

CN

One step S3 in container with Ceph

Ceph Nano (cn) helps you interact with an Amazon S3 API endpoint by leveraging container technologies.

It provides a CLI to perform operations on S3 as well as a dashboard.

cn also runs on Kubernetes!



The Project

- Lives at <https://github.com/ceph/cn>
- Written in Go, built around Docker's API
- Runs on Docker on your laptop (Linux and OSX with Docker for Mac)
- Relies on Ceph!
- Consumes official Ceph images from the <https://github.com/ceph/ceph-container> project
- Targeted for developers or demos
- Support from Ceph Jewel to Nautilus (current dev)

****Do not use cn in production!****

```
$ cn --help
Ceph Nano - One step S3 in container with Ceph.
```

```

      *((((((((((((
      (((((((((((((((
      ((((((((*      ,((((((((*
      ((((((          ((((((
*(((,              ,(((/
(((,      (((((/    *(((
(((      ((((((((((  (((
/(((      ((((((((((  (((
(((      ((((((      /(((/
(((      ((((((      *(((
.(((      ((((((
,((((((((*          /(((
.((((      (( (/ //  (((
      /(((      /((((      /((((
      .(((/ /
```

```
Usage:
  cn [command]
```

Available Commands:

- cluster Interact with a particular Ceph cluster
- s3 Interact with a particular S3 object server
- image Interact with cn's container image(s)
- version Print the version of cn
- kube Outputs cn kubernetes template (cn kube > kube-cn.yml)
- update-check Print cn current and latest version number

Flags:

- h, --help help for cn

Use "cn [command] --help" for more information about a command.



Your first cluster in less than 30s!

```
'cn cluster start my-cluster'
```

- Run container image from <https://hub.docker.com/r/ceph/daemon/tags/> or the one you just built on your laptop
- Storage:
 - Default uses Docker volumes
 - Dedicated directory (only on Linux)
 - Dedicated block device (only on Linux) or PVC on Kubernetes!
- Dedicated working directory to do uploads from the CLI

Interact with your cluster

cn supports the following actions:

- Start an object storage server
- Stat an object storage server
- Stop an object storage server
- Restart an object storage server
- Print an object storage server logs
- Purge an object storage server. DANGEROUS!
- Enter an object storage server

```
$ cn cluster status b
```

```
HEALTH_OK is the Ceph status  
Your working directory is: /tmp  
S3 access key is: 5G360TRRJ9FNyI7W1XWE  
S3 secret key is: sRA3QIZYjb7FyHiPIYwIRd0b5uVnw9MAkz3XxJMN  
S3 object server address is: http://30.40.53.39:8000  
Ceph Nano browser address is: http://30.40.53.39:5000
```

Your first S3 bucket

The CLI exposes many S3 calls:

- Make bucket
- Remove bucket
- List objects or buckets
- List all object in all buckets
- Put file into bucket
- Get file out of a bucket
- Delete file from bucket
- Disk usage by buckets
- Get various information about Buckets or Files
- Copy object
- Move object
- Synchronize a directory tree to S3

```
$ cn s3 mb my-app my-buck
Bucket 's3://my-buck/' created on cluster my-app

$ cn s3 put my-app /etc/passwd my-buck
upload: '/tmp/passwd' -> 's3://my-buck/passwd' [1 of 1]
, 4096 of 6774 60% in 0s 1237.62 kB/s_ 6774 of 6774 100% in 1s 4.72 kB/s 6774 of 6774 100%
in 1s 4.72 kB/s done on cluster my-app
```

Work with container images

- `cn` is able to list the container images available on the Docker Hub.
- Use `-a` to get the full list of tags (really verbose)
- Ability to update to the `'latest'` container image available
- Over 5K images!

```
$ cn image ls
ceph/daemon:latest-bis
ceph/daemon:latest-bis-luminous
ceph/daemon:latest
ceph/daemon:latest-mimic
ceph/daemon:latest-luminous
ceph/daemon:latest-master
ceph/daemon:master-17afd12-mimic-centos-7
ceph/daemon:master-17afd12-luminous-centos-7
ceph/daemon:master-17afd12-mimic-centos-7-aarch64
ceph/daemon:master-17afd12-luminous-centos-7-aarch64
```

Multi-cluster support

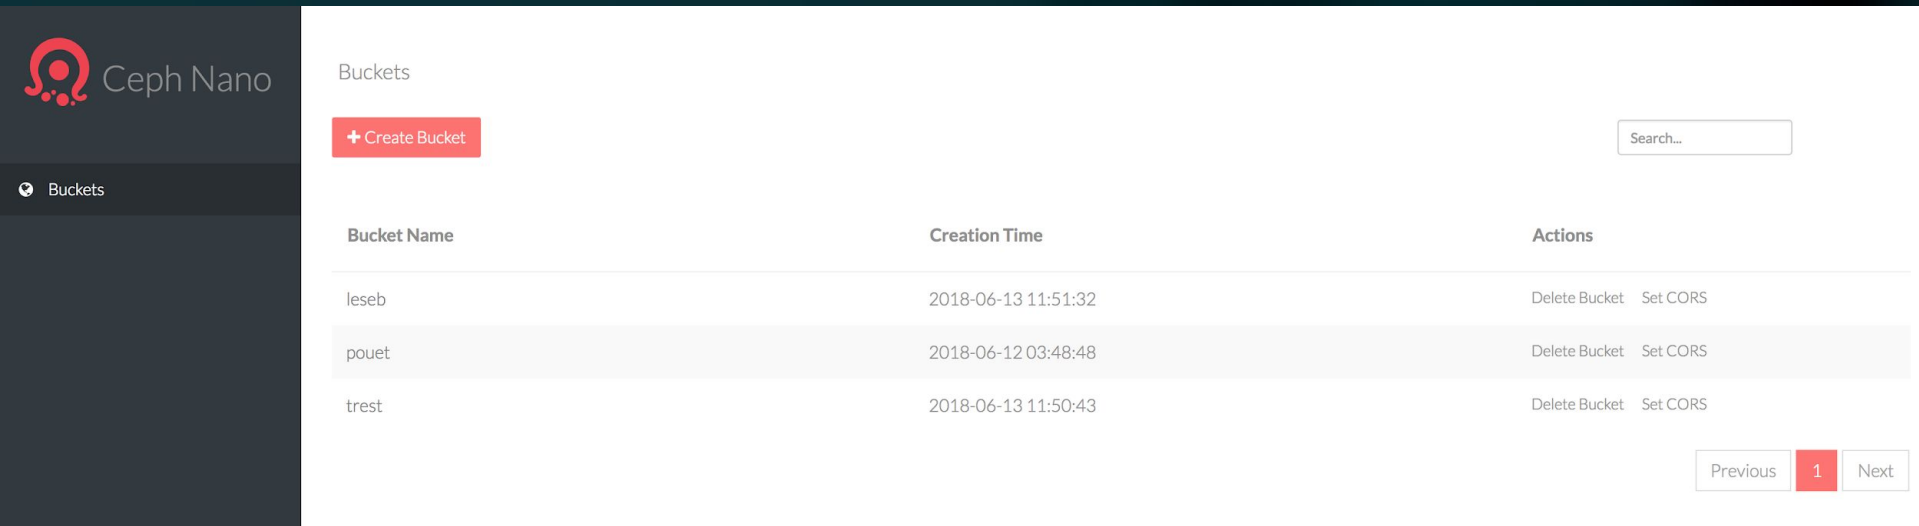
- 1 cluster = 1 container
- You can run as many clusters as you want on your laptop
- Each new cluster gets new ports allocated for both S3 endpoint and dashboard
- All the clusters are isolated from each others
- Each cluster can run a different version of Ceph

```
$ cn cluster ls
```

| NAME | STATUS | IMAGE | IMAGE RELEASE | IMAGE CREATION TIME |
|----------|---------|-----------------------------|----------------|--------------------------------|
| luminous | running | ceph/daemon:latest-luminous | master-7150c5c | 2018-09-05T13:12:42.171390511Z |
| mimic | running | ceph/daemon:latest-mimic | master-5d66b91 | 2018-11-02T00:16:00.112057779Z |
| my-app | running | ceph/daemon:latest | master-17afd12 | 2018-10-30T18:12:11.068242136Z |

Dashboard

- Create buckets and directories
- Upload and share objects URL



The screenshot displays the Ceph Nano dashboard interface. On the left, there is a dark sidebar with the Ceph Nano logo and a 'Buckets' menu item. The main content area is titled 'Buckets' and features a '+ Create Bucket' button and a search input field. Below these is a table with three columns: 'Bucket Name', 'Creation Time', and 'Actions'. The table lists three buckets: 'leseb', 'pouet', and 'trest', each with its creation time and 'Delete Bucket' and 'Set CORS' actions. At the bottom right, there are navigation buttons for 'Previous', '1', and 'Next'.

Buckets

[+ Create Bucket](#)

| Bucket Name | Creation Time | Actions |
|-------------|---------------------|------------------------|
| leseb | 2018-06-13 11:51:32 | Delete Bucket Set CORS |
| pouet | 2018-06-12 03:48:48 | Delete Bucket Set CORS |
| trest | 2018-06-13 11:50:43 | Delete Bucket Set CORS |

[Previous](#) [1](#) [Next](#)

Runs on Kubernetes

- Currently tested on 1.9
- 'cn kube' will output a Kubernetes template
- Edit the template and add PVC
- Dashboard not available

```
$ kubectl create -f tests/kube-cn.yml
service "ceph-nano-services" created
statefulset.apps "ceph-nano" created

$ kubectl get pods -l app=ceph
NAME          READY   STATUS    RESTARTS   AGE
ceph-nano-0   1/1     Running   0           27s

$ kubectl expose pod ceph-nano-0 --type=NodePort
service "ceph-nano-0" exposed

$ curl $(minikube service ceph-nano-0 --url)
<?xml version="1.0" encoding="UTF-8"?>
<ListAllMyBucketsResult xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
  <Owner><ID>anonymous</ID><DisplayName></DisplayName>
</Owner><Buckets></Buckets></ListAllMyBucketsResult>
```

Enter the void

- Curious? Want deeper understanding of Ceph?

```
$ cn cluster enter mimic

[root@ceph-nano-mimic-faa32aebf00b /]#
[root@ceph-nano-mimic-faa32aebf00b /]# ceph -s
cluster:
  id:          9d1cc61a-59ac-4f94-bdff-6202f1830e97
  health: HEALTH_OK

services:
  mon: 1 daemons, quorum ceph-nano-mimic-faa32aebf00b
  mgr: ceph-nano-mimic-faa32aebf00b(active)
  osd: 1 osds: 1 up, 1 in

data:
  pools:   5 pools, 40 pgs
  objects: 224 objects, 2.6 KiB
  usage:   1.0 GiB used, 9.0 GiB / 10 GiB avail
  pgs:    40 active+clean
```

DEMO

Roadmap

- Add cn as a backend for Rook (<https://rook.io/>)
- Add a config file to control different behaviour (new version available, different registry...)
- Enhancements coming in the ceph-container project so the Ceph version will be in the image tag
- Core Ceph developer mode
- A nicer website, any volunteer? :)
- Speed the bootstrap (who said under 20s?)

Wrap up

- cn is handy to interact with the S3 API from:
 - A CLI
 - A dashboard
 - Your own code
- You can run it from your laptop or on Kubernetes
- Useful to introspect any Ceph version's behaviour (easier than reading Ceph's code)
 - E.g: was this command available in 12.2.1 or 12.2.2?
 - E.g: how did the output of `ceph -s` look like in Jewel?
 - E.g: what was the default value of option X in version Y?

Merci !

seb@redhat.com