

LEVERAGING BTRFS, LINUX AND OPEN SOURCE



Suman Chakravartula

rockstor.com / [@rockstorinc](https://twitter.com/rockstorinc)

SDC 2015

Linux/BTRFS based Network Attached Storage(NAS) <http://www.rockstor.com/docs/contribute.html> — Edit

3,455 commits

2 branches

22 releases

9 contributors

Branch: master rockstor-core / +

Merge pull request #887 from schakrava/715_snapshots_smb_property

schakrava authored 22 hours ago latest commit f192658729

conf	don't allow deletion of cliapp api key.	a day ago
src/rockstor	update tooltip. #715	22 hours ago
test	moved webdriver tests to separate repo	2 years ago
.gitignore	Add mock to test_pools RAID0 crud in test_pools.py	5 months ago
AUTHORS	Initial transfer of code to github	3 years ago

Code

Issues 139

Pull requests 4

Wiki

Pulse

Graphs

Settings

PROJECT BRIEF

PROJECT PRINCIPLES

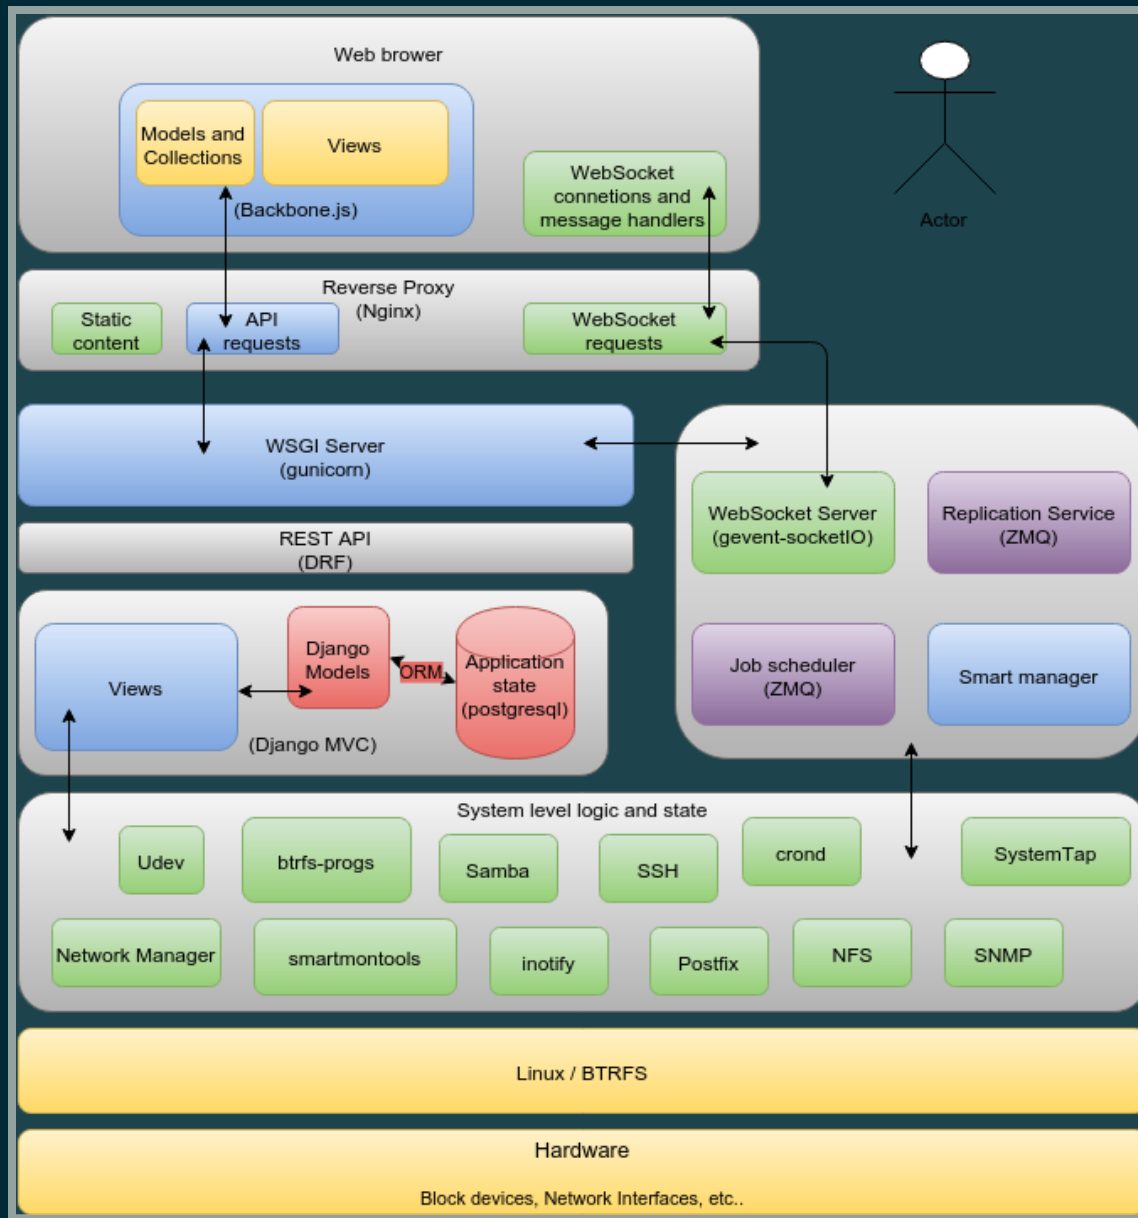


- Open Source
- Community focused
- Hardware agnostic
- Linux/BTRFS based
- Simple to install and manage
- Use boring technology whenever possible
- Contributor friendly

OVERVIEW



- Developing with BTRFS
- Server side API's with Django and DRF
- ZMQ for data replication
- Django and ZMQ for asynchronous APIs
- Backbone.js
- Gevents, SocketIO and websockets
- Docker



HIGH LEVEL ARCHITECTURE

BTRFS



Volume Management features

- mix and match drives
- add and remove drives
- change raid profiles

BTRFS



Copy on Write features

- Subvolumes
- Snapshots
- Clones
- Send/Receive Snapshots

BTRFS



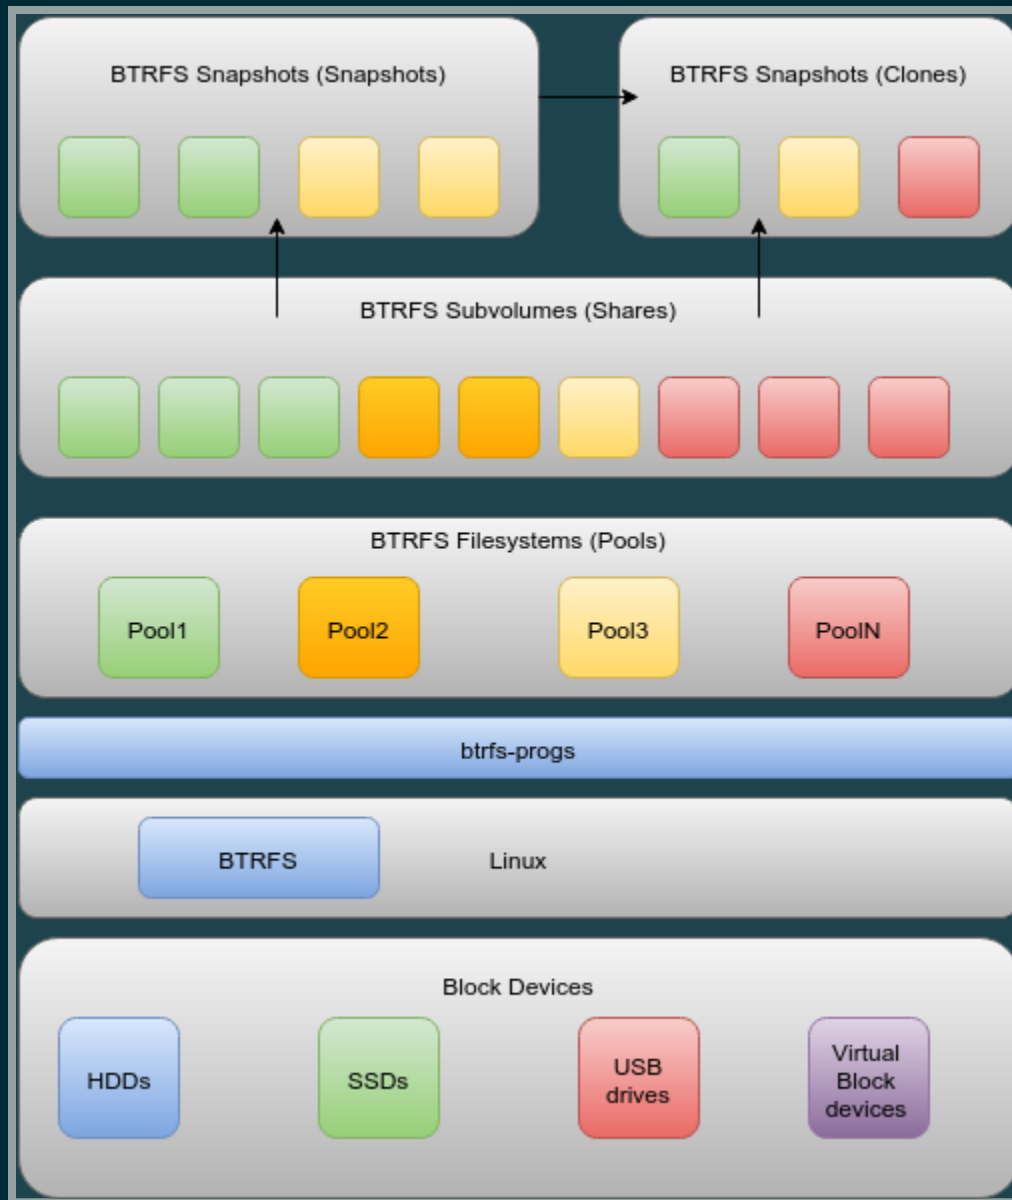
Other cool features

- Quota groups
- Compression
- Deduplication
- SSD awareness

DEVELOPING WITH BTRFS



- btrfs-progs user space utilities
- Cutting edge
- Not developer friendly yet.

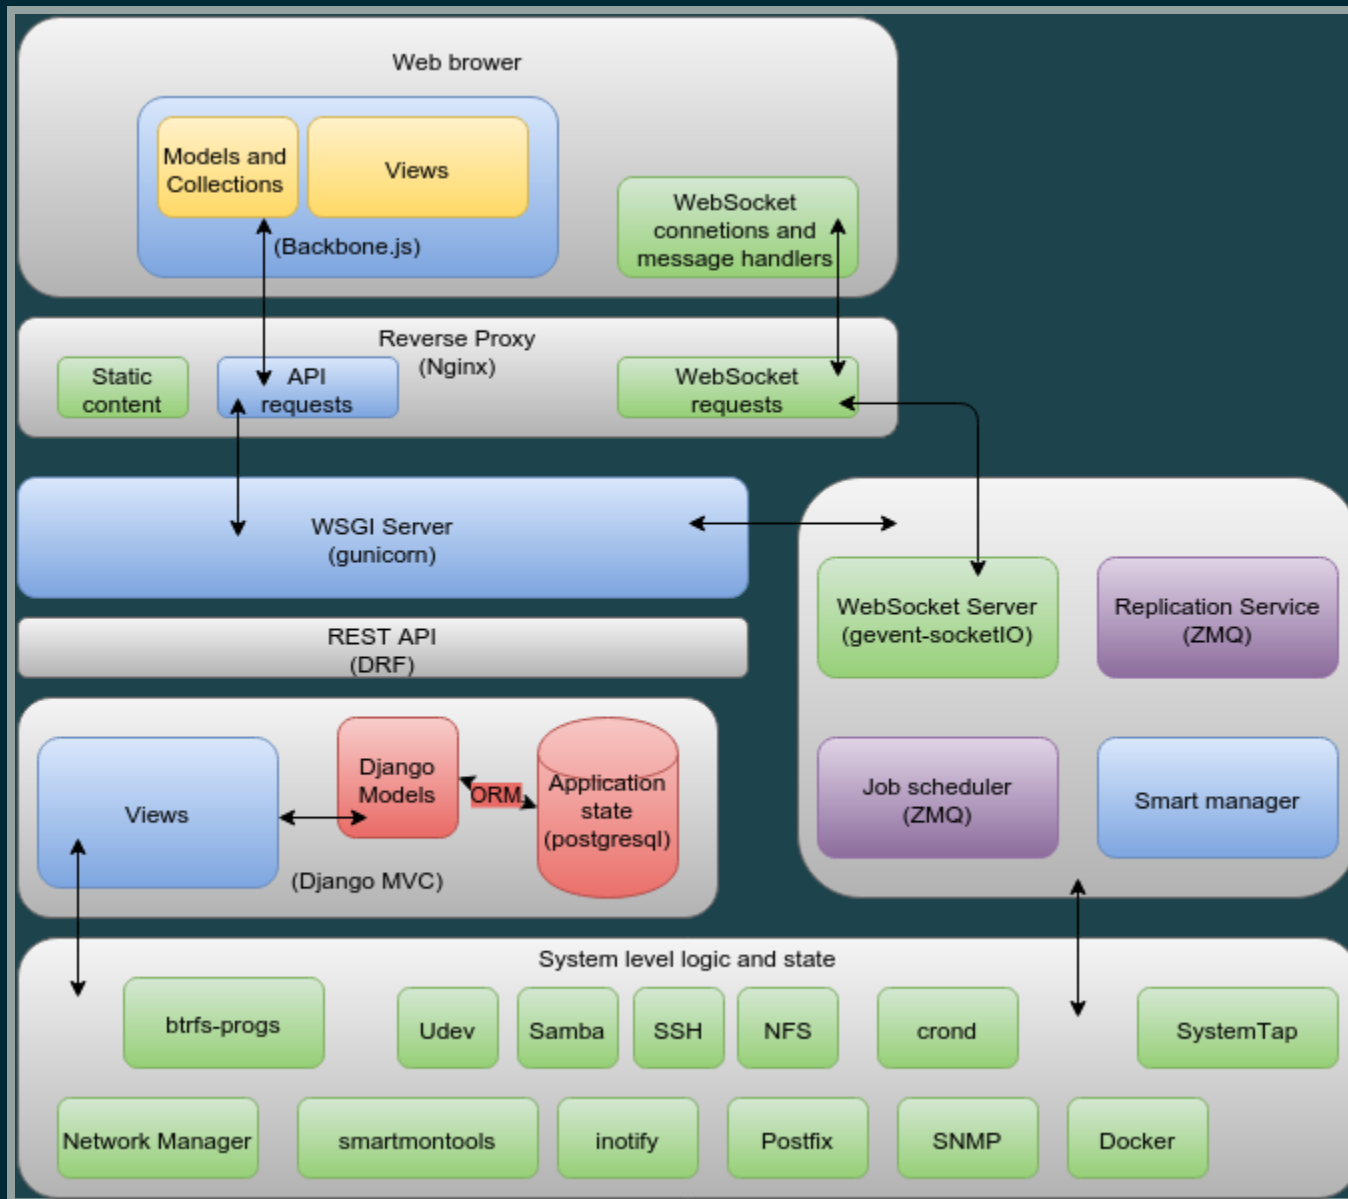


BTRFS LAYOUT

OTHER OS LEVEL TOOLS



- NetworkManager
- smartmontools
- udevadm
- SystemTap
- inotify
- and more...



APPLICATION ARCHITECTURE

DJANGO AND FRIENDS



- Django MVC
- Postgresql
- Django Rest Framework

STATE SYNCHRONIZATION



- System state is the fact
- Application state is the truth
- Minimal application state
- Eventual consistency

BACKBONE AND FRIENDS



- Backbone MVC
- JQuery
- D3

ZMQ

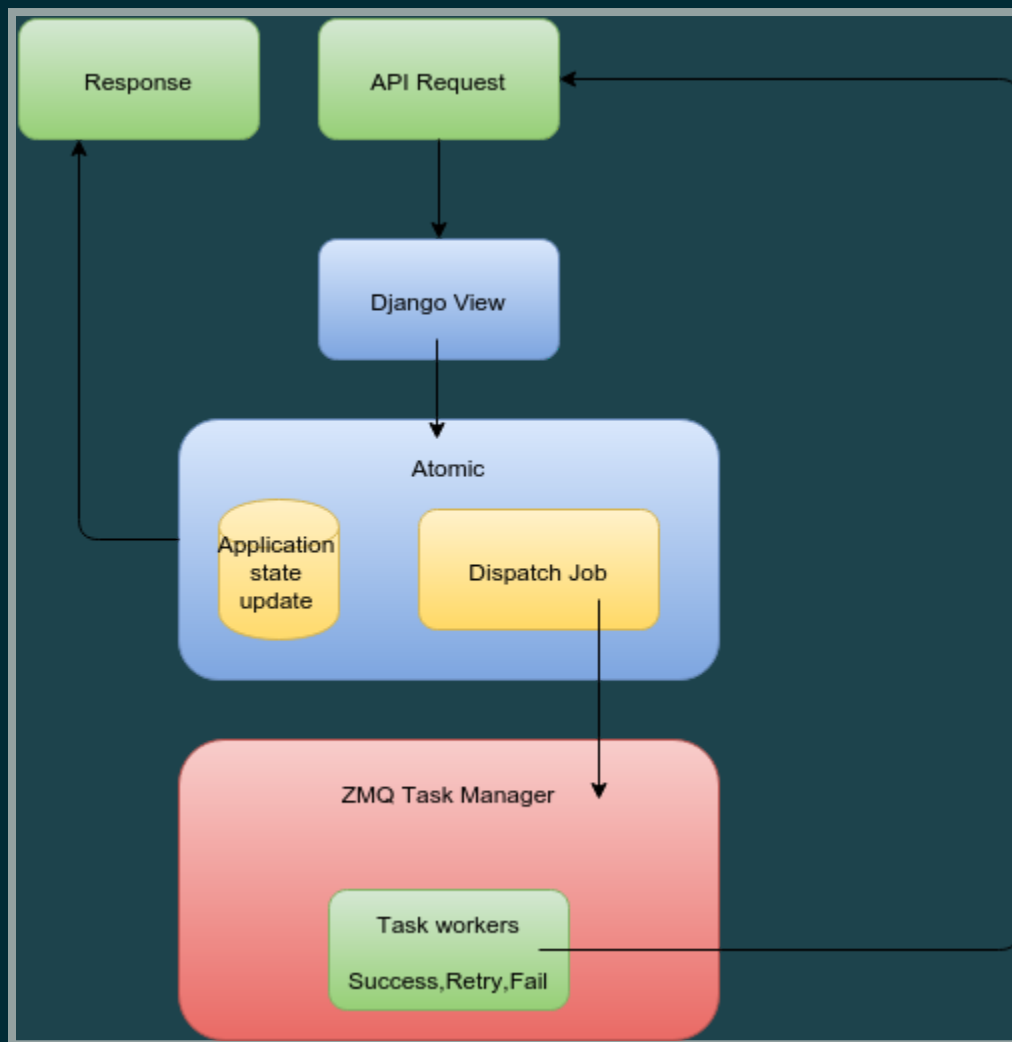


- Job scheduling
- Asynchronous replication of Shares
- SystemTap probes (On hold)

DJANGO + ZMQ = ASYNCHRONOUS



- Long running tasks: btrfs balance, scrub etc..
- Task manager uses the RESTful API



DJANGO + ZMQ

DOCKER + ZMQ



- Hosting Docker based applications
- ZMQ for app installs and updates.

ROCK-ON PLAY STORE



SYSTEMTAP AND ANALYTICS

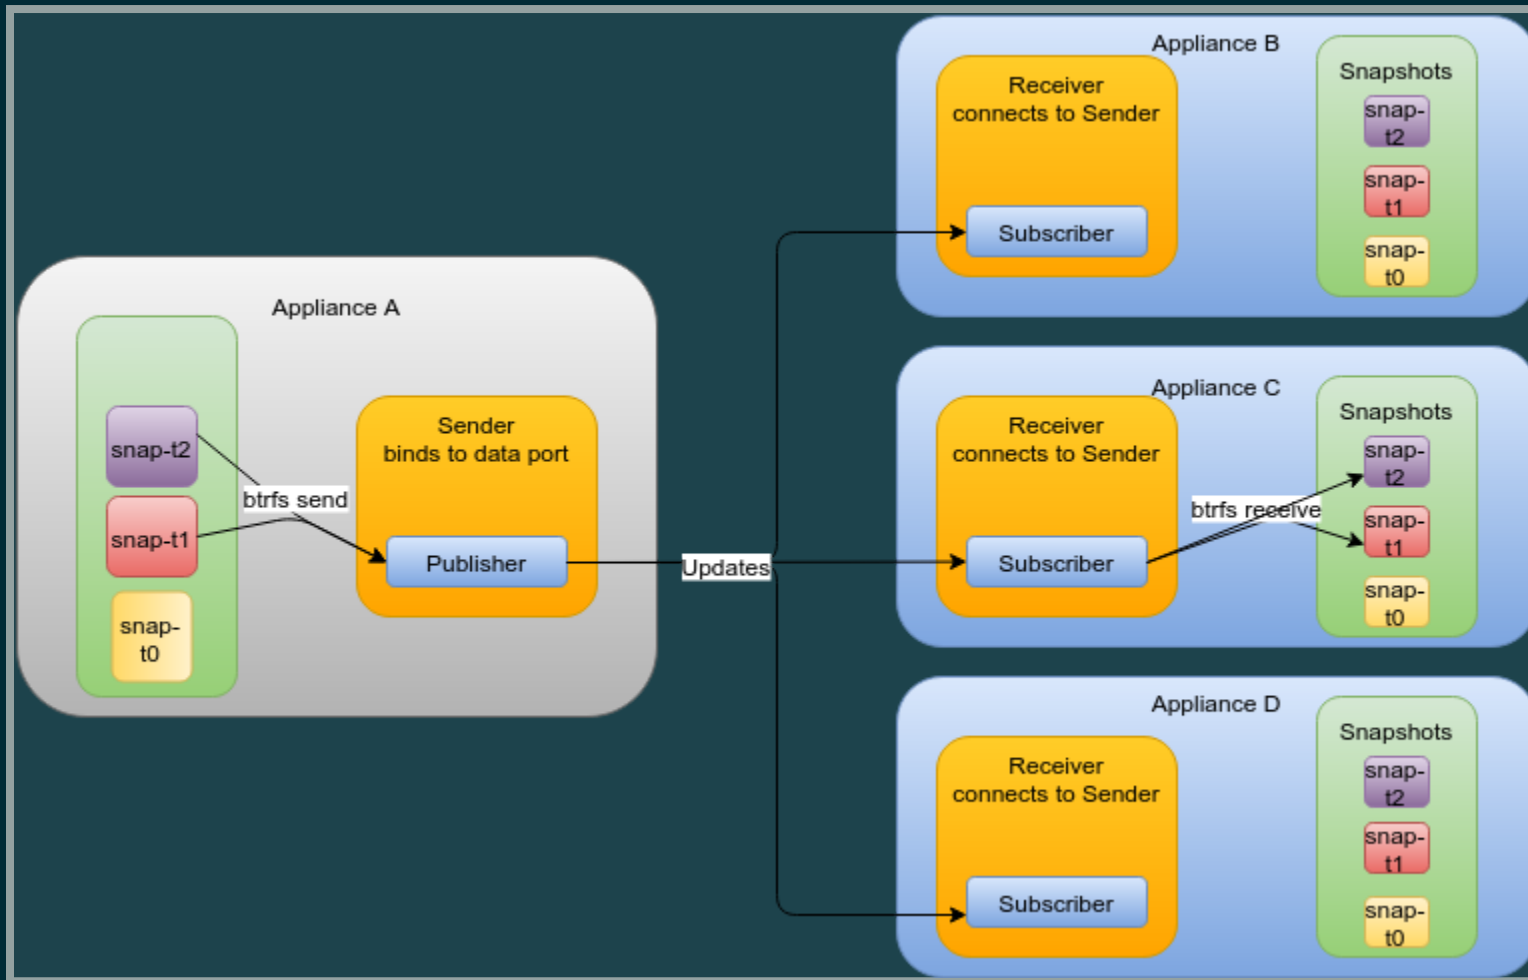


- ZMQ for scheduling stap scripts
- Probe management from the UI
- Reports using D3

ASYNCHRONOUS REPLICATION



- BTRFS send/receive
- LAN or WAN
- Replicate to multiple destinations (planned)
- Encryption (planned)



PUB-SUB FOR DATA TRANSFER

PUSH-PULL FOR META DATA



RESPONSIVE UI AND EVENT DRIVEN BACKEND



- Subscribe, not poll
- Event driven message passing
- Push notifications

GEVENT-SOCKETIO + WEBSOCKETS



SUMMARY



- developing with BTRFS
- server side API's with Django and DRF
- ZMQ for data replication
- Django and ZMQ for asynchronous APIs
- Backbone.js
- Gevents, SocketIO and websockets
- Docker

OPEN SOURCE SOFTWARE AND COMMUNITY FEEDBACK LOOP



- Community focused development
- Vision versus community wants

QUESTIONS OR COMMENTS?



Important links

- Website: rockstor.com
- Community: forum.rockstor.com
- Code: github.com/rockstor/rockstor-core
- Docs: rockstor.com/docs