

Fio

Christian Külker

2023-02-09

Contents

1	Installing Fio	1
1.1	As A Package	1
1.2	From Source	2
2	Running Fio	3
3	History	7
4	Disclaimer of Warranty	8
5	Limitation of Liability	8

1 Installing Fio

1.1 As A Package

Without **OFED** on some systems, installation from package works.

```
aptitude install fio
```

CentOS (and Red Hat Enterprise Linux) have more limited main repositories than most distributions. If you haven't done so already, you will need to add the EPEL repository to CentOS/RHEL to get fio.** However, this may conflict if you are using **Infini-band** with **OFED**. Use installation from source for this.

```
yum install epel-release -y
yum install fio
```

1.2 From Source

For older setups, download the 2.0.14 `tar` version from 2013, as **RPM** may conflict with the **OFED** stack:

```
# As user:
tar xvf fio-2.0.14.tar
cd fio-2.0.14
./configure
make
# As root:
make install
install -m 755 -d /usr/local/bin
install fio fio_generate_plots /usr/local/bin
install -m 755 -d /usr/local/man/man1
install -m 644 fio.1 /usr/local/man/man1
install -m 644 fio_generate_plots.1 /usr/local/man/man1
```

Version 3.19 of 2020 already contains more scripts and man pages:

```
# As user:
wget https://github.com/axboe/fio/archive/fio-3.19.tar.gz
tar xvzf fio-3.19.tar.gz
cd fio-fio-3.19/
./configure
make
# As root:
make install
install -m 755 -d /usr/local/bin
install fio t/fio-genzipf t/fio-btrace2fio t/fio-dedupe t/fio-verify-state
↵ \
./tools/fio_generate_plots ./tools/plot/fio2gnuplot ./tools/genfio \
./tools/fiologparser.py ./tools/hist/fiologparser_hist.py \
./tools/fio_jsonplus_clat2csv /usr/local/bin
install -m 755 -d /usr/local/man/man1
install -m 644 ./fio.1 /usr/local/man/man1
install -m 644 ./tools/fio_generate_plots.1 /usr/local/man/man1
install -m 644 ./tools/plot/fio2gnuplot.1 /usr/local/man/man1
install -m 644 ./tools/hist/fiologparser_hist.py.1 /usr/local/man/man1
install -m 755 -d /usr/local/share/fio
install -m 644 ./tools/plot/*gpm /usr/local/share/fio/
```

The above `wget` URL has changed, see the <https://github.com> repository tags for the new URL. (2022-05-19) The latest version 3.30 as of 2022-05-19:

```
# As user:
cd /tmp
wget https://github.com/axboe/fio/archive/refs/tags/fio-3.30.tar.gz
tar xvzf fio-3.30.tar.gz
cd fio-fio-3.30
./configure
make
# As root:
make install
install -m 755 -d /usr/local/bin
install fio t/fio-genzipf t/fio-btrace2fio t/fio-verify-state \
./tools/fio_generate_plots ./tools/plot/fio2gnuplot ./tools/genfio \
./tools/fiologparser.py ./tools/hist/fiologparser_hist.py \
./tools/hist/fio-histo-log-pctiles.py ./tools/fio_jsonplus_clat2csv \
/usr/local/bin
install -m 755 -d /usr/local/man/man1
install -m 644 ./fio.1 /usr/local/man/man1
install -m 644 ./tools/fio_generate_plots.1 /usr/local/man/man1
install -m 644 ./tools/plot/fio2gnuplot.1 /usr/local/man/man1
install -m 644 ./tools/hist/fiologparser_hist.py.1 /usr/local/man/man1
install -m 755 -d /usr/local/share/fio
install -m 644 ./tools/plot/*gpm /usr/local/share/fio/
```

Read the [documentation](#) for details

2 Running Fio

```
fio --name=random-write --ioengine=posixaio --rw=randwrite --bs=4k \
--numjobs=1 --size=4g --iodepth=1 --runtime=60 --time_based --end_fsync=1
```

The has a huge output, but most important is the write speed

```
1 [...]
2 Run status group 0 (all jobs):
3   WRITE: bw=127MiB/s (133MB/s), 127MiB/s-127MiB/s (133MB/s-133MB/s), \
4   io=8192MiB (8590MB), run=64602-64602msec
5 [...]
```

Reminder: MiB = Mebibyte (basis 1014), MB = Megabyte (basis 1000)

As we can see here the disk was fast, we requested 4g and it was able to run through its

approximately 2 times in 60 seconds. The most interesting value is **bw=127MiB/s**.

If you have the number of CPU cores and the bandwidth to the storage subsystem, run more jobs in parallel. Then repeat with different block sizes, such as 1m.

In contrast the version 3.19 from 2020 on a laptop from 2008:

```
fio --name=random-write --ioengine=posixaio --rw=randwrite --bs=4k \
--numjobs=1 --size=4g --iodepth=1 --runtime=60 --time_based --end_fsync=1
random-write: (g=0): rw=randwrite, bs=(R) 4096B-4096B, (W) 4096B-4096B,
↳ (T) \
4096B-4096B, ioengine=posixaio, iodepth=1
fio-3.19
Starting 1 process
random-write: Laying out IO file (1 file / 4096MiB)
Jobs: 1 (f=1): [F(1)][100.0%][w=23.4MiB/s][w=5992 IOPS][eta 00m:00s]
random-write: (groupid=0, jobs=1): err= 0: pid=2427: Tue May 19 11:55:53
↳ 2020
write: IOPS=10.1k, BW=39.3MiB/s (41.2MB/s)(2411MiB/61288msec); 0 zone
↳ resets
slat (usec): min=2, max=118728, avg=14.89, stdev=369.70
clat (nsec): min=1327, max=245121k, avg=71777.23, stdev=1507737.18
lat (usec): min=19, max=250225, avg=86.66, stdev=1556.20
clat percentiles (usec):
| 1.00th=[ 3], 5.00th=[ 3], 10.00th=[ 3], 20.00th=[ 23],
| 30.00th=[ 26], 40.00th=[ 29], 50.00th=[ 30], 60.00th=[ 30],
| 70.00th=[ 30], 80.00th=[ 37], 90.00th=[ 63], 95.00th=[ 91],
| 99.00th=[ 326], 99.50th=[ 799], 99.90th=[ 5866], 99.95th=[ 8717],
| 99.99th=[82314]
bw ( KiB/s): min=16406, max=66802, per=100.00%, avg=41176.64, \
stdev=8856.09, samples=118
iops      : min= 4101, max=16700, avg=10293.88, stdev=2214.04,
↳ samples=118
lat (usec) : 2=0.09%, 4=11.76%, 10=0.01%, 20=0.28%, 50=74.46%
lat (usec) : 100=8.98%, 250=3.08%, 500=0.64%, 750=0.18%, 1000=0.09%
lat (msec) : 2=0.16%, 4=0.12%, 10=0.12%, 20=0.02%, 50=0.01%
lat (msec) : 100=0.01%, 250=0.01%
cpu       : usr=10.78%, sys=15.06%, ctx=775939, majf=5, minf=29
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=0,617184,0,0 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=39.3MiB/s (41.2MB/s), 39.3MiB/s-39.3MiB/s (41.2MB/s-41.2MB/s), \
    ↪ \
io=2411MiB (2528MB), run=61288-61288msec

Disk stats (read/write):
  dm-1: ios=2381/631516, merge=0/0, ticks=9376/82883752,
    ↪ in_queue=82929236, \
util=80.66%, aggrrios=3246/655729, aggrmerge=0/0, aggrticks=10692/87229256,
    ↪ \
aggrin_queue=87239968, aggrutil=81.29%
  dm-0: ios=3246/655729, merge=0/0, ticks=10692/87229256,
    ↪ in_queue=87239968,\
util=81.29%, aggrrios=3235/653548, aggrmerge=11/2181,
    ↪ aggrticks=1945/108056, \
aggrin_queue=62344, aggrutil=31.25%
  sda: ios=3235/653548, merge=11/2181, ticks=1945/108056, in_queue=62344,
    ↪ \
util=31.25%

```

And again the line with the most interest:

```

1  WRITE: bw=39.3MiB/s (41.2MB/s), 39.3MiB/s-39.3MiB/s (41.2MB/s-41.2MB/s), \
2  io=2411MiB (2528MB), run=61288-61288msec

```

The laptop is much slower with **bw=39.3MiB/s** for 1 job

```

1  WRITE: bw=52.6MiB/s (55.2MB/s), 24.6MiB/s-28.0MiB/s (25.8MB/s-29.4MB/s), \
2  io=3232MiB (3389MB), run=61428-61428msec

```

And for 2 jobs **bw=52.6MiB/s**. Comparing the laptop from 2008 with a desktop from 2013, it holds quite well:

```

1  WRITE: bw=26.8MiB/s (28.1MB/s), 26.8MiB/s-26.8MiB/s (28.1MB/s-28.1MB/s), \
2  io=2581MiB (2706MB), run=96397-96397msec

```

Since the desktop freezes when `fio` is run on a core and performance is poor, this is in

some cases a matter of system design, not age.

In contrast the version 3.30 from 2022 on a laptop from 2008:

```
fio --name=random-write --ioengine=posixaio --rw=randwrite --bs=4k \
--numjobs=1 --size=4g --iodepth=1 --runtime=60 --time_based --end_fsync=1
random-write: (g=0): rw=randwrite, bs=(R) 4096B-4096B, (W) 4096B-4096B,
↳ (T) \
4096B-4096B, ioengine=posixaio, iodepth=1
fio-3.30
Starting 1 process
random-write: Laying out IO file (1 file / 4096MiB)
Jobs: 1 (f=1): [F(1)][100.0%][eta 00m:00s]
random-write: (groupid=0, jobs=1): err= 0: pid=28278: Thu May 19 15:19:14
↳ 2022
write: IOPS=3861, BW=15.1MiB/s (15.8MB/s)(1183MiB/78401msec); 0 zone
↳ resets
slat (usec): min=3, max=9870, avg= 7.60, stdev=59.69
clat (nsec): min=1257, max=21055M, avg=185432.51, stdev=39838405.45
lat (usec): min=22, max=21055k, avg=193.03, stdev=39838.46
clat percentiles (nsec):
| 1.00th=[ 1816], 5.00th=[ 20608], 10.00th=[ 21632],
| 20.00th=[ 22912], 30.00th=[ 24192], 40.00th=[ 24960],
| 50.00th=[ 25984], 60.00th=[ 27264], 70.00th=[ 28544],
| 80.00th=[ 30592], 90.00th=[ 34048], 95.00th=[ 64256],
| 99.00th=[ 978944], 99.50th=[ 5799936], 99.90th=[ 7634944],
| 99.95th=[ 8355840], 99.99th=[53215232]
bw ( KiB/s): min= 80, max=106970, per=100.00%, avg=36142.76, \
stdev=27917.69, samples=67
iops      : min= 20, max=26742, avg=9035.67, stdev=6979.41,
↳ samples=67
lat (usec) : 2=1.11%, 4=0.02%, 10=0.01%, 20=0.93%, 50=91.75%
lat (usec) : 100=2.75%, 250=1.98%, 500=0.28%, 750=0.10%, 1000=0.08%
lat (msec) : 2=0.22%, 4=0.20%, 10=0.55%, 20=0.02%, 50=0.01%
lat (msec) : 100=0.01%, 250=0.01%, 500=0.01%, 2000=0.01%, >=2000=0.01%
cpu        : usr=5.23%, sys=2.26%, ctx=380587, majf=0, minf=18
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=0,302771,0,0 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=15.1MiB/s (15.8MB/s), 15.1MiB/s-15.1MiB/s (15.8MB/s-15.8MB/s), \
    ↪ \
io=1183MiB (1240MB), run=78401-78401msec

Disk stats (read/write):
  dm-1: ios=88/278258, merge=0/0, ticks=325796/2519395188, \
in_queue=2521065440, util=97.99%, aggrios=91/279033, aggrmerge=0/0, \
aggrticks=325808/2519396324, aggrin_queue=2520489288, aggrutil=97.98%
  dm-0: ios=91/279033, merge=0/0, ticks=325808/2519396324, \
in_queue=2520489288, util=97.98%, aggrios=90/276897, aggrmerge=0/2133, \
aggrticks=1604/10742808, aggrin_queue=10744248, aggrutil=95.21%
  sda: ios=90/276897, merge=0/2133, ticks=1604/10742808,
    ↪ in_queue=10744248, \
util=95.21%
```

Compare `fio 3.19` with `fio 3.30` (2 years later) on the same 2008 hardware:

```
1 # fio 3.19 (redone 2022-05-19)
2 WRITE: bw=12.3MiB/s (12.9MB/s), 12.3MiB/s-12.3MiB/s (12.9MB/s-12.9MB/s), \
3 io=926MiB (971MB), run=75132-75132msec
4
5 # fio 3.30 (done 2022-05-19)
6 WRITE: bw=15.1MiB/s (15.8MB/s), 15.1MiB/s-15.1MiB/s (15.8MB/s-15.8MB/s), \
7 io=1183MiB (1240MB), run=78401-78401msec
```

So it seems that the newer version is faster. This has to be taken with a pinch of salt. However, it is recommended to average over several runs and stick to the same version of `fio`.

3 History

Version	Date	Notes
0.1.2	2023-02-09	Improve writing
0.1.1	2022-05-19	Change code blocks, +history, fio-3.30, comparison
0.1.0	2020-05-19	Initial release

4 Disclaimer of Warranty

THERE IS NO WARRANTY FOR THIS INFORMATION, DOCUMENTS AND PROGRAMS, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE INFORMATION, DOCUMENT OR THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE INFORMATION, DOCUMENTS AND PROGRAMS IS WITH YOU. SHOULD THE INFORMATION, DOCUMENTS OR PROGRAMS PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

5 Limitation of Liability

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS THE INFORMATION, DOCUMENTS OR PROGRAMS AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE INFORMATION, DOCUMENTS OR PROGRAMS (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE INFORMATION, DOCUMENTS OR PROGRAMS TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.