

# Guest performance

## Tips

The best results for Windows guest:

- HDD (no SSD)
  - probably CEPH better handles bunch of ordered requests (like for rotational drives)
- no DISCARD
  - very poor performance with DISCARD

## todo

Ext4 is extremely slow on RBD block device.

SSD emulation enabled, discard support enabled, SCSI single controller, IO thread.

Guest Kernel 6.0.

BTRFS detects SSD and use: single data, duplicated system and metadata.

Simple and fast test show huge difference between ext4 and btrfs.

```
fio --name=fiotest --directory=./2ext4 --size=256m --rw=randrw --bs=4K --
direct=1 --ioengine=libaio --iodepth=64 --numjobs=4 --runtime=10s --
time_based=1
```

Ext4 with unsafe options (barrier=0, data=writeback)

```
READ: bw=981KiB/s (1005kB/s), 234KiB/s-254KiB/s (240kB/s-260kB/s),
io=10.5MiB (11.0MB), run=10981-10991msec
WRITE: bw=1030KiB/s (1055kB/s), 251KiB/s-266KiB/s (257kB/s-273kB/s),
io=11.1MiB (11.6MB), run=10981-10991msec
```

Ext4 with default options:

```
READ: bw=997KiB/s (1021kB/s), 242KiB/s-267KiB/s (248kB/s-274kB/s),
io=10.6MiB (11.1MB), run=10583-10840msec
WRITE: bw=1048KiB/s (1073kB/s), 259KiB/s-271KiB/s (266kB/s-277kB/s),
io=11.1MiB (11.6MB), run=10583-10840msec
```

BTRFS with default options:

```
READ: bw=3281KiB/s (3359kB/s), 806KiB/s-855KiB/s (826kB/s-875kB/s),
io=33.6MiB (35.3MB), run=10496-10497msec
WRITE: bw=3431KiB/s (3513kB/s), 850KiB/s-873KiB/s (870kB/s-894kB/s),
```

```
io=35.2MiB (36.9MB), run=10496-10497msec
```

BTRFS with nodatacow:

```
READ: bw=2444KiB/s (2502kB/s), 604KiB/s-619KiB/s (618kB/s-634kB/s),  
io=25.1MiB (26.3MB), run=10469-10501msec  
WRITE: bw=2566KiB/s (2628kB/s), 629KiB/s-652KiB/s (644kB/s-668kB/s),  
io=26.3MiB (27.6MB), run=10469-10501msec
```

```
READ: bw=3109KiB/s (3184kB/s), 770KiB/s-796KiB/s (789kB/s-815kB/s),  
io=30.0MiB (32.5MB), run=10193-10195msec  
WRITE: bw=3253KiB/s (3331kB/s), 786KiB/s-847KiB/s (805kB/s-867kB/s),  
io=32.4MiB (33.0MB), run=10193-10195msec
```

BTRFS converted to SINGLE with command

```
btrfs balance start --force -sconvert=single -mconvert=single ./1btrfs/
```

```
READ: bw=3126KiB/s (3201kB/s), 755KiB/s-815KiB/s (773kB/s-835kB/s),  
io=31.4MiB (32.9MB), run=10283-10286msec  
WRITE: bw=3267KiB/s (3345kB/s), 793KiB/s-864KiB/s (812kB/s-884kB/s),  
io=32.8MiB (34.4MB), run=10283-10286msec
```

```
READ: bw=3861KiB/s (3953kB/s), 941KiB/s-984KiB/s (964kB/s-1008kB/s),  
io=39.2MiB (41.1MB), run=10341-10400msec  
WRITE: bw=3997KiB/s (4093kB/s), 971KiB/s-1034KiB/s (995kB/s-1059kB/s),  
io=40.6MiB (42.6MB), run=10341-10400msec
```

+ nodatacow:

```
READ: bw=3911KiB/s (4004kB/s), 959KiB/s-1004KiB/s (982kB/s-1028kB/s),  
io=39.3MiB (41.2MB), run=10290-10295msec  
WRITE: bw=4049KiB/s (4147kB/s), 993KiB/s-1039KiB/s (1016kB/s-1064kB/s),  
io=40.7MiB (42.7MB), run=10290-10295msec
```

```
READ: bw=3879KiB/s (3972kB/s), 941KiB/s-1001KiB/s (963kB/s-1025kB/s),  
io=38.3MiB (40.2MB), run=10110-10114msec  
WRITE: bw=4028KiB/s (4125kB/s), 976KiB/s-1041KiB/s (999kB/s-1066kB/s),  
io=39.8MiB (41.7MB), run=10110-10114msec
```

From:  
<https://niziak.spox.org/wiki/> - **niziak.spox.org**



Permanent link:  
[https://niziak.spox.org/wiki/vm:proxmox:ceph:guest\\_performance](https://niziak.spox.org/wiki/vm:proxmox:ceph:guest_performance)

Last update: **2023/06/01 14:24**