Ceph and Storage Management with openATTIC

FOSDEM 2017, Brussels, BE

Lenz Grimmer <lgrimmer@suse.com>
Engineering Team Lead
SUSE LINUX GmbH

openATTIC – Our Vision

Develop an alternative to proprietary storage management systems

“Traditional” unified storage (NAS/SAN)

Support Ceph for scale-out scenarios

Open Source, backed with professional support and services

OpenATTIC – Notable Changes

Removed Enterprise/Community Edition split

- Now fully under the GPLv2
- Removed requirement for CLA
- DCA (Signed-off-by) is all that’s needed to contribute

Public Issue tracker and roadmap (JIRA)

Public pull requests / code reviews on BitBucket

Monthly releases / nightly snapshot builds

Entire code base (backend/UI/Tests/Docs) in one branch

Part of SUSE since November 2016

openATTIC – Key Aspects

Focus on data center storage management

- Support both SAN and NAS functionality without limitations
- Ceph support

Fully Open Source (GPLv2)

- No arbitrary functional restrictions
- Low entrance barrier for adoption

Based on standard Linux / OSS tools and frameworks

Multiple Linux distributions (Debian/Ubuntu/Red Hat/SUSE)

- Well-established, mature technology stack
- Broad vendor support (e.g. device drivers)
- Broad user base

openATTIC – Open Source Storage Management

- Modern WebUI
- RESTful API (Software-Defined Storage)
- Unified Storage
  - NAS (NFS, CIFS, HTTP)
  - SAN (iSCSI, Fibre Channel)
- LVM, XFS, ZFS, Btrfs, ext3/4
- Volume mirroring (DRBD®)
- Multi-node support
- Monitoring (Nagios/Icinga) built-in
- Ceph management & monitoring
- Development sponsored by SUSE

openATTIC – Components

Backend

- Python (Django)
- Django REST Framework
- Linux tools for storage management (LVM, LIO, filesystem utilities, DRBD, etc.)
- Nagios/Icinga + PNP4Nagios (Monitoring and Graphing)
- Salt / DeepSea (Ceph deployment & management)
- Stick around for Jan Fajerski's talk about DeepSea at 15:00

Web Frontend

- AngularJS (JS framework)
- Bootstrap (HTML, CSS and JS framework)
- Uses REST API exclusively

Automated Test Suites

- Python/Django Unit Tests
- REST API Test framework (Gatling)
- WebUI E2E Tests (Protractor/Jasmine)

openATTIC – High Level Architecture

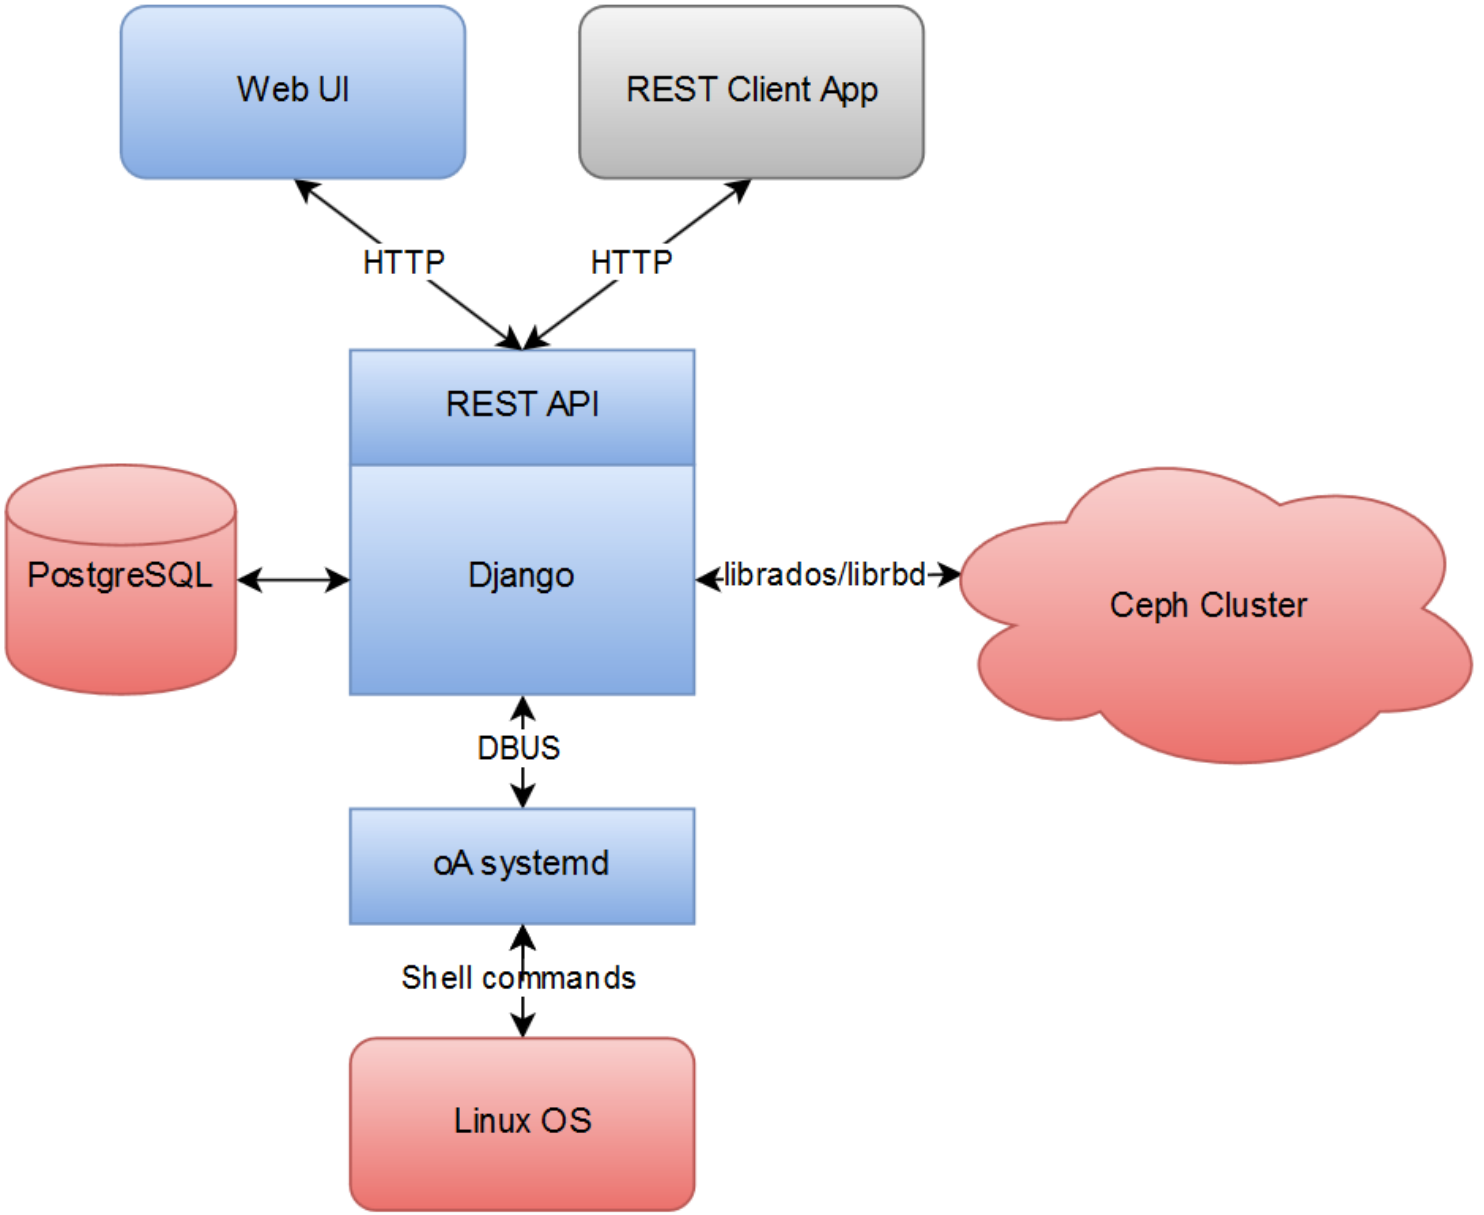

openATTIC – Storage Monitoring Infrastructure

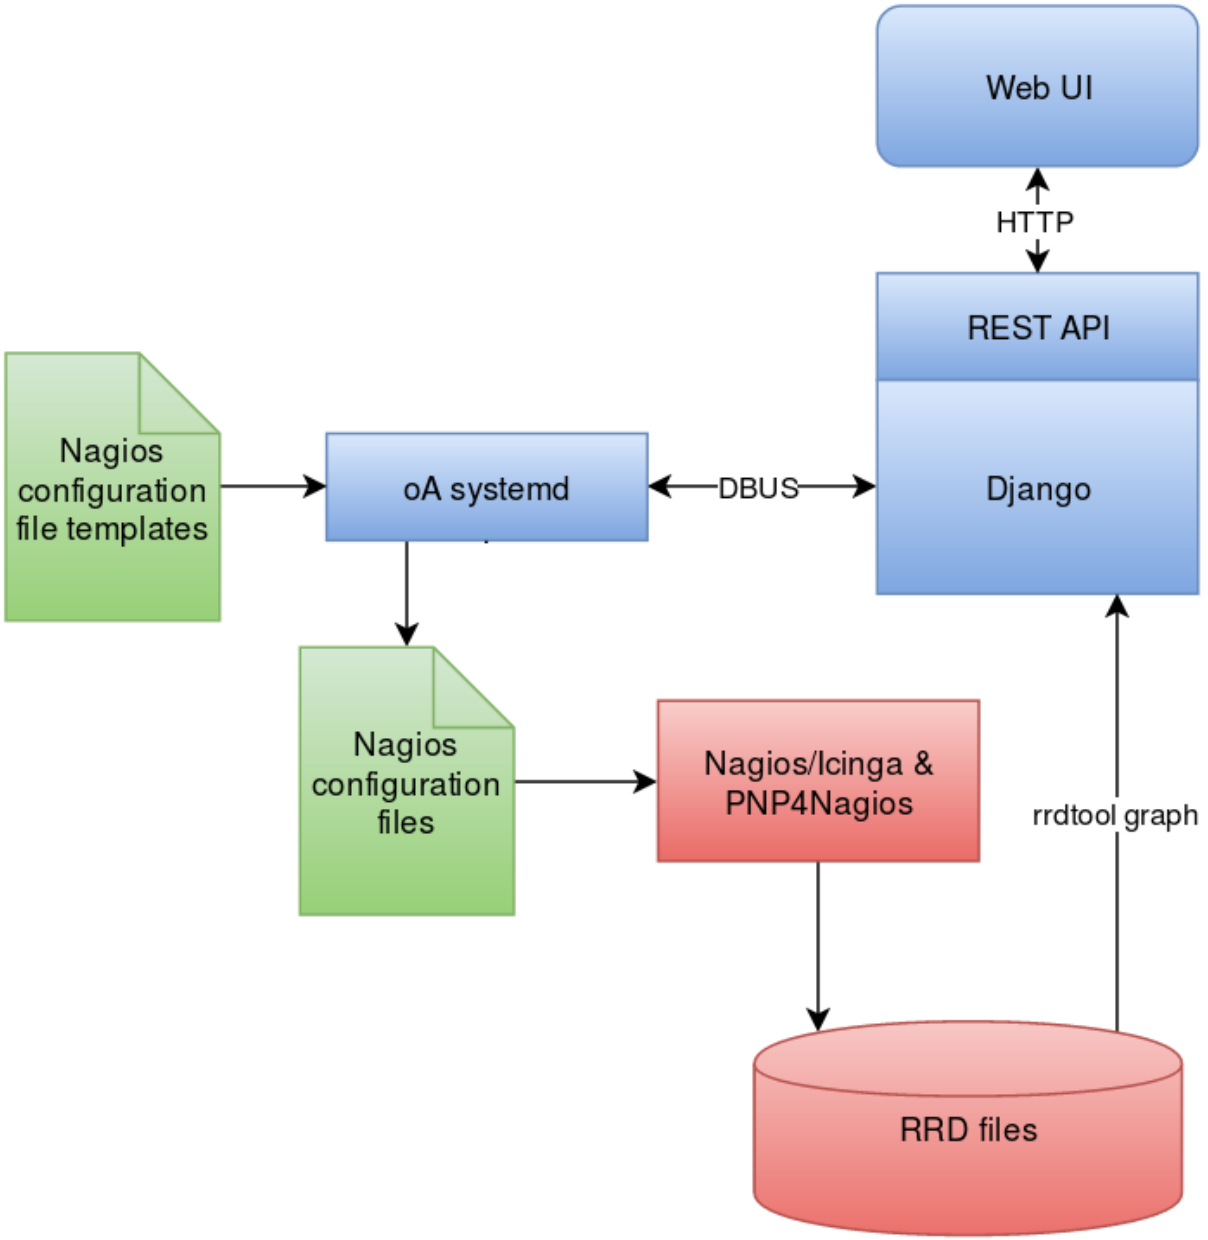

openATTIC – Ceph Monitoring Infrastructure

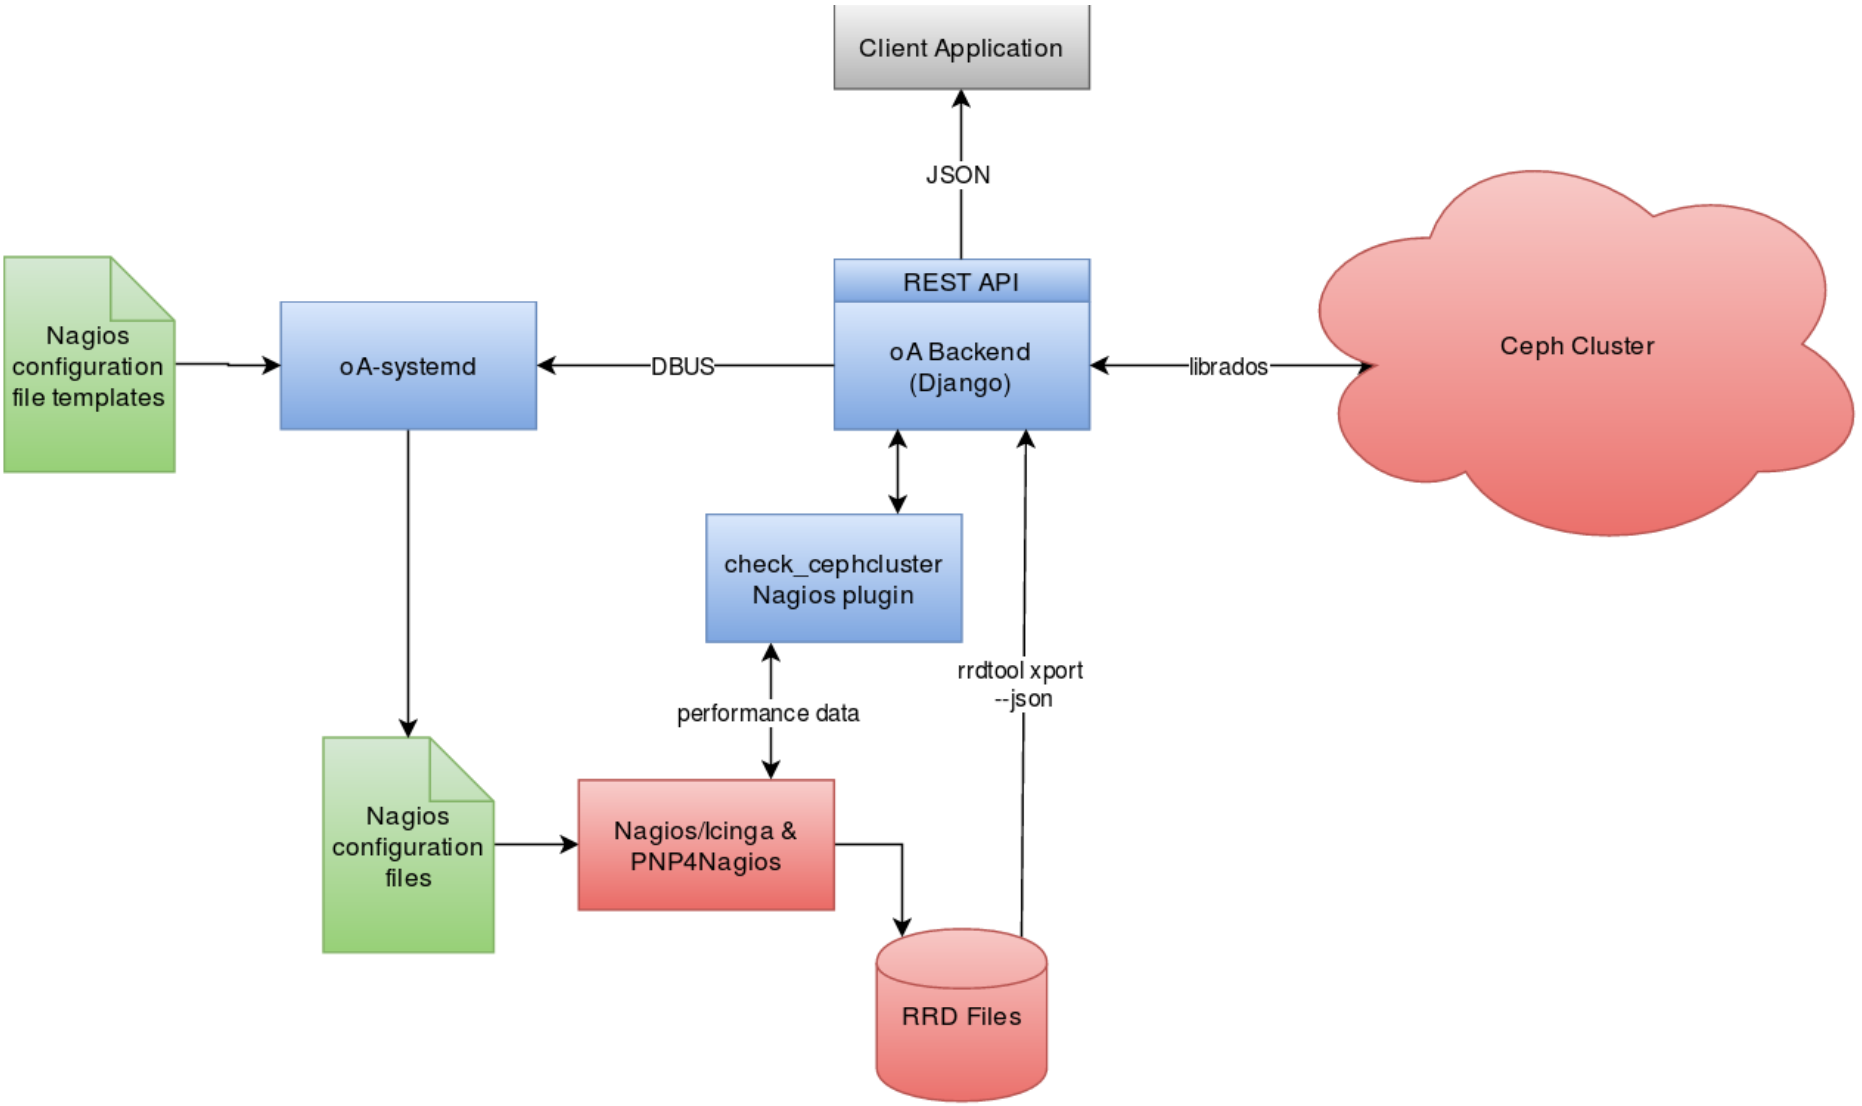

openATTIC – Storage Management Roadmap

Add DRBD volume mirroring to the WebUI (WIP)

Extend Disk and Storage Pool Management functionality

- Creating/Managing LVM Volume Groups
- Creating/Modifying Btrfs/ZFS Pools (incl. RAID setups)
- Automatic discovery of disks/pools (via udev)
- Disk health monitoring (SMART)

Extend SAN functionality (more iSCSI/FC features)

Public Roadmap on the openATTIC Jira/Wiki to solicit community feedback

openATTIC – Ceph Management Goals

Create a management & monitoring GUI tool

A tool that admins actually *want* to use

That scales without becoming overwhelming

Still should allow changes to be made elsewhere, without becoming inconsistent

openATTIC – Current Ceph Development Status

Ceph Cluster Status Dashboard (Performance Graphs, Health Status)

Pool management (view/create/delete)

Pool monitoring

Manage EC profiles

RBD management (view/create/delete/map)

RBD monitoring

OSD management (view)

CRUSH map editor

Support for managing multiple Ceph clusters

openATTIC – Ceph Development Roadmap

- Dashboard improvements (more metrics / graphs)
- Task Queue WebUI (WIP)
- Deployment, remote configuration / role assignment of Ceph nodes (via Salt Open & SUSE'S "DeepSea framework)
- iSCSI target management
- OSD Monitoring/Management
- RGW Management (e.g. users, buckets, keys) via RGW Admin Ops API
- Extend Pool Management
- CephFS Management / Monitoring
- Remote node monitoring (via Salt & collectd)

Screenshots

openATTIC – Storage Dashboard

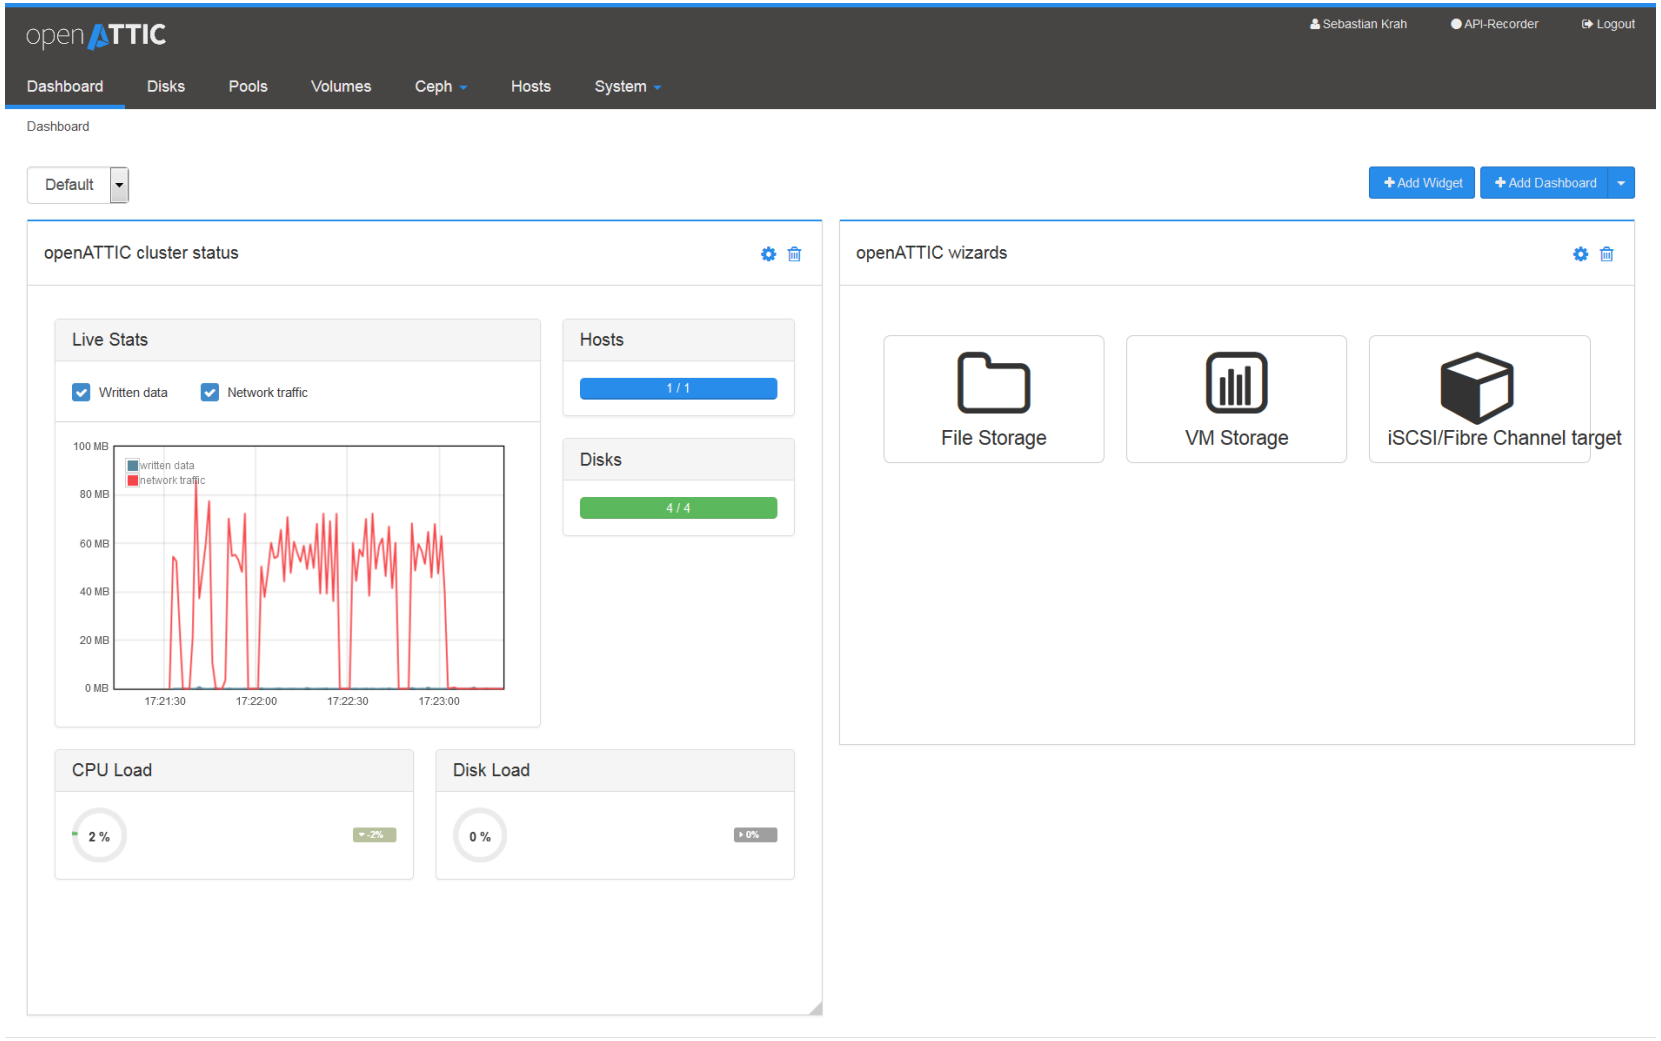

openATTIC – Volume Management

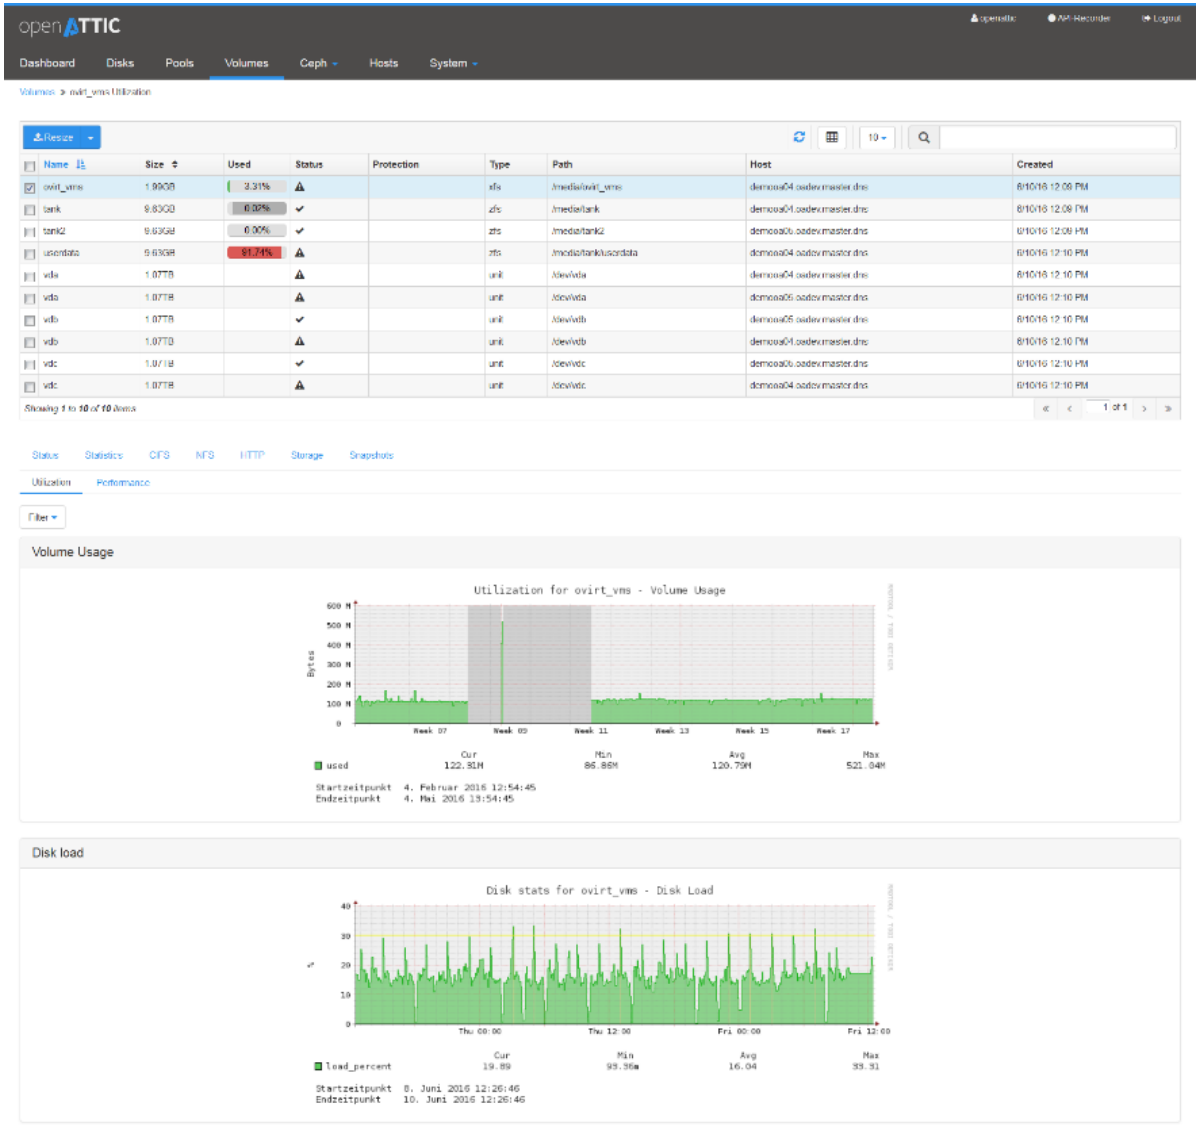

openATTIC – API Recorder

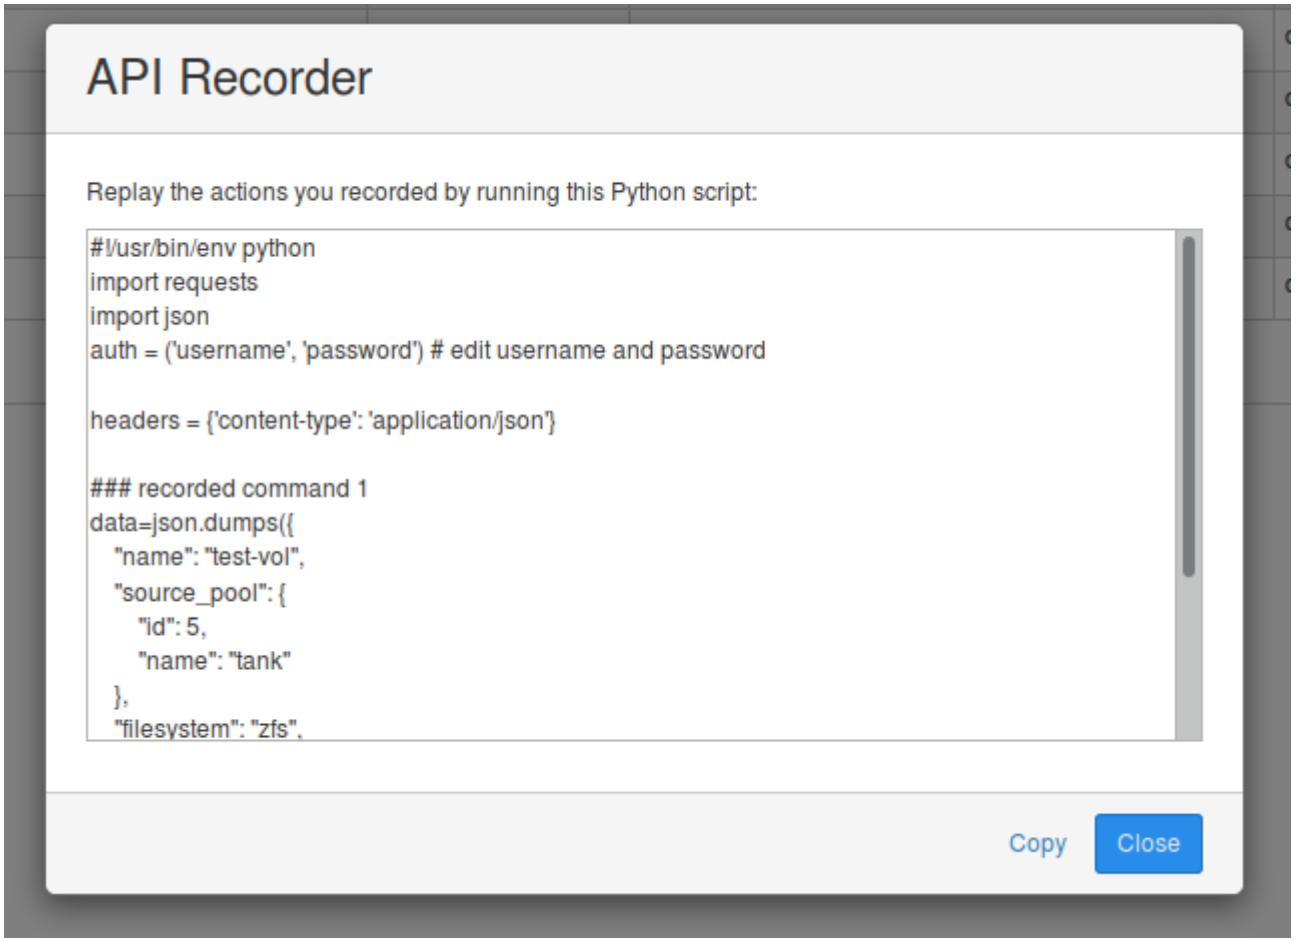

openATTIC – Ceph Cluster Dashboard

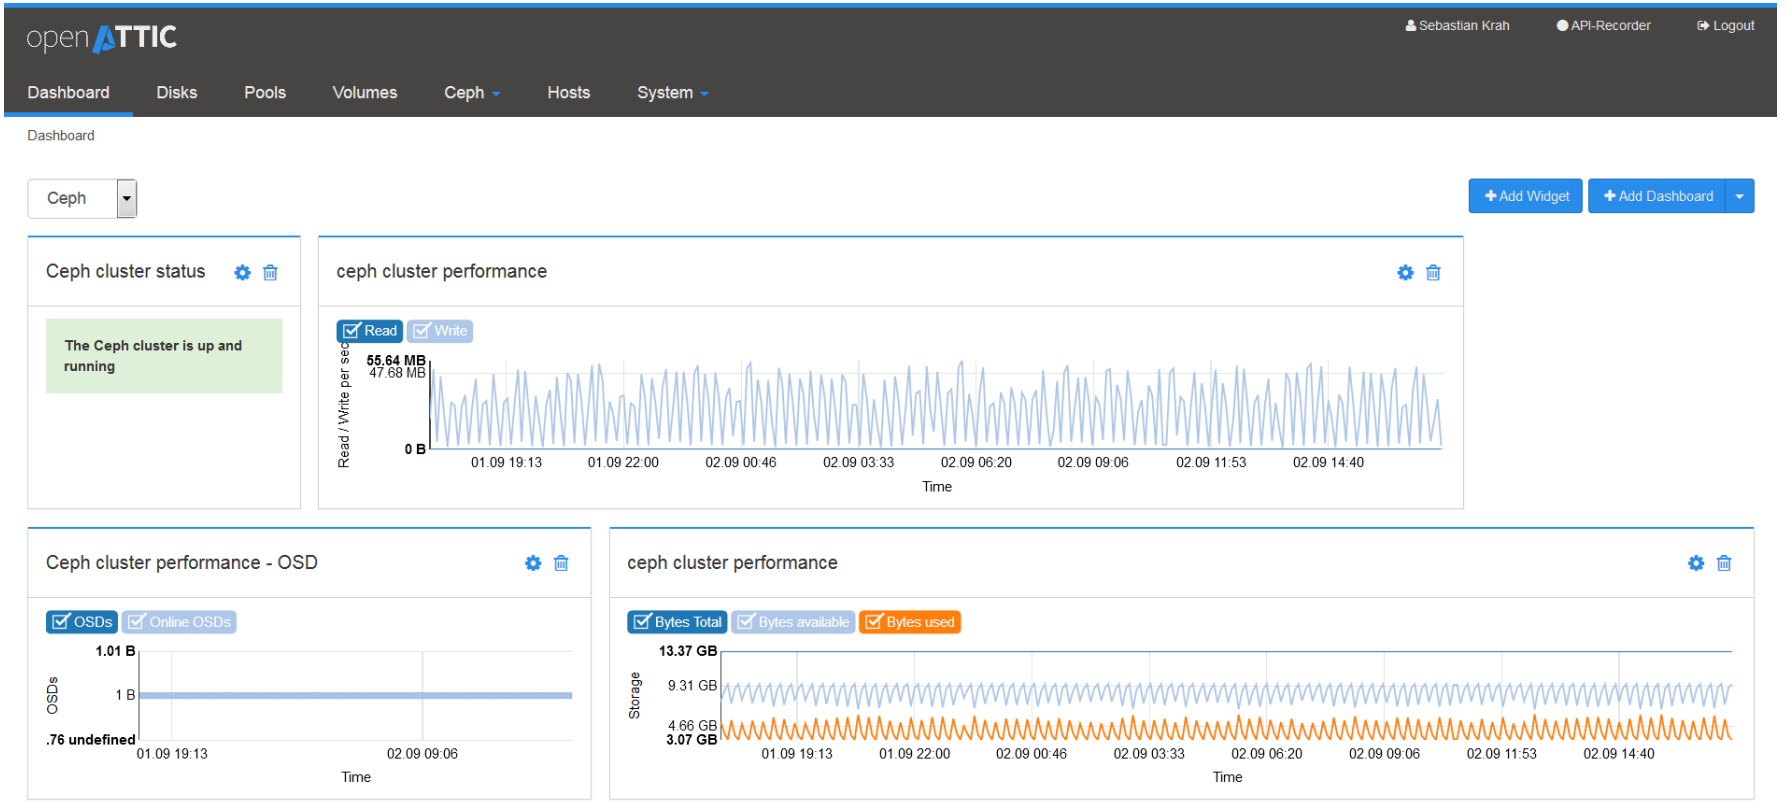

openATTIC – Ceph Pool List

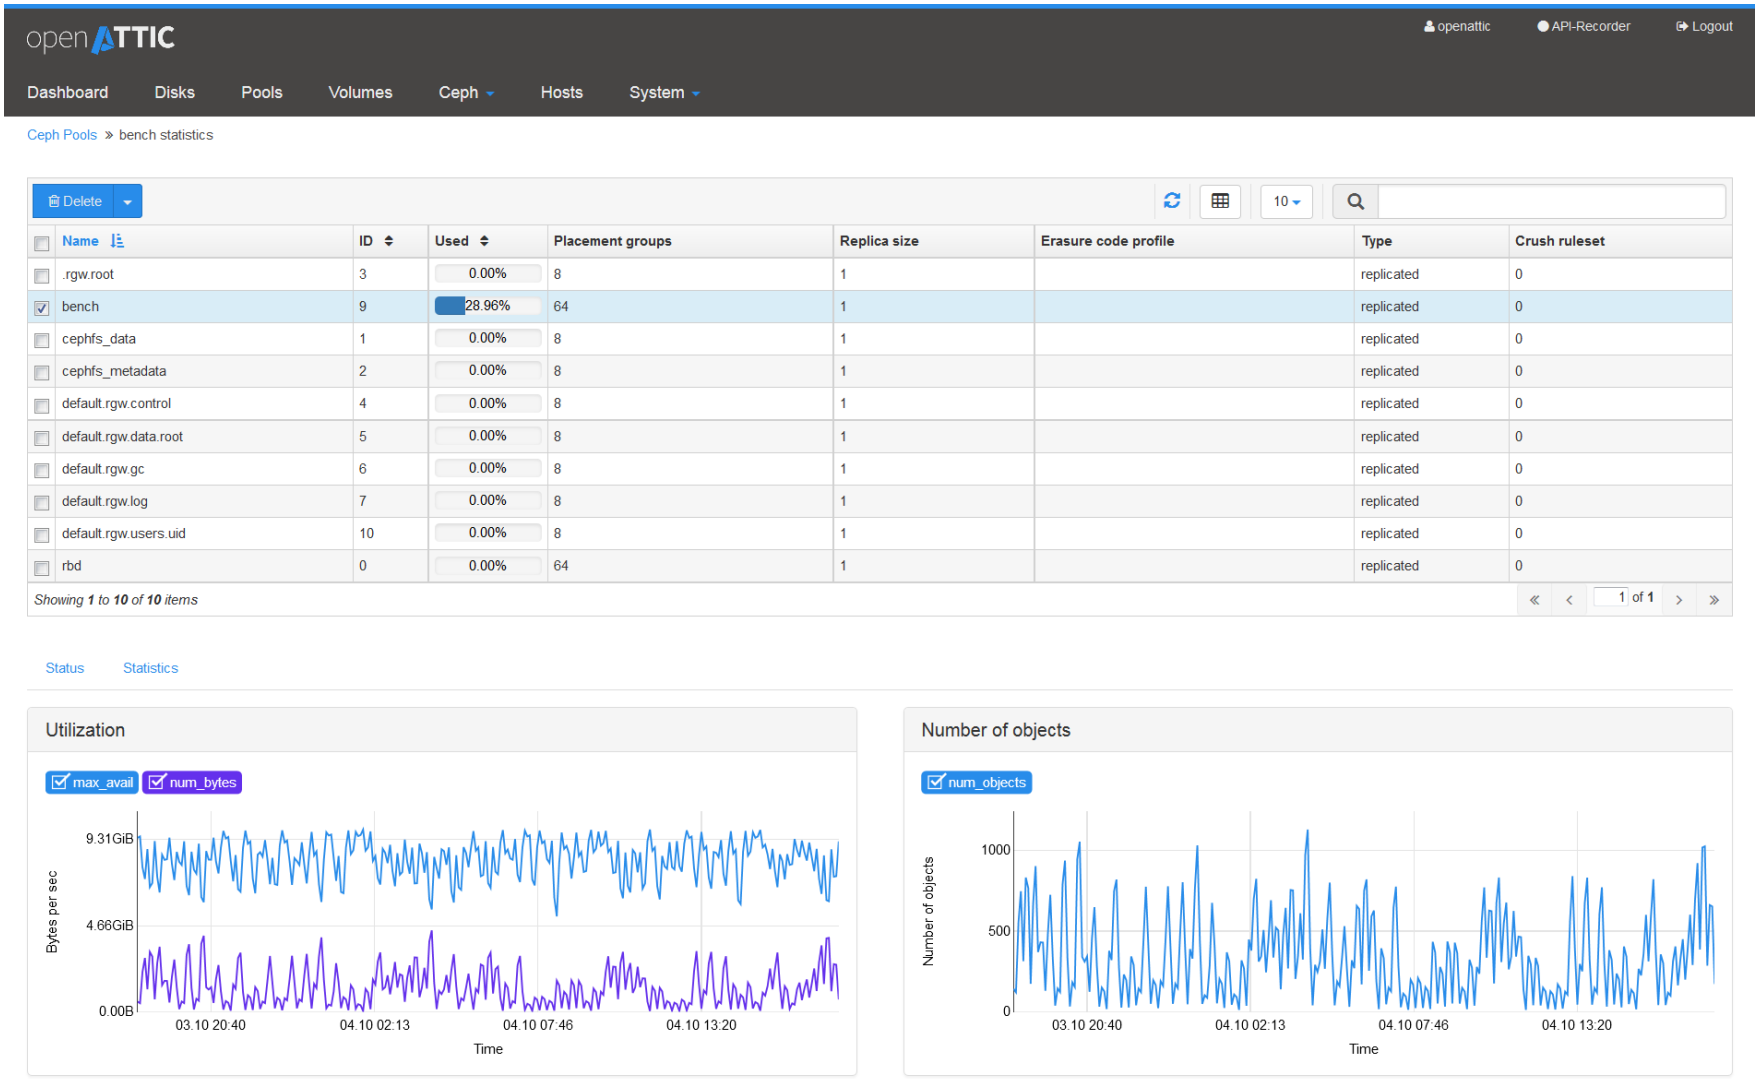

openATTIC – Ceph Pool Creation

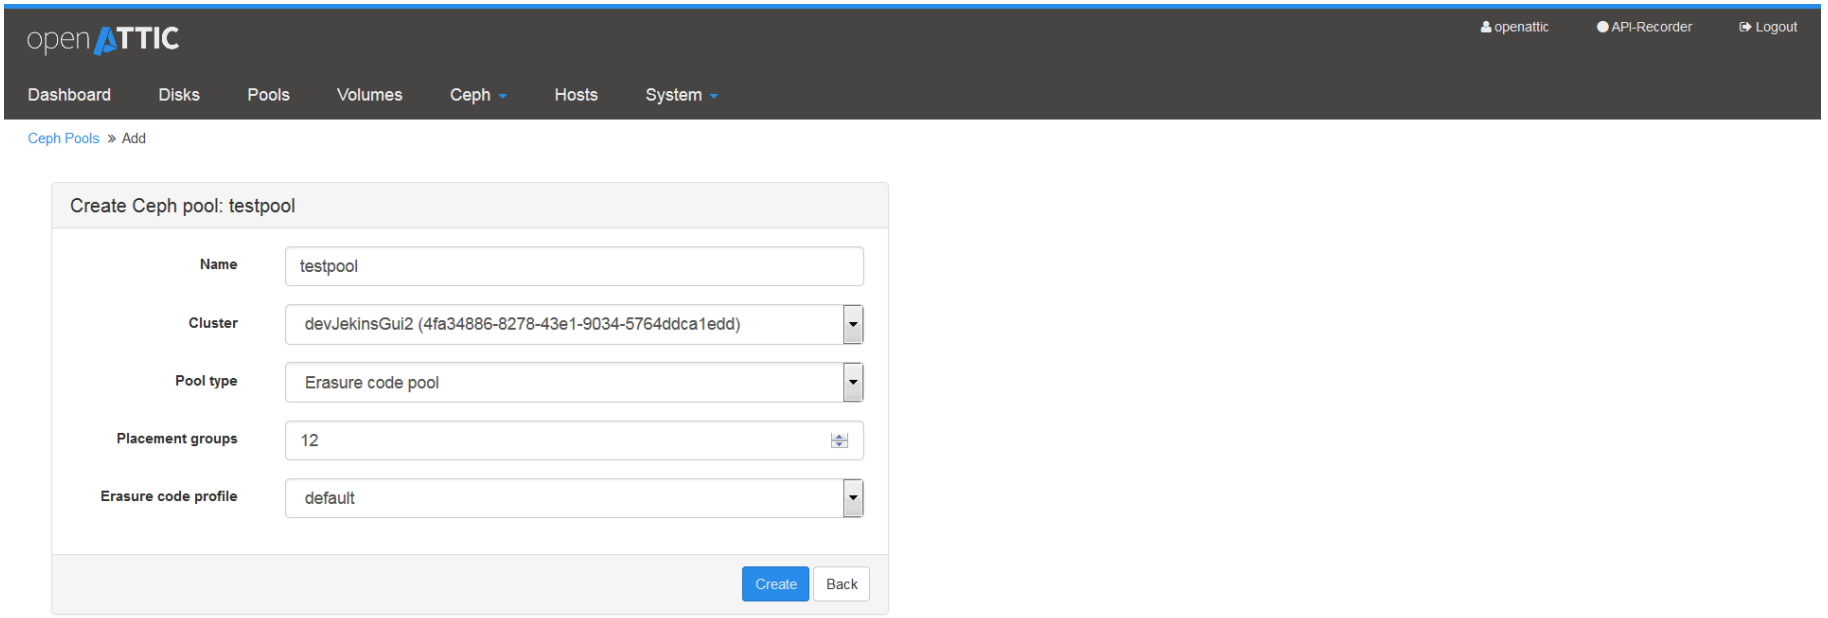

openATTIC – Ceph RBD List

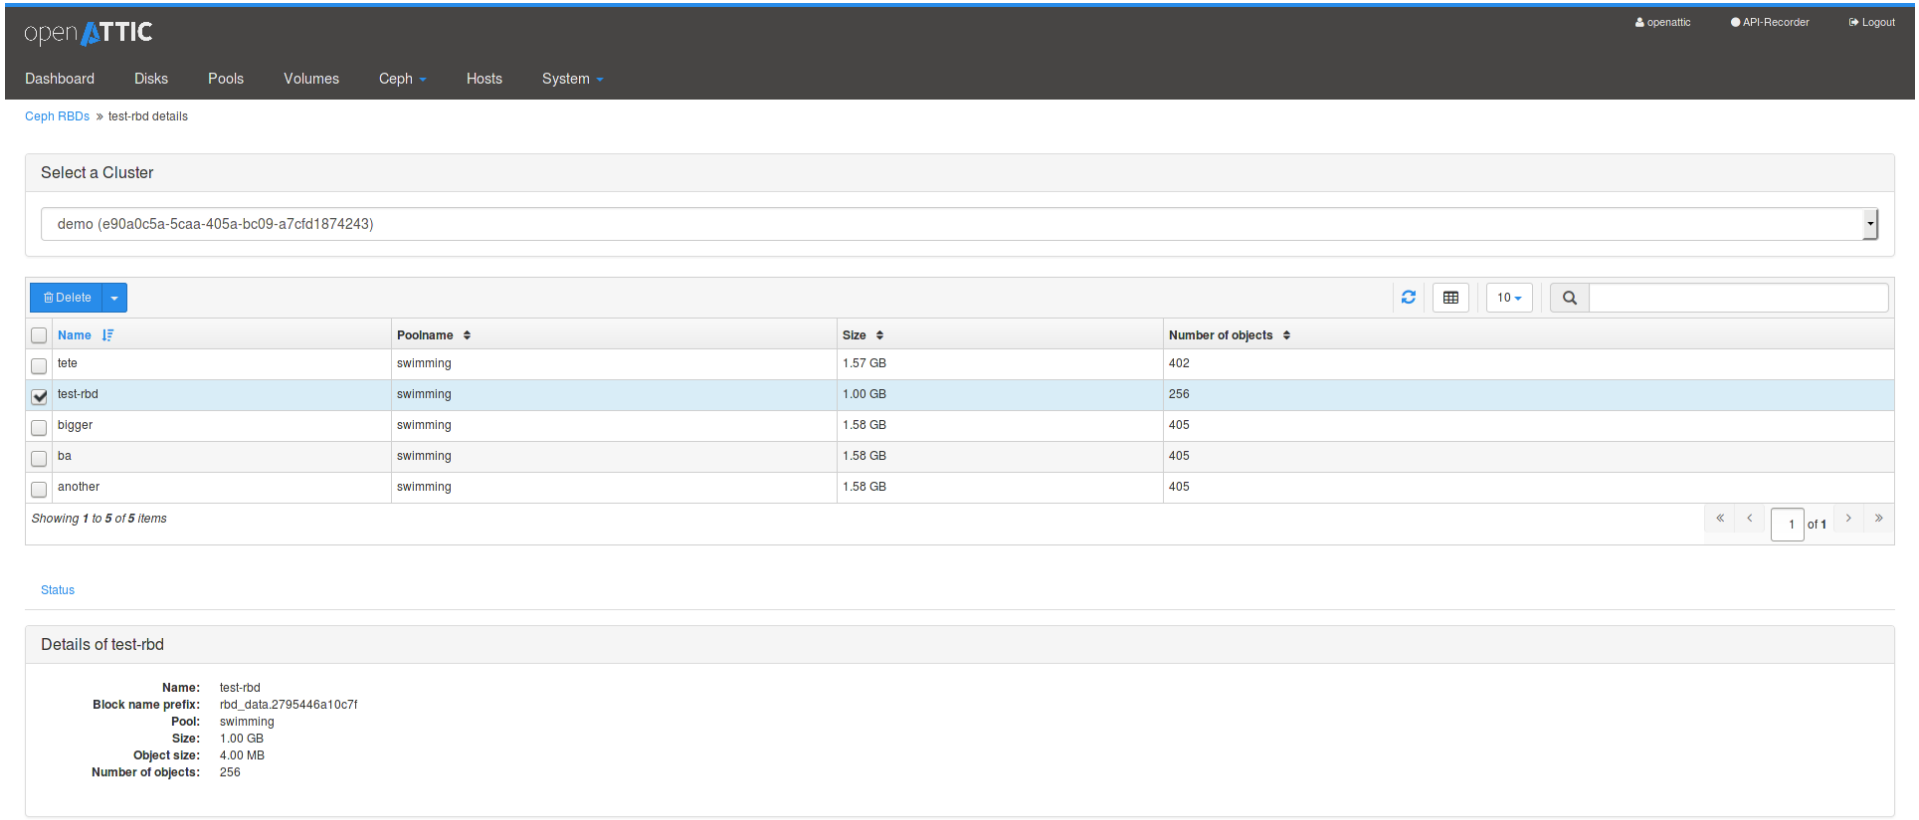

openATTIC – Ceph OSD List

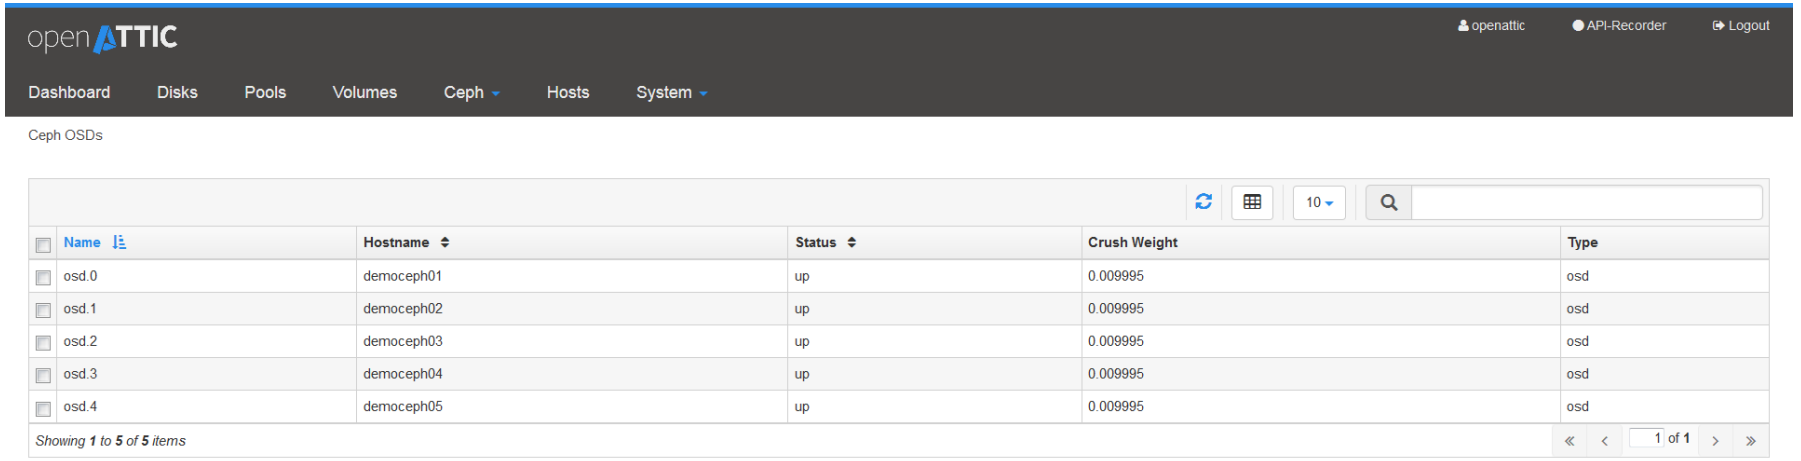

openATTIC – Ceph RBD List

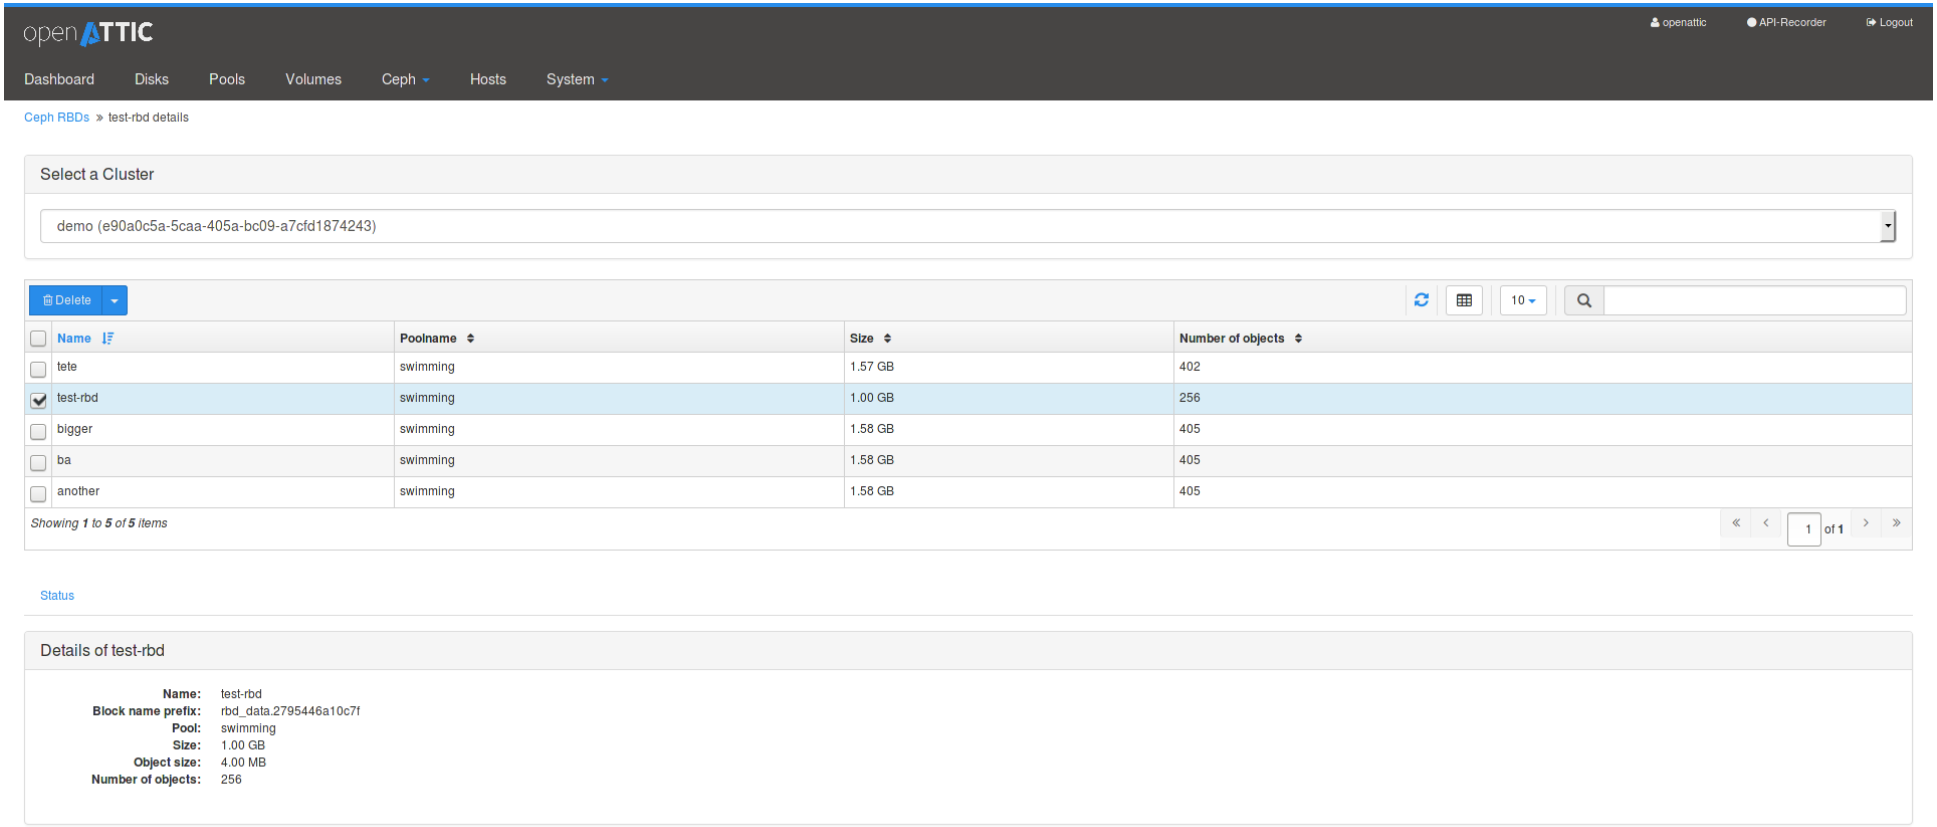

openATTIC – Ceph RBD Creation

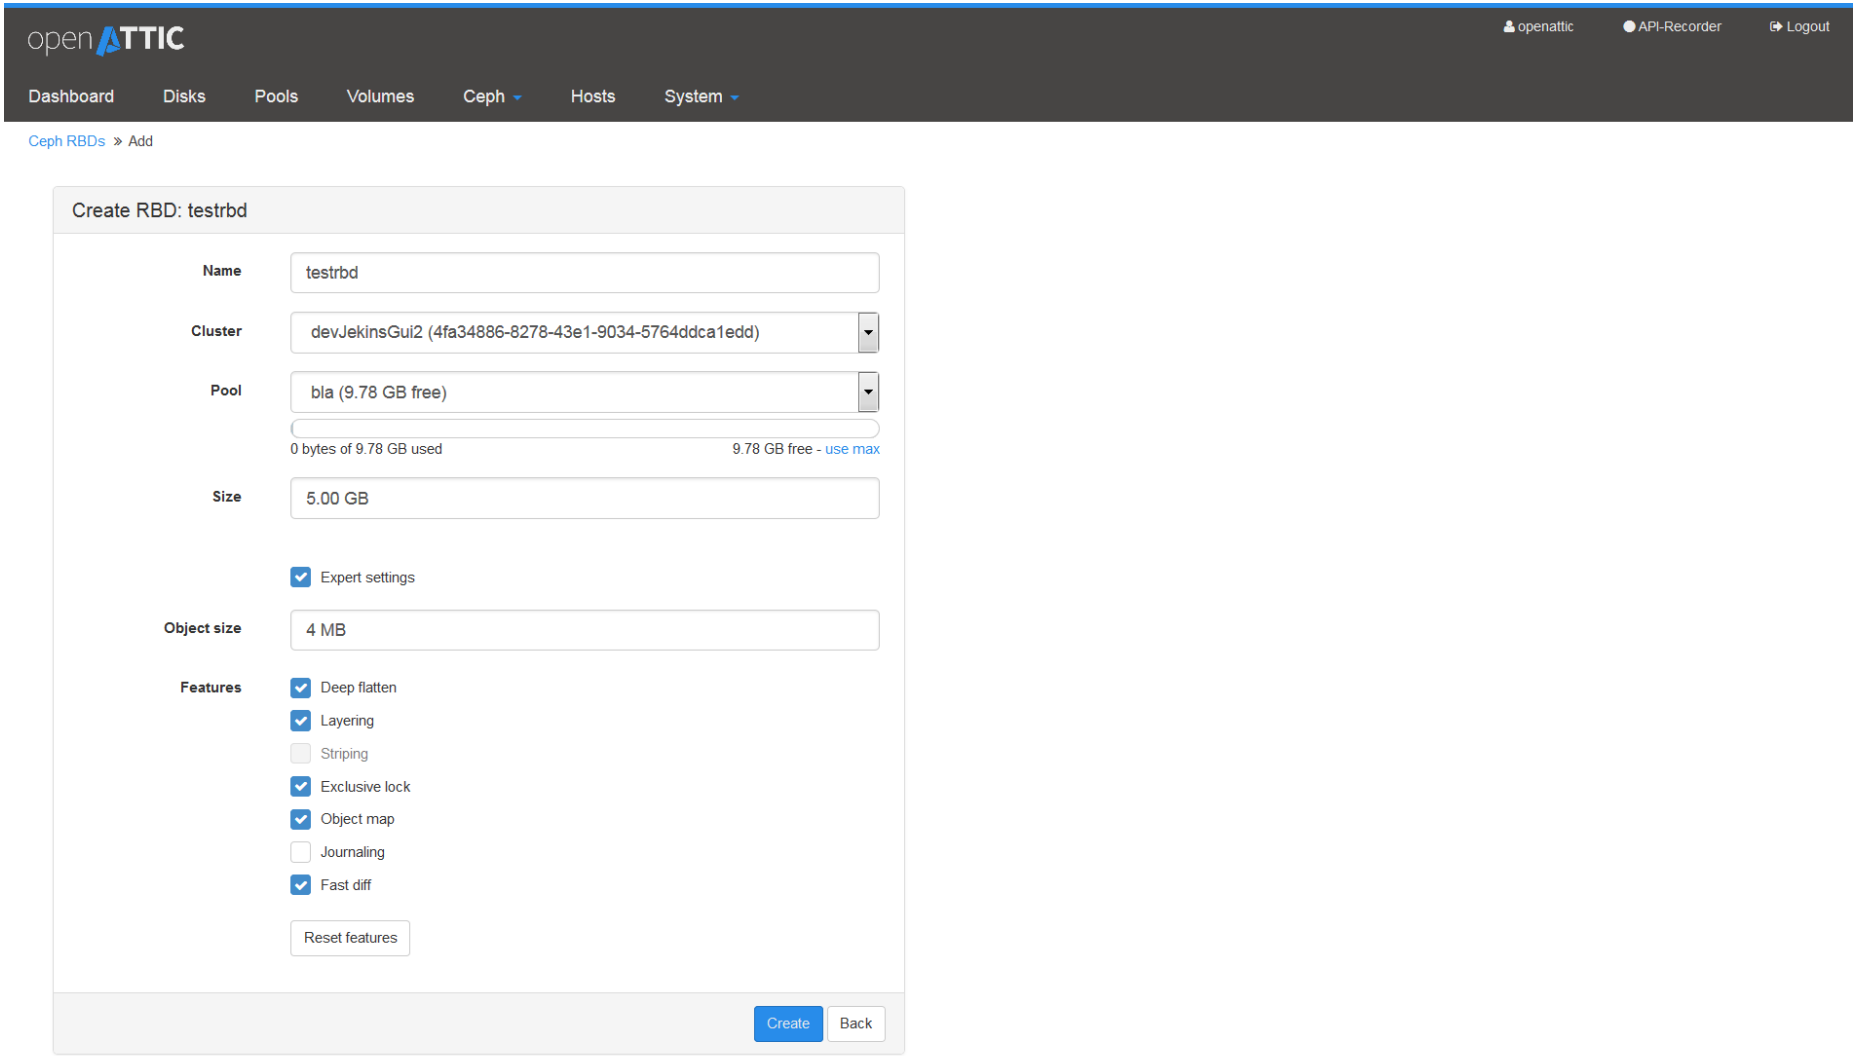

openATTIC – CRUSH Map Editor

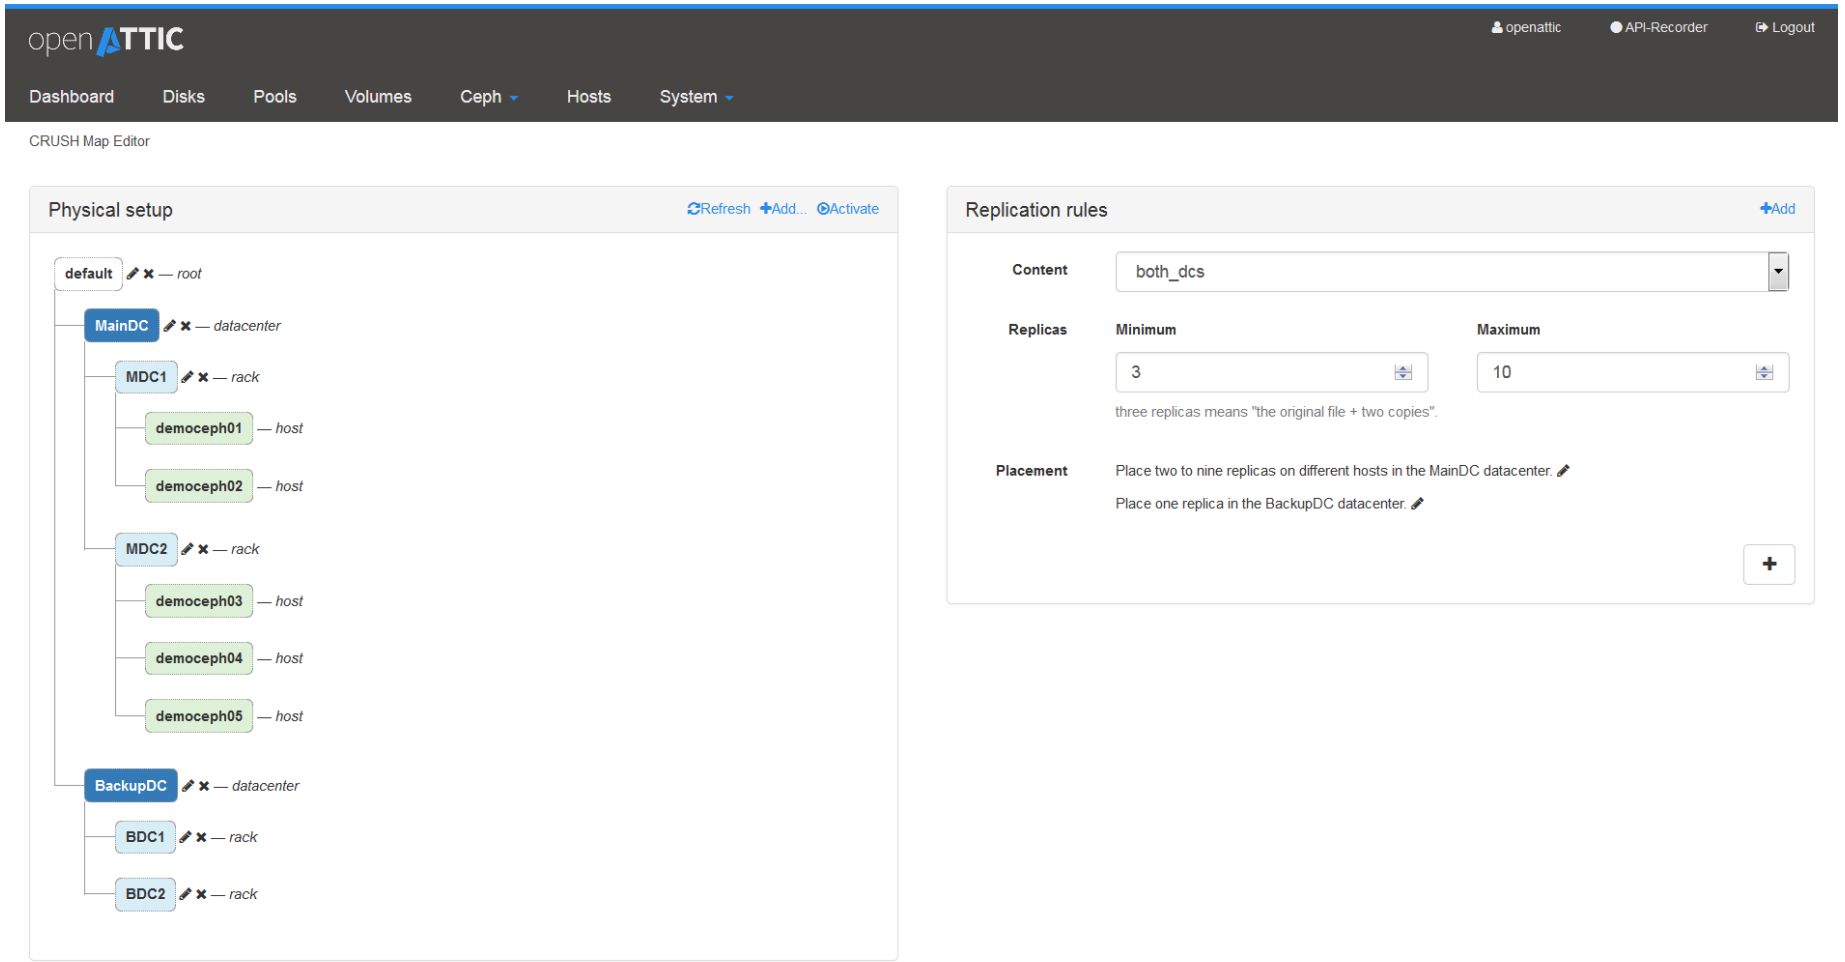

openATTIC – Resources

- www.openattic.org
- demo.openattic.org
- blog.openattic.org
- docs.openattic.org
- bitbucket.org/openattic
- tracker.openattic.org
- Twitter: @openattic
- G+: openATTIC

Questions?


SUSE®
We adapt. You succeed.