

# Proxmox VE

- [High Availability](#)
- [ZFS](#)
- [Useful additions](#)
- [Nested Virtualization Benchmarks](#)
- [Troubleshooting](#)
- [CEPH](#)

# High Availability

[Proxmox Documentation: High Availability](#)

Below are relevant excerpts from the above source materials for quick access to necessary commands for various tasks. Please consult the source material for more documentation related to the concepts and commands referenced there.

## Error Recovery

```
ha-manager set vm:100 --state disabled
```

## Node Maintenance

### Maintenance Mode

```
ha-manager status
```

```
ha-manager crm-command node-maintenance enable [NODENAME]
```

```
ha-manager crm-command node-maintenance disable [NODENAME]
```

# ZFS

## Upgrade single disk rpool to mirrored rpool

I added a second M.2 drive to a laptop already running Proxmox VE. From the `zpool-attach` man page:

“ If the existing device is a mirror or plain device (e.g. specified as "sda" or "mirror-7"), the new device will be mirrored with the existing device, a resilver will be initiated, and the new device will contribute to additional redundancy once the resilver completes.

Steps:

1. Use `gdisk` or `sgdisk` to copy the partition table from the existing drive to the new drive and change the GUID.
2. Make the second disk bootable
3. Attach the new disk to the existing disk using `zpool attach`.

```
# backup partition table from the existing disk nvme0n1
gdisk /dev/nvme0n1
b
gdisk.part.nvme0n1
q

# restore the backed up partition table to the new disk nvme1n1
gdisk /dev/nvme1n1
r
l
gdisk.part.nvme0n1
x
g
```

```
R
```

```
w
```

```
# same result as using gdisk above
# I have not actually used this method yet
sgdisk --replicate=/dev/nvme1n1 /dev/nvme0n1
sgdisk --randomize-guids /dev/nvme1n1
```

```
pve-efiboot-tool status
```

```
# Setup the UEFI bootloader
pve-efiboot-tool format /dev/nvme1n1p2 --force
pve-efiboot-tool init /dev/nvme1n1p2
```

```
pve-efiboot-tool status
```

```
#
# BEFORE
#
root@pve1:~# zpool status
  pool: rpool
  state: ONLINE
status: Some supported and requested features are not enabled on the pool.
       The pool can still be used, but some features are unavailable.
action: Enable all features using 'zpool upgrade'. Once this is done,
       the pool may no longer be accessible by software that does not support
       the features. See zpool-features(7) for details.
config:
```

	NAME	STATE	READ	WRITE
CKSUM				
0	rpool	ONLINE	0	
0	nvme-BC711_NVMe_SK_hynix_256GB____FYB1N066013901U4N-part3	ONLINE	0	
0				

```
zpool attach rpool nvme-BC711_NVMe_SK_hynix_256GB____FYB1N066013901U4N-part3 nvme-
KBG40ZNS256G_NVMe_KIOXIA_256GB_Y96PC1TDPTLL-part3
```

```

#
# AFTER
#
root@pvel:~# zpool status
  pool: rpool
  state: ONLINE
status: Some supported and requested features are not enabled on the pool.
       The pool can still be used, but some features are unavailable.
action: Enable all features using 'zpool upgrade'. Once this is done,
       the pool may no longer be accessible by software that does not support
       the features. See zpool-features(7) for details.
scan: resilvered 30.6G in 00:03:12 with 0 errors on Thu Aug  7 16:40:11 2025
config:

```

CKSUM	NAME	STATE	READ	WRITE
0	rpool	ONLINE	0	
0	mirror-0	ONLINE	0	
0	nvme-BC711_NVMe_SK_hynix_256GB____FYB1N066013901U4N-part3	ONLINE	0	
0	nvme-KBG40ZNS256G_NVMe_KIOXIA_256GB_Y96PC1TDPTLL-part3	ONLINE	0	

## Add SLOG

```

# identify the disk(s) you want to use
lsblk -o NAME,MODEL,SERIAL,SIZE,TYPE,MOUNTPOINT

# find the disk / partitions by id
find /dev/disk/by-id/ | grep -e VK0240GFDKF -e MTFDDAK960TCB

# add the log mirror
zpool add z-10k-SAS-DS-Mirror log mirror /dev/disk/by-id/ata-VK0240GFDKF_S269NX0H703269-part1
/dev/disk/by-id/ata-MTFDDAK960TCB_18181C3C1F61-part1

```

```
# verify
zpool status
```

```
# check usage
zpool iostat -v

# watch usage
zpool iostat -v 1

# or
watch zpool iostat -v
```

## Add cache (L2ARC)

```
zpool add z-10k-SAS-DS-Mirror cache /dev/disk/by-id/ata-MTFDDAK960TCB_18181C3C1F61-part2
```

## HPE Proliant

# Useful additions

## Packages to add

```
apt install lldpd snmpd
```

## snmpd.conf

Minimal configuration for `/etc/snmp/snmpd.conf` to enable monitoring physical host resources via snmp.

```
# file:/etc/snmp/snmpd.conf
sysLocation    Undisclosed location
sysContact     Ghostbusters! <ghost@busters.loc>
sysServices    72

# master agentx
agentaddress   udp:161

rocommunity   localmon 10.0.0.0/8
rocommunity   localmon 172.16.0.0/12
rocommunity   localmon 192.168.0.0/16
#rocommunity6 localmon fe80::/10

includeDir    /etc/snmp/snmpd.conf.d
```

```
cp /etc/snmp/snmpd.conf /etc/snmp/snmpd.backup.conf
```

```
cat << EOF > /etc/snmp/snmpd.conf
# file:/etc/snmp/snmpd.conf
sysLocation    Undisclosed location
sysContact     Ghostbusters! <ghost@busters.loc>
sysServices    72
```

```
# master agentx
agentaddress udp:161

rocommunity localmon 10.0.0.0/8
rocommunity localmon 172.16.0.0/12
rocommunity localmon 192.168.0.0/16
#rocommunity6 localmon fe80::/10

includeDir /etc/snmp/snmpd.conf.d
EOF
```

-end

# Nested Virtualization Benchmarks

## Disk benchmark tests

### fio simple write test

```
fio --name=sequential_write --rw=write --bs=1M --size=10G --filename=/root/testfile --direct=1
```

### fio random read job file

```
[global]
ioengine=libaio
direct=1
group_reporting
runtime=60
time_based

[random_read_job]
rw=randread
bs=4k
iodepth=64
filename=/root/testfile
size=10G
```

### fio sequential read job file

```
[global]
ioengine=libaio
direct=1
group_reporting
runtime=60
time_based

[random_read_job]
rw=read
```

```
bs=4k
iodepth=64
filename=/root/testfile
size=10G
```

BASH code to generate lots of different commands

```
for fio_sync in 0 1; do
  for fio_bs in 4KB 16KB 1MB; do
    for fio_rw in read write randwrite randread; do
      echo "fio --name=${fio_rw}_${fio_bs}_${fio_sync} --size=16g --bs=${fio_bs} --
rw=${fio_rw} --ioengine=libaio --sync=${fio_sync} --iodepth=32 --numjobs=1 --direct=1 --
end_fsync=1 --gtod_reduce=1 --time_based --runtime=60"
    done
  done
done
```

## Hardware: DELL Latitude 5520 (Laptop)

CPU: 11th Gen Intel(R) Core(TM) i5-1145G7 @ 2.60GHz

Disk: 2 x 256GB

Filesystem: ZFS Raid-1 Mirror

## PVE benchmark

```
# Disk information
root@pve:~# smartctl -a /dev/nvme0n1
Model Number:           Micron 2450 NVMe 256GB
Temperature:           43 Celsius
Percentage Used:        15%

root@pve:~# smartctl -a /dev/nvme1n1
Model Number:           PC SN520 NVMe WDC 256GB
Temperature:           45 Celsius
Percentage Used:        28%
```

```
# CPU Information
```

```
root@pve:~# cat /proc/cpuinfo | grep "model name"
```

```
model name      : 11th Gen Intel(R) Core(TM) i5-1145G7 @ 2.60GHz
model name      : 11th Gen Intel(R) Core(TM) i5-1145G7 @ 2.60GHz
model name      : 11th Gen Intel(R) Core(TM) i5-1145G7 @ 2.60GHz
model name      : 11th Gen Intel(R) Core(TM) i5-1145G7 @ 2.60GHz
model name      : 11th Gen Intel(R) Core(TM) i5-1145G7 @ 2.60GHz
model name      : 11th Gen Intel(R) Core(TM) i5-1145G7 @ 2.60GHz
model name      : 11th Gen Intel(R) Core(TM) i5-1145G7 @ 2.60GHz
model name      : 11th Gen Intel(R) Core(TM) i5-1145G7 @ 2.60GHz
```

```
# Benchmark - write
```

```
root@pve:~# fio --name=sequential_write --rw=write --bs=1M --size=10G --filename=testfile --
direct=1
```

```
sequential_write: (g=0): rw=write, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-1024KiB, (T) 1024KiB-
1024KiB, ioengine=psync, iodepth=1
```

```
fio-3.39
```

```
Starting 1 process
```

```
sequential_write: Laying out IO file (1 file / 10240MiB)
```

```
Jobs: 1 (f=1): [W(1)][100.0%][w=60.1MiB/s][w=60 IOPS][eta 00m:00s]
```

```
sequential_write: (groupid=0, jobs=1): err= 0: pid=1838716: Mon Aug 11 13:47:12 2025
```

```
write: IOPS=64, BW=64.1MiB/s (67.2MB/s)(10.0GiB/159737msec); 0 zone resets
```

```
clat (usec): min=1521, max=2857.4k, avg=15554.55, stdev=30350.48
```

```
lat (usec): min=1564, max=2857.4k, avg=15595.47, stdev=30349.98
```

```
clat percentiles (msec):
```

```
| 1.00th=[ 3], 5.00th=[ 3], 10.00th=[ 3], 20.00th=[ 15],
| 30.00th=[ 15], 40.00th=[ 16], 50.00th=[ 17], 60.00th=[ 17],
| 70.00th=[ 18], 80.00th=[ 18], 90.00th=[ 20], 95.00th=[ 22],
| 99.00th=[ 32], 99.50th=[ 37], 99.90th=[ 110], 99.95th=[ 226],
| 99.99th=[ 760]
```

```
bw ( KiB/s): min= 4096, max=448512, per=100.00%, avg=66923.14, stdev=44636.99, samples=313
```

```
iops      : min= 4, max= 438, avg=65.35, stdev=43.59, samples=313
```

```
lat (msec) : 2=0.50%, 4=12.31%, 10=1.23%, 20=77.88%, 50=7.84%
```

```
lat (msec) : 100=0.13%, 250=0.07%, 500=0.02%, 1000=0.01%, >=2000=0.01%
```

```
cpu       : usr=0.34%, sys=13.49%, ctx=11028, majf=5, minf=10
```

```
IO depths : 1=100.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%
```

```
submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
```

```
complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
```

```
issued rwts: total=0,10240,0,0 short=0,0,0,0 dropped=0,0,0,0
latency   : target=0, window=0, percentile=100.00%, depth=1
```

Run status group 0 (all jobs):

```
WRITE: bw=64.1MiB/s (67.2MB/s), 64.1MiB/s-64.1MiB/s (67.2MB/s-67.2MB/s), io=10.0GiB
(10.7GB), run=159737-159737msec
```

# Benchmark - read

```
root@pve:~# fio random_read_test.fio
```

```
random_read_job: (g=0): rw=randread, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B,
ioengine=libaio, iodepth=64
```

```
fio-3.39
```

Starting 1 process

```
Jobs: 1 (f=1): [r(1)][100.0%][r=10.3MiB/s][r=2626 IOPS][eta 00m:00s]
```

```
random_read_job: (groupid=0, jobs=1): err= 0: pid=1841119: Mon Aug 11 13:54:44 2025
```

```
read: IOPS=2638, BW=10.3MiB/s (10.8MB/s)(618MiB/60001msec)
```

```
slat (usec): min=167, max=6793, avg=377.03, stdev=155.55
```

```
clat (nsec): min=1728, max=36276k, avg=23873354.46, stdev=1833473.80
```

```
lat (usec): min=458, max=36588, avg=24250.38, stdev=1850.88
```

```
clat percentiles (usec):
```

```
| 1.00th=[20579], 5.00th=[21627], 10.00th=[21890], 20.00th=[22414],
| 30.00th=[22938], 40.00th=[23200], 50.00th=[23725], 60.00th=[23987],
| 70.00th=[24511], 80.00th=[25035], 90.00th=[26084], 95.00th=[27132],
| 99.00th=[29754], 99.50th=[31065], 99.90th=[33162], 99.95th=[33817],
| 99.99th=[34866]
```

```
bw ( KiB/s): min= 7944, max=11024, per=100.00%, avg=10553.40, stdev=363.11, samples=120
```

```
iops      : min= 1986, max= 2756, avg=2638.35, stdev=90.78, samples=120
```

```
lat (usec) : 2=0.01%, 500=0.01%, 750=0.01%
```

```
lat (msec) : 2=0.01%, 4=0.01%, 10=0.01%, 20=0.65%, 50=99.33%
```

```
cpu      : usr=0.64%, sys=6.66%, ctx=158463, majf=0, minf=72
```

```
IO depths : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.1%, 32=0.1%, >=64=100.0%
```

```
submit   : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
```

```
complete : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.1%, >=64=0.0%
```

```
issued rwts: total=158301,0,0,0 short=0,0,0,0 dropped=0,0,0,0
```

```
latency   : target=0, window=0, percentile=100.00%, depth=64
```

Run status group 0 (all jobs):

```
READ: bw=10.3MiB/s (10.8MB/s), 10.3MiB/s-10.3MiB/s (10.8MB/s-10.8MB/s), io=618MiB (648MB),
```

# Nested VM Benchmark

Ubuntu 24.04

```

root@nested-vm-l2:~$ cat /proc/cpuinfo | grep "model name"
model name      : 11th Gen Intel(R) Core(TM) i5-1145G7 @ 2.60GHz
model name      : 11th Gen Intel(R) Core(TM) i5-1145G7 @ 2.60GHz

root@nested-vm-l2:~$ fio --name=sequential_write --rw=write --bs=1M --size=10G --
filename=testfile --direct=1
sequential_write: (g=0): rw=write, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-1024KiB, (T) 1024KiB-
1024KiB, ioengine=psync, iodepth=1
fio-3.36
Starting 1 process
sequential_write: Laying out IO file (1 file / 10240MiB)
Jobs: 1 (f=1): [W(1)][97.1%][w=178MiB/s][w=177 IOPS][eta 00m:08s]
sequential_write: (groupid=0, jobs=1): err= 0: pid=1820: Mon Aug 11 18:43:38 2025
  write: IOPS=37, BW=37.9MiB/s (39.8MB/s)(10.0GiB/269886msec); 0 zone resets
    clat (nsec): min=0, max=44001M, avg=26294437.47, stdev=800866646.25
      lat (nsec): min=0, max=44001M, avg=26331483.73, stdev=800866317.34
    clat percentiles (usec):
      | 1.00th=[ 1270], 5.00th=[ 1385], 10.00th=[ 1434],
      | 20.00th=[ 1516], 30.00th=[ 1582], 40.00th=[ 1647],
      | 50.00th=[ 1745], 60.00th=[ 1860], 70.00th=[ 2057],
      | 80.00th=[ 2606], 90.00th=[ 4293], 95.00th=[ 5997],
      | 99.00th=[ 11600], 99.50th=[ 13960], 99.90th=[ 826278],
      | 99.95th=[17112761], 99.99th=[17112761]
    bw ( KiB/s): min= 2043, max=610304, per=100.00%, avg=332192.42, stdev=168523.35,
samples=62
    iops      : min= 1, max= 596, avg=324.08, stdev=164.57, samples=62
    lat (nsec) : 2=0.26%
    lat (usec) : 20=0.01%, 100=0.01%, 500=0.03%, 750=0.03%, 1000=0.02%
    lat (msec) : 2=67.67%, 4=20.74%, 10=9.56%, 20=1.42%, 50=0.10%
    lat (msec) : 100=0.04%, 250=0.01%, 1000=0.01%, >=2000=0.10%
    cpu        : usr=0.23%, sys=1.18%, ctx=10245, majf=0, minf=16
    IO depths  : 1=100.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%
    submit     : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
    complete   : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%

```

```
issued rwts: total=0,10240,0,0 short=0,0,0,0 dropped=0,0,0,0
latency   : target=0, window=0, percentile=100.00%, depth=1
```

Run status group 0 (all jobs):

```
WRITE: bw=37.9MiB/s (39.8MB/s), 37.9MiB/s-37.9MiB/s (39.8MB/s-39.8MB/s), io=10.0GiB
(10.7GB), run=269886-269886msec
```

Disk stats (read/write):

```
sda: ios=3017/10737, sectors=75792/21008312, merge=253/833, ticks=3237/526767,
in_queue=686667, util=96.37%
```

# Benchmark - read

```
root@nested-vm-l2:~# fio random_read_test.fio
```

```
random_read_job: (g=0): rw=randread, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B,
ioengine=libaio, iodepth=64
```

```
fio-3.36
```

Starting 1 process

```
random_read_job: Laying out IO file (1 file / 10240MiB)
```

```
Jobs: 1 (f=1): [r(1)][100.0%][r=20.0MiB/s][r=5112 IOPS][eta 00m:00s]
```

```
random_read_job: (groupid=0, jobs=1): err= 0: pid=3300: Mon Aug 11 19:02:19 2025
```

```
read: IOPS=4748, BW=18.5MiB/s (19.4MB/s)(1113MiB/60004msec)
```

```
slat (nsec): min=0, max=13840k, avg=116055.33, stdev=350395.40
```

```
clat (usec): min=2917, max=48588, avg=13307.65, stdev=5035.71
```

```
lat (usec): min=2958, max=48739, avg=13423.70, stdev=5063.84
```

```
clat percentiles (usec):
```

```
| 1.00th=[ 5735], 5.00th=[ 6980], 10.00th=[ 7832], 20.00th=[ 8848],
| 30.00th=[ 9896], 40.00th=[10945], 50.00th=[12125], 60.00th=[13566],
| 70.00th=[15533], 80.00th=[17695], 90.00th=[20317], 95.00th=[22938],
| 99.00th=[27657], 99.50th=[29754], 99.90th=[33817], 99.95th=[35390],
| 99.99th=[39584]
```

```
bw ( KiB/s): min=12645, max=34142, per=100.00%, avg=19000.53, stdev=4050.47, samples=119
```

```
iops      : min= 3161, max= 8535, avg=4749.84, stdev=1012.64, samples=119
```

```
lat (msec) : 4=0.01%, 10=30.78%, 20=58.31%, 50=10.90%
```

```
cpu       : usr=3.98%, sys=90.23%, ctx=2410, majf=0, minf=74
```

```
IO depths  : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.1%, 32=0.1%, >=64=100.0%
```

```
submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
```

```
complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.1%, >=64=0.0%
```

```
issued rwts: total=284908,0,0,0 short=0,0,0,0 dropped=0,0,0,0
```

```
latency : target=0, window=0, percentile=100.00%, depth=64
```

```
Run status group 0 (all jobs):
```

```
  READ: bw=18.5MiB/s (19.4MB/s), 18.5MiB/s-18.5MiB/s (19.4MB/s-19.4MB/s), io=1113MiB  
(1167MB), run=60004-60004msec
```

```
Disk stats (read/write):
```

```
  sda: ios=283830/8, sectors=2270640/72, merge=0/1, ticks=1580398/232, in_queue=1580794,  
  util=49.28%
```

# Hardware: HPE

## PVE - Mirrored 15K SAS

```
# CPU
```

```
root@btr-pve-c9ne1:~# grep "model name" /proc/cpuinfo
```

```
model name      : Intel(R) Xeon(R) CPU E5-2687W v4 @ 3.00GHz
```

```
# 2 Sockets, 24 Cores per socket, 48 total cores
```

```
# DISK: ZFS Mirrored HP EH0300JDYTH 15k SAS
```

```
# Benchmark - write
```

```
root@pve-hpe:~# fio --name=sequential_write --rw=write --bs=1M --size=10G --
```

```
filename=/root/testfile --direct=1
```

```
sequential_write: (g=0): rw=write, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-1024KiB, (T) 1024KiB-  
1024KiB, ioengine=psync, iodepth=1
```

```
fio-3.39
```

```
Starting 1 process
```

```
sequential_write: Laying out IO file (1 file / 10240MiB)
```

```
Jobs: 1 (f=1): [W(1)][100.0%][w=100MiB/s][w=100 IOPS][eta 00m:00s]
```

```
sequential_write: (groupid=0, jobs=1): err= 0: pid=1156898: Mon Aug 11 14:02:36 2025
```

```
  write: IOPS=93, BW=93.5MiB/s (98.0MB/s)(10.0GiB/109530msec); 0 zone resets
```

```
    clat (usec): min=8364, max=80908, avg=10618.01, stdev=3390.34
```

```
      lat (usec): min=8507, max=80994, avg=10691.62, stdev=3390.64
```

```
    clat percentiles (usec):
```

```
| 1.00th=[ 8717], 5.00th=[ 8848], 10.00th=[ 8848], 20.00th=[ 8979],
| 30.00th=[ 9110], 40.00th=[ 9503], 50.00th=[10028], 60.00th=[10552],
| 70.00th=[11076], 80.00th=[11469], 90.00th=[11863], 95.00th=[13042],
| 99.00th=[25035], 99.50th=[32113], 99.90th=[53216], 99.95th=[57934],
| 99.99th=[80217]
bw ( KiB/s): min=51200, max=104448, per=100.00%, avg=95758.09, stdev=11129.23, samples=218
iops      : min= 50, max= 102, avg=93.51, stdev=10.87, samples=218
lat (msec) : 10=50.61%, 20=47.46%, 50=1.81%, 100=0.13%
cpu       : usr=0.74%, sys=65.95%, ctx=10420, majf=5, minf=8
IO depths : 1=100.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%
submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
issued rwts: total=0,10240,0,0 short=0,0,0,0 dropped=0,0,0,0
latency   : target=0, window=0, percentile=100.00%, depth=1
```

Run status group 0 (all jobs):

```
WRITE: bw=93.5MiB/s (98.0MB/s), 93.5MiB/s-93.5MiB/s (98.0MB/s-98.0MB/s), io=10.0GiB
(10.7GB), run=109530-109530msec
```

# Benchmark - read

```
root@pve-hpe:~# fio random_read_test.fio
random_read_job: (g=0): rw=randread, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B,
ioengine=libaio, iodepth=64
fio-3.39
Starting 1 process
Jobs: 1 (f=1): [r(1)][100.0%][r=820KiB/s][r=205 IOPS][eta 00m:00s]
random_read_job: (groupid=0, jobs=1): err= 0: pid=1159109: Mon Aug 11 14:09:42 2025
read: IOPS=217, BW=871KiB/s (892kB/s)(51.0MiB/60001msec)
slat (usec): min=462, max=40890, avg=4580.11, stdev=1975.13
clat (usec): min=13, max=531922, avg=288716.77, stdev=40807.24
lat (msec): min=6, max=535, avg=293.30, stdev=41.16
clat percentiles (msec):
| 1.00th=[ 251], 5.00th=[ 262], 10.00th=[ 266], 20.00th=[ 271],
| 30.00th=[ 271], 40.00th=[ 275], 50.00th=[ 279], 60.00th=[ 284],
| 70.00th=[ 288], 80.00th=[ 292], 90.00th=[ 313], 95.00th=[ 393],
| 99.00th=[ 460], 99.50th=[ 481], 99.90th=[ 527], 99.95th=[ 531],
| 99.99th=[ 531]
```

```
bw ( KiB/s): min= 376, max= 1112, per=99.92%, avg=870.73, stdev=99.13, samples=120
iops       : min= 94, max= 278, avg=217.68, stdev=24.78, samples=120
lat (usec) : 20=0.01%
lat (msec) : 10=0.01%, 20=0.02%, 50=0.05%, 100=0.09%, 250=0.54%
lat (msec) : 500=98.97%, 750=0.32%
cpu        : usr=0.33%, sys=3.76%, ctx=13070, majf=0, minf=408
IO depths  : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.1%, 32=0.2%, >=64=99.5%
  submit   : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
  complete : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.1%, >=64=0.0%
  issued rwts: total=13061,0,0,0 short=0,0,0,0 dropped=0,0,0,0
  latency   : target=0, window=0, percentile=100.00%, depth=64
```

Run status group 0 (all jobs):

```
  READ: bw=871KiB/s (892kB/s), 871KiB/s-871KiB/s (892kB/s-892kB/s), io=51.0MiB (53.5MB),
run=60001-60001msec
```

## PVE - RAIDz3: 10 x 7200 rpm (MM1000GBKAL)

```
root@pve:/z-7200-SATA-Rz3# fio --name=sequential_write --rw=write --bs=1M --size=10G --
filename=testfile --direct=1
sequential_write: (g=0): rw=write, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-1024KiB, (T) 1024KiB-
1024KiB, ioengine=psync, iodepth=1
fio-3.39
Starting 1 process
sequential_write: Laying out IO file (1 file / 10240MiB)
Jobs: 1 (f=1): [W(1)][100.0%][w=37.0MiB/s][w=37 IOPS][eta 00m:00s]
sequential_write: (groupid=0, jobs=1): err= 0: pid=1546695: Tue Aug 12 10:18:08 2025
  write: IOPS=54, BW=55.0MiB/s (57.6MB/s)(10.0GiB/186319msec); 0 zone resets
    clat (msec): min=9, max=206, avg=18.11, stdev=10.06
      lat (msec): min=9, max=206, avg=18.19, stdev=10.06
    clat percentiles (msec):
      | 1.00th=[ 14], 5.00th=[ 15], 10.00th=[ 16], 20.00th=[ 16],
      | 30.00th=[ 17], 40.00th=[ 17], 50.00th=[ 17], 60.00th=[ 17],
      | 70.00th=[ 17], 80.00th=[ 18], 90.00th=[ 18], 95.00th=[ 21],
      | 99.00th=[ 77], 99.50th=[ 93], 99.90th=[ 127], 99.95th=[ 134],
      | 99.99th=[ 159]
    bw ( KiB/s): min=12288, max=67584, per=100.00%, avg=56300.17, stdev=11089.88, samples=372
    iops       : min= 12, max= 66, avg=54.98, stdev=10.83, samples=372
    lat (msec) : 10=0.03%, 20=94.98%, 50=2.80%, 100=1.90%, 250=0.28%
```

```
cpu          : usr=0.47%, sys=46.47%, ctx=11448, majf=0, minf=9
IO depths    : 1=100.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%
submit      : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
complete    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
issued rwts: total=0,10240,0,0 short=0,0,0,0 dropped=0,0,0,0
latency     : target=0, window=0, percentile=100.00%, depth=1
```

Run status group 0 (all jobs):

```
WRITE: bw=55.0MiB/s (57.6MB/s), 55.0MiB/s-55.0MiB/s (57.6MB/s-57.6MB/s), io=10.0GiB
(10.7GB), run=186319-186319msec
```

```
root@pve:/z-7200-SATA-Rz3# fio random_read_test.fio
```

```
random_read_job: (g=0): rw=randread, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B,
ioengine=libaio, iodepth=64
```

```
fio-3.39
```

```
Starting 1 process
```

```
Jobs: 1 (f=1): [r(1)][100.0%][r=436KiB/s][r=109 IOPS][eta 00m:00s]
```

```
random_read_job: (groupid=0, jobs=1): err= 0: pid=1578323: Tue Aug 12 10:46:14 2025
```

```
read: IOPS=105, BW=421KiB/s (431kB/s)(24.7MiB/60008msec)
```

```
slat (usec): min=519, max=100246, avg=9492.20, stdev=4271.24
```

```
clat (usec): min=14, max=1314.5k, avg=595573.00, stdev=91957.69
```

```
lat (msec): min=8, max=1323, avg=605.07, stdev=92.62
```

```
clat percentiles (msec):
```

```
| 1.00th=[ 550], 5.00th=[ 558], 10.00th=[ 567], 20.00th=[ 567],
| 30.00th=[ 575], 40.00th=[ 575], 50.00th=[ 575], 60.00th=[ 584],
| 70.00th=[ 584], 80.00th=[ 592], 90.00th=[ 617], 95.00th=[ 793],
| 99.00th=[ 927], 99.50th=[ 1301], 99.90th=[ 1318], 99.95th=[ 1318],
| 99.99th=[ 1318]
```

```
bw ( KiB/s): min= 112, max= 944, per=100.00%, avg=424.40, stdev=73.04, samples=119
```

```
iosps      : min= 28, max= 236, avg=106.10, stdev=18.26, samples=119
```

```
lat (usec) : 20=0.02%
```

```
lat (msec) : 10=0.02%, 20=0.02%, 50=0.05%, 100=0.08%, 250=0.27%
```

```
lat (msec) : 500=0.44%, 750=92.46%, 1000=5.69%, 2000=0.97%
```

```
cpu        : usr=0.16%, sys=2.54%, ctx=6315, majf=0, minf=82
```

```
IO depths  : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.3%, 32=0.5%, >=64=99.0%
```

```
submit     : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
```

```
complete   : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.1%, >=64=0.0%
```

```
issued rwts: total=6313,0,0,0 short=0,0,0,0 dropped=0,0,0,0
```

```
latency : target=0, window=0, percentile=100.00%, depth=64
```

```
Run status group 0 (all jobs):
```

```
READ: bw=421KiB/s (431kB/s), 421KiB/s-421KiB/s (431kB/s-431kB/s), io=24.7MiB (25.9MB),  
run=60008-60008msec
```

## PVE - draid3:6d:10c:1s-0: 10 x 10k SAS (EG1200JEMDA)

```
root@pve:/# cd z-10K-SAS-dRAID3/  
root@pve:/z-10K-SAS-dRAID3# fio --name=sequential_write --rw=write --bs=1M --size=10G --  
filename=/root/testfile --direct=1  
sequential_write: (g=0): rw=write, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-1024KiB, (T) 1024KiB-  
1024KiB, ioengine=psync, iodepth=1  
fio-3.39  
Starting 1 process  
sequential_write: Laying out IO file (1 file / 10240MiB)  
Jobs: 1 (f=1): [W(1)][100.0%][w=98.0MiB/s][w=98 IOPS][eta 00m:00s]  
sequential_write: (groupid=0, jobs=1): err= 0: pid=1550993: Tue Aug 12 10:20:41 2025  
write: IOPS=90, BW=90.8MiB/s (95.2MB/s)(10.0GiB/112818msec); 0 zone resets  
clat (usec): min=7765, max=67166, avg=10939.46, stdev=3638.62  
lat (usec): min=7858, max=67210, avg=11012.52, stdev=3639.09  
clat percentiles (usec):  
| 1.00th=[ 8160], 5.00th=[ 8356], 10.00th=[ 8455], 20.00th=[ 8979],  
| 30.00th=[ 9503], 40.00th=[10028], 50.00th=[10552], 60.00th=[10945],  
| 70.00th=[11469], 80.00th=[11863], 90.00th=[12125], 95.00th=[12387],  
| 99.00th=[28705], 99.50th=[36439], 99.90th=[49021], 99.95th=[56361],  
| 99.99th=[64226]  
bw ( KiB/s): min=45056, max=112640, per=100.00%, avg=93088.43, stdev=12260.36, samples=225  
iops      : min= 44, max= 110, avg=90.91, stdev=11.97, samples=225  
lat (msec) : 10=39.64%, 20=57.95%, 50=2.31%, 100=0.10%  
cpu      : usr=0.76%, sys=65.33%, ctx=10540, majf=0, minf=7  
IO depths : 1=100.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%  
submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%  
complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%  
issued rwts: total=0,10240,0,0 short=0,0,0,0 dropped=0,0,0,0  
latency   : target=0, window=0, percentile=100.00%, depth=1  
  
Run status group 0 (all jobs):  
WRITE: bw=90.8MiB/s (95.2MB/s), 90.8MiB/s-90.8MiB/s (95.2MB/s-95.2MB/s), io=10.0GiB  
(10.7GB), run=112818-112818msec
```

```

root@pve:/z-10K-SAS-dRAID3# fio random_read_test.fio
random_read_job: (g=0): rw=randread, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B,
ioengine=libaio, iodepth=64
fio-3.39
Starting 1 process
random_read_job: Laying out IO file (1 file / 10240MiB)
Jobs: 1 (f=1): [r(1)][100.0%][r=576KiB/s][r=144 IOPS][eta 00m:00s]
random_read_job: (groupid=0, jobs=1): err= 0: pid=1580837: Tue Aug 12 10:48:22 2025
  read: IOPS=163, BW=655KiB/s (671kB/s)(38.4MiB/60004msec)
    slat (usec): min=13, max=39844, avg=6090.61, stdev=2135.90
    clat (usec): min=12, max=523484, avg=383170.17, stdev=33200.56
      lat (msec): min=6, max=529, avg=389.26, stdev=33.38
    clat percentiles (msec):
      | 1.00th=[ 338], 5.00th=[ 355], 10.00th=[ 363], 20.00th=[ 368],
      | 30.00th=[ 372], 40.00th=[ 376], 50.00th=[ 380], 60.00th=[ 384],
      | 70.00th=[ 388], 80.00th=[ 393], 90.00th=[ 401], 95.00th=[ 409],
      | 99.00th=[ 510], 99.50th=[ 514], 99.90th=[ 523], 99.95th=[ 523],
      | 99.99th=[ 523]
    bw ( KiB/s): min= 112, max= 1104, per=99.95%, avg=655.40, stdev=75.42, samples=120
    iops       : min= 28, max= 276, avg=163.85, stdev=18.85, samples=120
    lat (usec) : 20=0.01%
    lat (msec) : 10=0.01%, 20=0.01%, 50=0.05%, 100=0.08%, 250=0.23%
    lat (msec) : 500=97.36%, 750=2.25%
    cpu        : usr=0.22%, sys=3.94%, ctx=9518, majf=0, minf=256
    IO depths  : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.2%, 32=0.3%, >=64=99.4%
      submit   : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
      complete : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.1%, >=64=0.0%
    issued rwts: total=9831,0,0,0 short=0,0,0,0 dropped=0,0,0,0
    latency    : target=0, window=0, percentile=100.00%, depth=64

Run status group 0 (all jobs):
  READ: bw=655KiB/s (671kB/s), 655KiB/s-655KiB/s (671kB/s-671kB/s), io=38.4MiB (40.3MB),
run=60004-60004msec

```

## PVE - ZFS Mirror + SLOG + L2ARC (cache)

2 x 10K SAS DS EH0300JDYTH + 2 x SSD SLOG + 1 x SSD L2ARC (cache)

All	3	1GiB	C: 36% (45/126GiB)	MB/s
	Read (MB/s)		Write (MB/s)	
SEQ1M Q8T1	2635.65		1572.68	
SEQ1M Q1T1	344.52		324.98	
RND4K Q32T1	18.92		17.18	
RND4K Q1T1	4.27		4.27	

All	3	1GiB	C: 36% (45/126GiB)
	Read (MB/s)		Write (MB/s)
SEQ1M Q1T1	459.03		507.95
RND4K Q1T1	4.03		3.88
RND4K (IOPS)	984.38		946.29
RND4K ( $\mu$ s)	968.37		1010.66

All	3	1GiB	Select Folder	MB/s
	Read (MB/s)		Write (MB/s)	
SEQ1M Q8T1	2681.86		128.34	
SEQ1M Q1T1	404.12		639.32	
RND4K Q32T1	20.18		16.86	
RND4K Q1T1	4.06		3.97	

All	3	1GiB	Select Folder
	Read (MB/s)		Write (MB/s)
SEQ1M Q1T1	481.12		623.66
RND4K Q1T1	4.56		3.98
RND4K (IOPS)	1113.53		971.19
RND4K ( $\mu$ s)	853.19		977.46

All	3	1GiB	Select Folder	MB/s
	Read (MB/s)		Write (MB/s)	
SEQ1M Q8T1	2515.36		1357.47	
SEQ1M Q1T1	393.23		10.06	
RND4K Q32T1	19.35		17.64	
RND4K Q1T1	4.29		4.38	

All	3	1GiB	Select Folder
	Read (MB/s)		Write (MB/s)
SEQ1M Q1T1	414.62		580.34
RND4K Q1T1	4.59		3.97
RND4K (IOPS)	1121.83		968.26
RND4K ( $\mu$ s)	847.31		980.28

```

root@pve:/# cd z-10k-SAS-DS-Mirror/
root@pve:/z-10k-SAS-DS-Mirror# fio --name=sequential_write --rw=write --bs=1M --size=10G --
filename=/root/testfile --direct=1
sequential_write: (g=0): rw=write, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-1024KiB, (T) 1024KiB-
1024KiB, ioengine=psync, iodepth=1
fio-3.39
Starting 1 process
Jobs: 1 (f=1): [W(1)][100.0%][w=69.0MiB/s][w=69 IOPS][eta 00m:00s]
sequential_write: (groupid=0, jobs=1): err= 0: pid=1553236: Tue Aug 12 10:22:50 2025
write: IOPS=89, BW=89.3MiB/s (93.7MB/s)(10.0GiB/114626msec); 0 zone resets
clat (msec): min=7, max=315, avg=11.11, stdev= 5.53
lat (msec): min=8, max=315, avg=11.18, stdev= 5.53
clat percentiles (msec):
| 1.00th=[ 9], 5.00th=[ 9], 10.00th=[ 9], 20.00th=[ 9],
| 30.00th=[ 10], 40.00th=[ 11], 50.00th=[ 11], 60.00th=[ 11],
| 70.00th=[ 12], 80.00th=[ 12], 90.00th=[ 13], 95.00th=[ 14],
| 99.00th=[ 33], 99.50th=[ 42], 99.90th=[ 68], 99.95th=[ 88],
| 99.99th=[ 116]

```

```
bw ( KiB/s): min=14336, max=108544, per=100.00%, avg=91516.09, stdev=14694.26, samples=229
iops      : min= 14, max= 106, avg=89.37, stdev=14.35, samples=229
lat (msec) : 10=39.91%, 20=57.56%, 50=2.22%, 100=0.28%, 250=0.02%
lat (msec) : 500=0.01%
cpu       : usr=0.71%, sys=64.60%, ctx=10487, majf=0, minf=7
IO depths : 1=100.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%
submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
issued rwts: total=0,10240,0,0 short=0,0,0,0 dropped=0,0,0,0
latency   : target=0, window=0, percentile=100.00%, depth=1
```

Run status group 0 (all jobs):

```
WRITE: bw=89.3MiB/s (93.7MB/s), 89.3MiB/s-89.3MiB/s (93.7MB/s-93.7MB/s), io=10.0GiB
(10.7GB), run=114626-114626msec
```

```
root@pve:/z-10k-SAS-DS-Mirror# fio random_read_test.fio
```

```
random_read_job: (g=0): rw=randread, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B,
ioengine=libaio, iodepth=64
```

```
fio-3.39
```

```
Starting 1 process
```

```
random_read_job: Laying out IO file (1 file / 10240MiB)
```

```
Jobs: 1 (f=1): [r(1)][100.0%][r=824KiB/s][r=206 IOPS][eta 00m:00s]
```

```
random_read_job: (groupid=0, jobs=1): err= 0: pid=1583195: Tue Aug 12 10:50:24 2025
```

```
read: IOPS=203, BW=812KiB/s (832kB/s)(47.6MiB/60003msec)
```

```
slat (usec): min=13, max=123917, avg=4911.77, stdev=3207.67
```

```
clat (usec): min=14, max=542655, avg=309167.09, stdev=35242.70
```

```
lat (msec): min=7, max=546, avg=314.08, stdev=35.48
```

```
clat percentiles (msec):
```

```
| 1.00th=[ 262], 5.00th=[ 279], 10.00th=[ 284], 20.00th=[ 292],
```

```
| 30.00th=[ 296], 40.00th=[ 300], 50.00th=[ 305], 60.00th=[ 309],
```

```
| 70.00th=[ 313], 80.00th=[ 321], 90.00th=[ 334], 95.00th=[ 359],
```

```
| 99.00th=[ 472], 99.50th=[ 498], 99.90th=[ 531], 99.95th=[ 535],
```

```
| 99.99th=[ 542]
```

```
bw ( KiB/s): min= 199, max= 1264, per=99.97%, avg=812.26, stdev=93.13, samples=120
```

```
iops      : min= 49, max= 316, avg=203.06, stdev=23.32, samples=120
```

```
lat (usec) : 20=0.01%
```

```
lat (msec) : 10=0.01%, 20=0.01%, 50=0.06%, 100=0.07%, 250=0.27%
```

```
lat (msec) : 500=99.21%, 750=0.37%
```

```
cpu          : usr=0.27%, sys=3.39%, ctx=11691, majf=0, minf=84
IO depths    : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.1%, 32=0.3%, >=64=99.5%
submit      : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
complete    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.1%, >=64=0.0%
issued rwts: total=12184,0,0,0 short=0,0,0,0 dropped=0,0,0,0
latency     : target=0, window=0, percentile=100.00%, depth=64
```

Run status group 0 (all jobs):

```
  READ: bw=812KiB/s (832kB/s), 812KiB/s-812KiB/s (832kB/s-832kB/s), io=47.6MiB (49.9MB),
run=60003-60003msec
```

## L1 VM - ZFS Mirrored SSD

```
# Benchmark - write
```

```
root@us-pve-101-testing:~# fio --name=sequential_write --rw=write --bs=1M --size=10G --
filename=/root/testfile --direct=1
sequential_write: (g=0): rw=write, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-1024KiB, (T) 1024KiB-
1024KiB, ioengine=psync, iodepth=1
fio-3.36
Starting 1 process
sequential_write: Laying out IO file (1 file / 10240MiB)
Jobs: 1 (f=1): [W(1)][96.6%][eta 00m:06s]
sequential_write: (groupid=0, jobs=1): err= 0: pid=7040: Mon Aug 11 22:41:53 2025
write: IOPS=58, BW=59.0MiB/s (61.8MB/s)(10.0GiB/173676msec); 0 zone resets
  clat (usec): min=647, max=45665k, avg=16898.79, stdev=707938.73
  lat (usec): min=656, max=45665k, avg=16955.22, stdev=707938.94
  clat percentiles (usec):
    | 1.00th=[ 988], 5.00th=[ 1352], 10.00th=[ 1450],
    | 20.00th=[ 1582], 30.00th=[ 1647], 40.00th=[ 1713],
    | 50.00th=[ 1762], 60.00th=[ 1795], 70.00th=[ 1827],
    | 80.00th=[ 1860], 90.00th=[ 1909], 95.00th=[ 1958],
    | 99.00th=[ 2278], 99.50th=[ 2409], 99.90th=[ 4359],
    | 99.95th=[11475616], 99.99th=[17112761]
  bw ( KiB/s): min=26677, max=819200, per=100.00%, avg=486353.63, stdev=170810.86,
samples=43
  iops          : min= 26, max= 800, avg=474.84, stdev=166.83, samples=43
```

```
lat (usec)   : 750=0.17%, 1000=0.89%
lat (msec)   : 2=95.51%, 4=3.31%, 10=0.06%, 750=0.01%, >=2000=0.06%
cpu          : usr=0.51%, sys=0.95%, ctx=10241, majf=0, minf=13
IO depths    : 1=100.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%
  submit     : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
  complete   : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
  issued rwts: total=0,10240,0,0 short=0,0,0,0 dropped=0,0,0,0
  latency    : target=0, window=0, percentile=100.00%, depth=1
```

Run status group 0 (all jobs):

```
WRITE: bw=59.0MiB/s (61.8MB/s), 59.0MiB/s-59.0MiB/s (61.8MB/s-61.8MB/s), io=10.0GiB
(10.7GB), run=173676-173676msec
```

Disk stats (read/write):

```
sda: ios=0/10537, sectors=0/21021040, merge=0/500, ticks=0/4724709, in_queue=4862347,
util=97.87%
```

# Benchmark - read

```
root@us-pve-101-testing:~# fio random_read_test.fio
```

```
random_read_job: (g=0): rw=randread, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B,
ioengine=libaio, iodepth=64
```

```
fio-3.36
```

```
Starting 1 process
```

```
Jobs: 1 (f=1): [r(1)][100.0%][r=308MiB/s][r=78.9k IOPS][eta 00m:00s]
```

```
random_read_job: (groupid=0, jobs=1): err= 0: pid=7047: Mon Aug 11 22:43:10 2025
```

```
read: IOPS=76.1k, BW=297MiB/s (312MB/s)(17.4GiB/60001msec)
```

```
slat (usec): min=2, max=5361, avg= 8.95, stdev= 8.40
```

```
clat (usec): min=115, max=8190, avg=831.02, stdev=330.81
```

```
lat (usec): min=158, max=8199, avg=839.97, stdev=334.39
```

```
clat percentiles (usec):
```

```
| 1.00th=[ 351], 5.00th=[ 388], 10.00th=[ 416], 20.00th=[ 486],
```

```
| 30.00th=[ 562], 40.00th=[ 619], 50.00th=[ 898], 60.00th=[ 1057],
```

```
| 70.00th=[ 1090], 80.00th=[ 1123], 90.00th=[ 1172], 95.00th=[ 1221],
```

```
| 99.00th=[ 1663], 99.50th=[ 1745], 99.90th=[ 1876], 99.95th=[ 2008],
```

```
| 99.99th=[ 2540]
```

```
bw ( KiB/s): min=213812, max=539568, per=100.00%, avg=305124.23, stdev=72008.27,
```

```
samples=119
```

```
iops      : min=53453, max=134892, avg=76281.10, stdev=18002.05, samples=119
lat (usec) : 250=0.01%, 500=21.93%, 750=26.28%, 1000=4.38%
lat (msec) : 2=47.36%, 4=0.05%, 10=0.01%
cpu       : usr=32.95%, sys=65.00%, ctx=11406, majf=0, minf=77
IO depths : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.1%, 32=0.1%, >=64=100.0%
  submit   : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
  complete : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.1%, >=64=0.0%
issued rwts: total=4564316,0,0,0 short=0,0,0,0 dropped=0,0,0,0
latency   : target=0, window=0, percentile=100.00%, depth=64
```

Run status group 0 (all jobs):

```
  READ: bw=297MiB/s (312MB/s), 297MiB/s-297MiB/s (312MB/s-312MB/s), io=17.4GiB (18.7GB),
run=60001-60001msec
```

Disk stats (read/write):

```
  sda: ios=4557508/19, sectors=36460064/176, merge=0/5, ticks=758830/37038, in_queue=832901,
util=68.19%
```

## L1 VM - ZFS Mirrored 10K SAS + SLOG + Cache

```
root@us-pve-101-testing:~# fio --name=sequential_write --rw=write --bs=1M --size=10G --
filename=/root/testfile --direct=1
sequential_write: (g=0): rw=write, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-1024KiB, (T) 1024KiB-
1024KiB, ioengine=psync, iodepth=1
fio-3.36
Starting 1 process
Jobs: 1 (f=1): [W(1)][100.0%][w=185MiB/s][w=185 IOPS][eta 00m:00s]
sequential_write: (groupid=0, jobs=1): err= 0: pid=997: Tue Aug 12 14:53:40 2025
  write: IOPS=253, BW=253MiB/s (265MB/s)(10.0GiB/40449msec); 0 zone resets
    clat (usec): min=1113, max=188603, avg=3871.65, stdev=2642.15
    lat (usec): min=1149, max=188642, avg=3943.94, stdev=2641.43
    clat percentiles (usec):
      | 1.00th=[ 1500], 5.00th=[ 1631], 10.00th=[ 1745], 20.00th=[ 1827],
      | 30.00th=[ 1942], 40.00th=[ 2311], 50.00th=[ 4817], 60.00th=[ 5211],
      | 70.00th=[ 5407], 80.00th=[ 5604], 90.00th=[ 5932], 95.00th=[ 6128],
      | 99.00th=[ 6652], 99.50th=[ 6849], 99.90th=[ 7177], 99.95th=[ 7898],
      | 99.99th=[64750]
    bw ( KiB/s): min=149504, max=610304, per=100.00%, avg=260124.17, stdev=140385.29,
samples=80
```

```
iops      : min= 146, max= 596, avg=253.86, stdev=137.09, samples=80
lat (msec) : 2=32.85%, 4=15.03%, 10=52.07%, 20=0.03%, 100=0.01%
lat (msec) : 250=0.01%
cpu       : usr=3.08%, sys=3.15%, ctx=10243, majf=0, minf=13
IO depths : 1=100.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%
submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
issued rwts: total=0,10240,0,0 short=0,0,0,0 dropped=0,0,0,0
latency   : target=0, window=0, percentile=100.00%, depth=1
```

Run status group 0 (all jobs):

```
WRITE: bw=253MiB/s (265MB/s), 253MiB/s-253MiB/s (265MB/s-265MB/s), io=10.0GiB (10.7GB),
run=40449-40449msec
```

Disk stats (read/write):

```
sda: ios=1/10259, sectors=8/20959881, merge=0/12, ticks=3/58119, in_queue=77002, util=93.49%
```

```
root@us-pve-101-testing:~# fio random_read_test.fio
```

```
random_read_job: (g=0): rw=randread, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B,
ioengine=libaio, iodepth=64
fio-3.36
```

Starting 1 process

```
Jobs: 1 (f=1): [r(1)][100.0%][r=486MiB/s][r=124k IOPS][eta 00m:00s]
```

```
random_read_job: (groupid=0, jobs=1): err= 0: pid=1000: Tue Aug 12 14:54:52 2025
```

```
read: IOPS=93.0k, BW=363MiB/s (381MB/s)(21.3GiB/60001msec)
```

```
slat (usec): min=2, max=525, avg= 7.23, stdev= 7.05
```

```
clat (usec): min=106, max=4144, avg=679.83, stdev=343.27
```

```
lat (usec): min=166, max=4152, avg=687.06, stdev=346.99
```

```
clat percentiles (usec):
```

```
| 1.00th=[ 355], 5.00th=[ 383], 10.00th=[ 400], 20.00th=[ 420],
| 30.00th=[ 437], 40.00th=[ 457], 50.00th=[ 486], 60.00th=[ 537],
| 70.00th=[ 832], 80.00th=[ 1139], 90.00th=[ 1188], 95.00th=[ 1254],
| 99.00th=[ 1532], 99.50th=[ 1762], 99.90th=[ 1876], 99.95th=[ 1909],
| 99.99th=[ 2089]
```

```
bw ( KiB/s): min=211328, max=580672, per=99.82%, avg=371243.65, stdev=97421.62,
samples=119
```

```
iops      : min=52832, max=145168, avg=92810.87, stdev=24355.37, samples=119
```

```
lat (usec) : 250=0.01%, 500=54.35%, 750=14.73%, 1000=2.17%
```

```
lat (msec) : 2=28.73%, 4=0.02%, 10=0.01%
cpu        : usr=35.72%, sys=62.92%, ctx=11510, majf=0, minf=76
IO depths  : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.1%, 32=0.1%, >=64=100.0%
  submit   : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
  complete : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.1%, >=64=0.0%
  issued rwts: total=5578560,0,0,0 short=0,0,0,0 dropped=0,0,0,0
  latency  : target=0, window=0, percentile=100.00%, depth=64
```

Run status group 0 (all jobs):

```
  READ: bw=363MiB/s (381MB/s), 363MiB/s-363MiB/s (381MB/s-381MB/s), io=21.3GiB (22.8GB),
run=60001-60001msec
```

Disk stats (read/write):

```
  sda: ios=5562235/9, sectors=44497880/80, merge=0/2, ticks=1046679/49508, in_queue=1124064,
util=53.76%
```

## Nested VM Benchmark - ZFS Mirrored SSD

```
# Benchmark - Write
```

```
root@us-hv-nested:~# fio --name=sequential_write --rw=write --bs=1M --size=10G --
filename=/root/testfile --direct=1
sequential_write: (g=0): rw=write, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-1024KiB, (T) 1024KiB-
1024KiB, ioengine=psync, iodepth=1
fio-3.36
Starting 1 process
sequential_write: Laying out IO file (1 file / 10240MiB)
Jobs: 1 (f=1): [W(1)][99.1%][w=165MiB/s][w=165 IOPS][eta 00m:02s]
sequential_write: (groupid=0, jobs=1): err= 0: pid=1439: Mon Aug 11 22:26:29 2025
  write: IOPS=44, BW=44.5MiB/s (46.6MB/s)(10.0GiB/230323msec); 0 zone resets
    clat (usec): min=9, max=28370k, avg=22023.80, stdev=535433.32
    lat (usec): min=9, max=28370k, avg=22285.21, stdev=535432.32
    clat percentiles (msec):
      | 1.00th=[  4], 5.00th=[  5], 10.00th=[  6], 20.00th=[  6],
      | 30.00th=[  6], 40.00th=[  6], 50.00th=[  7], 60.00th=[  7],
      | 70.00th=[  7], 80.00th=[  8], 90.00th=[  9], 95.00th=[ 10],
```

```
| 99.00th=[ 16], 99.50th=[ 25], 99.90th=[ 1838], 99.95th=[14429],
| 99.99th=[17113]
bw ( KiB/s): min= 4079, max=176761, per=100.00%, avg=132012.25, stdev=31904.93,
samples=157
iops      : min= 3, max= 172, avg=128.38, stdev=31.21, samples=157
lat (usec) : 10=0.01%, 500=0.01%
lat (msec) : 2=0.06%, 4=0.99%, 10=94.77%, 20=3.56%, 50=0.33%
lat (msec) : 100=0.06%, 250=0.08%, 500=0.01%, 750=0.01%, 2000=0.02%
lat (msec) : >=2000=0.10%
cpu       : usr=0.91%, sys=8.89%, ctx=12455, majf=0, minf=16
IO depths : 1=100.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%
submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
issued rwts: total=0,10240,0,0 short=0,0,0,0 dropped=0,0,0,0
latency   : target=0, window=0, percentile=100.00%, depth=1
```

Run status group 0 (all jobs):

```
WRITE: bw=44.5MiB/s (46.6MB/s), 44.5MiB/s-44.5MiB/s (46.6MB/s-46.6MB/s), io=10.0GiB
(10.7GB), run=230323-230323msec
```

Disk stats (read/write):

```
sda: ios=111/10313, sectors=3728/20900664, merge=90/235, ticks=2312/790448, in_queue=937669,
util=89.85%
```

# Benchmark - Read

```
root@us-hv-nested:~# fio random_read_test.fio
random_read_job: (g=0): rw=randread, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B,
ioengine=libaio, iodepth=64
fio-3.36
Starting 1 process
Jobs: 1 (f=1): [r(1)][100.0%][r=3498KiB/s][r=874 IOPS][eta 00m:00s]
random_read_job: (groupid=0, jobs=1): err= 0: pid=1454: Mon Aug 11 22:27:43 2025
read: IOPS=1151, BW=4608KiB/s (4718kB/s)(270MiB/60009msec)
slat (nsec): min=1200, max=30439k, avg=506941.31, stdev=1047300.59
clat (msec): min=2, max=335, avg=54.75, stdev=34.49
lat (msec): min=3, max=335, avg=55.26, stdev=34.75
```

clat percentiles (msec):

```
| 1.00th=[ 15], 5.00th=[ 16], 10.00th=[ 17], 20.00th=[ 19],  
| 30.00th=[ 24], 40.00th=[ 37], 50.00th=[ 61], 60.00th=[ 68],  
| 70.00th=[ 73], 80.00th=[ 80], 90.00th=[ 93], 95.00th=[ 111],  
| 99.00th=[ 157], 99.50th=[ 176], 99.90th=[ 215], 99.95th=[ 288],  
| 99.99th=[ 326]
```

```
bw ( KiB/s): min= 2003, max=11768, per=100.00%, avg=4613.21, stdev=1889.89, samples=118  
iops       : min= 500, max= 2942, avg=1152.94, stdev=472.50, samples=118  
lat (msec) : 4=0.01%, 10=0.01%, 20=24.47%, 50=21.04%, 100=46.64%  
lat (msec) : 250=7.75%, 500=0.08%  
cpu        : usr=5.15%, sys=81.55%, ctx=7548, majf=0, minf=76  
IO depths  : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.1%, 32=0.1%, >=64=99.9%  
submit     : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%  
complete   : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.1%, >=64=0.0%  
issued rwts: total=69124,0,0,0 short=0,0,0,0 dropped=0,0,0,0  
latency    : target=0, window=0, percentile=100.00%, depth=64
```

Run status group 0 (all jobs):

```
READ: bw=4608KiB/s (4718kB/s), 4608KiB/s-4608KiB/s (4718kB/s-4718kB/s), io=270MiB (283MB),  
run=60009-60009msec
```

Disk stats (read/write):

```
sda: ios=69121/20, sectors=552968/288, merge=0/5, ticks=583586/133, in_queue=583767,  
util=67.38%
```

## Hardware: PVE: MeLE Overclock 4C

# sequential\_write

```
bw ( KiB/s): min=397312, max=444416, per=100.00%, avg=429994.67, stdev=10951.34, samples=48  
iops       : min= 388, max= 434, avg=419.92, stdev=10.69, samples=48
```

# random\_read

```
bw ( KiB/s): min=15336, max=16192, per=100.00%, avg=15984.67, stdev=114.94, samples=120  
iops       : min= 3834, max= 4048, avg=3996.17, stdev=28.74, samples=120
```

# sequential\_read

```
bw ( KiB/s): min=10184, max=24472, per=100.00%, avg=19554.20, stdev=1018.25, samples=120  
iops       : min= 2546, max= 6118, avg=4888.53, stdev=254.56, samples=120
```

```
root@pmvh2501:~# fio --name=sequential_write --rw=write --bs=1M --size=10G --
filename=/root/testfile --direct=1
sequential_write: (g=0): rw=write, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-1024KiB, (T) 1024KiB-
1024KiB, ioengine=psync, iodepth=1
fio-3.39
Starting 1 process
sequential_write: Laying out IO file (1 file / 10240MiB)
Jobs: 1 (f=1): [W(1)][100.0%][w=418MiB/s][w=418 IOPS][eta 00m:00s]
sequential_write: (groupid=0, jobs=1): err= 0: pid=2040323: Tue Aug 12 21:03:13 2025
  write: IOPS=419, BW=420MiB/s (440MB/s)(10.0GiB/24394msec); 0 zone resets
    clat (usec): min=1719, max=48467, avg=2336.25, stdev=896.61
    lat (usec): min=1764, max=48512, avg=2380.72, stdev=896.69
  clat percentiles (usec):
    | 1.00th=[ 2147], 5.00th=[ 2180], 10.00th=[ 2212], 20.00th=[ 2278],
    | 30.00th=[ 2278], 40.00th=[ 2311], 50.00th=[ 2311], 60.00th=[ 2311],
    | 70.00th=[ 2311], 80.00th=[ 2311], 90.00th=[ 2376], 95.00th=[ 2409],
    | 99.00th=[ 2573], 99.50th=[ 2802], 99.90th=[17957], 99.95th=[21365],
    | 99.99th=[38536]
  bw ( KiB/s): min=397312, max=444416, per=100.00%, avg=429994.67, stdev=10951.34,
samples=48
  iops      : min= 388, max= 434, avg=419.92, stdev=10.69, samples=48
  lat (msec) : 2=0.04%, 4=99.70%, 10=0.10%, 20=0.10%, 50=0.07%
  cpu       : usr=2.16%, sys=92.04%, ctx=12136, majf=5, minf=10
  IO depths  : 1=100.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%
  submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
  complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
  issued rwts: total=0,10240,0,0 short=0,0,0,0 dropped=0,0,0,0
  latency   : target=0, window=0, percentile=100.00%, depth=1
```

Run status group 0 (all jobs):

```
  WRITE: bw=420MiB/s (440MB/s), 420MiB/s-420MiB/s (440MB/s-440MB/s), io=10.0GiB (10.7GB),
run=24394-24394msec
```

```
root@pmvh2501:~# fio random_read_test.fio
```

```
random_read_job: (g=0): rw=randread, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B,
ioengine=libaio, iodepth=64
```

```
fio-3.39
```

```
Starting 1 process
```

```
Jobs: 1 (f=1): [r(1)][100.0%][r=15.6MiB/s][r=3989 IOPS][eta 00m:00s]
random_read_job: (groupid=0, jobs=1): err= 0: pid=2040401: Tue Aug 12 21:04:13 2025
read: IOPS=3996, BW=15.6MiB/s (16.4MB/s)(937MiB/60001msec)
  slat (usec): min=151, max=4918, avg=248.40, stdev=43.31
  clat (usec): min=2, max=21443, avg=15762.61, stdev=450.75
    lat (usec): min=257, max=21696, avg=16011.01, stdev=454.68
  clat percentiles (usec):
    | 1.00th=[15008], 5.00th=[15139], 10.00th=[15270], 20.00th=[15401],
    | 30.00th=[15533], 40.00th=[15664], 50.00th=[15664], 60.00th=[15795],
    | 70.00th=[15926], 80.00th=[16057], 90.00th=[16188], 95.00th=[16450],
    | 99.00th=[17433], 99.50th=[17695], 99.90th=[18220], 99.95th=[18744],
    | 99.99th=[21103]
bw ( KiB/s): min=15336, max=16192, per=100.00%, avg=15984.67, stdev=114.94, samples=120
iops      : min= 3834, max= 4048, avg=3996.17, stdev=28.74, samples=120
lat (usec) : 4=0.01%, 500=0.01%, 750=0.01%
lat (msec)  : 2=0.01%, 4=0.01%, 10=0.01%, 20=99.96%, 50=0.02%
cpu         : usr=0.72%, sys=9.61%, ctx=240318, majf=0, minf=72
IO depths   : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.1%, 32=0.1%, >=64=100.0%
  submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
  complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.1%, >=64=0.0%
  issued rwts: total=239770,0,0,0 short=0,0,0,0 dropped=0,0,0,0
  latency   : target=0, window=0, percentile=100.00%, depth=64
```

Run status group 0 (all jobs):

```
  READ: bw=15.6MiB/s (16.4MB/s), 15.6MiB/s-15.6MiB/s (16.4MB/s-16.4MB/s), io=937MiB (982MB),
run=60001-60001msec
```

```
random_read_job: (g=0): rw=read, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B,
ioengine=libaio, iodepth=64
```

fio-3.39

Starting 1 process

```
Jobs: 1 (f=1): [R(1)][100.0%][r=18.8MiB/s][r=4820 IOPS][eta 00m:00s]
```

```
random_read_job: (groupid=0, jobs=1): err= 0: pid=2041137: Tue Aug 12 21:07:41 2025
```

```
read: IOPS=4888, BW=19.1MiB/s (20.0MB/s)(1146MiB/60001msec)
  slat (usec): min=135, max=32402, avg=202.85, stdev=82.92
  clat (usec): min=2, max=39730, avg=12873.56, stdev=959.97
    lat (usec): min=206, max=67804, avg=13076.41, stdev=978.59
  clat percentiles (usec):
```

```
| 1.00th=[10159], 5.00th=[12387], 10.00th=[12518], 20.00th=[12518],
| 30.00th=[12649], 40.00th=[12780], 50.00th=[12780], 60.00th=[12911],
| 70.00th=[13042], 80.00th=[13042], 90.00th=[13304], 95.00th=[13566],
| 99.00th=[15008], 99.50th=[15795], 99.90th=[30278], 99.95th=[34866],
| 99.99th=[38011]

bw ( KiB/s): min=10184, max=24472, per=100.00%, avg=19554.20, stdev=1018.25, samples=120
iops        : min= 2546, max= 6118, avg=4888.53, stdev=254.56, samples=120
lat (usec)  : 4=0.01%, 250=0.01%, 500=0.01%, 750=0.01%, 1000=0.01%
lat (msec)  : 2=0.01%, 4=0.01%, 10=0.87%, 20=98.95%, 50=0.17%
cpu         : usr=0.86%, sys=11.28%, ctx=293962, majf=0, minf=71
IO depths   : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.1%, 32=0.1%, >=64=100.0%
submit      : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
complete    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.1%, >=64=0.0%
issued rwts: total=293313,0,0,0 short=0,0,0,0 dropped=0,0,0,0
latency     : target=0, window=0, percentile=100.00%, depth=64
```

Run status group 0 (all jobs):

```
READ: bw=19.1MiB/s (20.0MB/s), 19.1MiB/s-19.1MiB/s (20.0MB/s-20.0MB/s), io=1146MiB
(1201MB), run=60001-60001msec
```

## Hardware: Nutanix NX-8155-G8

2 x SSD + 4 HDD

Benchmark was run during business hours with normal production workloads occurring.

CrystalDiskMark 9.0.1 x64 20250812 - Windows Server 2022 VM on Nutanix  
NX-8155-G8 - 2x SSD 4x HDD

File Settings Profile The

All 3 1GiB C: 44% (55/126GiB) MB/s


	Read (MB/s)	Write (MB/s)
SEQ1M Q8T1	554.43	309.30
SEQ1M Q1T1	406.66	204.46
RND4K Q32T1	42.77	80.46
RND4K Q1T1	8.98	8.30

CrystalDiskMark 9.0.1 x64 [Admin]

File Settings Profile Theme Help Language

All 3 1GiB C: 44% (55/126GiB)

	Read (MB/s)	Write (MB/s)
SEQ1M Q1T1	413.01	216.23
RND4K Q1T1	12.23	5.26
RND4K (IOPS)	2985.60	1284.67
RND4K (μs)	334.43	776.88



```

root@nutanix-l1-vm:~# fio --name=sequential_write --rw=write --bs=1M --size=10G --
filename=/root/testfile --direct=1
sequential_write: (g=0): rw=write, bs=(R) 1024KiB-1024KiB, (W) 1024KiB-1024KiB, (T) 1024KiB-
1024KiB, ioengine=psync, iodepth=1
fio-3.28
Starting 1 process
sequential_write: Laying out IO file (1 file / 10240MiB)
Jobs: 1 (f=1): [W(1)][100.0%][w=143MiB/s][w=143 IOPS][eta 00m:00s]
sequential_write: (groupid=0, jobs=1): err= 0: pid=129336: Tue Aug 12 10:51:08 2025
write: IOPS=155, BW=155MiB/s (163MB/s)(10.0GiB/66001msec); 0 zone resets
clat (usec): min=4496, max=56501, avg=6377.43, stdev=2162.18
lat (usec): min=4542, max=56536, avg=6438.83, stdev=2164.59
clat percentiles (usec):
| 1.00th=[ 4686], 5.00th=[ 4752], 10.00th=[ 4817], 20.00th=[ 4948],
| 30.00th=[ 5080], 40.00th=[ 5276], 50.00th=[ 5669], 60.00th=[ 6128],
| 70.00th=[ 6718], 80.00th=[ 7373], 90.00th=[ 8717], 95.00th=[10290],
| 99.00th=[14222], 99.50th=[16057], 99.90th=[21365], 99.95th=[25035],
| 99.99th=[54789]
bw ( KiB/s): min=120832, max=198656, per=100.00%, avg=159134.29, stdev=25598.04,
samples=131
iops      : min= 118, max= 194, avg=155.40, stdev=25.00, samples=131
lat (msec) : 10=94.29%, 20=5.58%, 50=0.12%, 100=0.02%
cpu       : usr=1.30%, sys=1.36%, ctx=10368, majf=0, minf=11
IO depths : 1=100.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%

```

```
submit      : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
complete   : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
issued rwts: total=0,10240,0,0 short=0,0,0,0 dropped=0,0,0,0
latency    : target=0, window=0, percentile=100.00%, depth=1
```

Run status group 0 (all jobs):

```
WRITE: bw=155MiB/s (163MB/s), 155MiB/s-155MiB/s (163MB/s-163MB/s), io=10.0GiB (10.7GB),
run=66001-66001msec
```

Disk stats (read/write):

```
sda: ios=0/20672, merge=0/278, ticks=0/122765, in_queue=122764, util=99.93%
```

```
root@nutanix-l1-vm:~# fio random_read_test.fio
```

```
random_read_job: (g=0): rw=randread, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T) 4096B-4096B,
ioengine=libaio, iodepth=64
```

```
fio-3.28
```

```
Starting 1 process
```

```
Jobs: 1 (f=1): [r(1)][100.0%][r=43.5MiB/s][r=11.1k IOPS][eta 00m:00s]
```

```
random_read_job: (groupid=0, jobs=1): err= 0: pid=129344: Tue Aug 12 10:53:31 2025
```

```
read: IOPS=15.3k, BW=59.7MiB/s (62.6MB/s)(3582MiB/60009msec)
```

```
slat (nsec): min=1483, max=352698, avg=2964.05, stdev=2532.94
```

```
clat (usec): min=256, max=236442, avg=4184.75, stdev=4942.01
```

```
lat (usec): min=258, max=236447, avg=4187.82, stdev=4942.20
```

```
clat percentiles (usec):
```

```
| 1.00th=[ 537], 5.00th=[ 832], 10.00th=[ 1090], 20.00th=[ 1467],
| 30.00th=[ 1778], 40.00th=[ 2114], 50.00th=[ 2606], 60.00th=[ 3326],
| 70.00th=[ 4424], 80.00th=[ 5997], 90.00th=[ 8979], 95.00th=[ 12387],
| 99.00th=[ 21365], 99.50th=[ 25822], 99.90th=[ 47973], 99.95th=[ 62653],
| 99.99th=[110625]
```

```
bw ( KiB/s): min=11144, max=151008, per=100.00%, avg=61258.53, stdev=26522.54, samples=119
```

```
iops      : min= 2786, max=37752, avg=15314.62, stdev=6630.65, samples=119
```

```
lat (usec) : 500=0.72%, 750=3.00%, 1000=4.50%
```

```
lat (msec) : 2=28.65%, 4=29.85%, 10=25.19%, 20=6.85%, 50=1.16%
```

```
lat (msec) : 100=0.08%, 250=0.01%
```

```
cpu      : usr=2.75%, sys=8.08%, ctx=475071, majf=0, minf=75
```

```
I/O depths : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.1%, 32=0.1%, >=64=100.0%
```

```
submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
```

```
complete : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.1%, >=64=0.0%
issued rwts: total=916907,0,0,0 short=0,0,0,0 dropped=0,0,0,0
latency   : target=0, window=0, percentile=100.00%, depth=64
```

Run status group 0 (all jobs):

READ: bw=59.7MiB/s (62.6MB/s), 59.7MiB/s-59.7MiB/s (62.6MB/s-62.6MB/s), io=3582MiB  
(3756MB), run=60009-60009msec


Disk stats (read/write):

sda: ios=915429/91, merge=0/14, ticks=3817950/1460, in\_queue=3819410, util=99.91%

## Hardware: Azure VM

Standard D2as v6 (2 vcpus, 8 GiB memory)

Storage: Premium SSD LRS



-end

# Troubleshooting

Various notes for random troubleshooting

## List all block devices with model

```
lsblk -o NAME,MODEL,SERIAL,SIZE,TYPE,MOUNTPOINT
```

## Backup to NAS mount via SMB - stale file handle

Testing backups to a QNAP via SMB resulted in failures with the message **stale file handle**. Setting the option below may have been what fixed the glitch, or it was just unmounting and letting PVE remount. Further testing needed.

```
pvesm set <storage_ID> --options noserverino
umount /mnt/pve/<storage_ID>
```

## Notifications via Discord

If your test notifications fail with a 400/500 message, make sure you've added the header Content-Type with value application/json. Below is the Body value I'm using currently.

Discord has a limit to the length of messages that can be sent, and PVE likes to send long messages sometimes. If your test messages are still failing with a status 400 message, they're probably too long. Need a way to limit the number of characters.

```
{
  "content": null,
  "embeds": [
    {
      "title": "{{ escape title }}",
```

```
    "url": "https://[IP]:8006",
    "description": "`#{ escape message }`",
    "color": null,
    "author": {
      "name": "Proxmox Cluster / Node - [cluster / node name]",
      "url": "https://[IP]:8006"
    },
    "footer": {
      "text": "Proxmox - [cluster/node name]"
    }
  }
],
"attachments": []
}
```

-end

# CEPH

## Useful commands

```
# quick status check
ceph status

# summary of storage pools, raw usage, and estimated capacity
ceph df

# utilization per osd
ceph osd df
```

## Adding CEPH Manager Dashboard

### [Source](#)

```
# run this on all service manager nodes
apt install ceph-mgr-dashboard

# run this to configure ceph mgr

echo "admin-password" > ~/ceph-dashboard-admin-password.txt

ceph mgr module enable dashboard
ceph dashboard ac-user-create cephdash administrator -i ~/ceph-dashboard-admin-password.txt
ceph config-key set mgr/dashboard/server_addr ::
ceph dashboard create-self-signed-cert
```

```
ceph mgr module disable dashboard
ceph mgr module enable dashboard

systemctl restart ceph-mgr@[servername].service

# Then https://[IP or FQDN]:8443 or http://[IP or FQDN]:8080
```

# Shutting down entire Proxmox cluster

If you need to shutdown your entire Proxmox cluster for any reason (maintenance, power issues, moving to a new site, whatever):

```
# prep for shutdown
ceph osd set noout
ceph osd set nobackfill
ceph osd set norecover

# shutdown all VMs and containers
```

```
# prep for service resumption
ceph osd unset norecover
ceph osd unset nobackfill
ceph osd unset noout

# start desired VMs and containers
```

# Upgrading Proxmox running CEPH

This has not yet been tested and verified, but adding it here for future reference.

[Source](#)

```
# Node maintenance

# stop and wait for scrub and deep-scrub operations

ceph osd set noscrub
ceph osd set nodeep-scrub

ceph status

# set cluster in maintenance mode with :

ceph osd set noout

# for node 1..N

# migrate VMs and CTs off node

# (GUI or CLI)

# run updates

apt update && pveupgrade

reboot # if required, e.g., kernel update

# wait for node to come back on line and quorate

# next N

# restore

ceph osd unset noout

# when all the PGs are active, re-enable the scrub and deep-scrub operations

ceph status

ceph osd unset noscrub
ceph osd unset nodeep-scrub
```

# done