

# CEPH performance monitoring

## basic info

### ceph

```
ceph -s
ceph -w
ceph df
ceph osd tree
ceph osd df tree
```

### rados

```
rados df
```

Where:

- **USED COMPR:** amount of space allocated for compressed data (i.e. this includes compressed data plus all the allocation, replication and erasure coding overhead).
- **UNDER COMPR:** amount of data passed through compression (summed over all replicas) and beneficial enough to be stored in a compressed form.

### RBD Rados Block Device

```
rbd ls
rbd du
```

### perf

```
ceph iostat

ceph osd perf

rbd perf image iostat # and wait 30 sec
rbd perf image iotop # and wait 30 sec

rbd du
```

<https://ceph.io/community/new-mimic-iostat-plugin/>

## rados benchmark

Hints from [https://yourcmc.ru/wiki/Ceph\\_performance#Test\\_your\\_Ceph\\_cluster](https://yourcmc.ru/wiki/Ceph_performance#Test_your_Ceph_cluster):

- Don't use `rados bench`. It creates a small number of objects (1-2 for a thread) so all of them always reside in cache and improve the results far beyond they should be.
- You can use `rbd bench`, but fio is better.

```
ceph osd pool create test
rados -p test bench 10 write --no-cleanup
rados -p test bench 10 seq
rados -p test bench -t 4 10 seq
rados -p test bench 10 rand

rados bench -p test 60 write -b 4M -t 16 --no-cleanup
rados bench -p test 60 seq -t 16
rados bench -p test 60 rand -t 16

rados -p test cleanup
```

## rbd benchmark

### Chapter 9. Benchmarking Performance

```
# Image 1G:
rbd create test/myimage --size 1024

#TBD ...
#rbd device map test/myimage --cluster
#rbd device unmap /dev/rbdX

rbd rm test/myimage
```

From:

<https://niziak.spoX.org/wiki/> - niziak.spoX.org

Permanent link:

[https://niziak.spoX.org/wiki/vm:proxmox:ceph:performance\\_monitoring](https://niziak.spoX.org/wiki/vm:proxmox:ceph:performance_monitoring)

Last update: **2021/04/08 11:44**

