

fio - тестирование скорости дисковой подсистемы

Тут неплохой список тестов: <https://docs.oracle.com/en-us/iaas/Content/Block/References/samplefiocommandslinux.htm>

Команды:

```
#IOPS
```

```
fio --filename=/file --size=5GB --direct=1 --rw=randrw --bs=4k \
--ioengine=libaio --iodepth=256 --runtime=120 --numjobs=4 --time_based \
--group_reporting --name=iops-test-job --eta-newline=1
```

```
# B/W
```

```
fio --filename=/file --size=5GB --direct=1 --rw=randrw --bs=64k \
--ioengine=libaio --iodepth=64 --runtime=120 --numjobs=4 --time_based \
--group_reporting --name=throughput-test-job --eta-newline=1
```

```
# Latency
```

```
fio --filename=/file --size=5GB --direct=1 --rw=randrw --bs=4k \
--ioengine=libaio --iodepth=1 --numjobs=1 --time_based --group_reporting \
--name=rwlatency-test-job --runtime=120 --eta-newline=1
```

Вот пример random r/w iops на десктопном Toshiba HDD:

```
iops-test-job: (groupid=0, jobs=4): err= 0: pid=251421: Wed Jul 17 10:37:16 2024
read: IOPS=76, BW=307KiB/s (314kB/s)(36.1MiB/120567msec)
```

slat (usec): min=3, max=1056.9k, avg=25388.41, stdev=70969.08

clat (msec): min=285, max=10659, avg=6440.46, stdev=987.92

lat (msec): min=500, max=10763, avg=6465.85, stdev=990.85

clat percentiles (msec):

| 1.00th=[1653], 5.00th=[5671], 10.00th=[6007], 20.00th=[6141],
| 30.00th=[6275], 40.00th=[6342], 50.00th=[6477], 60.00th=[6544],
| 70.00th=[6678], 80.00th=[6812], 90.00th=[7148], 95.00th=[7819],
| 99.00th=[8926], 99.50th=[9060], 99.90th=[9866], 99.95th=[10268],
| 99.99th=[10671]

bw (KiB/s): min= 40, max= 608, per=100.00%, avg=307.70, stdev=26.41, samples=907

iosps : min= 10, max= 152, avg=76.92, stdev= 6.60, samples=907

write: IOPS=79, BW=319KiB/s (326kB/s)(37.5MiB/120567msec); 0 zone resets

slat (usec): min=3, max=975286, avg=25577.27, stdev=70511.46

clat (msec): min=469, max=9979, avg=6323.55, stdev=931.04

lat (msec): min=500, max=9979, avg=6349.13, stdev=933.87

clat percentiles (msec):

| 1.00th=[1569], 5.00th=[5604], 10.00th=[5940], 20.00th=[6074],
| 30.00th=[6208], 40.00th=[6275], 50.00th=[6342], 60.00th=[6409],
| 70.00th=[6544], 80.00th=[6611], 90.00th=[6879], 95.00th=[7684],
| 99.00th=[8658], 99.50th=[8792], 99.90th=[9194], 99.95th=[9597],
| 99.99th=[10000]

bw (KiB/s): min= 64, max= 632, per=100.00%, avg=320.56, stdev=26.98, samples=909

iosps : min= 16, max= 158, avg=80.14, stdev= 6.74, samples=909

lat (msec) : 500=0.02%, 750=0.25%, 1000=0.26%, 2000=0.76%, >=2000=98.71%

cpu : usr=0.03%, sys=0.04%, ctx=11061, majf=0, minf=43

IO depths : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.2%, 16=0.3%, 32=0.7%, >=64=98.7%

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=0.1%

issued rwts: total=9243,9607,0,0 short=0,0,0,0 dropped=0,0,0,0

latency : target=0, window=0, percentile=100.00%, depth=256

Run status group 0 (all jobs):

READ: bw=307KiB/s (314kB/s), 307KiB/s-307KiB/s (314kB/s-314kB/s), io=36.1MiB (37.9MB), run=120567-120567msec

WRITE: bw=319KiB/s (326kB/s), 319KiB/s-319KiB/s (326kB/s-326kB/s), io=37.5MiB (39.3MB), run=120567-120567msec

Disk stats (read/write):

sda: ios=9257/9607, merge=0/0, ticks=4143068/3114506, in_queue=7257575, util=99.37%

Блоками по 512 картина еще унылее.

На 8к вцелом плюс-минус то же

```
iops-test-job: (groupid=0, jobs=4): err= 0: pid=251505: Wed Jul 17 10:42:18 2024
read: IOPS=77, BW=620KiB/s (635kB/s)(73.0MiB/120481msec)
    slat (usec): min=3, max=387273, avg=25978.96, stdev=69175.81
    clat (msec): min=320, max=8886, avg=6372.20, stdev=876.49
    lat (msec): min=486, max=8886, avg=6398.18, stdev=878.93
    clat percentiles (msec):
    | 1.00th=[ 1552], 5.00th=[ 5738], 10.00th=[ 6007], 20.00th=[ 6208],
    | 30.00th=[ 6275], 40.00th=[ 6409], 50.00th=[ 6477], 60.00th=[ 6544],
    | 70.00th=[ 6678], 80.00th=[ 6812], 90.00th=[ 7013], 95.00th=[ 7215],
    | 99.00th=[ 7684], 99.50th=[ 7819], 99.90th=[ 8288], 99.95th=[ 8423],
    | 99.99th=[ 8926]
bw ( KiB/s): min= 176, max= 1344, per=99.33%, avg=616.50, stdev=53.06, samples=915
iops      : min= 22, max= 168, avg=77.06, stdev= 6.63, samples=915
write: IOPS=80, BW=644KiB/s (660kB/s)(75.8MiB/120481msec); 0 zone resets
    slat (usec): min=3, max=417654, avg=24495.72, stdev=67714.31
    clat (msec): min=305, max=7858, avg=6266.56, stdev=797.47
    lat (msec): min=486, max=7903, avg=6291.05, stdev=799.85
    clat percentiles (msec):
    | 1.00th=[ 1670], 5.00th=[ 5738], 10.00th=[ 5940], 20.00th=[ 6141],
    | 30.00th=[ 6208], 40.00th=[ 6275], 50.00th=[ 6409], 60.00th=[ 6477],
    | 70.00th=[ 6544], 80.00th=[ 6611], 90.00th=[ 6745], 95.00th=[ 6879],
    | 99.00th=[ 7215], 99.50th=[ 7416], 99.90th=[ 7752], 99.95th=[ 7752],
    | 99.99th=[ 7886]
bw ( KiB/s): min= 96, max= 1264, per=99.48%, avg=641.88, stdev=54.29, samples=918
iops      : min= 12, max= 158, avg=80.24, stdev= 6.79, samples=918
lat (msec) : 500=0.06%, 750=0.25%, 1000=0.17%, 2000=0.88%, >=2000=98.63%
cpu        : usr=0.03%, sys=0.04%, ctx=11609, majf=0, minf=44
IO depths  : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.2%, 16=0.3%, 32=0.7%, >=64=98.7%
    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=0.1%
```

issued rwts: total=9340,9704,0,0 short=0,0,0,0 dropped=0,0,0,0

latency : target=0, window=0, percentile=100.00%, depth=256

Run status group 0 (all jobs):

READ: bw=620KiB/s (635kB/s), 620KiB/s-620KiB/s (635kB/s-635kB/s), io=73.0MiB (76.5MB), run=120481-120481msec

WRITE: bw=644KiB/s (660kB/s), 644KiB/s-644KiB/s (660kB/s-660kB/s), io=75.8MiB (79.5MB), run=120481-120481msec

Disk stats (read/write):

sda: ios=9328/9704, merge=0/0, ticks=4175537/3088888, in_queue=7264425, util=99.67%

Иопсы те же, но скорость выросла.

WD RED на 5400 оборотов. Несмотря на то что он на 5400, скорость выше чем у 7200 Toshiba HDD

iops-test-job: (groupid=0, jobs=4): err= 0: pid=251543: Wed Jul 17 10:46:32 2024

read: IOPS=86, BW=348KiB/s (356kB/s)(41.1MiB/120812msec)

slat (usec): min=2, max=652947, avg=22065.33, stdev=65152.95

clat (msec): min=411, max=8507, avg=5712.90, stdev=758.82

lat (msec): min=699, max=8507, avg=5734.97, stdev=760.79

clat percentiles (msec):

| 1.00th=[2106], 5.00th=[4799], 10.00th=[5134], 20.00th=[5336],
| 30.00th=[5537], 40.00th=[5604], 50.00th=[5738], 60.00th=[5873],
| 70.00th=[6007], 80.00th=[6208], 90.00th=[6477], 95.00th=[6678],
| 99.00th=[7215], 99.50th=[7416], 99.90th=[7819], 99.95th=[7886],
| 99.99th=[8154]

bw (KiB/s): min= 56, max= 784, per=100.00%, avg=350.09, stdev=32.50, samples=915

iops : min= 14, max= 196, avg=87.52, stdev= 8.12, samples=915

write: IOPS=90, BW=362KiB/s (371kB/s)(42.7MiB/120812msec); 0 zone resets

slat (usec): min=2, max=612526, avg=22706.93, stdev=65998.81

clat (msec): min=542, max=7598, avg=5518.32, stdev=699.72

lat (msec): min=699, max=7598, avg=5541.03, stdev=702.26

clat percentiles (msec):

| 1.00th=[1989], 5.00th=[4665], 10.00th=[5000], 20.00th=[5201],

```
| 30.00th=[ 5336], 40.00th=[ 5470], 50.00th=[ 5604], 60.00th=[ 5671],
| 70.00th=[ 5805], 80.00th=[ 5940], 90.00th=[ 6141], 95.00th=[ 6342],
| 99.00th=[ 6678], 99.50th=[ 6812], 99.90th=[ 7215], 99.95th=[ 7282],
| 99.99th=[ 7483]
bw ( KiB/s): min= 56, max= 848, per=100.00%, avg=364.32, stdev=33.56, samples=914
iops      : min= 14, max= 212, avg=91.08, stdev= 8.39, samples=914
lat (msec) : 500=0.01%, 750=0.05%, 1000=0.08%, 2000=0.85%, >=2000=99.01%
cpu       : usr=0.04%, sys=0.04%, ctx=13052, majf=0, minf=47
IO depths  : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.3%, 32=0.6%, >=64=98.8%
  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=0.1%
  issued rwts: total=10509,10929,0,0 short=0,0,0,0 dropped=0,0,0,0
  latency   : target=0, window=0, percentile=100.00%, depth=256
```

Run status group 0 (all jobs):

```
  READ: bw=348KiB/s (356kB/s), 348KiB/s-348KiB/s (356kB/s-356kB/s), io=41.1MiB (43.0MB), run=120812-
120812msec
  WRITE: bw=362KiB/s (371kB/s), 362KiB/s-362KiB/s (371kB/s-371kB/s), io=42.7MiB (44.8MB), run=120812-
120812msec
```

Disk stats (read/write):

```
sdc: ios=10532/10929, merge=0/0, ticks=4648821/2648107, in_queue=7296927, util=99.67%
```

Seagate Barracuda 7200.14

```
iops-test-job: (groupid=0, jobs=4): err= 0: pid=251590: Wed Jul 17 10:49:41 2024
```

```
read: IOPS=27, BW=111KiB/s (114kB/s)(13.1MiB/120986msec)
```

```
slat (usec): min=3, max=1103.8k, avg=72458.05, stdev=194381.62
```

```
clat (msec): min=488, max=20448, avg=17135.15, stdev=3593.05
```

```
lat (msec): min=1045, max=20537, avg=17207.61, stdev=3594.34
```

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

```
| 1.00th=[ 2333], 5.00th=[ 7215], 10.00th=[13758], 20.00th=[17113],
| 30.00th=[17113], 40.00th=[17113], 50.00th=[17113], 60.00th=[17113],
| 70.00th=[17113], 80.00th=[17113], 90.00th=[17113], 95.00th=[17113],
| 99.00th=[17113], 99.50th=[17113], 99.90th=[17113], 99.95th=[17113],
| 99.99th=[17113]
```

```
bw ( KiB/s): min= 32, max= 336, per=100.00%, avg=131.32, stdev=12.72, samples=697
iops      : min= 8, max= 84, avg=32.83, stdev= 3.18, samples=697
write: IOPS=29, BW=116KiB/s (119kB/s)(13.7MiB/120986msec); 0 zone resets
slat (usec): min=3, max=1236.0k, avg=67783.98, stdev=189096.67
clat (msec): min=471, max=22212, avg=16727.15, stdev=3523.66
lat (msec): min=1045, max=22508, avg=16794.94, stdev=3528.58
clat percentiles (msec):
| 1.00th=[ 2265], 5.00th=[ 7215], 10.00th=[13221], 20.00th=[16979],
| 30.00th=[17113], 40.00th=[17113], 50.00th=[17113], 60.00th=[17113],
| 70.00th=[17113], 80.00th=[17113], 90.00th=[17113], 95.00th=[17113],
| 99.00th=[17113], 99.50th=[17113], 99.90th=[17113], 99.95th=[17113],
| 99.99th=[17113]
bw ( KiB/s): min= 32, max= 368, per=100.00%, avg=135.84, stdev=13.41, samples=704
iops      : min= 8, max= 92, avg=33.96, stdev= 3.35, samples=704
lat (msec) : 500=0.03%, 1000=0.03%, 2000=0.82%, >=2000=99.13%
cpu        : usr=0.02%, sys=0.02%, ctx=3369, majf=0, minf=44
IO depths  : 1=0.1%, 2=0.1%, 4=0.2%, 8=0.5%, 16=0.9%, 32=1.9%, >=64=96.3%
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=99.9%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.1%
issued rwts: total=3358,3509,0,0 short=0,0,0,0 dropped=0,0,0,0
latency    : target=0, window=0, percentile=100.00%, depth=256
```

Run status group 0 (all jobs):

```
READ: bw=111KiB/s (114kB/s), 111KiB/s-111KiB/s (114kB/s-114kB/s), io=13.1MiB (13.8MB), run=120986-
120986msec
WRITE: bw=116KiB/s (119kB/s), 116KiB/s-116KiB/s (119kB/s-119kB/s), io=13.7MiB (14.4MB), run=120986-
120986msec
```

Disk stats (read/write):

```
sdd: ios=3373/3509, merge=0/0, ticks=4327686/2960324, in_queue=7288009, util=99.87%
```

Revision #4

Created 16 July 2024 19:40:57 by Ivan

Updated 4 March 2025 10:05:59 by Ivan