expected data durability is high enough so that the possibility of a lost object (and thus the need to run any full checks) can be neglected over long periods of time.

FILE SYSTEM CREATION

A S3QL file system is created with the **mkfs.s3ql** command. It has the following syntax:

```
mkfs.s3ql [options] <storage url>
```

This command accepts the following options:

- cachedir <path>** Store cached data in this directory (default: `~/ .s3ql`)
- log <target>** Destination for log messages. Specify `none` for standard output or `syslog` for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: `None`
- debug-modules <modules>** Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the `--log` option.
- debug** Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the `--log` option.
- quiet** be really quiet
- backend-options <options>** Backend specific options (separate by commas). See back-end documentation for available options.
- version** just print program version and exit
- authfile <path>** Read authentication credentials from this file (default: `~/ .s3ql/authinfo2`)
- L <name>** Filesystem label
- max-obj-size <size>** Maximum size of storage objects in KiB. Files bigger than this will be spread over multiple objects in the storage backend. Default: 10240 KiB.
- plain** Create unencrypted file system.

Unless you have specified the `--plain` option, **mkfs.s3ql** will ask you to enter an encryption password. This password will *not* be read from an authentication file specified with the `--authfile` option to prevent accidental creation of an encrypted file system.

Note that:

- All data that is stored under the given storage url is assumed to be managed exclusively by S3QL. Trying to manually save additional objects (or remove or manipulate existing objects) will lead to file system corruption, and **fsck.s3ql** may delete objects that do not belong to the file system.

- With most storage backends, slashes in the storage url prefix do not have special meaning. For example, the storage urls `s3://mybucket/myprefix/` and `s3://mybucket/myprefix` are distinct. In the first case, the prefix is `myprefix/`, while in the second it is `myprefix`.
- S3QL file systems can not be “stacked”, i.e. you cannot have one file system stored at `s3://bucketname/outerprefix` and a second one at `s3://bucketname/outerprefix/innerprefix`.

MANAGING FILE SYSTEMS

The `s3qladm` command performs various operations on *unmounted* S3QL file systems. The file system *must not be mounted* when using `s3qladm` or things will go wrong badly.

The syntax is

```
s3qladm [options] <action> <storage-url>
```

where `action` may be either of **passphrase**, **upgrade**, **clear** or **download-metadata**.

The **s3qladm** accepts the following general options, no matter what specific action is being invoked:

- authfile <path>** Read authentication credentials from this file (default: `~/.s3ql/authinfo2`)
- debug-modules <modules>** Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the `--log` option.
- debug** Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the `--log` option.
- quiet** be really quiet
- log <target>** Destination for log messages. Specify `none` for standard output or `syslog` for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: `None`
- backend-options <options>** Backend specific options (separate by commas). See back-end documentation for available options.
- cachedir <path>** Store cached data in this directory (default: `~/.s3ql`)
- version** just print program version and exit

Hint: run `s3qladm <action> --help` to get help on the additional arguments that the different actions take.

6.1 Changing the Passphrase

To change the passphrase of a file system, use the `passphrase` subcommand:

```
s3qladm passphrase <storage url>
```

That this will not actually re-encrypt all the stored data with the new passphrase. Rather, S3QL generates an internal “master key” when the file system is created and uses this key to encrypt all stored data. The user-supplied passphrase

is used to encrypt/decrypt only the master key, and can therefore be changed. This means, however, that if a passphrase has been disclosed to someone untrusted, changing it afterwards will **not revoke access to people who had access to the old passphrase** (because they could have decrypted the master key and retained a copy).

6.2 Upgrading the file system

If you have installed a new version of S3QL, it may sometimes be necessary to upgrade the file system metadata as well. Note that in this case the file system can no longer be accessed with older versions of S3QL after the upgrade.

During the upgrade you have to make sure that the command is not interrupted, and that no one else tries to mount, check or upgrade the file system at the same time.

To upgrade a file system from the previous to the current revision, execute

```
s3qladm upgrade <storage url>
```

6.3 Deleting a file system

A file system can be deleted with:

```
s3qladm clear <storage url>
```

This physically deletes all the data and file system structures.

6.4 Restoring Metadata Backups

If the most-recent copy of the file system metadata has been damaged irreparably, it is possible to restore one of the automatically created backup copies.

The command

```
s3qladm download-metadata <storage url>
```

will give you a list of the available metadata backups and allow you to download them. This will create two new files in the current directory, ending in `.db` and `.params`. To actually use the downloaded backup, you need to move these files into the `~/ .s3ql/` directory and run `fsck .s3ql`.

Warning: You should probably not use this functionality without having asked for help on the mailing list first (see *Further Resources / Getting Help*).

MOUNTING

A S3QL file system is mounted with the **mount.s3ql** command. It has the following syntax:

```
mount.s3ql [options] <storage url> <mountpoint>
```

Note: S3QL is not a network file system like **NFS** or **CIFS**. It can only be mounted on one computer at a time.

This command accepts the following options:

- log <target>** Destination for log messages. Specify `none` for standard output or `syslog` for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: `~/ .s3ql/mount.log`
- cachedir <path>** Store cached data in this directory (default: `~/ .s3ql`)
- debug-modules <modules>** Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the `--log` option.
- debug** Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the `--log` option.
- quiet** be really quiet
- backend-options <options>** Backend specific options (separate by commas). See back-end documentation for available options.
- version** just print program version and exit
- authfile <path>** Read authentication credentials from this file (default: `~/ .s3ql/authinfo2`)
- compress <algorithm-lvl>** Compression algorithm and compression level to use when storing new data. *algorithm* may be any of `lzma`, `bzip2`, `zlib`, or `none`. *lvl* may be any integer from 0 (fastest) to 9 (slowest). Default: `lzma-6`
- cachesize <size>** Cache size in KiB (default: autodetect).
- max-cache-entries <num>** Maximum number of entries in cache (default: autodetect). Each cache entry requires one file descriptor, so if you increase this number you have to make sure that your process file descriptor limit (as set with `ulimit -n`) is high enough (at least the number of cache entries + 100).

--keep-cache	Do not purge locally cached files on exit.
--allow-other	Normally, only the user who called <code>mount.s3ql</code> can access the mount point. This user then also has full access to it, independent of individual file permissions. If the <code>--allow-other</code> option is specified, other users can access the mount point as well and individual file permissions are taken into account for all users.
--allow-root	Like <code>--allow-other</code> , but restrict access to the mounting user and the root user.
--fg	Do not daemonize, stay in foreground
--fs-name FS_NAME	Mount name passed to fuse, the name will be shown in the first column of the system mount command output. If not specified your storage url is used.
--systemd	Run as systemd unit. Consider specifying <code>-log none</code> as well to make use of journald.
--metadata-upload-interval <seconds>	Interval in seconds between complete metadata uploads. Set to 0 to disable. Default: 24h.
--threads <no>	Number of parallel upload threads to use (default: auto).
--nfs	Enable some optimizations for exporting the file system over NFS. (default: False)

7.1 Permission Checking

If the file system is mounted with neither the `allow-root` nor `allow-other` option, the mounting user has full permissions on the S3QL file system (he is effectively root). If one (or both) of the options is used, standard unix permission checks apply, i.e. only the real root user has full access and all other users (including the mounting user) are subject to permission checks.

7.2 Compression Algorithms

S3QL supports three compression algorithms, LZMA, Bzip2 and zlib (with LZMA being the default). The compression algorithm can be specified freely whenever the file system is mounted, since it affects only the compression of new data blocks.

Roughly speaking, LZMA is slower but achieves better compression ratios than Bzip2, while Bzip2 in turn is slower but achieves better compression ratios than zlib.

For maximum file system performance, the best algorithm therefore depends on your network connection speed: the compression algorithm should be fast enough to saturate your network connection.

To find the optimal algorithm and number of parallel compression threads for your system, S3QL ships with a program called `benchmark.py` in the `contrib` directory. You should run this program on a file that has a size that is roughly equal to the block size of your file system and has similar contents. It will then determine the compression speeds for the different algorithms and the upload speeds for the specified backend and recommend the best algorithm that is fast enough to saturate your network connection.

Obviously you should make sure that there is little other system load when you run `benchmark.py` (i.e., don't compile software or encode videos at the same time).

7.3 Notes about Caching

S3QL maintains a local cache of the file system data to speed up access. The cache is block based, so it is possible that only parts of a file are in the cache.

7.3.1 Maximum Number of Cache Entries

The maximum size of the cache can be configured with the `--cachesize` option. In addition to that, the maximum number of objects in the cache is limited by the `--max-cache-entries` option, so it is possible that the cache does not grow up to the maximum cache size because the maximum number of cache elements has been reached. The reason for this limit is that each cache entry requires one open file descriptor, and Linux distributions usually limit the total number of file descriptors per process to about a thousand.

If you specify a value for `--max-cache-entries`, you should therefore make sure to also configure your system to increase the maximum number of open file handles. This can be done temporarily with the `ulimit -n` command. The method to permanently change this limit system-wide depends on your distribution.

7.3.2 Cache Flushing and Expiration

S3QL flushes changed blocks in the cache to the backend whenever a block has not been accessed for at least 10 seconds. Note that when a block is flushed, it still remains in the cache.

Cache expiration (i.e., removal of blocks from the cache) is only done when the maximum cache size is reached. S3QL always expires the least recently used blocks first.

7.4 NFS Support

S3QL filesystems can be exported over NFS. The `--nfs` option is recommended to improve performance when NFS is used, but no harm will occur when it is not specified.

NFS supports persistence of client mounts across server restarts. This means that if a client has mounted an S3QL file system over NFS, the server may unmount and remount the S3QL filesystem (or even reboot) without the client being affected beyond temporarily becoming unavailable. This poses several challenges, but is supported by S3QL as long as no `fsck.s3ql` operation is run:

Warning: If `fsck.s3ql` modifies a file system in any way, all NFS clients must unmount and re-mount the NFS share before the S3QL file system is re-mounted on the server.

7.5 Failure Modes

Once an S3QL file system has been mounted, there is a multitude of problems that can occur when communicating with the remote server. Generally, `mount.s3ql` always tries to keep the file system as accessible as possible under the circumstances. That means that if network connectivity is lost, data can still be written as long as there is space in the local cache. Attempts to read data not already present in the cache, however, will block until connection is re-established. If any sort of data corruption is detected, the file system will switch to read-only mode. Attempting to read files that are affected by the corruption will return an input/output error (*errno* set to `EIO`).

In case of other unexpected or fatal problems, `mount.s3ql` terminates, but does not unmount the file system. Any attempt to access the mountpoint will result in a “Transport endpoint not connected” error (*errno* set to `ESHUTDOWN`).

This ensures that a mountpoint whose `mount.s3ql` process has terminated can not be confused with a mountpoint containing an empty file system (which would be fatal if e.g. the mountpoint is automatically mirrored). When this has happened, the mountpoint can be cleared by using the `fusermount` command (provided by FUSE) with the `-u` parameter.

`mount.s3ql` will automatically try to re-establish the connection to the server if network connectivity is lost, and retry sending a request when the connection is established but the remote server signals a temporary problem. These attempts will be made at increasing intervals for a period up to 24 hours, with retry intervals starting at 20 ms and increasing up to 5 minutes. After 24 hours, `mount.s3ql` will give up and terminate, leaving the mountpoint inaccessible as described above.

Generally, `mount.s3ql` will also emit log messages for any unusual conditions that it encounters. The destination for these messages can be set with the `--log` parameter. It is highly recommended to periodically check these logs, for example with a tool like `logcheck`. Many potential issues that `mount.s3ql` may encounter do not justify restricting access to the file system, but should nevertheless be investigated if they occur. Checking the log messages is the only way to find out about them.

7.6 Automatic Mounting

If you want to mount and unmount an S3QL file system automatically at system startup and shutdown, you should do so with a dedicated S3QL init job (instead of using `/etc/fstab`). When using `systemd`, `mount.s3ql` can be started with `--systemd` to run as a `systemd` service of type `notify`.

Note: In principle, it is also possible to automatically mount an S3QL file system with an appropriate entry in `/etc/fstab`. However, this is not recommended for several reasons:

- file systems mounted in `/etc/fstab` will be unmounted with the `umount` command, so your system will not wait until all data has been uploaded but shutdown (or restart) immediately (this is a FUSE limitation, cf <https://github.com/libfuse/libfuse/issues/1>).
 - There is no way to tell the system that mounting S3QL requires a Python interpreter to be available, so it may attempt to run `mount.s3ql` before it has mounted the volume containing the Python interpreter.
 - There is no standard way to tell the system that internet connection has to be up before the S3QL file system can be mounted.
-

ADVANCED S3QL FEATURES

8.1 Snapshotting and Copy-on-Write

The command `s3qlcp` can be used to duplicate a directory tree without physically copying the file contents. This is made possible by the data de-duplication feature of S3QL.

The syntax of `s3qlcp` is:

```
s3qlcp [options] <src> <target>
```

This will replicate the contents of the directory `<src>` in the directory `<target>`. `<src>` has to be an existing directory and `<target>` must not exist. Moreover, both directories have to be within the same S3QL file system.

The replication will not take any additional space. Only if one of directories is modified later on, the modified data will take additional storage space.

`s3qlcp` can only be called by the user that mounted the file system and (if the file system was mounted with `--allow-other` or `--allow-root`) the root user.

Note that:

- After the replication, both source and target directory will still be completely ordinary directories. You can regard `<src>` as a snapshot of `<target>` or vice versa. However, the most common usage of `s3qlcp` is to regularly duplicate the same source directory, say `documents`, to different target directories. For a e.g. monthly replication, the target directories would typically be named something like `documents_January` for the replication in January, `documents_February` for the replication in February etc. In this case it is clear that the target directories should be regarded as snapshots of the source directory.
- Exactly the same effect could be achieved by an ordinary copy program like `cp -a`. However, this procedure would be orders of magnitude slower, because `cp` would have to read every file completely (so that S3QL had to fetch all the data over the network from the backend) before writing them into the destination folder.

8.1.1 Snapshotting vs Hardlinking

Snapshot support in S3QL is inspired by the hardlinking feature that is offered by programs like `rsync` or `storeBackup`. These programs can create a hardlink instead of copying a file if an identical file already exists in the backup. However, using hardlinks has two large disadvantages:

- backups and restores always have to be made with a special program that takes care of the hardlinking. The backup must not be touched by any other programs (they may make changes that inadvertently affect other hardlinked files)
- special care needs to be taken to handle files which are already hardlinked (the restore program needs to know that the hardlink was not just introduced by the backup program to save space)

S3QL snapshots do not have these problems, and they can be used with any backup program.

8.2 Getting Statistics

You can get more information about a mounted S3QL file system with the `s3qlstat` command. It has the following syntax:

```
s3qlstat [options] <mountpoint>
```

This will print out something like this

```
Directory entries: 1488068
Inodes:           1482991
Data blocks:      87948
Total data size:  400 GiB
After de-duplication: 51 GiB (12.98% of total)
After compression: 43 GiB (10.85% of total, 83.60% of de-duplicated)
Database size:    172 MiB (uncompressed)
(some values do not take into account not-yet-uploaded dirty blocks in cache)
```

Probably the most interesting numbers are the total size of your data, the total size after duplication, and the final size after de-duplication and compression.

`s3qlstat` can only be called by the user that mounted the file system and (if the file system was mounted with `--allow-other` or `--allow-root`) the root user.

For a full list of available options, run `s3qlstat --help`.

8.3 Immutable Trees

The command `s3qllock` can be used to make a directory tree immutable. Immutable trees can no longer be changed in any way whatsoever. You can not add new files or directories and you can not change or delete existing files and directories. The only way to get rid of an immutable tree is to use the `s3qlrm` command (see below).

For example, to make the directory tree beneath the directory `2010-04-21` immutable, execute

```
s3qllock 2010-04-21
```

Immutability is a feature designed for backups. Traditionally, backups have been made on external tape drives. Once a backup was made, the tape drive was removed and locked somewhere in a shelf. This has the great advantage that the contents of the backup are now permanently fixed. Nothing (short of physical destruction) can change or delete files in the backup.

In contrast, when backing up into an online storage system like S3QL, all backups are available every time the file system is mounted. Nothing prevents a file in an old backup from being changed again later on. In the worst case, this may make your entire backup system worthless. Imagine that your system gets infected by a nasty virus that simply deletes all files it can find – if the virus is active while the backup file system is mounted, the virus will destroy all your old backups as well!

Even if the possibility of a malicious virus or trojan horse is excluded, being able to change a backup after it has been made is generally not a good idea. A common S3QL use case is to keep the file system mounted at all times and periodically create backups with `rsync -a`. This allows every user to recover her files from a backup without having to call the system administrator. However, this also allows every user to accidentally change or delete files in one of the old backups.

Making a backup immutable protects you against all these problems. Unless you happen to run into a virus that was specifically programmed to attack S3QL file systems, backups can be neither deleted nor changed after they have been made immutable.

8.4 Fast Recursive Removal

The `s3qlrm` command can be used to recursively delete files and directories on an S3QL file system. Although `s3qlrm` is faster than using e.g. `rm -r`, the main reason for its existence is that it allows you to delete immutable trees as well. The syntax is rather simple:

```
s3qlrm <directory>
```

Be warned that there is no additional confirmation. The directory will be removed entirely and immediately.

8.5 Runtime Configuration

The `s3qlctrl` can be used to control a mounted S3QL file system. Its syntax is

```
s3qlctrl [options] <action> <mountpoint> ...
```

<mountpoint> must be the location of a mounted S3QL file system. For a list of valid options, run `s3qlctrl --help`. <action> may be either of:

- flushcache** Flush file system cache. The command blocks until the cache has been flushed.
- dropcache** Flush, and then drop file system cache. The command blocks until the cache has been flushed and dropped.
- log** Change log level.
- cachesize** Change file system cache size.
- upload-meta** Trigger a metadata upload.

UNMOUNTING

To unmount an S3QL file system, use the command:

```
umount.s3ql [options] <mountpoint>
```

This will block until all data has been written to the backend.

Only the user who mounted the file system with **mount.s3ql** is able to unmount it again. If you are root and want to unmount an S3QL file system mounted by an ordinary user, you have to use the **fusermount -u** or **umount** command instead. Note that these commands do not block until all data has been uploaded, so if you use them instead of **umount.s3ql** then you should manually wait for the **mount.s3ql** process to terminate before shutting down the system.

The **umount.s3ql** command accepts the following options:

- log <target>** Destination for log messages. Specify *none* for standard output or *syslog* for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: *None*
- debug-modules <modules>** Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the **--log** option.
- debug** Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the **--log** option.
- quiet** be really quiet
- version** just print program version and exit
- lazy, -z** Lazy unmount. Detaches the file system immediately, even if there are still open files. The data will be uploaded in the background once all open files have been closed.

If, for some reason, the **umount.s3ql** command does not work, the file system can also be unmounted with **fusermount -u -z**. Note that this command will return immediately and the file system may continue to upload data in the background for a while longer.

CHECKING FOR ERRORS

It is recommended to periodically run the `fsck.s3ql` and `s3ql_verify` commands (in this order) to ensure that the file system is consistent, and that there has been no data corruption or data loss in the storage backend.

`fsck.s3ql` is intended to detect and correct problems with the internal file system structure, caused by e.g. a file system crash or a bug in S3QL. It assumes that the storage backend can be fully trusted, i.e. if the backend reports that a specific storage object exists, `fsck.s3ql` takes that as proof that the data is present and intact.

In contrast to that, the `s3ql_verify` command is intended to check the consistency of the storage backend. It assumes that the internal file system data is correct, and verifies that all data can actually be retrieved from the backend. Running `s3ql_verify` may therefore take much longer than running `fsck.s3ql`.

10.1 Checking and repairing internal file system errors

`fsck.s3ql` checks that the internal file system structure is consistent and attempts to correct any problems it finds. If an S3QL file system has not been unmounted correctly for any reason, you need to run `fsck.s3ql` before you can mount the file system again.

The `fsck.s3ql` command has the following syntax:

```
fsck.s3ql [options] <storage url>
```

This command accepts the following options:

- log <target>** Destination for log messages. Specify `none` for standard output or `syslog` for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: `~/.s3ql/fsck.log`
- cachedir <path>** Store cached data in this directory (default: `~/.s3ql`)
- debug-modules <modules>** Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the `--log` option.
- debug** Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the `--log` option.
- quiet** be really quiet
- backend-options <options>** Backend specific options (separate by commas). See back-end documentation for available options.
- version** just print program version and exit

--authfile <path>	Read authentication credentials from this file (default: <code>~/ .s3ql/ authinfo2</code>)
--compress <algorithm-lvl>	Compression algorithm and compression level to use when storing new data. <i>algorithm</i> may be any of <code>lzma</code> , <code>bzip2</code> , <code>zlib</code> , or <code>none</code> . <i>lvl</i> may be any integer from 0 (fastest) to 9 (slowest). Default: <code>lzma-6</code>
--keep-cache	Do not purge locally cached files on exit.
--batch	If user input is required, exit without prompting.
--force	Force checking even if file system is marked clean.
--force-remote	Force use of remote metadata even when this would likely result in data loss.

10.2 Detecting and handling backend data corruption

The `s3ql_verify` command verifies all data in the file system. In contrast to `fsck.s3ql`, `s3ql_verify` does not trust the object listing returned by the backend, but actually attempts to retrieve every object. By default, `s3ql_verify` will attempt to retrieve just the metadata for every object (for e.g. the S3-compatible or Google Storage backends this corresponds to a HEAD request for each object), which is generally sufficient to determine if the object still exists. When specifying the `--data` option, `s3ql_verify` will instead read every object entirely. To determine how much data will be transmitted in total when using `--data`, look at the *After compression* row in the `s3qlstat` output.

`s3ql_verify` is not able to correct any data corruption that it finds. Instead, a list of the corrupted and/or missing objects is written to a file and the decision about the proper course of action is left to the user. If you have administrative access to the backend server, you may want to investigate the cause of the corruption or check if the missing/corrupted objects can be restored from backups. If you believe that the missing/corrupted objects are indeed lost irrevocably, you can use the `remove_objects.py` script (from the `contrib` directory of the S3QL distribution) to explicitly delete the objects from the storage backend. After that, you should run `fsck.s3ql`. Since the (now explicitly deleted) objects should now no longer be included in the object index reported by the backend, `fsck.s3ql` will identify the objects as missing, update the internal file system structures accordingly, and move the affected files into the `lost+found` directory.

The `s3ql_verify` command has the following syntax:

```
s3ql_verify [options] <storage url>
```

This command accepts the following options:

--log <target>	Destination for log messages. Specify <code>none</code> for standard output or <code>syslog</code> for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: <code>None</code>
--debug-modules <modules>	Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the <code>--log</code> option.
--debug	Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the <code>--log</code> option.
--quiet	be really quiet
--version	just print program version and exit

- cachedir <path>** Store cached data in this directory (default: `~/ .s3ql`)
- backend-options <options>** Backend specific options (separate by commas). See backend documentation for available options.
- authfile <path>** Read authentication credentials from this file (default: `~/ .s3ql/ authinfo2`)
- missing-file <name>** File to store keys of missing objects.
- corrupted-file <name>** File to store keys of corrupted objects.
- data** Read every object completely, instead of checking just the metadata.
- parallel PARALLEL** Number of connections to use in parallel.
- start-with <n>** Skip over first `<n>` objects and with verifying object `<n>+1`.

STORING BACKEND INFORMATION AND CREDENTIALS

Normally, S3QL reads username and password for the backend as well as an encryption passphrase for the file system from the terminal. Alternatively, these (and some other backend parameters) may be specified in the `~/.s3ql/authinfo2` file, or a file specified by the `--authfile` parameter.

The `authinfo` file consists of sections, led by a `[section]` header and followed by `name: value` entries. The section headers themselves are not used by S3QL but have to be unique within the file.

In each section, the following entries can be defined:

storage-url Specifies the storage url to which this section applies. If a storage url starts with the value of this entry, the section is considered applicable.

backend-login Specifies the username to use for authentication with the backend.

backend-password Specifies the password to use for authentication with the backend.

fs-passphrase Specifies the passphrase to use to decrypt the file system (if it is encrypted).

Furthermore, command line parameters that are shared between all S3QL commands that work with storage URL may be specified in this file as well by omitting the leading dashes, e.g. the `--backend-options notls` parameter may be written into the `authinfo` file as `backend-options: notls`.

When reading the `authinfo` file, S3QL considers every applicable section in order and uses the last value that it found for each entry. For example, consider the following authentication file:

```
[s3]
storage-url: s3://
backend-login: joe
backend-password: notquitesecret

[fs1]
storage-url: s3://joes-first-bucket
fs-passphrase: neitheristhis

[fs2]
storage-url: s3://joes-second-bucket
fs-passphrase: swordfish

[fs3]
storage-url: s3://joes-second-bucket/with-prefix
backend-login: bill
backend-password: bi2311
fs-passphrase: 1123bi
```

With this authentication file, S3QL would try to log in as “joe” whenever the s3 backend is used, except when accessing a storage url that begins with “s3://joes-second-bucket/with-prefix”. In that case, the last section becomes active and

S3QL would use the “bill” credentials. Furthermore, file system encryption passphrases will be used for storage urls that start with “s3://joes-first-bucket” or “s3://joes-second-bucket”.

The authentication file is parsed by the [Python ConfigParser module](#).

CONTRIBUTED PROGRAMS

S3QL comes with a few contributed programs that are not part of the core distribution (and are therefore not installed automatically by default), but which may nevertheless be useful. These programs are in the `contrib` directory of the source distribution or in `/usr/share/doc/s3ql/contrib` if you installed S3QL from a package.

12.1 benchmark.py

This program measures S3QL write performance, uplink bandwidth and compression speed to determine the limiting factor. It also gives recommendation for compression algorithm and number of upload threads to achieve maximum performance.

12.2 clone_fs.py

This program physically clones an S3QL file system from one backend into another, without recompressing or reencrypting. It can be used to migrate S3 buckets to a different storage region or storage class (standard or reduced redundancy).

12.3 pcp.py

`pcp.py` is a wrapper program that starts several `rsync` processes to copy directory trees in parallel. This is important because transferring files in parallel significantly enhances performance when copying data from an S3QL file system (see *Improving copy performance* for details).

To recursively copy the directory `/mnt/home-backup` into `/home/joe` using 8 parallel processes and preserving permissions, you would execute

```
pcp.py -a --processes=8 /mnt/home-backup/ /home/joe
```

12.4 s3ql_backup.sh

This is an example script that demonstrates how to set up a simple but powerful backup solution using S3QL and `rsync`.

The `s3ql_backup.sh` script automates the following steps:

1. Mount the file system

2. Replicate the previous backup with *s3qlcp*
3. Update the new copy with the data from the backup source using *rsync*
4. Make the new backup immutable with *s3qllock*
5. Delete old backups that are no longer needed
6. Unmount the file system

The backups are stored in directories of the form `YYYY-MM-DD_HH:mm:ss` and the *expire_backups.py* command is used to delete old backups.

12.5 expire_backups.py

expire_backups.py is a program to intelligently remove old backups that are no longer needed.

To define what backups you want to keep for how long, you define a number of *age ranges*. **expire_backups** ensures that you will have at least one backup in each age range at all times. It will keep exactly as many backups as are required for that and delete any backups that become redundant.

Age ranges are specified by giving a list of range boundaries in terms of backup cycles. Every time you create a new backup, the existing backups age by one cycle.

Example: when **expire_backups** is called with the age range definition `1 3 7 14 31`, it will guarantee that you always have the following backups available:

1. A backup that is 0 to 1 cycles old (i.e, the most recent backup)
2. A backup that is 1 to 3 cycles old
3. A backup that is 3 to 7 cycles old
4. A backup that is 7 to 14 cycles old
5. A backup that is 14 to 31 cycles old

Note: If you do backups in fixed intervals, then one cycle will be equivalent to the backup interval. The advantage of specifying the age ranges in terms of backup cycles rather than days or weeks is that it allows you to gracefully handle irregular backup intervals. Imagine that for some reason you do not turn on your computer for one month. Now all your backups are at least a month old, and if you had specified the above backup strategy in terms of absolute ages, they would all be deleted! Specifying age ranges in terms of backup cycles avoids these sort of problems.

expire_backups usage is simple. It requires backups to be stored in directories of the form `year-month-day_hour:minute:seconds (YYYY-MM-DD_HH:mm:ss)` and works on all backups in the current directory. So for the above backup strategy, the correct invocation would be:

```
expire_backups.py 1 3 7 14 31
```

When storing your backups on an S3QL file system, you probably want to specify the `--use-s3qlrm` option as well. This tells **expire_backups** to use the *s3qlrm* command to delete directories.

expire_backups uses a “state file” to keep track which backups are how many cycles old (since this cannot be inferred from the dates contained in the directory names). The standard name for this state file is `.expire_backups.dat`. If this file gets damaged or deleted, **expire_backups** no longer knows the ages of the backups and refuses to work. In this case you can use the `--reconstruct-state` option to try to reconstruct the state from the backup dates. However, the accuracy of this reconstruction depends strongly on how rigorous you have been with making backups (it is only completely correct if the time between subsequent backups has always been exactly the same), so it’s generally a good idea not to tamper with the state file.

For a full list of available options, run `expire_backups.py --help`.

12.6 `remove_objects.py`

`remove_objects.py` is a program to remove a list of objects from a storage backend. Since it acts on the backend-level, the backend need not contain an S3QL file system.

TIPS & TRICKS

13.1 SSH Backend

By combining S3QL's local backend with `sshfs`, it is possible to store an S3QL file system on arbitrary SSH servers: first mount the remote target directory into the local filesystem,

```
sshfs user@my.server.com:/mnt/s3ql /mnt/sshfs
```

and then give the mountpoint to S3QL as a local destination:

```
mount.s3ql local:///mnt/sshfs/myfsdata /mnt/s3ql
```

13.2 Permanently mounted backup file system

If you use S3QL as a backup file system, it can be useful to mount the file system permanently (rather than just mounting it for a backup and unmounting it afterwards). Especially if your file system becomes large, this saves you long mount- and unmount times if you only want to restore a single file.

If you decide to do so, you should make sure to

- Use `s3qllock` to ensure that backups are immutable after they have been made.
- Call `s3qlctrl upload-meta` right after a every backup to make sure that the newest metadata is stored safely (if you do backups often enough, this may also allow you to set the `--metadata-upload-interval` option of `mount.s3ql` to zero).

13.3 Improving copy performance

Note: The following applies only when copying data **from** an S3QL file system, **not** when copying data **to** an S3QL file system.

If you want to copy a lot of smaller files *from* an S3QL file system (e.g. for a system restore) you will probably notice that the performance is rather bad.

The reason for this is intrinsic to the way S3QL works. Whenever you read a file, S3QL first has to retrieve this file over the network from the backend. This takes a minimum amount of time (the network latency), no matter how big or small the file is. So when you copy lots of small files, 99% of the time is actually spend waiting for network data.

Theoretically, this problem is easy to solve: you just have to copy several files at the same time. In practice, however, almost all unix utilities (`cp`, `rsync`, `tar` and friends) insist on copying data one file at a time. This makes a lot of sense when copying data on the local hard disk, but in case of S3QL this is really unfortunate.

The best workaround that has been found so far is to copy files by starting several `rsync` processes at once and use exclusion rules to make sure that they work on different sets of files.

For example, the following script will start 3 `rsync` instances. The first instance handles all filenames starting with `a-f`, the second the filenames from `g-l` and the third covers the rest. The `+ */` rule ensures that every instance looks into all directories.

```
#!/bin/bash

RSYNC_ARGS="-aHv /mnt/s3ql/ /home/restore/"

rsync -f "+ */" -f "--! [a-f]*" $RSYNC_ARGS &
rsync -f "+ */" -f "--! [g-l]*" $RSYNC_ARGS &
rsync -f "+ */" -f "-- [a-l]*" $RSYNC_ARGS &

wait
```

The optimum number of parallel processes depends on your network connection and the size of the files that you want to transfer. However, starting about 10 processes seems to be a good compromise that increases performance dramatically in almost all situations.

S3QL comes with a script named `pcp.py` in the `contrib` directory that can be used to transfer files in parallel without having to write an explicit script first. See the description of [pcp.py](#) for details.

FREQUENTLY ASKED QUESTIONS

14.1 What does “python-apsw must be linked dynamically to sqlite3” mean?

This error occurs if your python-apsw module is statically linked against SQLite. You need to recompile and install the module with the proper settings for dynamic linking.

If you are using Ubuntu or Debian, make sure that you are **not** using the apsw PPA at <https://launchpad.net/~ubuntu-rogerbinns/+archive/apsw>, since this PPA provides statically linked packages. Instead, use the *python-apsw* packages shipped with Debian/Ubuntu. You can reinstall the package (after removing the apsw PPA from `/etc/apt/sources.list` and `/etc/apt/sources.list.d/*` with:

```
# dpkg --purge --force-depends python-apsw
# apt-get install python-apsw
```

14.2 How can I improve the file system throughput?

There are three possible limiting factors for file system throughput:

1. Throughput from process through kernel to S3QL
2. Compression speed
3. Upload speed

To figure out what’s the limiting factor in your case, you can use the script `contrib/benchmark.py` in the S3QL tarball. It will determine the individual throughputs, and also give a recommendation for compression algorithm and number of upload threads.

14.3 Why does `df` show 1 TB of free space?

The file system size of S3QL is unlimited. However, most Unix utilities (including `df`) have no concept of an infinite size, so S3QL has to provide some numerical value. Therefore, S3QL always reports that the file system is 50% free, but it also reports at least 1 TB. This means that if you store 1.5 TB you will have another 1.5 TB free, but if you store just 2 MB, you will still have TB free rather than just another 2 MB.

14.4 Which operating systems are supported?

S3QL is developed on Linux, but should in principle work on all FUSE supported operating systems that have the required Python modules. The `{umount}.s3ql` program does some tricks which may not work on non-Linux systems, but you can always umount with the `fusermount -u` command provided by FUSE itself (the only difference is that it will not block until all data has been uploaded to S3 but return immediately). Please report your success (or failure) with other operating systems on the mailing list, so that this information can be extended.

14.5 Is there a file size limit?

No, S3QL puts no limits on the size of your files. Any limits on maximum object size imposed by the storage service do not matter for S3QL, since files are automatically and transparently split into smaller blocks.

14.6 Suppose I want to make a small change in a very large file. Will S3QL download and re-upload the entire file?

No, S3QL splits files in blocks of a configurable size (default: 10 MB). So to make a small change in a big file, only one block needs to be transferred and not the entire file.

14.7 I don't quite understand this de-duplication feature...

The great thing about data de-duplication is that you don't need to know anything about it and will still get all the benefits. Towards you, S3QL will behave and look exactly like a file system without data duplication. When talking to the storage backend, however, the S3QL will try to store identical data only once and thus occupy less space than an ordinary file system would.

14.8 What does the “Transport endpoint not connected” error mean?

It means that the file system has crashed. Please check `mount.log` for a more useful error message and report a bug if appropriate. If you can't find any errors in `mount.log`, the mount process may have “segfaulted”. To confirm this, look for a corresponding message in the `dmesg` output. If the mount process segfaulted, please try to obtain a C backtrace (see [[Providing Debugging Info]]) of the crash and file a bug report.

To make the mountpoint available again (i.e., unmount the crashed file system), use the `fusermount -u` command.

Before reporting a bug, please make sure that you're not just using the most recent S3QL version, but also the most-recent version of the most important dependencies (`python-llfuse`, `python-apsw`, `python-lzma`).

14.9 What does “Backend reports that fs is still mounted elsewhere, aborting” mean?

It means that either the file systems has not been unmounted cleanly, or the data from the most-recent mount is not yet available from the backend. In the former case you should try to run `fsck` on the computer where the file system has been mounted most recently. In the later case, waiting may fix the problem: many storage providers have a so-called “consistency window” for which recently uploaded data may not yet be available for download. If you know that your

backend is fully consistent, or the consistency window has passed, you can also run `fsck` in this situation. S3QL will then assume that whatever data is missing at this point is not going to show up again in the future. Affected files will be moved to `/lost+found`.

14.10 Can I access an S3QL file system on multiple computers simultaneously?

To share a file system between several computers, you need to have a method for these computers to communicate, so that they can tell each other when they are changing data. However, S3QL is designed to store data in cloud storage services which can not be used for this kind of communication, because different computers may end up talking to different servers in the cloud (and there is no way for S3QL to force the cloud servers to synchronize quickly and often enough).

Therefore, the only way to share an S3QL file system among different computers is to have one “master” computer that runs `mount.s3ql`, and then shares the mountpoint over a network file system like NFS, CIFS or sshfs. If the participating computers are connected over an insecure network, NFS and CIFS should be combined with VPN software like `Tinc` (which is very easy to set up for a few computers) or OpenVPN (which is a lot more complicated).

In principle, both VPN and NFS/CIFS-alike functionality could be integrated into S3QL to allow simultaneous mounting using `mount.s3ql` directly. However, consensus among the S3QL developers is that this is not worth the increased complexity.

14.11 What maximum object size should I use?

The `--max-obj-size` option of the `mkfs.s3ql` command determines the maximum size of the data chunks that S3QL exchanges with the backend server.

For files smaller than the maximum object size, this option has no effect. Files larger than the maximum object size are split into multiple chunks. Whenever you upload or download data from the backend, this is done in complete chunks. So if you have configured a maximum object size of 10 MB, and want to read (or write) 5 bytes in a 100 MB file, you will still download (or upload) 10 MB of data. If you decreased the maximum object size to 5 MB, you'd download/upload 5 MB.

On the other hand, if you want to read the whole 100 MB, and have configured an object size of 10 MB, S3QL will have to send 10 separate requests to the backend. With a maximum object size of 1 MB, there'd be 100 separate requests. The larger the number of requests, the more inefficient this becomes, because there is a fixed time associated with the processing of each request. Also, many storage providers charge a fixed amount for each request, so downloading 100 MB in one request of 100 MB is cheaper than downloading it with 100 requests of 1 MB.

When choosing a maximum object size, you have to find a balance between these two effects. The bigger the size, the less requests will be used, but the more data is transferred uselessly. The smaller the size, the more requests will be used, but less traffic will be wasted.

Generally you should go with the default unless you have a good reason to change it. When adjusting the value, keep in mind that it only affects files larger than the maximum object size. So if most of your files are less than 1 MB, decreasing the maximum object size from the default 10 MB to 1 MB will have almost no effect.

14.12 Is there a way to make the cache persistent / access the file system offline?

No, there is no way to do this. However, for most use-cases this feature is actually not required. If you want to ensure that all data in a directory is stored in the cloud *and* available offline, all you need to do is store your data in a local directory that you periodically synchronize to an S3QL mountpoint. For example, if you would like to have an S3QL mountpoint with persistent cache at `/mnt/data`, you get can this as follows:

1. Mount a local block device at `/mnt/data`, e.g. using `ext4` or `btrfs`
2. Mount an S3QL file system at `/mnt/data_online` using a small cache size and number of threads (eg. `--threads 2 --max-cache-entries 10`).
3. Keep `/mnt/data_online` synchronized with `/mnt/data` by either
 - Periodically running `rsync`, e.g. by calling `rsync -aHA --delete-during --partial` from `cron`, or
 - Continuously synchronizing using e.g. `watch.sh` or `lsyncd` (these programs use `inotify` to constantly monitor the source directory and immediately copy over any changes)
4. Now use `/mnt/data` in the same way as you would use an S3QL file system with a persistent cache.

14.13 I would like to use S3QL with Hubic, but...

HubiC has a terribly designed API, effectively no customer service, and the servers are unreliable and produce sporadic weird errors. Don't expect any help if you encounter problems with S3QL and hubiC.

14.14 What's a reasonable metadata upload interval?

The metadata upload interval can be specified using the `--metadata-upload-interval` option of `mount.s3ql` and a reasonable value to use is (surprise!) the default of 24 hours.

To understand why intervals smaller than a few hours do not make sense, consider what happens during a metadata upload:

1. The file system is frozen (all requests will block)
2. The entire SQLite database holding the metadata is dumped into a more space efficient format.
3. The file system is unfrozen
4. The dumped database is uploaded
5. The previous metadata is backed up, and backups are rotated

In other words, uploading metadata is a very expensive operation that typically takes several seconds to several minutes.

On the other hand, consider the situation that the periodic upload is intended to prevent:

1. Your computer crashes, and the SQLite database in your S3QL cache directory is **irreparably** corrupted or lost.

This is a *very* unlikely situation. When `mount.s3ql` crashes (but your computer keeps running), the local metadata will not get corrupted. If your entire computer crashes (e.g. because of a power outage), `fsck.s3ql` is almost always still able to recover the local metadata. You have to be exceedingly unlucky to depend on the periodically uploaded metadata.

In other words, if you set the metadata upload interval to x hours, what you are saying is that you consider your system so unstable that there is a good chance that your local hard disk will be irreparably corrupted sometime in the next x hours. In such a situation, you should thus be **backing up your entire system every x hours** (not just the S3QL metadata). If the metadata upload interval coincides with the interval at which you're doing full system backups (not necessarily using S3QL), you are using the right value. If the metadata upload interval is smaller (or if you're not doing full system backups at all), you are probably uploading metadata too often.

KNOWN ISSUES

- S3QL de-duplicates data blocks based solely only on SHA256 checksums, without doing a byte-by-byte comparison of the blocks. Since it is possible for two data blocks to have the same checksum despite having different contents, this can lead to problems. If two such blocks are stored in an S3QL file system, the data in one block will be lost and replaced by the data in the other block. However, the chances of this occurring for any two blocks are about 1 in 10^{77} (2^{256}). For a file system that holds a total of 10^{34} blocks, the chances of a collision increase to about 1 in 10^9 . Storing more than 10^{34} blocks (or about 10^{25} TB with an (extremely small) block size of 4 kB) is therefore not recommended. Being exceptionally unlucky may also be a disadvantage.
- S3QL does not support Access Control Lists (ACLs). There is nothing fundamentally that prevents this, someone just has to write the necessary code.
- As of Linux kernel 3.5 S3QL file systems do not implement the “write protect” bit on directories. In other words, even if a directory has the write protect bit set, the owner of the directory can delete any files and (empty) subdirectories inside it. This is a bug in the FUSE kernel module (cf. <https://github.com/libfuse/libfuse/issues/23>) and needs to be fixed in the kernel. Unfortunately it does not look as if this is going to be fixed anytime soon (as of 2016/2/28).
- S3QL is rather slow when an application tries to write data in unreasonably small chunks. If a 1 MiB file is copied in chunks of 1 KB, this will take more than 10 times as long as when it’s copied with the (recommended) chunk size of 128 KiB.

This is a limitation of the FUSE library (which does not yet support write caching) which will hopefully be addressed in some future FUSE version.

Most applications, including e.g. GNU `cp` and `rsync`, use reasonably large buffers and are therefore not affected by this problem and perform very efficient on S3QL file systems.

However, if you encounter unexpectedly slow performance with a specific program, this might be due to the program using very small write buffers. Although this is not really a bug in the program, it might be worth to ask the program’s authors for help.

- S3QL always updates file and directory access times as if the `relatime` mount option has been specified: the access time (“`atime`”) is only updated if it is currently earlier than either the status change time (“`ctime`”) or modification time (“`mtime`”).
- S3QL directories always have an `st_nlink` value of 1. This may confuse programs that rely on directories having `st_nlink` values of $(2 + \text{number of sub directories})$.

Note that this is not a bug in S3QL. Including sub directories in the `st_nlink` value is a Unix convention, but by no means a requirement. If an application blindly relies on this convention being followed, then this is a bug in the application.

A prominent example are early versions of GNU `find`, which required the `--noleaf` option to work correctly on S3QL file systems. This bug has already been fixed in recent `find` versions.

- The `umount` and `fusermount -u` commands will *not* block until all data has been uploaded to the backend. (this is a FUSE limitation, cf. <https://github.com/libfuse/libfuse/issues/1>). If you use either command to unmount an S3QL file system, you have to take care to explicitly wait for the `mount.s3ql` process to terminate before you shut down or restart the system. Therefore it is generally not a good idea to mount an S3QL file system in `/etc/fstab` (you should use a dedicated init script instead).
- S3QL relies on the backends not to run out of space. This is a given for big storage providers like Amazon S3 or Google Storage, but you may stumble upon this if you use your own server or smaller providers.

If there is no space left in the backend, attempts to write more data into the S3QL file system will fail and the file system will be in an inconsistent state and require a file system check (and you should make sure to make space available in the backend before running the check).

Unfortunately, there is no way to handle insufficient space in the backend without leaving the file system inconsistent. Since S3QL first writes data into the cache, it can no longer return an error when it later turns out that the cache can not be committed to the backend.

- When using `python-dugong` versions 3.3 or earlier, S3QL supports only CONNECT-style proxying, which may cause issues with some proxy servers when using plain HTTP. Upgrading to `python-dugong` 3.4 or newer removes this limitation.

MANPAGES

The man pages are installed with S3QL on your system and can be viewed with the **man** command. For reference, they are also included here in the User's Guide.

16.1 The `mkfs.s3ql` command

16.1.1 Synopsis

```
mkfs.s3ql [options] <storage url>
```

16.1.2 Description

The `mkfs.s3ql` command creates a new file system in the location specified by *storage url*. The storage url depends on the backend that is used. The S3QL User's Guide should be consulted for a description of the available backends.

Unless you have specified the `--plain` option, `mkfs.s3ql` will ask you to enter an encryption password. This password will *not* be read from an authentication file specified with the `--authfile` option to prevent accidental creation of an encrypted file system.

16.1.3 Options

The `mkfs.s3ql` command accepts the following options.

- `--cachedir <path>` Store cached data in this directory (default: `~/ .s3ql`)
- `--log <target>` Destination for log messages. Specify `none` for standard output or `syslog` for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: `None`
- `--debug-modules <modules>` Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the `--log` option.
- `--debug` Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the `--log` option.
- `--quiet` be really quiet
- `--backend-options <options>` Backend specific options (separate by commas). See backend documentation for available options.

--version	just print program version and exit
--authfile <path>	Read authentication credentials from this file (default: <code>~/ .s3ql/ authinfo2</code>)
-L <name>	Filesystem label
--max-obj-size <size>	Maximum size of storage objects in KiB. Files bigger than this will be spread over multiple objects in the storage backend. Default: 10240 KiB.
--plain	Create unencrypted file system.

16.1.4 Exit Codes

`mkfs.s3ql` may terminate with the following exit codes:

- 0 Everything went well.
- 1 An unexpected error occurred. This may indicate a bug in the program.
- 2 Invalid command line argument or configuration file key.
- 3 Invalid backend option.
- 11 No such backend.
- 12 Authentication file has insecure permissions.
- 13 Unable to parse proxy settings.
- 14 Invalid credentials (Authentication failed).
- 15 No permission to access backend (Authorization denied).
- 16 Invalid storage URL, specified location does not exist in backend.
- 19 Unable to connect to backend, can't resolve hostname.
- 45 Unable to access cache directory.

16.1.5 See Also

The S3QL homepage is at <https://github.com/s3ql/s3ql/>.

The full S3QL documentation should also be installed somewhere on your system, common locations are `/usr/share/doc/s3ql` or `/usr/local/doc/s3ql`.

16.2 The `s3qladm` command

16.2.1 Synopsis

```
s3qladm [options] <action> <storage url>
```

where `action` may be either of **passphrase**, **upgrade**, **delete** or **download-metadata**.

16.2.2 Description

The **s3qladm** command performs various operations on *unmounted* S3QL file systems. The file system *must not be mounted* when using **s3qladm** or things will go wrong badly.

The storage url depends on the backend that is used. The S3QL User's Guide should be consulted for a description of the available backends.

16.2.3 Options

The **s3qladm** command accepts the following options.

- authfile <path>** Read authentication credentials from this file (default: `~/ .s3ql/ authinfo2`)
- debug-modules <modules>** Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the `--log` option.
- debug** Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the `--log` option.
- quiet** be really quiet
- log <target>** Destination for log messages. Specify `none` for standard output or `syslog` for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: `None`
- backend-options <options>** Backend specific options (separate by commas). See back-end documentation for available options.
- cachedir <path>** Store cached data in this directory (default: `~/ .s3ql`)
- version** just print program version and exit

Hint: run `s3qladm <action> --help` to get help on the additional arguments that the different actions take.

16.2.4 Actions

The following actions may be specified:

passphrase Changes the encryption passphrase of the file system.

upgrade Upgrade the file system to the newest revision.

clear Delete the file system with all the stored data.

download-metadata Interactively download backups of the file system metadata.

16.2.5 Exit Codes

s3qladm may terminate with the following exit codes:

- 0** Everything went well.
- 1** An unexpected error occurred. This may indicate a bug in the program.
- 2** Invalid command line argument or configuration file key.

- 3 Invalid backend option.
- 10 Could not open log file for writing.
- 11 No such backend.
- 12 Authentication file has insecure permissions.
- 13 Unable to parse proxy settings.
- 14 Invalid credentials (Authentication failed).
- 15 No permission to access backend (Authorization denied).
- 16 Invalid storage URL, specified location does not exist in backend.
- 17 Wrong file system passphrase.
- 18 No S3QL file system found at given storage URL.
- 19 Unable to connect to backend, can't resolve hostname.
- 45 Unable to access cache directory.

16.2.6 See Also

The S3QL homepage is at <https://github.com/s3ql/s3ql/>.

The full S3QL documentation should also be installed somewhere on your system, common locations are `/usr/share/doc/s3ql` or `/usr/local/doc/s3ql`.

16.3 The `mount .s3ql` command

16.3.1 Synopsis

```
mount.s3ql [options] <storage url> <mount point>
```

16.3.2 Description

The `mount .s3ql` command mounts the S3QL file system stored in *storage url* in the directory *mount point*. The storage url depends on the backend that is used. The S3QL User's Guide should be consulted for a description of the available backends.

16.3.3 Options

The `mount .s3ql` command accepts the following options.

- `--log <target>` Destination for log messages. Specify `none` for standard output or `syslog` for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: `~/ .s3ql/mount.log`
- `--cachedir <path>` Store cached data in this directory (default: `~/ .s3ql`)

- debug-modules <modules>** Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the `--log` option.
- debug** Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the `--log` option.
- quiet** be really quiet
- backend-options <options>** Backend specific options (separate by commas). See back-end documentation for available options.
- version** just print program version and exit
- authfile <path>** Read authentication credentials from this file (default: `~/ .s3ql/ authinfo2`)
- compress <algorithm-lvl>** Compression algorithm and compression level to use when storing new data. *algorithm* may be any of `lzma`, `bzip2`, `zlib`, or `none`. *lvl* may be any integer from 0 (fastest) to 9 (slowest). Default: `lzma-6`
- cachesize <size>** Cache size in KiB (default: autodetect).
- max-cache-entries <num>** Maximum number of entries in cache (default: autodetect). Each cache entry requires one file descriptor, so if you increase this number you have to make sure that your process file descriptor limit (as set with `ulimit -n`) is high enough (at least the number of cache entries + 100).
- keep-cache** Do not purge locally cached files on exit.
- allow-other** Normally, only the user who called `mount.s3ql` can access the mount point. This user then also has full access to it, independent of individual file permissions. If the `--allow-other` option is specified, other users can access the mount point as well and individual file permissions are taken into account for all users.
- allow-root** Like `--allow-other`, but restrict access to the mounting user and the root user.
- fg** Do not daemonize, stay in foreground
- fs-name FS_NAME** Mount name passed to fuse, the name will be shown in the first column of the system mount command output. If not specified your storage url is used.
- systemd** Run as systemd unit. Consider specifying `--log none` as well to make use of journald.
- metadata-upload-interval <seconds>** Interval in seconds between complete metadata uploads. Set to 0 to disable. Default: 24h.
- threads <no>** Number of parallel upload threads to use (default: auto).
- nfs** Enable some optimizations for exporting the file system over NFS. (default: False)

16.3.4 Exit Codes

`mount.s3ql` may terminate with the following exit codes:

- 0 Everything went well.
- 1 An unexpected error occurred. This may indicate a bug in the program.
- 2 Invalid command line argument or configuration file key.
- 3 Invalid backend option.
- 10 Could not open log file for writing.
- 11 No such backend.
- 12 Authentication file has insecure permissions.
- 13 Unable to parse proxy settings.
- 14 Invalid credentials (Authentication failed).
- 15 No permission to access backend (Authorization denied).
- 16 Invalid storage URL, specified location does not exist in backend.
- 17 Wrong file system passphrase.
- 18 No S3QL file system found at given storage URL.
- 19 Unable to connect to backend, can't resolve hostname.
- 30 File system was not unmounted cleanly.
- 31 File system appears to be mounted elsewhere.
- 32 Unsupported file system revision (too old).
- 33 Unsupported file system revision (too new).
- 34 Insufficient free nodes, need to run `fsck.s3ql`.
- 35 Attempted to mount read-only, this is not supported.
- 36 Mountpoint does not exist.
- 37 Not enough available file descriptors.
- 39 Unable to bind file system to mountpoint.
- 45 Unable to access cache directory.

16.3.5 See Also

The S3QL homepage is at <https://github.com/s3ql/s3ql/>.

The full S3QL documentation should also be installed somewhere on your system, common locations are `/usr/share/doc/s3ql` or `/usr/local/doc/s3ql`.

16.4 The `s3qlstat` command

16.4.1 Synopsis

```
s3qlstat [options] <mountpoint>
```

16.4.2 Description

The **s3qlstat** command prints statistics about the S3QL file system mounted at `mountpoint`.

s3qlstat can only be called by the user that mounted the file system and (if the file system was mounted with `--allow-other` or `--allow-root`) the root user.

16.4.3 Options

The **s3qlstat** command accepts the following options:

--log <target>	Destination for log messages. Specify <code>none</code> for standard output or <code>syslog</code> for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: <code>None</code>
--debug-modules <modules>	Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the <code>--log</code> option.
--debug	Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the <code>--log</code> option.
--quiet	be really quiet
--version	just print program version and exit
--raw	Do not pretty-print numbers

16.4.4 Exit Codes

s3qlstat may terminate with the following exit codes:

- 0 Everything went well.
- 1 An unexpected error occurred. This may indicate a bug in the program.
- 2 Invalid command line argument or configuration file key.

16.4.5 See Also

The S3QL homepage is at <https://github.com/s3ql/s3ql/>.

The full S3QL documentation should also be installed somewhere on your system, common locations are `/usr/share/doc/s3ql` or `/usr/local/doc/s3ql`.

16.5 The **s3qlctrl** command

16.5.1 Synopsis

```
s3qlctrl [options] <action> <mountpoint> ...
```

where `action` may be either of **flushcache**, **upload-meta**, **cache-size** or **log-metadata**.

16.5.2 Description

The **s3qlctrl** command performs various actions on the S3QL file system mounted in `mountpoint`.

s3qlctrl can only be called by the user that mounted the file system and (if the file system was mounted with `--allow-other` or `--allow-root`) the root user.

The following actions may be specified:

flushcache Flush file system cache, i.e. upload all changed file data to the backend. The command blocks until the cache has been flushed.

dropcache Flush, and then drop file system cache. The command blocks until the cache has been flushed and dropped.

upload-meta Upload metadata to the backend. All file system operations will block while a snapshot of the metadata is prepared for upload.

cachesize Changes the cache size of the file system. This action requires an additional argument that specifies the new cache size in KiB, so the complete command line is:

```
s3qlctrl [options] cachesize <mountpoint> <new-cache-size>
```

log Change the amount of information that is logged into `~/ .s3ql/mount.log` file. The complete syntax is:

```
s3qlctrl [options] log <mountpoint> <level> [<module> [<module> ...]]
```

here `level` is the desired new log level and may be either of *debug*, *info* or *warn*. One or more `module` may only be specified with the *debug* level and allow to restrict the debug output to just the listed modules.

16.5.3 Options

The **s3qlctrl** command also accepts the following options, no matter what specific action is being invoked:

- log <target>** Destination for log messages. Specify *none* for standard output or *syslog* for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: *None*
- debug-modules <modules>** Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the `--log` option.
- debug** Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the `--log` option.
- quiet** be really quiet
- version** just print program version and exit

Hint: run `s3qlctrl <action> --help` to get help on the additional arguments that the different actions take.

16.5.4 Exit Codes

s3qlctrl may terminate with the following exit codes:

- 0** Everything went well.
- 1** An unexpected error occurred. This may indicate a bug in the program.
- 2** Invalid command line argument or configuration file key.

16.5.5 See Also

The S3QL homepage is at <https://github.com/s3ql/s3ql/>.

The full S3QL documentation should also be installed somewhere on your system, common locations are `/usr/share/doc/s3ql` or `/usr/local/doc/s3ql`.

16.6 The `s3qlcp` command

16.6.1 Synopsis

```
s3qlcp [options] <source-dir> <dest-dir>
```

16.6.2 Description

The `s3qlcp` command duplicates the directory tree `source-dir` into `dest-dir` without physically copying the file contents. Both source and destination must lie inside the same S3QL file system.

The replication will not take any additional space. Only if one of directories is modified later on, the modified data will take additional storage space.

`s3qlcp` can only be called by the user that mounted the file system and (if the file system was mounted with `--allow-other` or `--allow-root`) the root user.

Note that:

- After the replication, both source and target directory will still be completely ordinary directories. You can regard `<src>` as a snapshot of `<target>` or vice versa. However, the most common usage of `s3qlcp` is to regularly duplicate the same source directory, say `documents`, to different target directories. For a e.g. monthly replication, the target directories would typically be named something like `documents_January` for the replication in January, `documents_February` for the replication in February etc. In this case it is clear that the target directories should be regarded as snapshots of the source directory.
- Exactly the same effect could be achieved by an ordinary copy program like `cp -a`. However, this procedure would be orders of magnitude slower, because `cp` would have to read every file completely (so that S3QL had to fetch all the data over the network from the backend) before writing them into the destination folder.

Snapshotting vs Hardlinking

Snapshot support in S3QL is inspired by the hardlinking feature that is offered by programs like `rsync` or `storeBackup`. These programs can create a hardlink instead of copying a file if an identical file already exists in the backup. However, using hardlinks has two large disadvantages:

- backups and restores always have to be made with a special program that takes care of the hardlinking. The backup must not be touched by any other programs (they may make changes that inadvertently affect other hardlinked files)
- special care needs to be taken to handle files which are already hardlinked (the restore program needs to know that the hardlink was not just introduced by the backup program to save space)

S3QL snapshots do not have these problems, and they can be used with any backup program.

16.6.3 Options

The `s3qlcp` command accepts the following options:

<code>--log <target></code>	Destination for log messages. Specify <code>none</code> for standard output or <code>syslog</code> for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: <code>None</code>
<code>--debug-modules <modules></code>	Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the <code>--log</code> option.
<code>--debug</code>	Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the <code>--log</code> option.
<code>--quiet</code>	be really quiet
<code>--version</code>	just print program version and exit

16.6.4 Exit Codes

`s3qlcp` may terminate with the following exit codes:

- 0 Everything went well.
- 1 An unexpected error occurred. This may indicate a bug in the program.
- 2 Invalid command line argument or configuration file key.

16.6.5 See Also

The S3QL homepage is at <https://github.com/s3ql/s3ql/>.

The full S3QL documentation should also be installed somewhere on your system, common locations are `/usr/share/doc/s3ql` or `/usr/local/doc/s3ql`.

16.7 The `s3qlrm` command

16.7.1 Synopsis

```
s3qlrm [options] <directory>
```

16.7.2 Description

The `s3qlrm` command recursively deletes files and directories on an S3QL file system. Although `s3qlrm` is faster than using e.g. `rm -r``, the main reason for its existence is that it allows you to delete immutable trees (which can be created with `s3qllock`) as well.

Be warned that there is no additional confirmation. The directory will be removed entirely and immediately.

`s3qlrm` can only be called by the user that mounted the file system and (if the file system was mounted with `--allow-other` or `--allow-root`) the root user.

16.7.3 Options

The **s3qlrm** command accepts the following options:

--log <target>	Destination for log messages. Specify <code>none</code> for standard output or <code>syslog</code> for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: <code>None</code>
--debug-modules <modules>	Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the <code>--log</code> option.
--debug	Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the <code>--log</code> option.
--quiet	be really quiet
--version	just print program version and exit

16.7.4 Exit Codes

s3qlrm may terminate with the following exit codes:

- 0 Everything went well.
- 1 An unexpected error occurred. This may indicate a bug in the program.
- 2 Invalid command line argument or configuration file key.

16.7.5 See Also

The S3QL homepage is at <https://github.com/s3ql/s3ql/>.

The full S3QL documentation should also be installed somewhere on your system, common locations are `/usr/share/doc/s3ql` or `/usr/local/doc/s3ql`.

16.8 The **s3qllock** command

16.8.1 Synopsis

```
s3qllock [options] <directory>
```

16.8.2 Description

The **s3qllock** command makes a directory tree in an S3QL file system immutable. Immutable trees can no longer be changed in any way whatsoever. You can not add new files or directories and you can not change or delete existing files and directories. The only way to get rid of an immutable tree is to use the **s3qlrm** command.

s3qllock can only be called by the user that mounted the file system and (if the file system was mounted with `--allow-other` or `--allow-root`) the root user.

16.8.3 Rationale

Immutability is a feature designed for backups. Traditionally, backups have been made on external tape drives. Once a backup was made, the tape drive was removed and locked somewhere in a shelf. This has the great advantage that the contents of the backup are now permanently fixed. Nothing (short of physical destruction) can change or delete files in the backup.

In contrast, when backing up into an online storage system like S3QL, all backups are available every time the file system is mounted. Nothing prevents a file in an old backup from being changed again later on. In the worst case, this may make your entire backup system worthless. Imagine that your system gets infected by a nasty virus that simply deletes all files it can find – if the virus is active while the backup file system is mounted, the virus will destroy all your old backups as well!

Even if the possibility of a malicious virus or trojan horse is excluded, being able to change a backup after it has been made is generally not a good idea. A common S3QL use case is to keep the file system mounted at all times and periodically create backups with `rsync -a`. This allows every user to recover her files from a backup without having to call the system administrator. However, this also allows every user to accidentally change or delete files *in* one of the old backups.

Making a backup immutable protects you against all these problems. Unless you happen to run into a virus that was specifically programmed to attack S3QL file systems, backups can be neither deleted nor changed after they have been made immutable.

16.8.4 Options

The `s3qllock` command accepts the following options:

<code>--log <target></code>	Destination for log messages. Specify <code>none</code> for standard output or <code>syslog</code> for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: <code>None</code>
<code>--debug-modules <modules></code>	Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the <code>--log</code> option.
<code>--debug</code>	Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the <code>--log</code> option.
<code>--quiet</code>	be really quiet
<code>--version</code>	just print program version and exit

16.8.5 Exit Codes

`s3qllock` may terminate with the following exit codes:

- 0 Everything went well.
- 1 An unexpected error occurred. This may indicate a bug in the program.
- 2 Invalid command line argument or configuration file key.

16.8.6 See Also

The S3QL homepage is at <https://github.com/s3ql/s3ql/>.

The full S3QL documentation should also be installed somewhere on your system, common locations are `/usr/share/doc/s3ql` or `/usr/local/doc/s3ql`.

16.9 The `umount .s3ql` command

16.9.1 Synopsis

```
umount.s3ql [options] <mount point>
```

16.9.2 Description

The `umount .s3ql` command unmounts the S3QL file system mounted in the directory *mount point* and blocks until all data has been uploaded to the storage backend.

Only the user who mounted the file system with `mount .s3ql` is able to unmount it with `umount .s3ql`. If you are root and want to unmount an S3QL file system mounted by an ordinary user, you have to use the `fusermount -u` or `umount` command instead. Note that these commands do not block until all data has been uploaded, so if you use them instead of `umount .s3ql` then you should manually wait for the `mount .s3ql` process to terminate before shutting down the system.

16.9.3 Options

The `umount .s3ql` command accepts the following options.

<code>--log <target></code>	Destination for log messages. Specify <code>none</code> for standard output or <code>syslog</code> for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: <code>None</code>
<code>--debug-modules <modules></code>	Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the <code>--log</code> option.
<code>--debug</code>	Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the <code>--log</code> option.
<code>--quiet</code>	be really quiet
<code>--version</code>	just print program version and exit
<code>--lazy, -z</code>	Lazy unmount. Detaches the file system immediately, even if there are still open files. The data will be uploaded in the background once all open files have been closed.

16.9.4 Exit Codes

`umount .s3ql` may terminate with the following exit codes:

- 0 Everything went well.
- 1 An unexpected error occurred. This may indicate a bug in the program.
- 2 Invalid command line argument or configuration file key.

16.9.5 See Also

The S3QL homepage is at <https://github.com/s3ql/s3ql/>.

The full S3QL documentation should also be installed somewhere on your system, common locations are `/usr/share/doc/s3ql` or `/usr/local/doc/s3ql`.

16.10 The `fsck.s3ql` command

16.10.1 Synopsis

```
fsck.s3ql [options] <storage url>
```

16.10.2 Description

The `fsck.s3ql` command checks the file system in the location specified by *storage url* for errors and attempts to repair any problems. The storage url depends on the backend that is used. The S3QL User's Guide should be consulted for a description of the available backends.

16.10.3 Options

The `fsck.s3ql` command accepts the following options.

- log <target>** Destination for log messages. Specify `none` for standard output or `syslog` for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: `~/ .s3ql/fsck.log`
- cachedir <path>** Store cached data in this directory (default: `~/ .s3ql`)
- debug-modules <modules>** Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the `--log` option.
- debug** Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the `--log` option.
- quiet** be really quiet
- backend-options <options>** Backend specific options (separate by commas). See back-end documentation for available options.
- version** just print program version and exit
- authfile <path>** Read authentication credentials from this file (default: `~/ .s3ql/authinfo2`)
- compress <algorithm-lvl>** Compression algorithm and compression level to use when storing new data. *algorithm* may be any of `lzma`, `bzip2`, `zlib`, or `none`. *lvl* may be any integer from 0 (fastest) to 9 (slowest). Default: `lzma-6`
- keep-cache** Do not purge locally cached files on exit.

--batch	If user input is required, exit without prompting.
--force	Force checking even if file system is marked clean.
--force-remote	Force use of remote metadata even when this would likely result in data loss.

16.10.4 Exit Codes

If **fsck.s3ql** found any file system errors (no matter if they were corrected or not), the exit code will be 128 plus one of the codes listed below. If no errors were found, the following exit codes are used as-is:

- 0** Everything went well.
- 1** An unexpected error occurred. This may indicate a bug in the program.
- 2** Invalid command line argument or configuration file key.
- 3** Invalid backend option.
- 10** Could not open log file for writing.
- 11** No such backend.
- 12** Authentication file has insecure permissions.
- 13** Unable to parse proxy settings.
- 14** Invalid credentials (Authentication failed).
- 15** No permission to access backend (Authorization denied).
- 16** Invalid storage URL, specified location does not exist in backend.
- 17** Wrong file system passphrase.
- 18** No S3QL file system found at given storage URL.
- 19** Unable to connect to backend, can't resolve hostname.
- 32** Unsupported file system revision (too old).
- 33** Unsupported file system revision (too new).
- 40** Cannot check mounted file system.
- 41** User input required, but running in batch mode.
- 42** File system check aborted by user.
- 43** Local metadata is corrupted.
- 44** Uncorrectable errors found.
- 45** Unable to access cache directory.
- 128** This error code will be *added* to one of the codes above if any file system errors have been found (no matter if they were corrected or not).

16.10.5 See Also

The S3QL homepage is at <https://github.com/s3ql/s3ql/>.

The full S3QL documentation should also be installed somewhere on your system, common locations are `/usr/share/doc/s3ql` or `/usr/local/doc/s3ql`.

16.11 The `s3ql_oauth_client` command

16.11.1 Synopsis

```
s3ql_oauth_client [options]
```

16.11.2 Description

The `s3ql_oauth_client` command may be used to obtain OAuth2 authentication tokens for use with Google Storage. It requests “user code” from Google which has to be pasted into the browser to complete the authentication process interactively. Once authentication in the browser has been completed, `s3ql_oauth_client` displays the OAuth2 refresh token.

When combined with the special username `oauth2`, the refresh token can be used as a backend passphrase when using the Google Storage S3QL backend.

16.11.3 Options

The `s3ql_oauth_client` command accepts the following options:

- | | |
|--|---|
| <code>--log <target></code> | Destination for log messages. Specify <code>none</code> for standard output or <code>syslog</code> for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: <code>None</code> |
| <code>--debug-modules <modules></code> | Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the <code>--log</code> option. |
| <code>--debug</code> | Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the <code>--log</code> option. |
| <code>--quiet</code> | be really quiet |
| <code>--version</code> | just print program version and exit |

16.11.4 Exit Codes

`s3ql_oauth_client` may terminate with the following exit codes:

- 0 Everything went well.
- 1 An unexpected error occurred. This may indicate a bug in the program.
- 2 Invalid command line argument or configuration file key.

16.11.5 See Also

The S3QL homepage is at <https://github.com/s3ql/s3ql/>.

The full S3QL documentation should also be installed somewhere on your system, common locations are `/usr/share/doc/s3ql` or `/usr/local/doc/s3ql`.

16.12 The `s3ql_verify` command

16.12.1 Synopsis

```
s3ql_verify [options] <storage url>
```

16.12.2 Description

The `s3ql_verify` command verifies all data in the file system. In contrast to `fsck.s3ql`, `s3ql_verify` does not trust the object listing returned by the backend, but actually attempts to retrieve every object. It therefore takes a lot longer.

The format of `<storage url>` depends on the backend that is used. The S3QL User's Guide should be consulted for a description of the available backends.

16.12.3 Options

The `s3ql_verify` command accepts the following options.

- `--log <target>` Destination for log messages. Specify `none` for standard output or `syslog` for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: `None`
- `--debug-modules <modules>` Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the `--log` option.
- `--debug` Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the `--log` option.
- `--quiet` be really quiet
- `--version` just print program version and exit
- `--cachedir <path>` Store cached data in this directory (default: `~/ .s3ql`)
- `--backend-options <options>` Backend specific options (separate by commas). See back-end documentation for available options.
- `--authfile <path>` Read authentication credentials from this file (default: `~/ .s3ql/ authinfo2`)
- `--missing-file <name>` File to store keys of missing objects.
- `--corrupted-file <name>` File to store keys of corrupted objects.
- `--data` Read every object completely, instead of checking just the metadata.
- `--parallel PARALLEL` Number of connections to use in parallel.
- `--start-with <n>` Skip over first `<n>` objects and with verifying object `<n>+1`.

16.12.4 Exit Codes

`s3ql_verify` may terminate with the following exit codes:

- 0 Everything went well.
- 1 An unexpected error occurred. This may indicate a bug in the program.
- 2 Invalid command line argument or configuration file key.
- 3 Invalid backend option.
- 10 Could not open log file for writing.
- 11 No such backend.
- 12 Authentication file has insecure permissions.
- 13 Unable to parse proxy settings.
- 14 Invalid credentials (Authentication failed).
- 15 No permission to access backend (Authorization denied).
- 16 Invalid storage URL, specified location does not exist in backend.
- 17 Wrong file system passphrase.
- 18 No S3QL file system found at given storage URL.
- 19 Unable to connect to backend, can't resolve hostname.
- 32 Unsupported file system revision (too old).
- 33 Unsupported file system revision (too new).
- 45 Unable to access cache directory.
- 46 The file system data was verified, and some objects were found to be missing or corrupted.

16.12.5 See Also

The S3QL homepage is at <https://github.com/s3ql/s3ql/>.

The full S3QL documentation should also be installed somewhere on your system, common locations are `/usr/share/doc/s3ql` or `/usr/local/doc/s3ql`.

16.13 The `pcp` command

16.13.1 Synopsis

```
pcp [options] <source> [<source> ...] <destination>
```

16.13.2 Description

The `pcp` command is a wrapper that starts several **sync** processes to copy directory trees in parallel. This is allows much better copying performance on file system that have relatively high latency when retrieving individual files like S3QL.

Note: Using this program only improves performance when copying *from* an S3QL file system. When copying *to* an S3QL file system, using `pcp` is more likely to *decrease* performance.

16.13.3 Options

The **pcp** command accepts the following options:

--log <target>	Destination for log messages. Specify <code>none</code> for standard output or <code>syslog</code> for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: <code>None</code>
--quiet	be really quiet
--debug-modules <modules>	Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the <code>--log</code> option.
--debug	Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the <code>--log</code> option.
--version	just print program version and exit
-a	Pass <code>-aHAX</code> option to <code>rsync</code> .
--processes <no>	Number of <code>rsync</code> processes to use (default: 10).

16.13.4 Exit Codes

pcp may terminate with the following exit codes:

- 0 Everything went well.
- 1 An unexpected error occurred. This may indicate a bug in the program.
- 2 Invalid command line argument or configuration file key.

16.13.5 See Also

pcp is shipped as part of S3QL, <https://github.com/s3ql/s3ql/>.

16.14 The `expire_backups` command

16.14.1 Synopsis

```
expire_backups [options] <age> [<age> ...]
```

16.14.2 Description

The **expire_backups** command intelligently remove old backups that are no longer needed.

To define what backups you want to keep for how long, you define a number of *age ranges*. **expire_backups** ensures that you will have at least one backup in each age range at all times. It will keep exactly as many backups as are required for that and delete any backups that become redundant.

Age ranges are specified by giving a list of range boundaries in terms of backup cycles. Every time you create a new backup, the existing backups age by one cycle.

Example: when `expire_backups` is called with the age range definition `1 3 7 14 31`, it will guarantee that you always have the following backups available:

1. A backup that is 0 to 1 cycles old (i.e, the most recent backup)
2. A backup that is 1 to 3 cycles old
3. A backup that is 3 to 7 cycles old
4. A backup that is 7 to 14 cycles old
5. A backup that is 14 to 31 cycles old

Note: If you do backups in fixed intervals, then one cycle will be equivalent to the backup interval. The advantage of specifying the age ranges in terms of backup cycles rather than days or weeks is that it allows you to gracefully handle irregular backup intervals. Imagine that for some reason you do not turn on your computer for one month. Now all your backups are at least a month old, and if you had specified the above backup strategy in terms of absolute ages, they would all be deleted! Specifying age ranges in terms of backup cycles avoids these sort of problems.

`expire_backups` usage is simple. It requires backups to be stored in directories of the form `year-month-day_hour:minute:seconds (YYYY-MM-DD_HH:mm:ss)` and works on all backups in the current directory. So for the above backup strategy, the correct invocation would be:

```
expire_backups.py 1 3 7 14 31
```

When storing your backups on an S3QL file system, you probably want to specify the `--use-s3qlrm` option as well. This tells `expire_backups` to use the `s3qlrm` command to delete directories.

`expire_backups` uses a “state file” to keep track which backups are how many cycles old (since this cannot be inferred from the dates contained in the directory names). The standard name for this state file is `.expire_backups.dat`. If this file gets damaged or deleted, `expire_backups` no longer knows the ages of the backups and refuses to work. In this case you can use the `--reconstruct-state` option to try to reconstruct the state from the backup dates. However, the accuracy of this reconstruction depends strongly on how rigorous you have been with making backups (it is only completely correct if the time between subsequent backups has always been exactly the same), so it’s generally a good idea not to tamper with the state file.

16.14.3 Options

The `expire_backups` command accepts the following options:

- | | |
|--|---|
| <code>--quiet</code> | be really quiet |
| <code>--log <target></code> | Destination for log messages. Specify <code>none</code> for standard output or <code>syslog</code> for the system logging daemon. Anything else will be interpreted as a file name. Log files will be rotated when they reach 1 MiB, and at most 5 old log files will be kept. Default: <code>None</code> |
| <code>--debug-modules <modules></code> | Activate debugging output from specified modules (use commas to separate multiple modules). Debug messages will be written to the target specified by the <code>--log</code> option. |
| <code>--debug</code> | Activate debugging output from all S3QL modules. Debug messages will be written to the target specified by the <code>--log</code> option. |
| <code>--version</code> | just print program version and exit |
| <code>--state <file></code> | File to save state information in (default: <code>“.expire_backups.dat”</code>) |
| <code>-n</code> | Dry run. Just show which backups would be deleted. |

- reconstruct-state** Try to reconstruct a missing state file from backup dates.
- use-s3qlrm** Use `s3qlrm` command to delete backups.

16.14.4 Exit Codes

`expire_backups` may terminate with the following exit codes:

- 0** Everything went well.
- 1** An unexpected error occurred. This may indicate a bug in the program.
- 2** Invalid command line argument or configuration file key.

16.14.5 See Also

`expire_backups` is shipped as part of S3QL, <https://github.com/s3ql/s3ql/>.

FURTHER RESOURCES / GETTING HELP

If you have questions or problems with S3QL that you weren't able to resolve with this manual, you might want to consider the following other resources:

- The [S3QL Wiki](#)
- The [S3QL Mailing List](#). You can subscribe by sending a mail to s3ql+subscribe@googlegroups.com.

Please report any bugs you may encounter in the [Issue Tracker](#).

IMPLEMENTATION DETAILS

This section provides some background information on how S3QL works internally. Reading this section is not necessary to use S3QL.

18.1 Metadata Storage

Like most unix filesystems, S3QL has a concept of inodes.

The contents of directory inodes (aka the names and inodes of the files and sub directories contained in a directory) are stored directly in an [SQLite](#) database. This database is stored in a special storage object that is downloaded when the file system is mounted and uploaded periodically in the background and when the file system is unmounted. This has two implications:

1. The entire file system tree can be read from the database. Fetching/storing storage objects from/in the storage backend is only required to access the contents of files (or, more precisely, inodes). This makes most file system operations very fast because no data has to be send over the network.
2. An S3QL filesystem can only be mounted on one computer at a time, using a single `mount.s3ql` process. Otherwise changes made in one mountpoint will invariably be overwritten when the second mount point is unmounted.

Sockets, FIFOs and character devices do not need any additional storage, all information about them is contained in the database.

18.2 Data Storage

The contents of file inodes are split into individual blocks. The maximum size of a block is specified when the file system is created and cannot be changed afterwards. Every block is stored as an individual object in the backend, and the mapping from inodes to blocks and from blocks to objects is stored in the database.

While the file system is mounted, blocks are cached locally.

Blocks can also be compressed and encrypted before they are stored in the storage backend. This happens during upload, i.e. the cached data is unencrypted and uncompressed.

If some files have blocks with identical contents, the blocks will be stored in the same backend object (i.e., the data is only stored once).

18.3 Data De-Duplication

Instead of uploading every block, S3QL first computes a checksum (a SHA256 hash) to check if an identical blocks has already been stored in an backend object. If that is the case, the new block will be linked to the existing object instead of being uploaded.

This procedure is invisible for the user and the contents of the block can still be changed. If several blocks share a backend object and one of the blocks is changed, the changed block is automatically stored in a new object (so that the contents of the other block remain unchanged).

18.4 Caching

When an application tries to read or write from a file, S3QL determines the block that contains the required part of the file and retrieves it from the backend or creates it if it does not yet exist. The block is then held in the cache directory. It is committed to S3 when it has not been accessed for more than a few seconds. Blocks are removed from the cache only when the maximum cache size is reached.

When the file system is unmounted, all modified blocks are written to the backend and the cache is cleaned.

18.5 Eventual Consistency Handling

S3QL has to take into account that with some storage providers, changes in objects do not propagate immediately. For example, when an Amazon S3 object is uploaded and immediately downloaded again, the downloaded data might not yet reflect the changes done in the upload (see also <http://developer.amazonwebservices.com/connect/message.jspa?messageID=38538>)

For the data blocks this is not a problem because a data blocks always get a new object ID when they are updated.

For the metadata however, S3QL has to make sure that it always downloads the most recent copy of the database when mounting the file system.

To that end, metadata versions are numbered, and the most recent version number is stored as part of the object id of a very small “marker” object. When S3QL has downloaded the metadata it checks the version number against the marker object and, if the two do not agree, waits for the most recent metadata to become available. Once the current metadata is available, the version number is increased and the marker object updated.

18.6 Encryption

When the file system is created, `mkfs.s3ql` generates a 256 bit master key by reading from `/dev/random`. The master key is encrypted with the passphrase that is entered by the user, and then stored with the rest of the file system data. Since the passphrase is only used to access the master key (which is used to encrypt the actual file system data), the passphrase can easily be changed.

Data is encrypted with a new session key for each object and each upload. The session key is generated by appending a nonce to the master key and then calculating the SHA256 hash. The nonce is generated by concatenating the object id and the current UTC time as a 32 bit float. The precision of the time is given by the Python `time()` function and usually at least 1 millisecond. The SHA256 implementation is included in the Python standard library.

Once the session key has been calculated, a SHA256 HMAC is calculated over the data that is to be uploaded. Afterwards, the data is compressed (unless `--compress none` was passed to `mount.s3ql`) and the HMAC inserted at the beginning. Both HMAC and compressed data are then encrypted using 256 bit AES in CTR mode using `PyCrypto`.

Finally, the nonce is inserted in front of the encrypted data and HMAC, and the packet is send to the backend as a new S3 object.