
WiZFS Documentation

Release 0.1.0.dev0

Ben Cole

April 07, 2015

1	Contents	3
1.1	Installation	3
1.2	Developer Documentation	4
2	Indices and tables	7

Django ZFS NAS management system. Primarily for browsing snapshots for file restoration

Contents

1.1 Installation

Before building out the project, first make sure the prerequisite system packages are installed

As this makes use of a ZFS library, it either needs to be run as root (dangerous) or you need to allow `/dev/zfs` to be accessed by non-root users

See [this issue on github](#).

1.1.1 Prerequisites

- ZFS
- SPL
- PostgreSQL
- Redis

1.1.2 Other dependencies

These are installed by buildout

- Celery (task queue)

Celery should be configured to run as a service. This is left as an exercise to the user :) (for now)

Todo

- Document how to set up Celery as a service (systemd, supervisor)
-

1.1.3 Installation

Clone the project, then from inside the project directory

```
virtualenv -p /usr/bin/python2 .  
bin/pip install cffi  
bin/pip install -r requirements/local.txt
```

Run `mange.py` with the virtualenv interpreter

```
bin/python manage.py runserver
```

1.1.4 Database Initialisation

```
su postgres
createuser wizfs --interactive
createdb -U wizfs wizfs
psql
psql> alter role wizfs unencrypted password '{your password here}';
```

Don't forget to update the settings.py with your database password.

Then run the migrations:

```
bin/python manage.py migrate
```

1.2 Developer Documentation

This covers code style, planned features, and architectural decisions for reference

1.2.1 Code Style

- PEP8 always

1.2.2 File recovery feature

The primary function of this project, is to allow convenient recovery of files from ZFS snapshots.

- Search for files across “live” filesystem(s) and snapshots
 - Done for live filesystem
- Browse filesystems and snapshot
 - Done for live filesystem
- Recover files from snapshots to live filesystem

Todo

- [Indexing snapshots](#)
-

Indexing snapshots

- WiZFS will eventually handle snapshot creation and rotation. A call to a view will create a snapshot, optionally recursively and rotate old ones.
- When a snapshot is created, we fire a task to index it. Likewise when one is deleted, remove it from the index.
- Indexing doesn't have to do anything fancy like hash files.

1.2.3 Planned Features

As the primary use case is for ZFS based storage servers there are many other useful possibilities that could be developed later:

- Disk space usage analyser - find large files/folders
- Duplicate Finder
- System health dashboard
- Backup to multiple locations

1.2.4 Duplicate Finder

A good way to free up space is to find duplicate files. There are numerous ways to do this; hashing, filesize comparison, filename comparison etc.

I intend to do something a little unique. My use case is remove duplicate HD video files (left over from transcoding etc) These are likely to have different filesize and hashes (as the content will differ). But they should have at least similar metadata.

The plan therefore is to do the following:

- Extract metadata from files whilst indexing them
- Store the metadata in Xapian faceted
- Use facets to determine duplicates

1.2.5 System health dashboard

As ZFS storage servers tend to be tucked away they're often forgotten about until something breaks.

There aren't many (any?) good, simple monitoring tools "that just work" out the box

- Monitor hardware - temp, disk health etc
- Monitor pool(s) - disk space, faulted disks
- Notifications

1.2.6 Backup to multiple locations

Using zfs send and receive commands pools can be easily backed up whilst online. There are other tools that do this, but as wizfs will manage snapshots it makes sense for it to manage backups.

Possible backup options include:

- Mirroring pool to local or remote pool
- Incremental backup
- A customisable history of snapshots
- Backups piped to xz/tar for compression
- Backups piped to GnuPG or OpenSSL for encryption
- Backup to a file to be stored on non-zfs filesystem
 - Cloud, DVD, Bluray, dumb hard drive, etc.

- Combinations of the above

Perhaps even having access to cloud providers within the app to do this automatically.

Indices and tables

- *genindex*
- *modindex*
- *search*