www.phoronix-test-suite.com

ZFS SSD Benchmark: RAW IMAGE vs QCOW2 vs ZVOL for KVM

This benchmark showâ€™s the performance of a zfs pool
providing storage to a kvm virtual machine with three different formats:

raw image files on plain dataset
qcow2 image file on plain dataset
zvol

For each of those several filesystem benchmarks from the phoronix-benchmark-suite are run with two different recordsizes:

8k
64k

Host Details:

cpu: IntelÂ® XeonÂ® CPU E5320 @ 1.86GHz
ram: 6GB DDR2
sata: Intel Corporation 631xESB/632xESB SATA AHCI Controller (rev 09
os: Debian GNU/Linux 9.4 (stretch), fully upgraded as of 2018-03-15
zfs: 0.7.6-1~bpo9+1 (from stretch-backports)
ssd: 2x samsung 840 pro 256GB

KVM machine:

cpu: host
ram: 1024MB
disk: bus=scsi,cache=writeback,io=threads
os: Debian GNU/Linux 9.1 (stretch)
fs: ext4

The zfs pool used consists of a single mirrored vdev with samsung 840 pro ssdâ€™s.
Each benchmark is run like this:

drop all caches
trim whole ssdâ€™s
create zfs pool (ashift=13,compression=lz4,atime=on,relatime=on
benchmark,devices=off,exec=off,setuid=off,xattr=off)
create disk:

qemu-img create -f raw debian9.raw 50G
qemu-img create -f qcow2 -o cluster_size=8k,preallocation=metadata,compat=1.1,lazy_refcounts=on debian9.qcow2 50G
zfs create -o volblocksize=8k -V 50G benchmark/kvm/debian9

create kvm machine
take timestamp
let debian9 install automatically
save install time
install phoronix-test-suite and needed dependencies in the vm
run the pts
trim the filesystem and save the time it took
copy the results to the kvm host
destroy vm and zpool

other test results:

8k_raw - install time: 619s, trim: 39.401s
64k raw - install time: 589s, trim: 0m9.593s
8k qcow2 - install time: 615s, trim: 1m45.421s
64k qcow2 - install time: 590s, trim: 0m7.598s
8k zvol - install time: 621s, trim: 0m11.353s
64k zvol - install time: 605s

Automated Executive Summary

zfs ssd mirror, ashift=13, recordsize 64k, raw image had the most wins, coming in first place for 67% of the tests.

Based on the geometric mean of all complete results, the fastest (zfs ssd mirror, ashift=13, recordsize 64k, raw image) was 1.662x the speed of the slowest (zfs ssd mirror, ashift=13, recordsize 8k, zvol). zfs ssd mirror, ashift=13, recordsize

64k, qcow2 image was 0.983x the speed of zfs ssd mirror, ashift=13, recordsize 64k, raw image, zfs ssd mirror, ashift=13, recordsize 8k, raw image was 0.688x the speed of zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image, zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image was 0.968x the speed of zfs ssd mirror, ashift=13, recordsize 8k, raw image, zfs ssd mirror, ashift=13, recordsize 64k, zvol was 0.989x the speed of zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image, zfs ssd mirror, ashift=13, recordsize 8k, zvol was 0.929x the speed of zfs ssd mirror, ashift=13, recordsize 64k, zvol.

The results with the greatest spread from best to worst included:

Flexible IO Tester (Type: Sequential Write - IO Engine: Linux AIO - Buffered: No - Direct: Yes - Block Size: 64KB - Disk Target: Default Test Directory) at 3.816x
Flexible IO Tester (Type: Random Write - IO Engine: Linux AIO - Buffered: No - Direct: Yes - Block Size: 1MB - Disk Target: Default Test Directory) at 3.472x
Flexible IO Tester (Type: Sequential Write - IO Engine: Linux AIO - Buffered: No - Direct: Yes - Block Size: 1MB - Disk Target: Default Test Directory) at 3.353x
Flexible IO Tester (Type: Random Write - IO Engine: Linux AIO - Buffered: No - Direct: Yes - Block Size: 64KB - Disk Target: Default Test Directory) at 3.045x
IOzone (Record Size: 64Kb - File Size: 8GB - Disk Test: Write Performance) at 2.877x
IOzone (Record Size: 4Kb - File Size: 8GB - Disk Test: Write Performance) at 2.843x
IOzone (Record Size: 1MB - File Size: 8GB - Disk Test: Read Performance) at 2.691x
IOzone (Record Size: 64Kb - File Size: 8GB - Disk Test: Read Performance) at 2.668x
IOzone (Record Size: 4Kb - File Size: 8GB - Disk Test: Read Performance) at 2.646x
IOzone (Record Size: 1MB - File Size: 8GB - Disk Test: Write Performance) at 2.63x.

Test Systems:

zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image

zfs ssd mirror, ashift=13, recordsize 64k, raw image

zfs ssd mirror, ashift=13, recordsize 64k, zvol

zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image

zfs ssd mirror, ashift=13, recordsize 8k, raw image

zfs ssd mirror, ashift=13, recordsize 8k, zvol

Processor: 8 x Intel Celeron_4x0 (Conroe/Merom Class Core 2) @ 1.86GHz (8 Cores), Motherboard: QEMU Standard PC (i440FX + PIIX 1996), Chipset: Intel 440FX- 82441FX PMC, Memory: 1 x 1024 MB RAM QEMU, Disk: 54GB QEMU HDD, Graphics: Cirrus Logic GD 5446, Network: Realtek RTL-8100/8101L/8139

OS: Debian 9.1, Kernel: 4.9.0-4-amd64 (x86_64), Compiler: GCC 6.3.0 20170516, File-System: ext4, Screen Resolution: 1024x768, System Layer: qemu

ZFS SSD Benchmark: RAW IMAGE vs QCOW2 vs ZVOL for KVM

Compiler Notes: --build=x86_64-linux-gnu --disable-browser-plugin --disable-vtable-verify --enable-checking=release --enable-clocale=gnu --enable-default-pie --enable-gnu-unique-object --enable-gtk-cairo --enable-java-awt=gtk --enable-java-home --enable-languages=c,ada,c++,java,go,d,fortran,objc,obj-c++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-plugin --enable-shared --enable-threads=posix --host=x86_64-linux-gnu --program-prefix=x86_64-linux-gnu- --target=x86_64-linux-gnu --with-abi=m64 --with-arch-32=i686 --with-arch-directory=amd64 --with-default-libstdcxx-abi=new --with-multilib-list=m32,m64,mx32 --with-target-system-zlib --with-tune=generic -v
Disk Notes: CFQ / data=ordered,errors=remount-ro,relatime,rw
System Notes: Python 2.7.13.

	zfs ssd mirror, ashift=13, recordsize 64k, qcow2	zfs ssd mirror, ashift=13, recordsize 64k, raw	zfs ssd mirror, ashift=13, recordsize 64k, zvol	zfs ssd mirror, ashift=13, recordsize 8k, qcow2	zfs ssd mirror, ashift=13, recordsize 8k, raw image	zfs ssd mirror, ashift=13, recordsize 8k, zvol
AIO-Stress - Rand Write (MB/s)	125.74	**131.81**	111.33	**93.70**	98.01	100.87
Normalized	95.39%	100%	84.46%	71.09%	74.36%	76.53%
Standard Deviation	1.1%	1.2%	0.8%	2.2%	0.9%	7.3%
Flexible IO Tester - Rand Read - Linux AIO - No - Yes - 1MB (MB/s)	292	294	454	**228**	230	**473**
Normalized	61.73%	62.16%	95.98%	48.2%	48.63%	100%
Standard Deviation	0.7%	1.2%	5.6%	0.8%	0.8%	3.9%
Flexible IO Tester - Rand Read - Linux AIO - No - Yes - 8KB (MB/s)	38.60	38.93	**40.73**	40.65	**36.80**	39.20
Normalized	94.77%	95.58%	100%	99.8%	90.35%	96.24%
Standard Deviation	0.4%	0.1%	2.2%	3.2%	1.9%	1.7%
Flexible IO Tester - Rand Read - Linux AIO - No - Yes - 64KB (MB/s)	237	**239**	230	232	**229**	230
Normalized	99.16%	100%	96.23%	97.07%	95.82%	96.23%
Standard Deviation		1.1%	0.4%	1.9%		0.7%
Flexible IO Tester - Rand Write - Linux AIO - No - Yes - 1MB (MB/s)	143	**146**	54.75	70.23	71.20	**42.05**
Normalized	97.95%	100%	37.5%	48.1%	48.77%	28.8%
Standard Deviation	0.8%	0.7%	9.4%	0.2%	0.7%	6.3%
Flexible IO Tester - Rand Write - Linux AIO - No - Yes - 8KB (MB/s)	15.97	17.20	**12.05**	25.73	**29.20**	20.70
Normalized	54.69%	58.9%	41.27%	88.12%	100%	70.89%
Standard Deviation	2.6%	2.7%	3.5%	1.2%	1.2%	0.8%
Flexible IO Tester - Rand Write - Linux AIO - No - Yes - 64KB (MB/s)	113	**118**	66.40	59.40	61.33	**38.75**
Normalized	95.76%	100%	56.27%	50.34%	51.97%	32.84%
Standard Deviation	1%	1.3%	2.1%	0.7%	0.1%	20.9%

ZFS SSD Benchmark: RAW IMAGE vs QCOW2 vs ZVOL for KVM

Flexible IO Tester - Seq Read - Linux AIO - No - Yes - 1MB (MB/s)	291	289	458	226	229	472
Normalized	61.65%	61.23%	97.03%	47.88%	48.52%	100%
Standard Deviation	0.7%	0.4%	2.1%	0.5%		5.4%
Flexible IO Tester - Seq Read - Linux AIO - No - Yes - 8KB (MB/s)	54.17	55.50	44.73	45.85	58.32	45.02
Normalized	92.88%	95.16%	76.7%	78.62%	100%	77.19%
Standard Deviation	6%	0.7%	1.2%	8.1%	8.9%	6%
Flexible IO Tester - Seq Read - Linux AIO - No - Yes - 64KB (MB/s)	306	303	347	244	252	353
Normalized	86.69%	85.84%	98.3%	69.12%	71.39%	100%
Standard Deviation	0.8%	1%	5.3%	0.4%	0.9%	2.2%
Flexible IO Tester - Seq Write - Linux AIO - No - Yes - 1MB (MB/s)	147	152	66.28	72.03	74.37	45.33
Normalized	96.71%	100%	43.61%	47.39%	48.93%	29.82%
Standard Deviation	2.8%		7.2%	2.6%	2.2%	34.7%
Flexible IO Tester - Seq Write - Linux AIO - No - Yes - 8KB (MB/s)	46.30	47.10	29.77	32.70	41.83	24.15
Normalized	98.3%	100%	63.21%	69.43%	88.81%	51.27%
Standard Deviation	9.1%	9.4%	5.3%	13.5%	10.4%	5.8%
Flexible IO Tester - Seq Write - Linux AIO - No - Yes - 64KB (MB/s)	139	138	71.37	70.77	74.27	36.43
Normalized	100%	99.28%	51.35%	50.91%	53.43%	26.21%
Standard Deviation	0.8%		1.4%	0.1%	1%	22.3%
Dbench - 1 (MB/s)	42.33	42.60	42.72	42.20	42.18	40.87
Normalized	99.09%	99.72%	100%	98.78%	98.74%	95.67%
Standard Deviation	0.3%	0.4%	0.1%	0.1%	0.2%	0.6%
Dbench - 6 (MB/s)	102.98	105.25	88.54	89.99	91.96	86.79
Normalized	97.84%	100%	84.12%	85.5%	87.37%	82.46%
Standard Deviation	1.5%	1.8%	2.5%	1.7%	2.6%	0.3%
Dbench - 12 (MB/s)	162.59	166.12	133.55	137.50	137.66	128.76
Normalized	97.88%	100%	80.39%	82.77%	82.87%	77.51%
Standard Deviation	0.8%	0.6%	0.5%	1.9%	1.9%	1.1%
Dbench - 48 (MB/s)	294.03	292.22	177.89	176.41	181.38	153.63
Normalized	100%	99.38%	60.5%	60%	61.69%	52.25%
Standard Deviation	2.6%	2.5%	8.9%	9.8%	9.6%	10.9%
Dbench - 128 (MB/s)	203.76	207.32	102.07	89.75	97.79	88.41
Normalized	98.28%	100%	49.23%	43.29%	47.17%	42.64%
Standard Deviation	0.8%	1.7%	11.9%	15.7%	14.6%	12.5%
Dbench - 256 (MB/s)	118.70	120.18	62.00	55.57	55.77	58.47
Normalized	98.77%	100%	51.59%	46.24%	46.41%	48.65%
Standard Deviation	0.3%	0.6%	7.9%	2.6%	9.3%	5.8%
IOzone - 1MB - 8GB - Read Performance (MB/s)	215.34	216.63	86.07	165.90	167.07	80.50
Normalized	99.4%	100%	39.73%	76.58%	77.12%	37.16%
Standard Deviation	0.6%	0.5%	1.2%	1.3%	1.3%	2.7%

ZFS SSD Benchmark: RAW IMAGE vs QCOW2 vs ZVOL for KVM

IOzone - 4Kb - 8GB - Read Performance (MB/s)	223.31	**224.74**	94.99	173.68	174.63	**84.95**
Normalized	99.36%	100%	42.27%	77.28%	77.7%	37.8%
Standard Deviation	0.7%	1.6%	2.3%	0.1%	0.8%	6.3%
IOzone - 1MB - 8GB - Write Performance (MB/s)	156.30	**157.79**	66.79	65.22	**60.00**	72.07
Normalized	99.06%	100%	42.33%	41.33%	38.03%	45.67%
Standard Deviation	0.9%	0.9%	1.3%	7.6%	5.8%	6.9%
IOzone - 4Kb - 8GB - Write Performance (MB/s)	161.59	**162.04**	63.10	61.43	**57.00**	74.06
Normalized	99.72%	100%	38.94%	37.91%	35.18%	45.7%
Standard Deviation	1.9%	1.1%	3%	4.1%	6.9%	2.6%
IOzone - 64Kb - 8GB - Read Performance (MB/s)	225.08	**226.37**	88.90	174.56	172.38	**84.86**
Normalized	99.43%	100%	39.27%	77.11%	76.15%	37.49%
Standard Deviation	0.6%	0.3%	3.1%	0.5%	0.5%	6%
IOzone - 64Kb - 8GB - Write Performance (MB/s)	156.04	**160.28**	60.03	58.87	**55.72**	67.62
Normalized	97.35%	100%	37.45%	36.73%	34.76%	42.19%
Standard Deviation	4.3%	4.5%	7.2%	3.8%	1.2%	11.4%
Compile Bench - Compile (MB/s)	142.88	**147.18**	81.59	**61.65**	72.13	91.69
Normalized	97.08%	100%	55.44%	41.89%	49.01%	62.3%
Standard Deviation	8.6%	8.7%	3%	19.8%	7.2%	7.3%
Compile Bench - Initial Create (MB/s)	82.40	**86.85**	50.11	**42.54**	48.40	49.32
Normalized	94.88%	100%	57.7%	48.98%	55.73%	56.79%
Standard Deviation	13.5%	6.1%	20.6%	33.9%	6.8%	12.6%
Compile Bench - Read Compiled Tree (MB/s)	64.69	**64.83**	60.72	56.72	55.72	**46.58**
Normalized	99.78%	100%	93.66%	87.49%	85.95%	71.85%
Standard Deviation	0.4%	0.9%	7.2%	1.6%	0.5%	2.4%

AIO-Stress 0.21

Test: Random Write

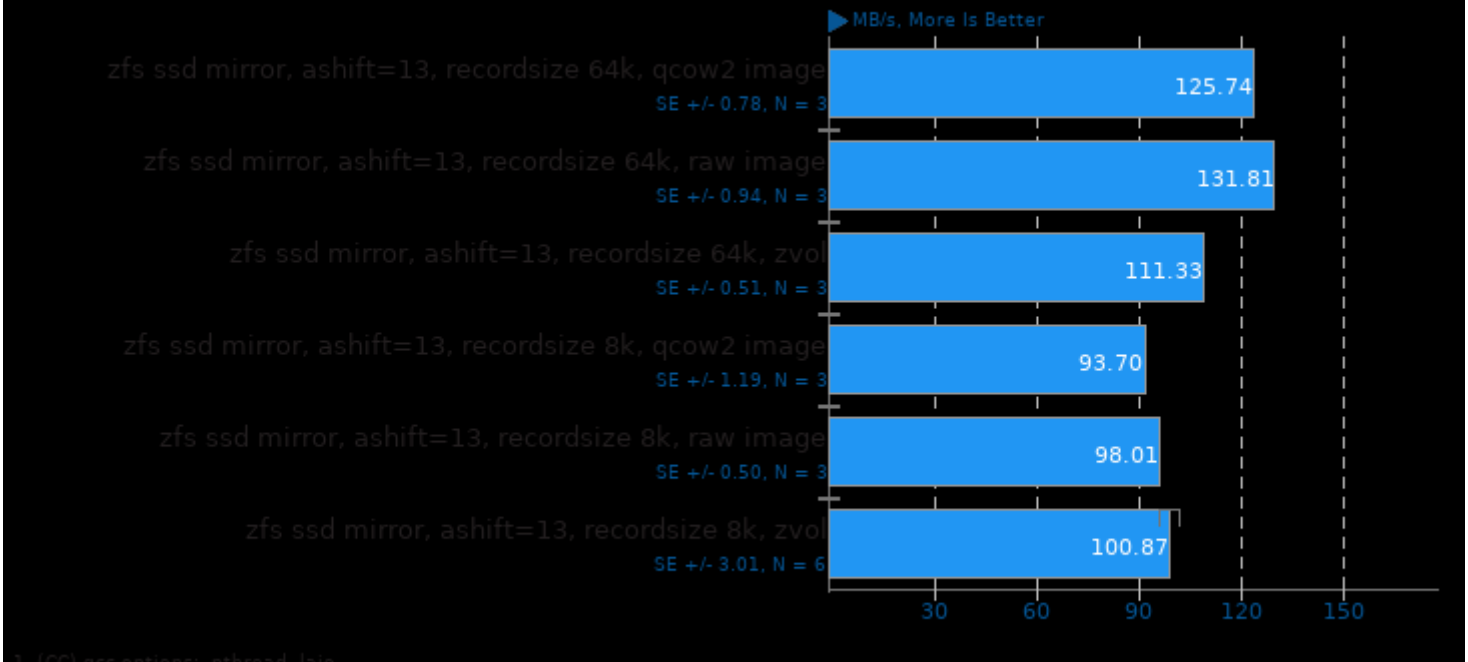

1. (CC) gcc options: -pthread -laio

Flexible IO Tester 3.1

Type: Random Read - IO Engine: Linux AIO - Buffered: No - Direct: Yes - Block Size: 1MB - Disk Target: Default Test Directory

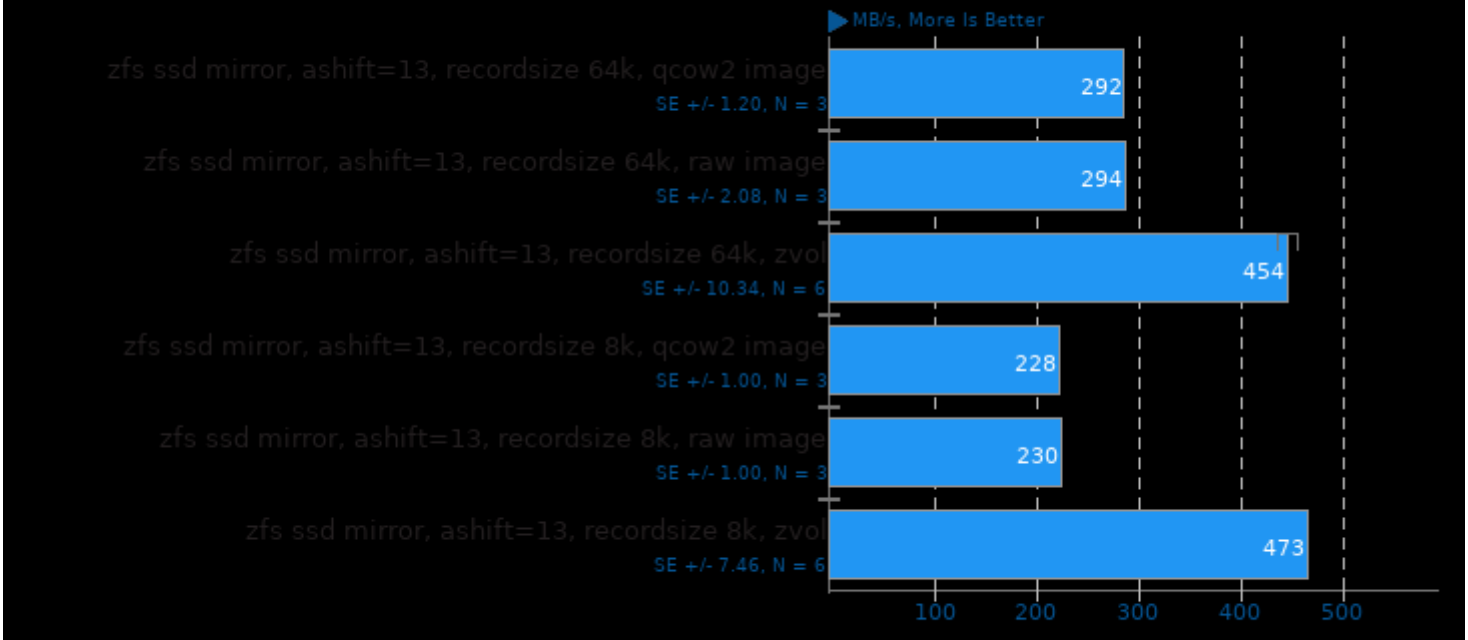

1. (CC) gcc options: -rdynamic -std=gnu99 -ffast-math -include -O3 -U_FORTIFY_SOURCE -lrt -laio -lm -lpthread -ldl

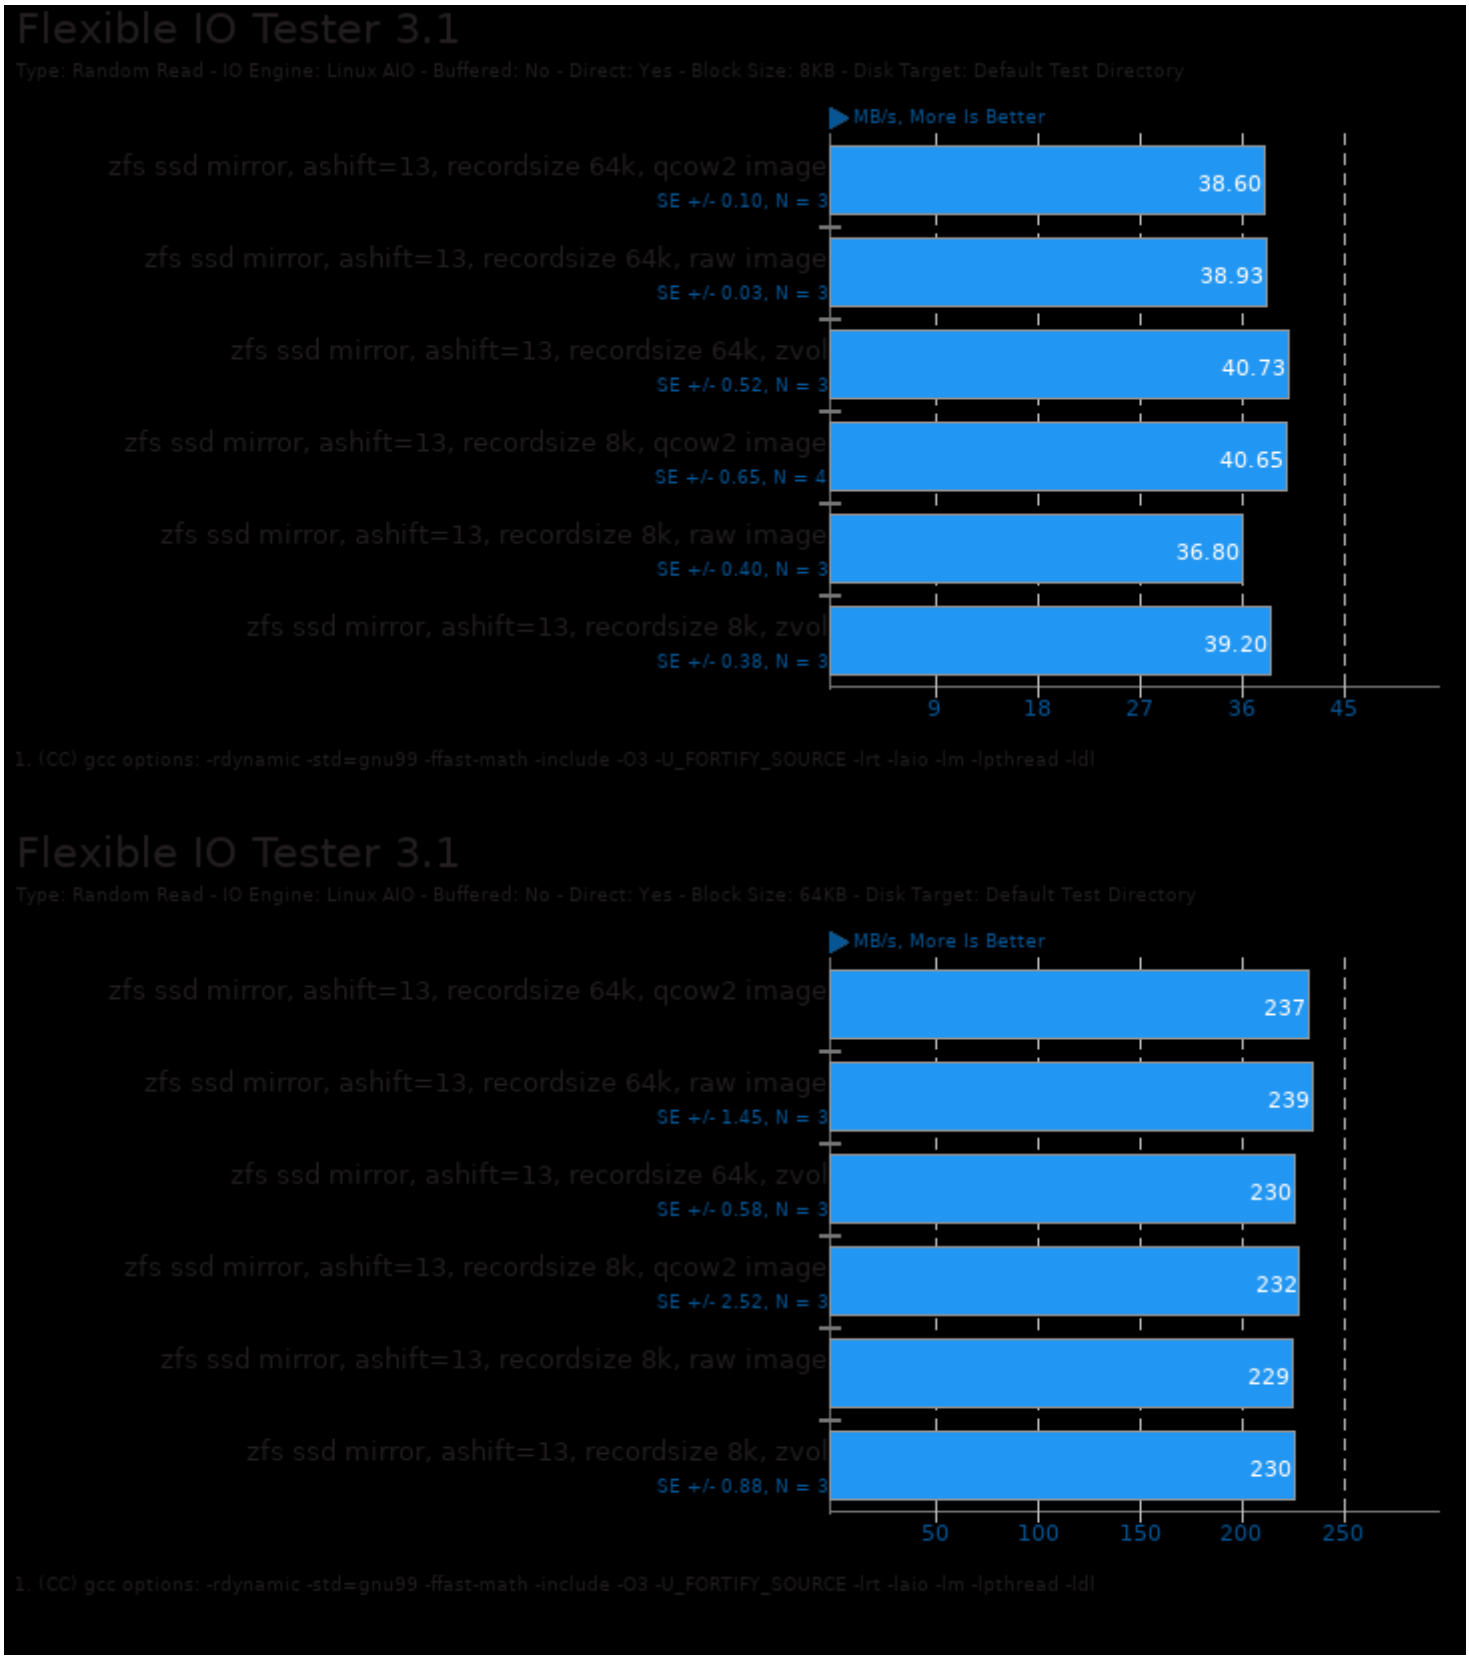
Flexible IO Tester 3.1
Type: Random Read - IO Engine: Linux AIO - Buffered: No - Direct: Yes - Block Size: 8KB - Disk Target: Default Test Directory
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 0.10, N = 3
38.60
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 0.03, N = 3
38.93
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 0.52, N = 3
40.73
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 0.65, N = 4
40.65
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 0.40, N = 3
36.80
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 0.38, N = 3
39.20
9
18
27
36
45
1. (CC) gcc options: -rdynamic -std=gnu99 -ffast-math -include -O3 -U_FORTIFY_SOURCE -lrt -laio -lm -lpthread -ldl
Flexible IO Tester 3.1
Type: Random Read - IO Engine: Linux AIO - Buffered: No - Direct: Yes - Block Size: 64KB - Disk Target: Default Test Directory
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
237
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 1.45, N = 3
239
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 0.58, N = 3
230
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 2.52, N = 3
232
zfs ssd mirror, ashift=13, recordsize 8k, raw image
229
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 0.88, N = 3
230
50
100
150
200
250
1. (CC) gcc options: -rdynamic -std=gnu99 -ffast-math -include -O3 -U_FORTIFY_SOURCE -lrt -laio -lm -lpthread -ldl

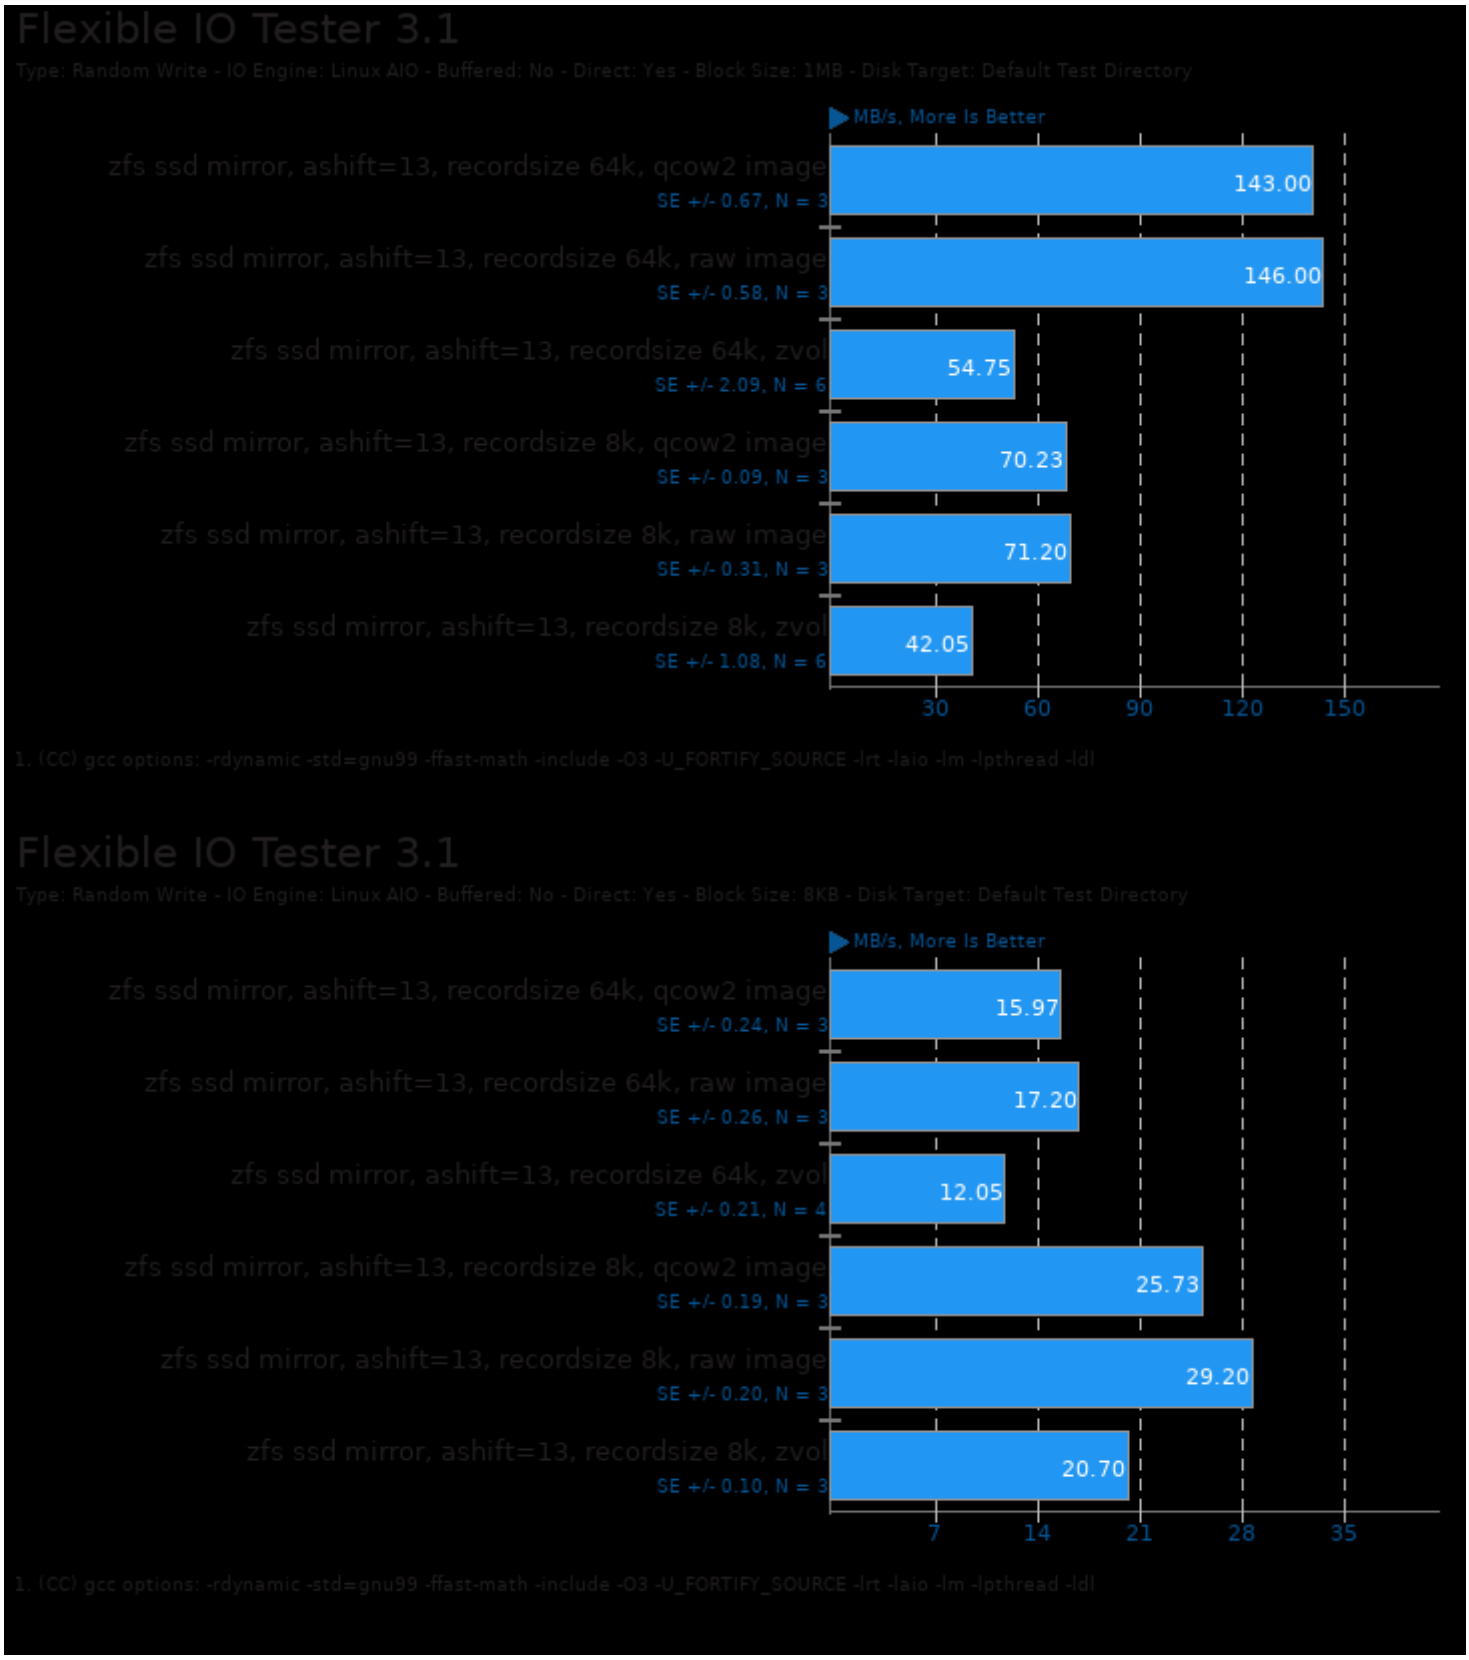
Flexible IO Tester 3.1
Type: Random Write - IO Engine: Linux AIO - Buffered: No - Direct: Yes - Block Size: 1MB - Disk Target: Default Test Directory
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 0.67, N = 3
143.00
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 0.58, N = 3
146.00
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 2.09, N = 6
54.75
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 0.09, N = 3
70.23
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 0.31, N = 3
71.20
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 1.08, N = 6
42.05
30
60
90
120
150
1. (CC) gcc options: -rdynamic -std=gnu99 -ffast-math -include -O3 -U_FORTIFY_SOURCE -lrt -laio -lm -lpthread -ldl
Flexible IO Tester 3.1
Type: Random Write - IO Engine: Linux AIO - Buffered: No - Direct: Yes - Block Size: 8KB - Disk Target: Default Test Directory
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 0.24, N = 3
15.97
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 0.26, N = 3
17.20
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 0.21, N = 4
12.05
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 0.19, N = 3
25.73
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 0.20, N = 3
29.20
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 0.10, N = 3
20.70
7
14
21
28
35
1. (CC) gcc options: -rdynamic -std=gnu99 -ffast-math -include -O3 -U_FORTIFY_SOURCE -lrt -laio -lm -lpthread -ldl

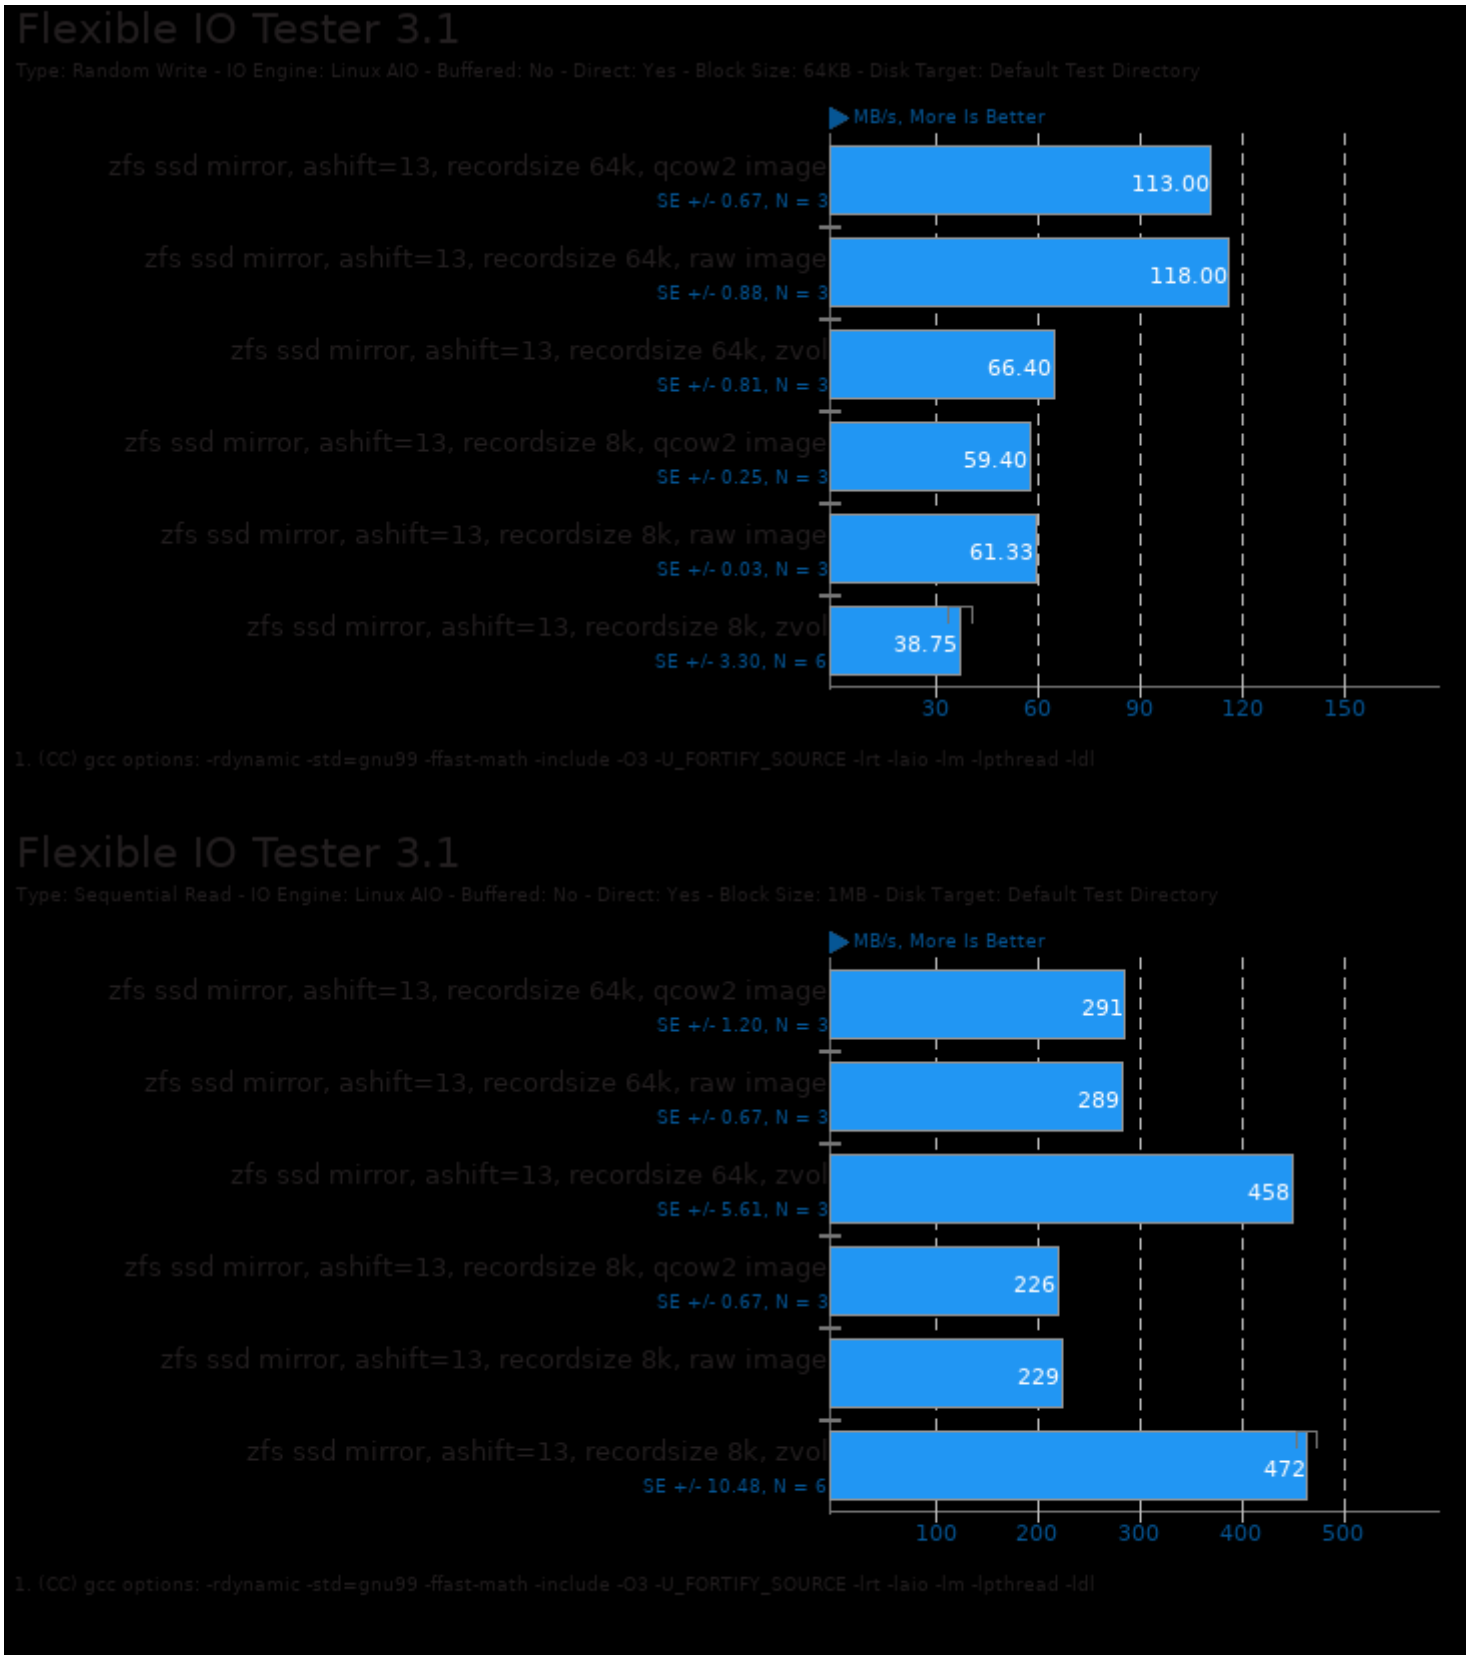
Flexible IO Tester 3.1
Type: Random Write - IO Engine: Linux AIO - Buffered: No - Direct: Yes - Block Size: 64KB - Disk Target: Default Test Directory
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 0.67, N = 3
113.00
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 0.88, N = 3
118.00
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 0.81, N = 3
66.40
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 0.25, N = 3
59.40
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 0.03, N = 3
61.33
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 3.30, N = 6
38.75
30
60
90
120
150
1. (CC) gcc options: -rdynamic -std=gnu99 -ffast-math -include -O3 -U_FORTIFY_SOURCE -lrt -laio -lm -lpthread -ldl
Flexible IO Tester 3.1
Type: Sequential Read - IO Engine: Linux AIO - Buffered: No - Direct: Yes - Block Size: 1MB - Disk Target: Default Test Directory
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 1.20, N = 3
291
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 0.67, N = 3
289
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 5.61, N = 3
458
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 0.67, N = 3
226
zfs ssd mirror, ashift=13, recordsize 8k, raw image
229
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 10.48, N = 6
472
100
200
300
400
500
1. (CC) gcc options: -rdynamic -std=gnu99 -ffast-math -include -O3 -U_FORTIFY_SOURCE -lrt -laio -lm -lpthread -ldl

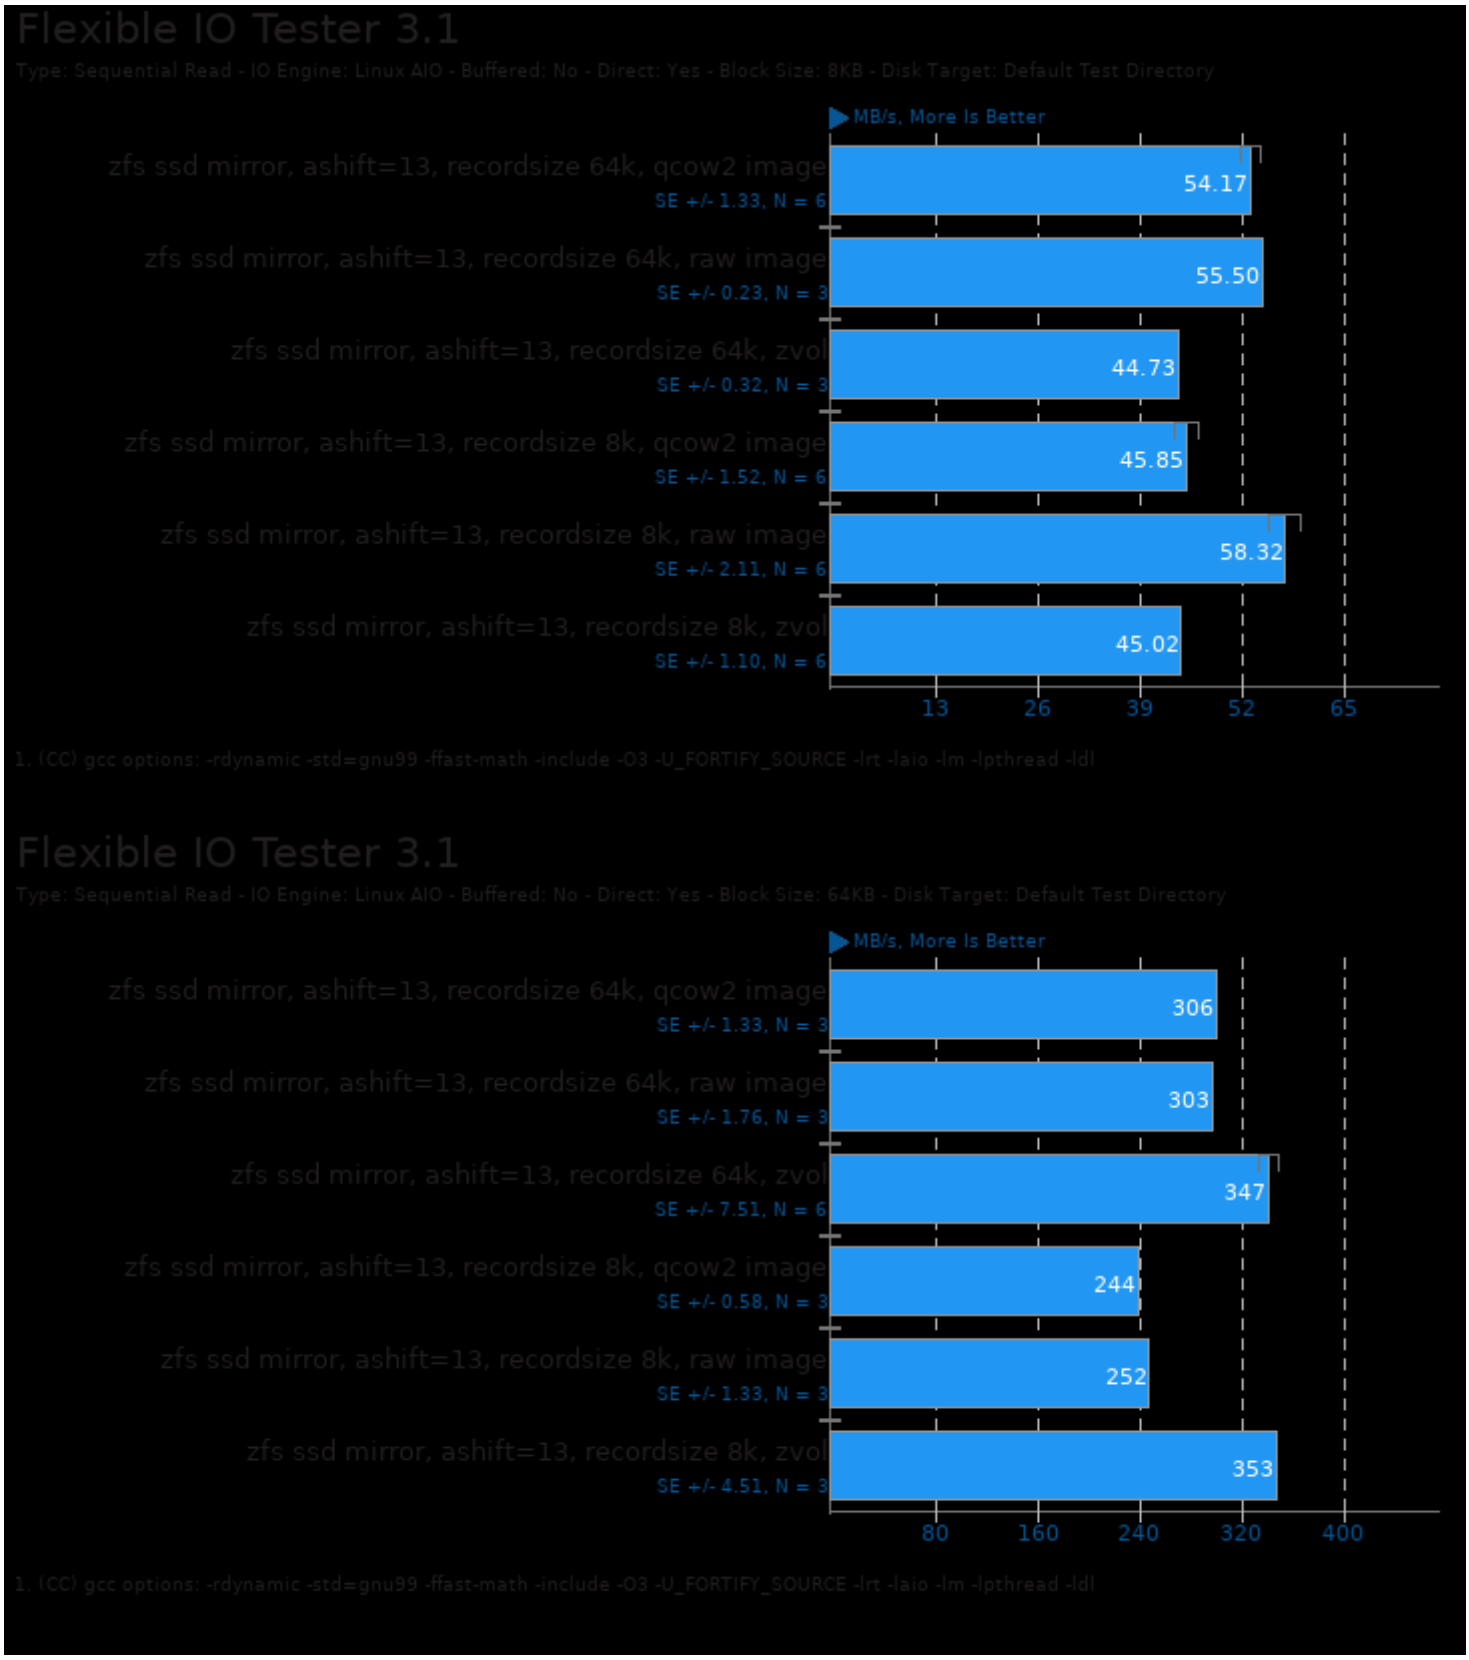
Flexible IO Tester 3.1
Type: Sequential Read - IO Engine: Linux AIO - Buffered: No - Direct: Yes - Block Size: 8KB - Disk Target: Default Test Directory
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 1.33, N = 6
54.17
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 0.23, N = 3
55.50
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 0.32, N = 3
44.73
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 1.52, N = 6
45.85
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 2.11, N = 6
58.32
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 1.10, N = 6
45.02
13
26
39
52
65
1. (CC) gcc options: -rdynamic -std=gnu99 -ffast-math -include -O3 -U_FORTIFY_SOURCE -lrt -laio -lm -lpthread -ldl

Flexible IO Tester 3.1
Type: Sequential Read - IO Engine: Linux AIO - Buffered: No - Direct: Yes - Block Size: 64KB - Disk Target: Default Test Directory
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 1.33, N = 3
306
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 1.76, N = 3
303
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 7.51, N = 6
347
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 0.58, N = 3
244
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 1.33, N = 3
252
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 4.51, N = 3
353
80
160
240
320
400
1. (CC) gcc options: -rdynamic -std=gnu99 -ffast-math -include -O3 -U_FORTIFY_SOURCE -lrt -laio -lm -lpthread -ldl

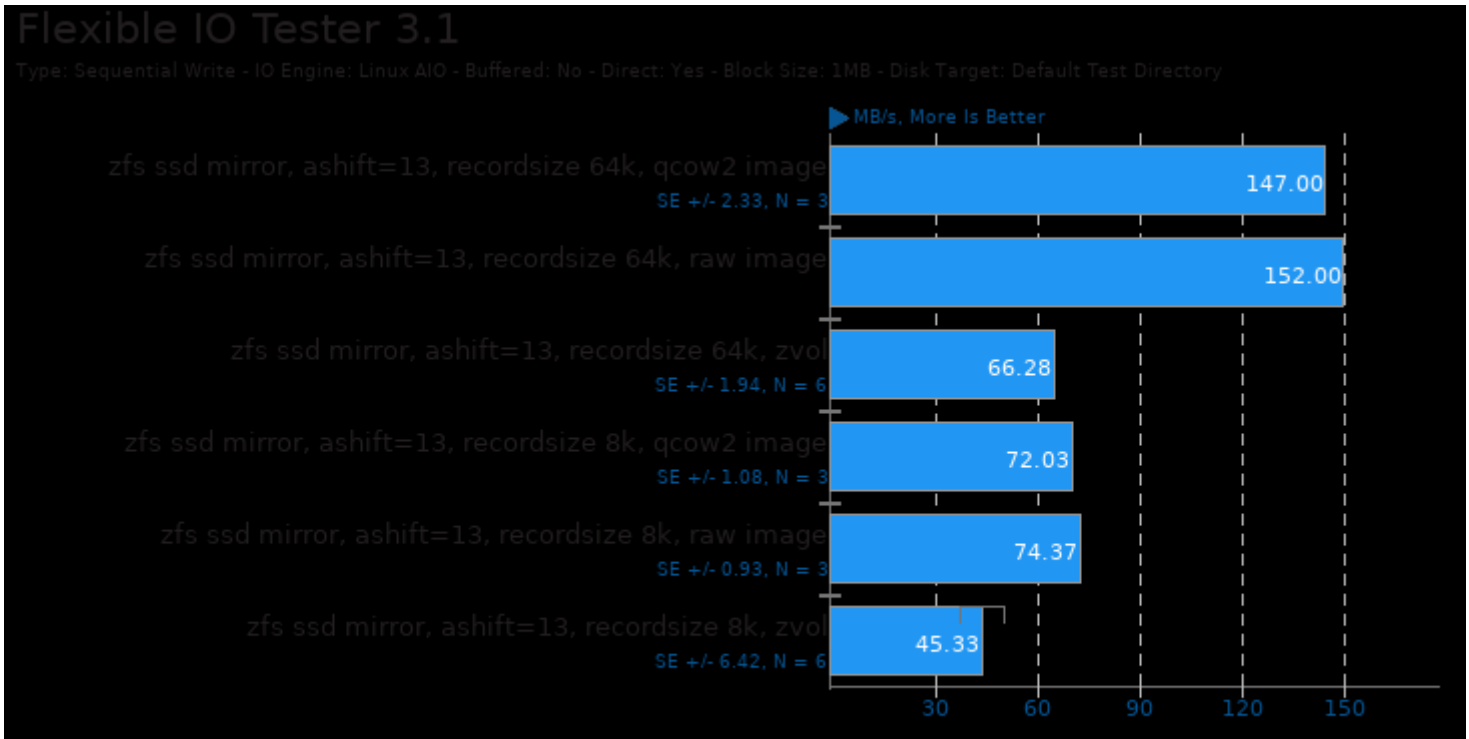
Flexible IO Tester 3.1
Type: Sequential Write - IO Engine: Linux AIO - Buffered: No - Direct: Yes - Block Size: 1MB - Disk Target: Default Test Directory
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 2.33, N = 3
147.00
zfs ssd mirror, ashift=13, recordsize 64k, raw image
152.00
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 1.94, N = 6
66.28
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 1.08, N = 3
72.03
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 0.93, N = 3
74.37
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 6.42, N = 6
45.33
30
60
90
120
150
1. (CC) gcc options: -rdynamic -std=gnu99 -ffast-math -include -O3 -U_FORTIFY_SOURCE -lrt -laio -lm -lpthread -ldl

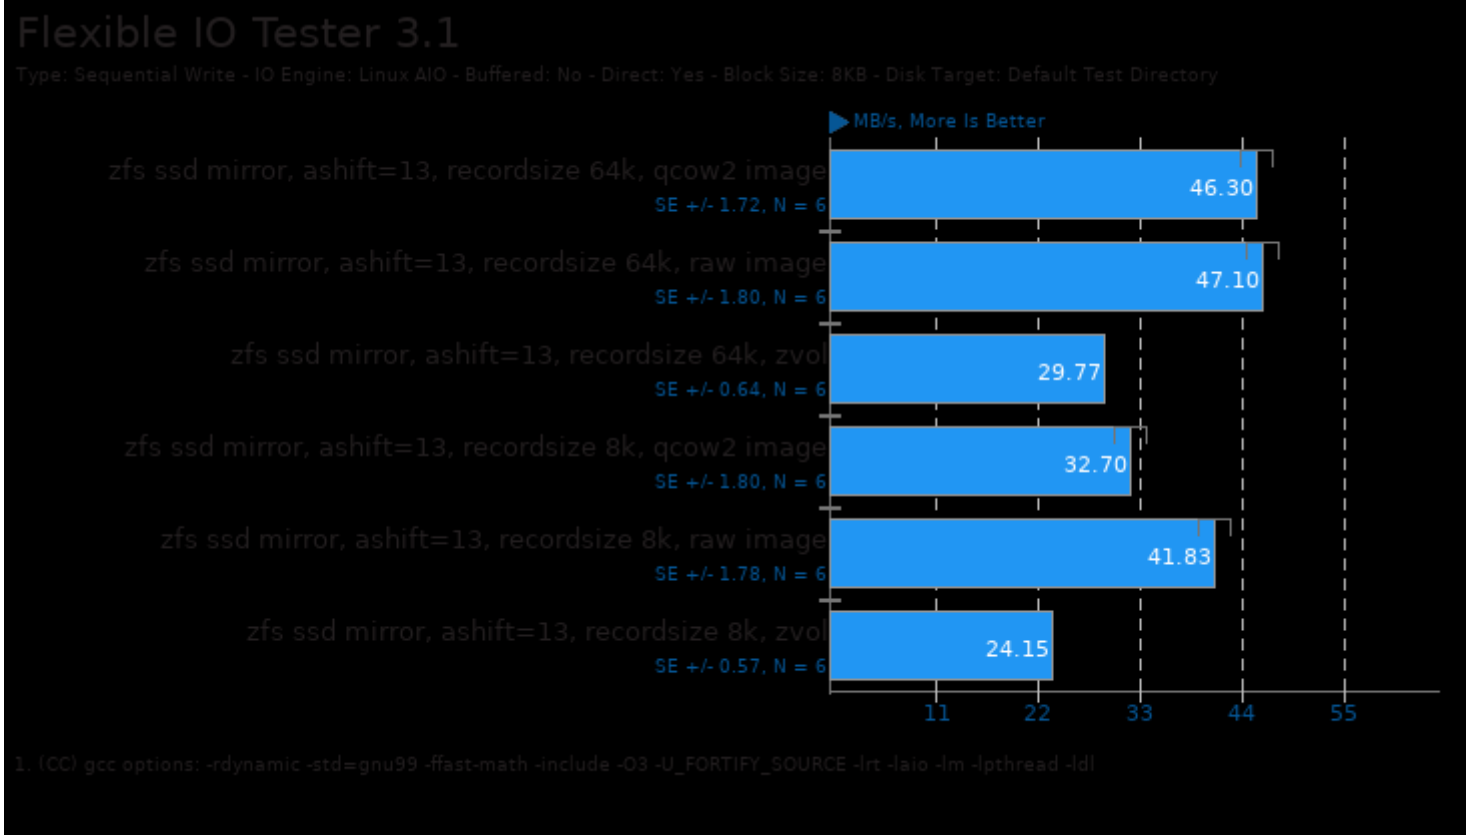
Flexible IO Tester 3.1
Type: Sequential Write - IO Engine: Linux AIO - Buffered: No - Direct: Yes - Block Size: 8KB - Disk Target: Default Test Directory
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 1.72, N = 6
46.30
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 1.80, N = 6
47.10
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 0.64, N = 6
29.77
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 1.80, N = 6
32.70
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 1.78, N = 6
41.83
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 0.57, N = 6
24.15
11
22
33
44
55
1. (CC) gcc options: -rdynamic -std=gnu99 -ffast-math -include -O3 -U_FORTIFY_SOURCE -lrt -laio -lm -lpthread -ldl

Flexible IO Tester 3.1

Type: Sequential Write - IO Engine: Linux AIO - Buffered: No - Direct: Yes - Block Size: 64KB - Disk Target: Default Test Directory

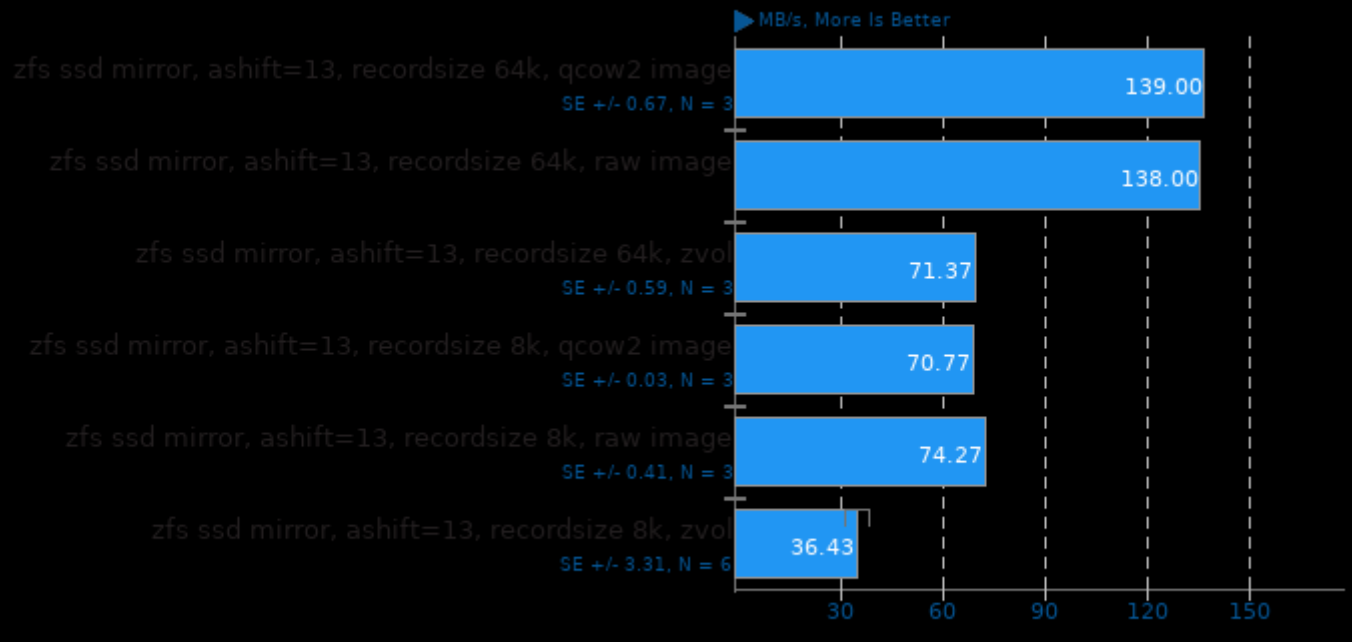

1. (CC) gcc options: -rdynamic -std=gnu99 -ffast-math -include -O3 -U_FORTIFY_SOURCE -lrt -laio -lm -lpthread -ldl

Dbench 4.0

Client Count: 1

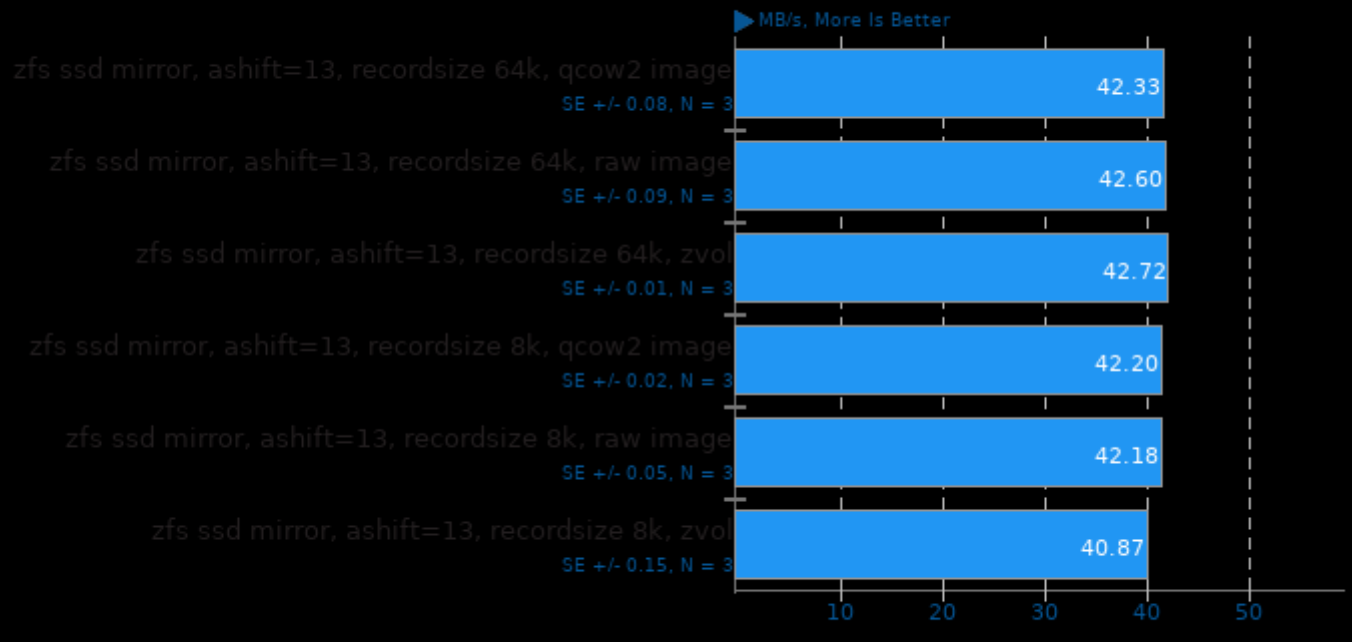

1. (CC) gcc options: -lpopt -O2

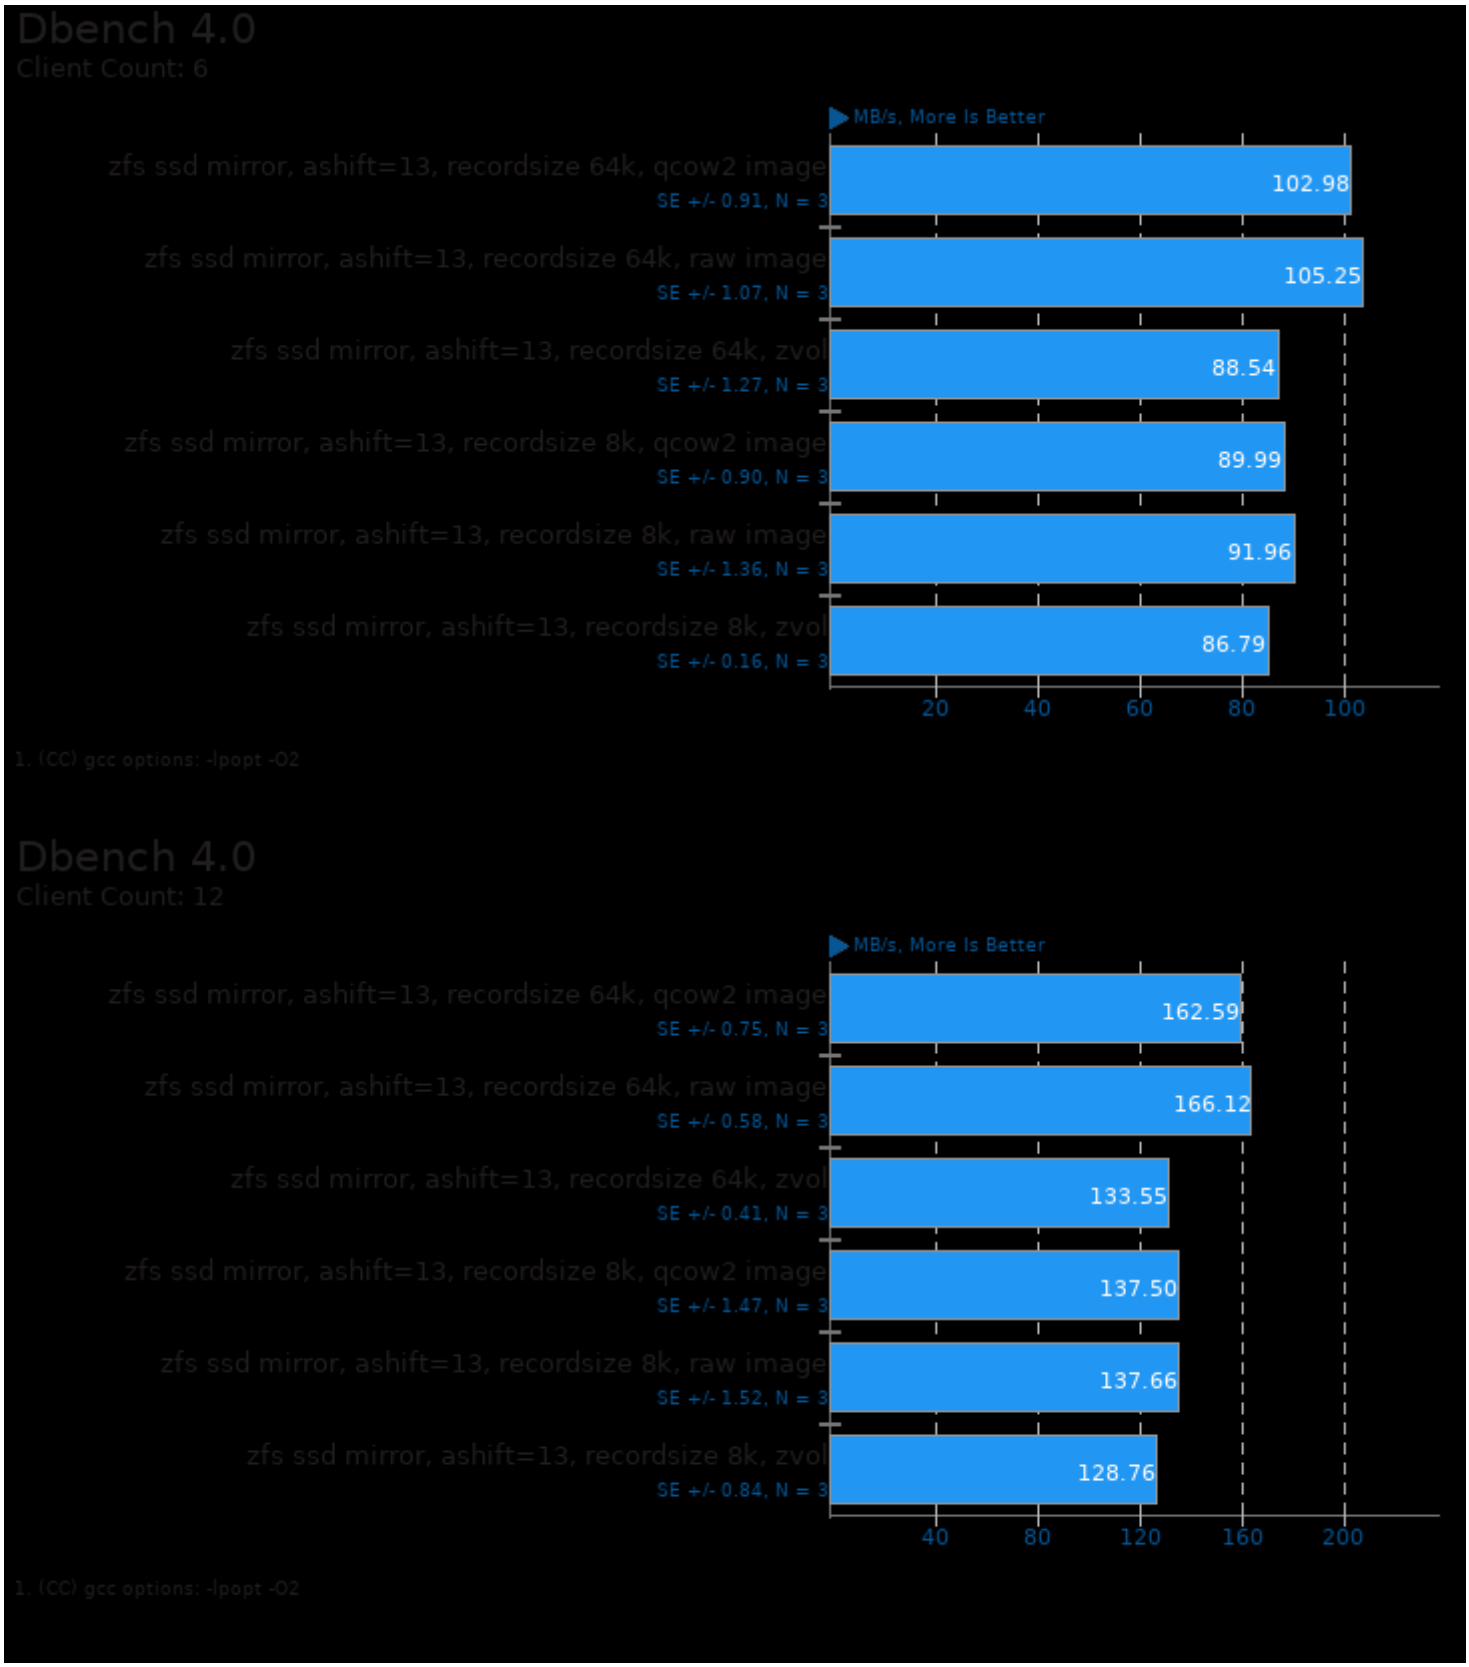
Dbench 4.0
Client Count: 6
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 0.91, N = 3
102.98
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 1.07, N = 3
105.25
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 1.27, N = 3
88.54
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 0.90, N = 3
89.99
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 1.36, N = 3
91.96
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 0.16, N = 3
86.79
20
40
60
80
100
1. (CC) gcc options: -lpopt -O2
Dbench 4.0
Client Count: 12
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 0.75, N = 3
162.59
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 0.58, N = 3
166.12
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 0.41, N = 3
133.55
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 1.47, N = 3
137.50
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 1.52, N = 3
137.66
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 0.84, N = 3
128.76
40
80
120
160
200
1. (CC) gcc options: -lpopt -O2

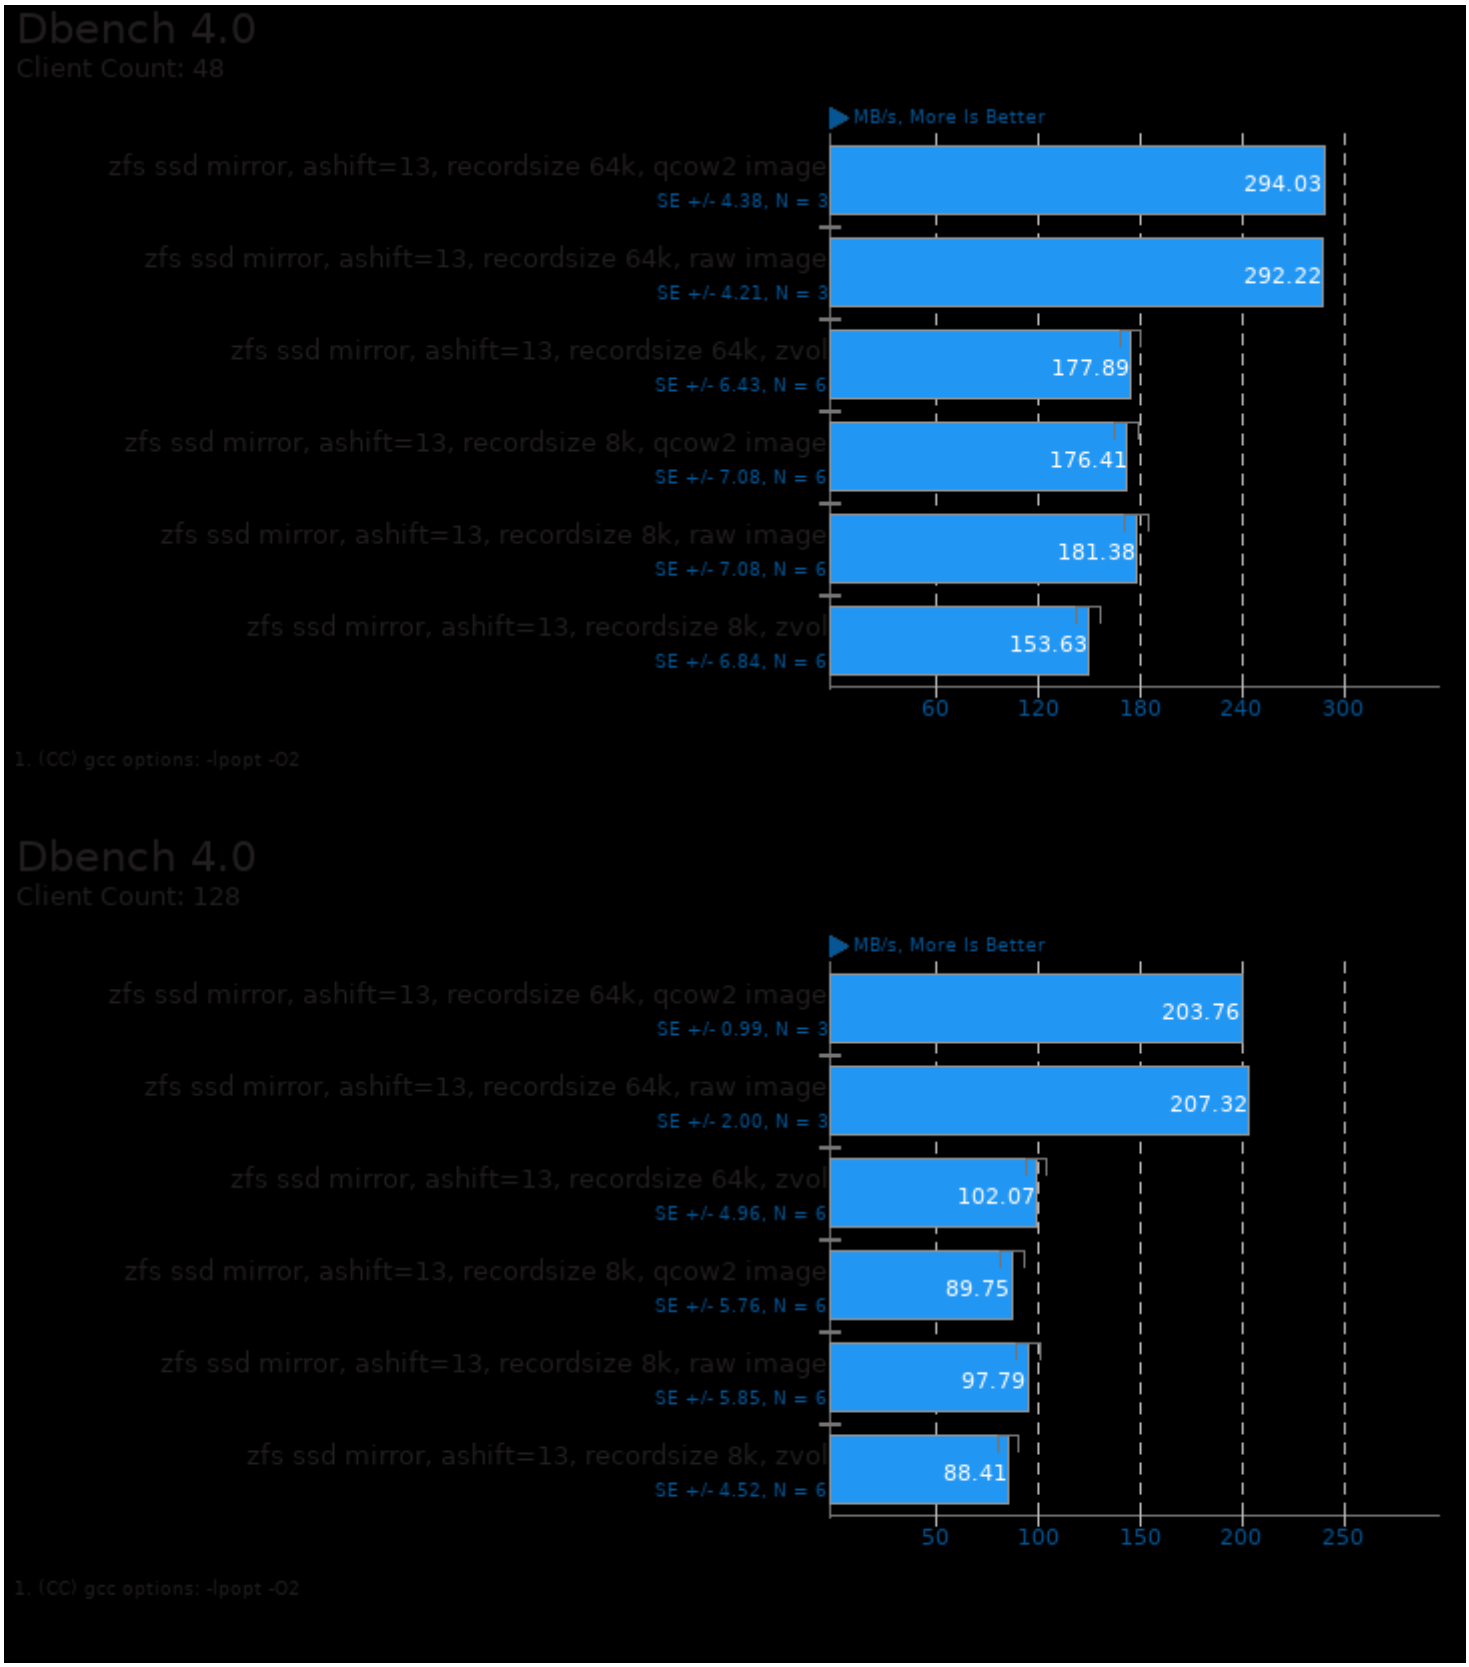
Dbench 4.0
Client Count: 48
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 4.38, N = 3
294.03
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 4.21, N = 3
292.22
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 6.43, N = 6
177.89
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 7.08, N = 6
176.41
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 7.08, N = 6
181.38
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 6.84, N = 6
153.63
60
120
180
240
300
1. (CC) gcc options: -lpopt -O2
Dbench 4.0
Client Count: 128
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 0.99, N = 3
203.76
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 2.00, N = 3
207.32
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 4.96, N = 6
102.07
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 5.76, N = 6
89.75
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 5.85, N = 6
97.79
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 4.52, N = 6
88.41
50
100
150
200
250
1. (CC) gcc options: -lpopt -O2

Dbench 4.0

Client Count: 256

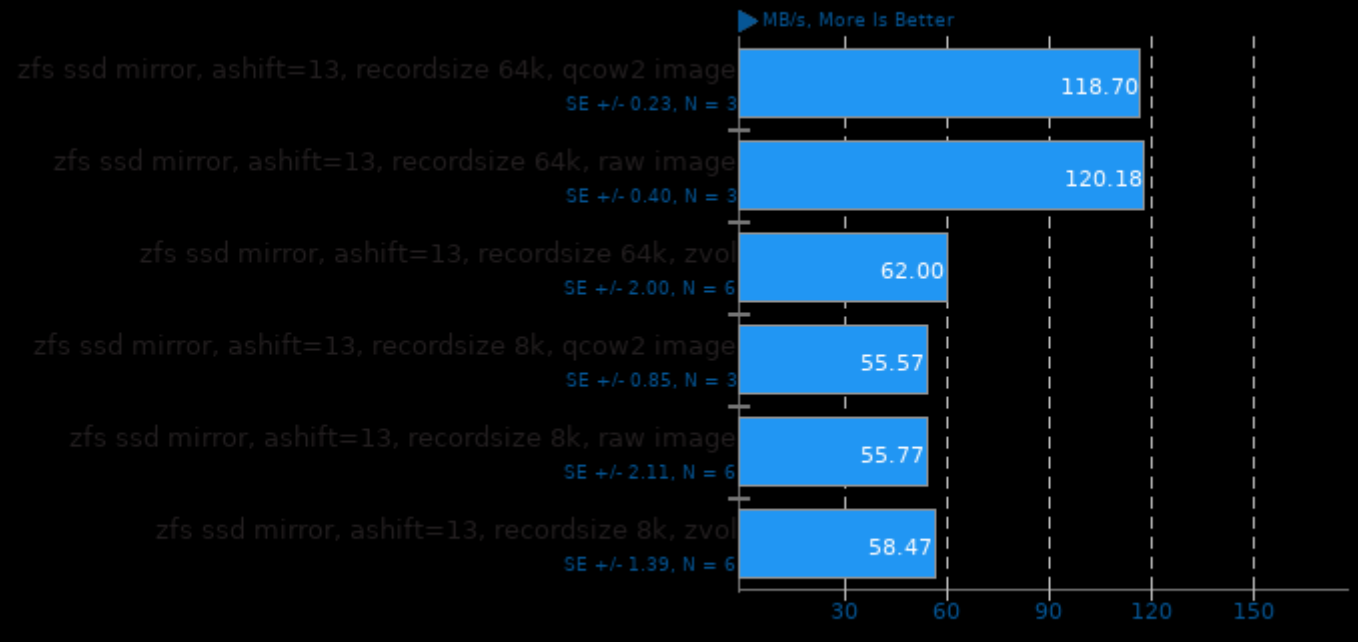

1. (CC) gcc options: -lpopt -O2

IOzone 3.465

Record Size: 1MB - File Size: 8GB - Disk Test: Read Performance

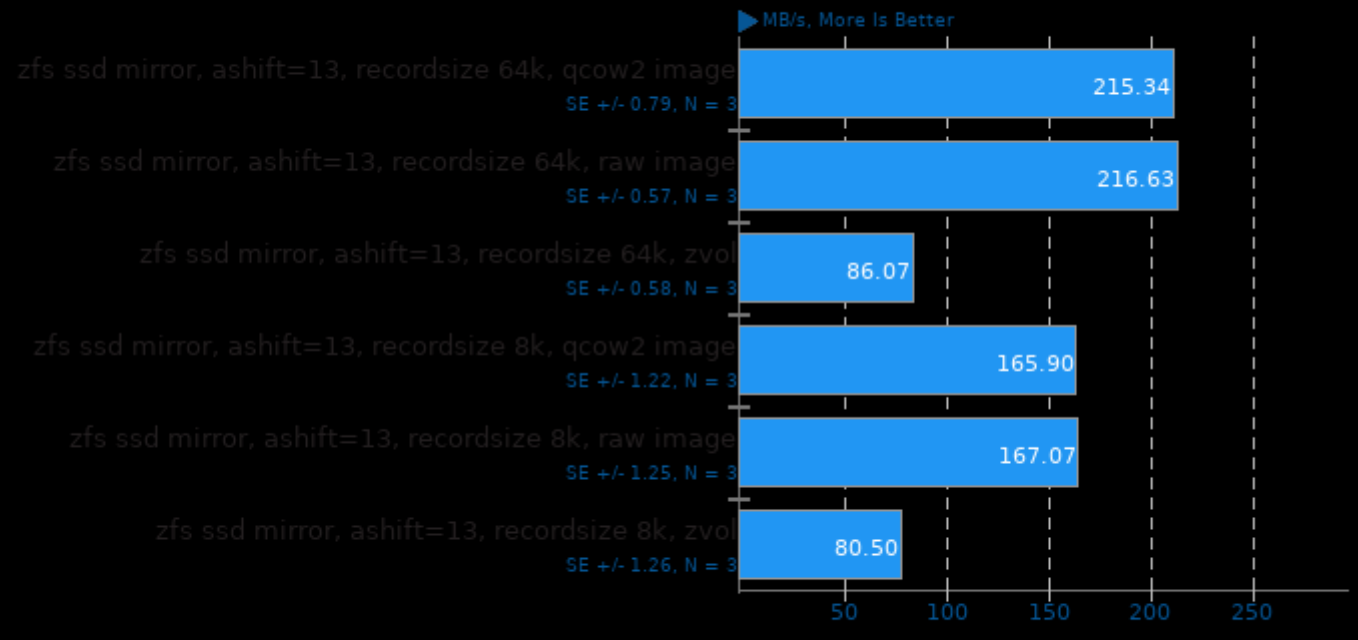

1. (CC) gcc options: -O3

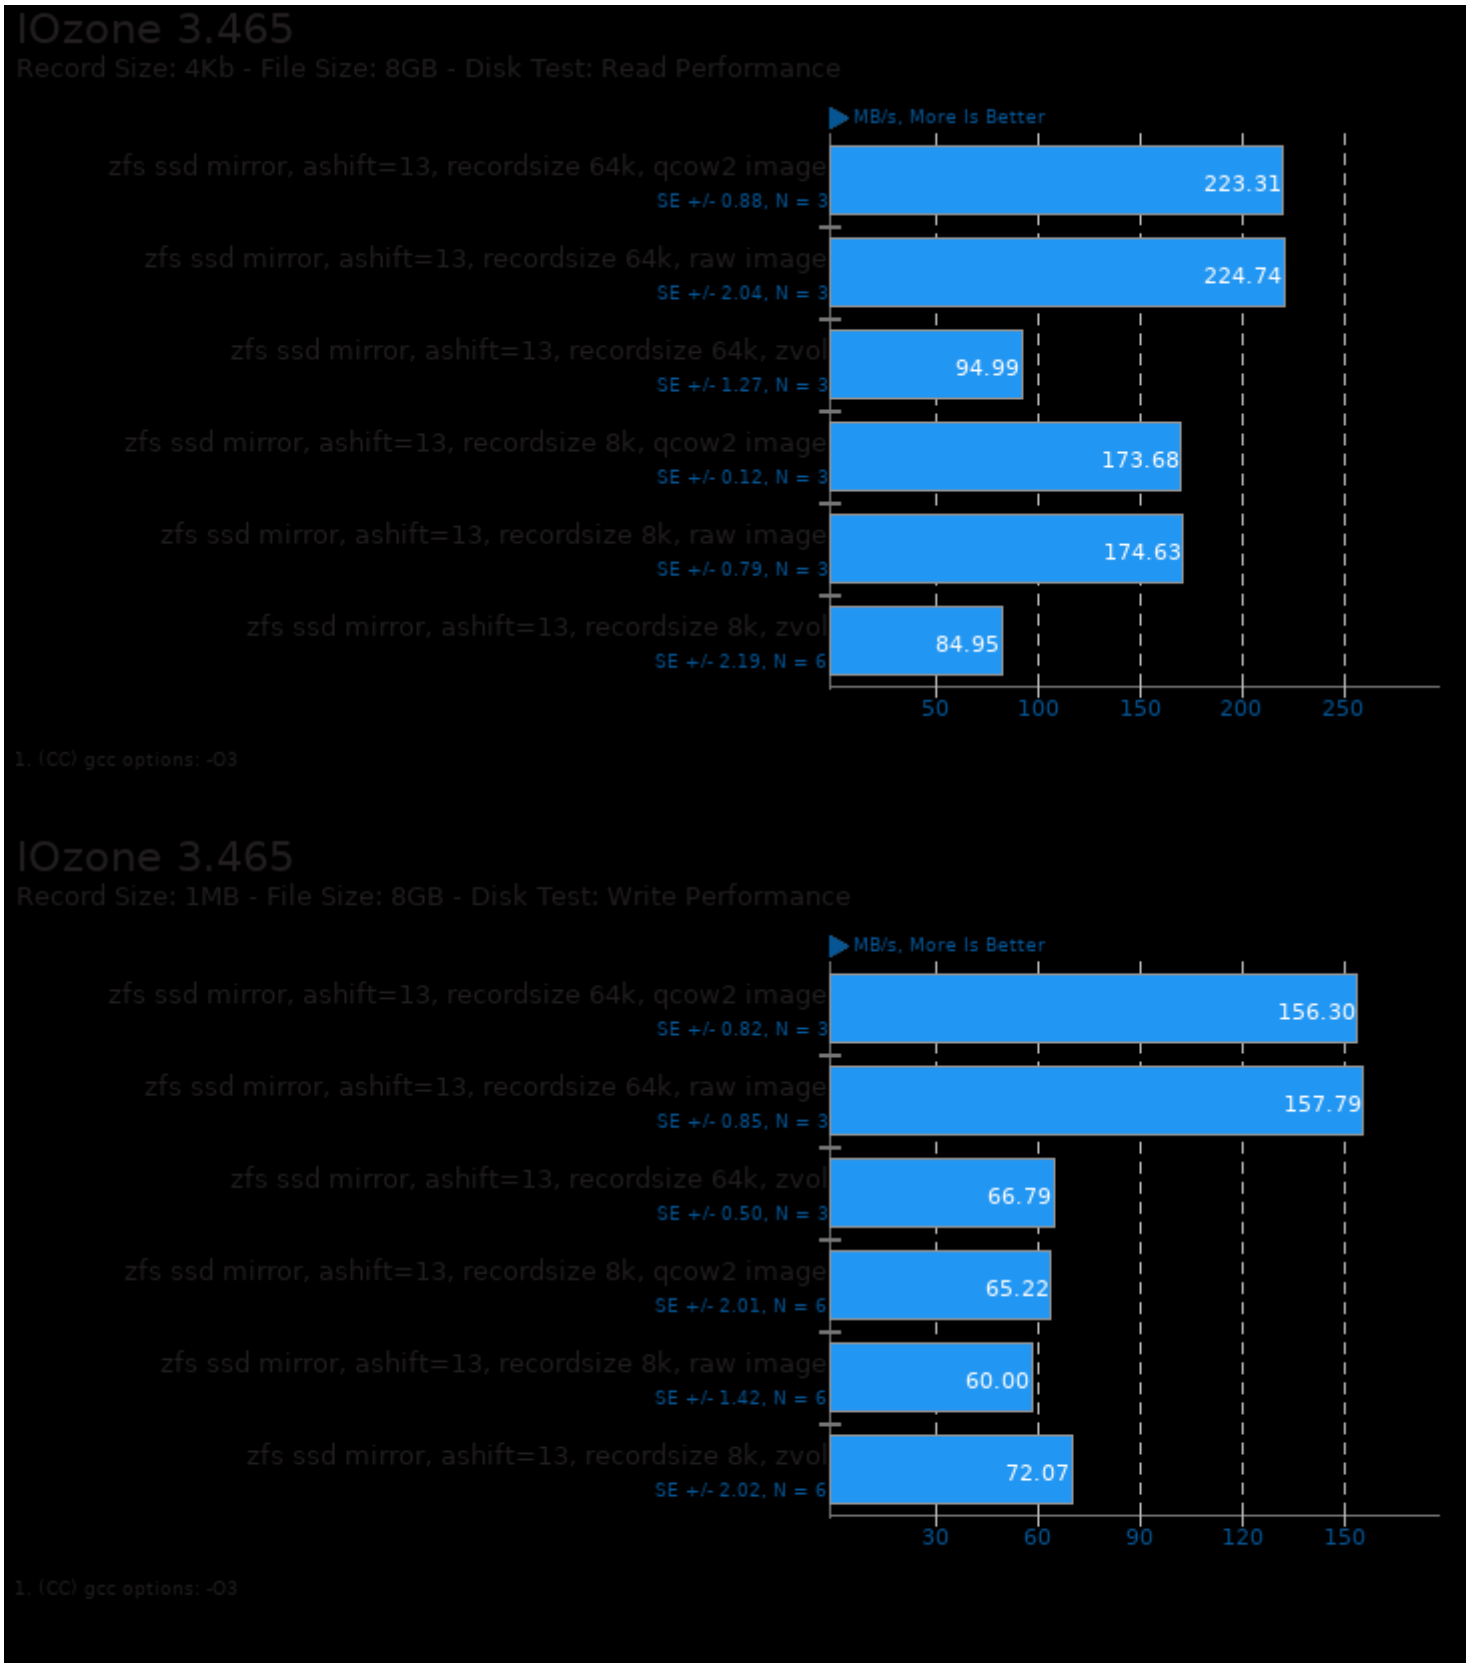
IOzone 3.465
Record Size: 4Kb - File Size: 8GB - Disk Test: Read Performance
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 0.88, N = 3
223.31
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 2.04, N = 3
224.74
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 1.27, N = 3
94.99
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 0.12, N = 3
173.68
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 0.79, N = 3
174.63
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 2.19, N = 6
84.95
50
100
150
200
250
1. (CC) gcc options: -O3
IOzone 3.465
Record Size: 1MB - File Size: 8GB - Disk Test: Write Performance
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 0.82, N = 3
156.30
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 0.85, N = 3
157.79
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 0.50, N = 3
66.79
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 2.01, N = 6
65.22
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 1.42, N = 6
60.00
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 2.02, N = 6
72.07
30
60
90
120
150
1. (CC) gcc options: -O3

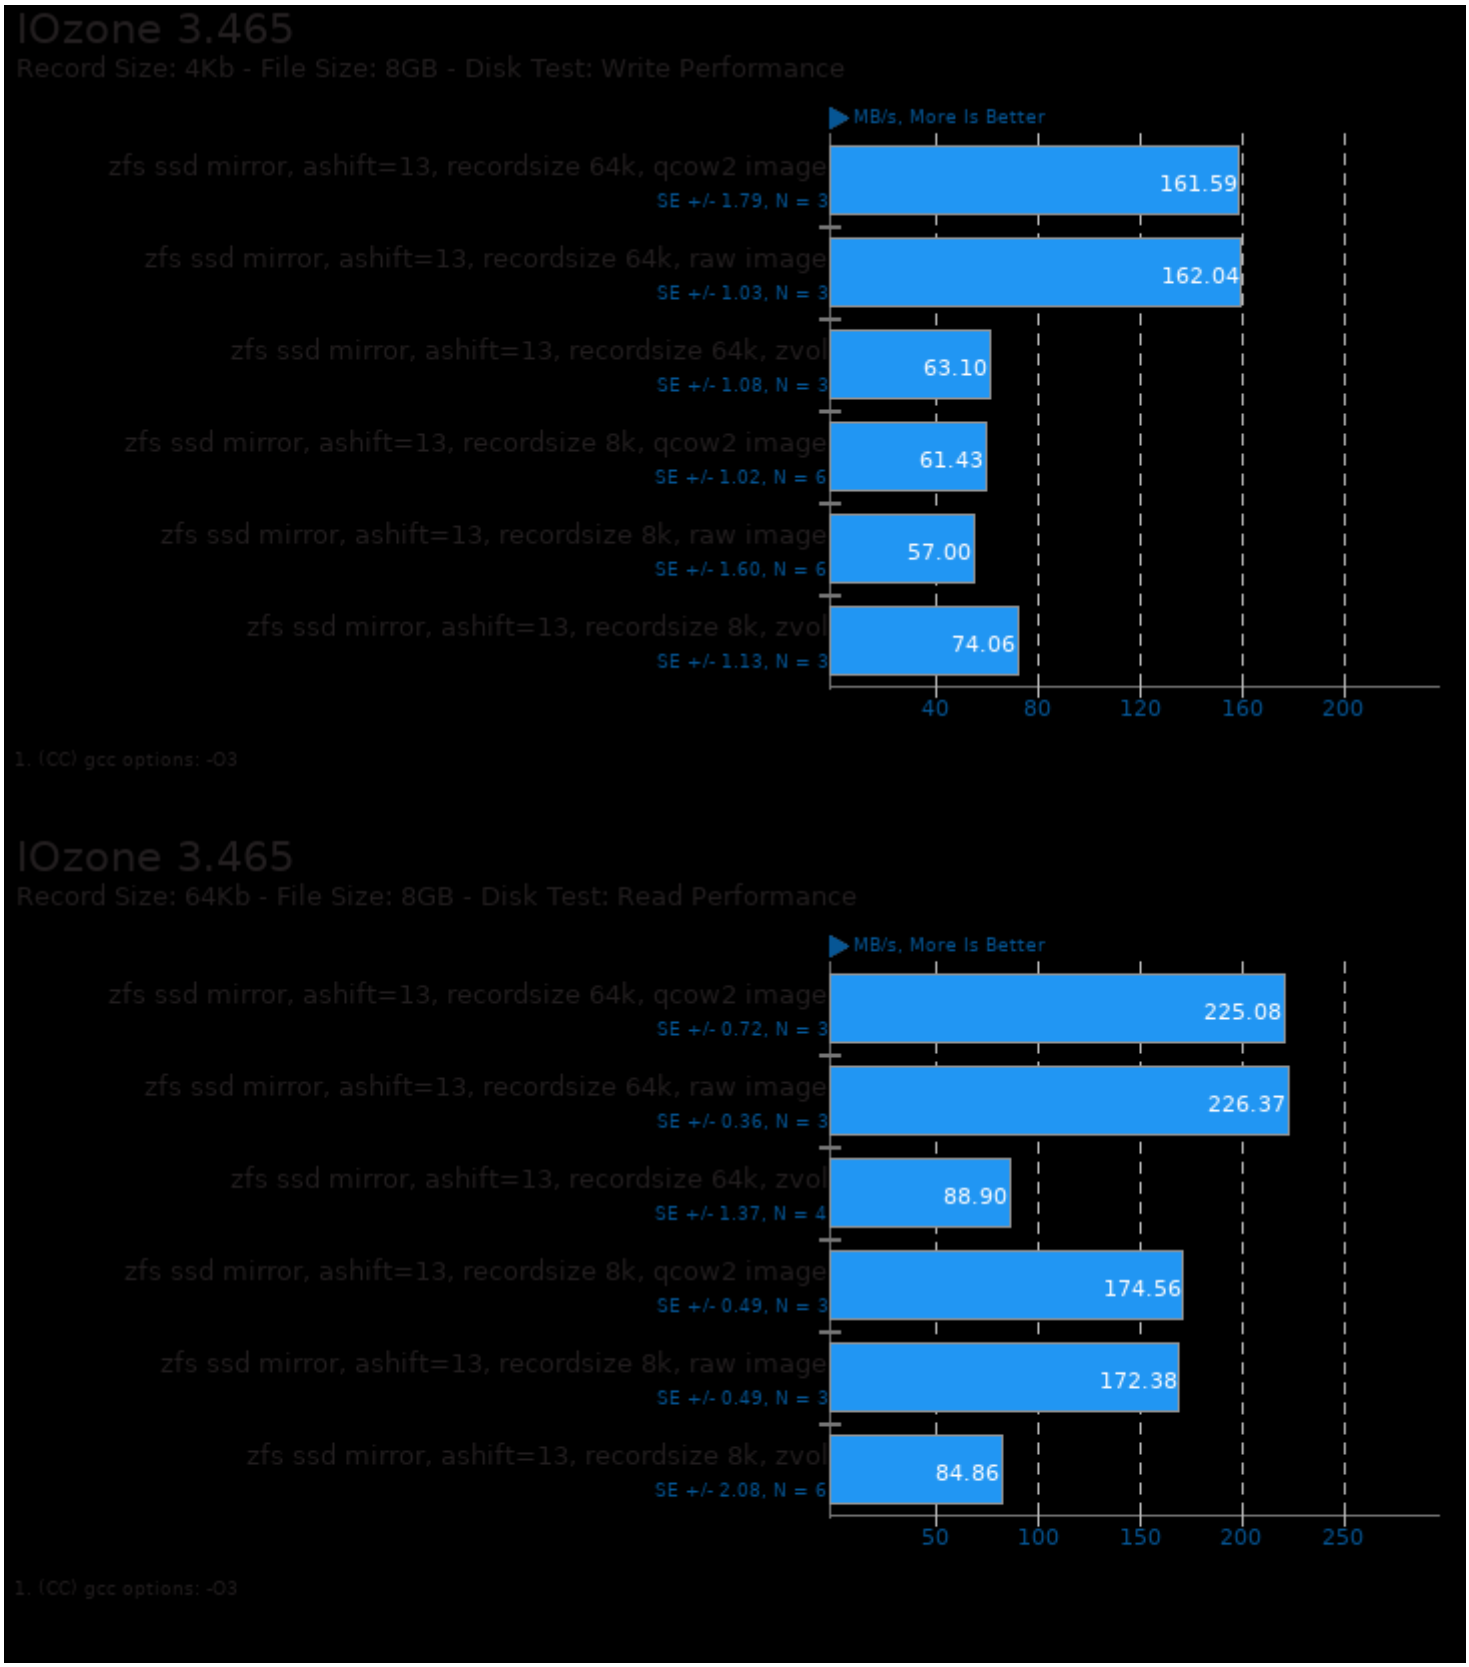
IOzone 3.465
Record Size: 4Kb - File Size: 8GB - Disk Test: Write Performance
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 1.79, N = 3
161.59
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 1.03, N = 3
162.04
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 1.08, N = 3
63.10
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 1.02, N = 6
61.43
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 1.60, N = 6
57.00
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 1.13, N = 3
74.06
40
80
120
160
200
1. (CC) gcc options: -O3
IOzone 3.465
Record Size: 64Kb - File Size: 8GB - Disk Test: Read Performance
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 0.72, N = 3
225.08
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 0.36, N = 3
226.37
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 1.37, N = 4
88.90
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 0.49, N = 3
174.56
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 0.49, N = 3
172.38
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 2.08, N = 6
84.86
50
100
150
200
250
1. (CC) gcc options: -O3

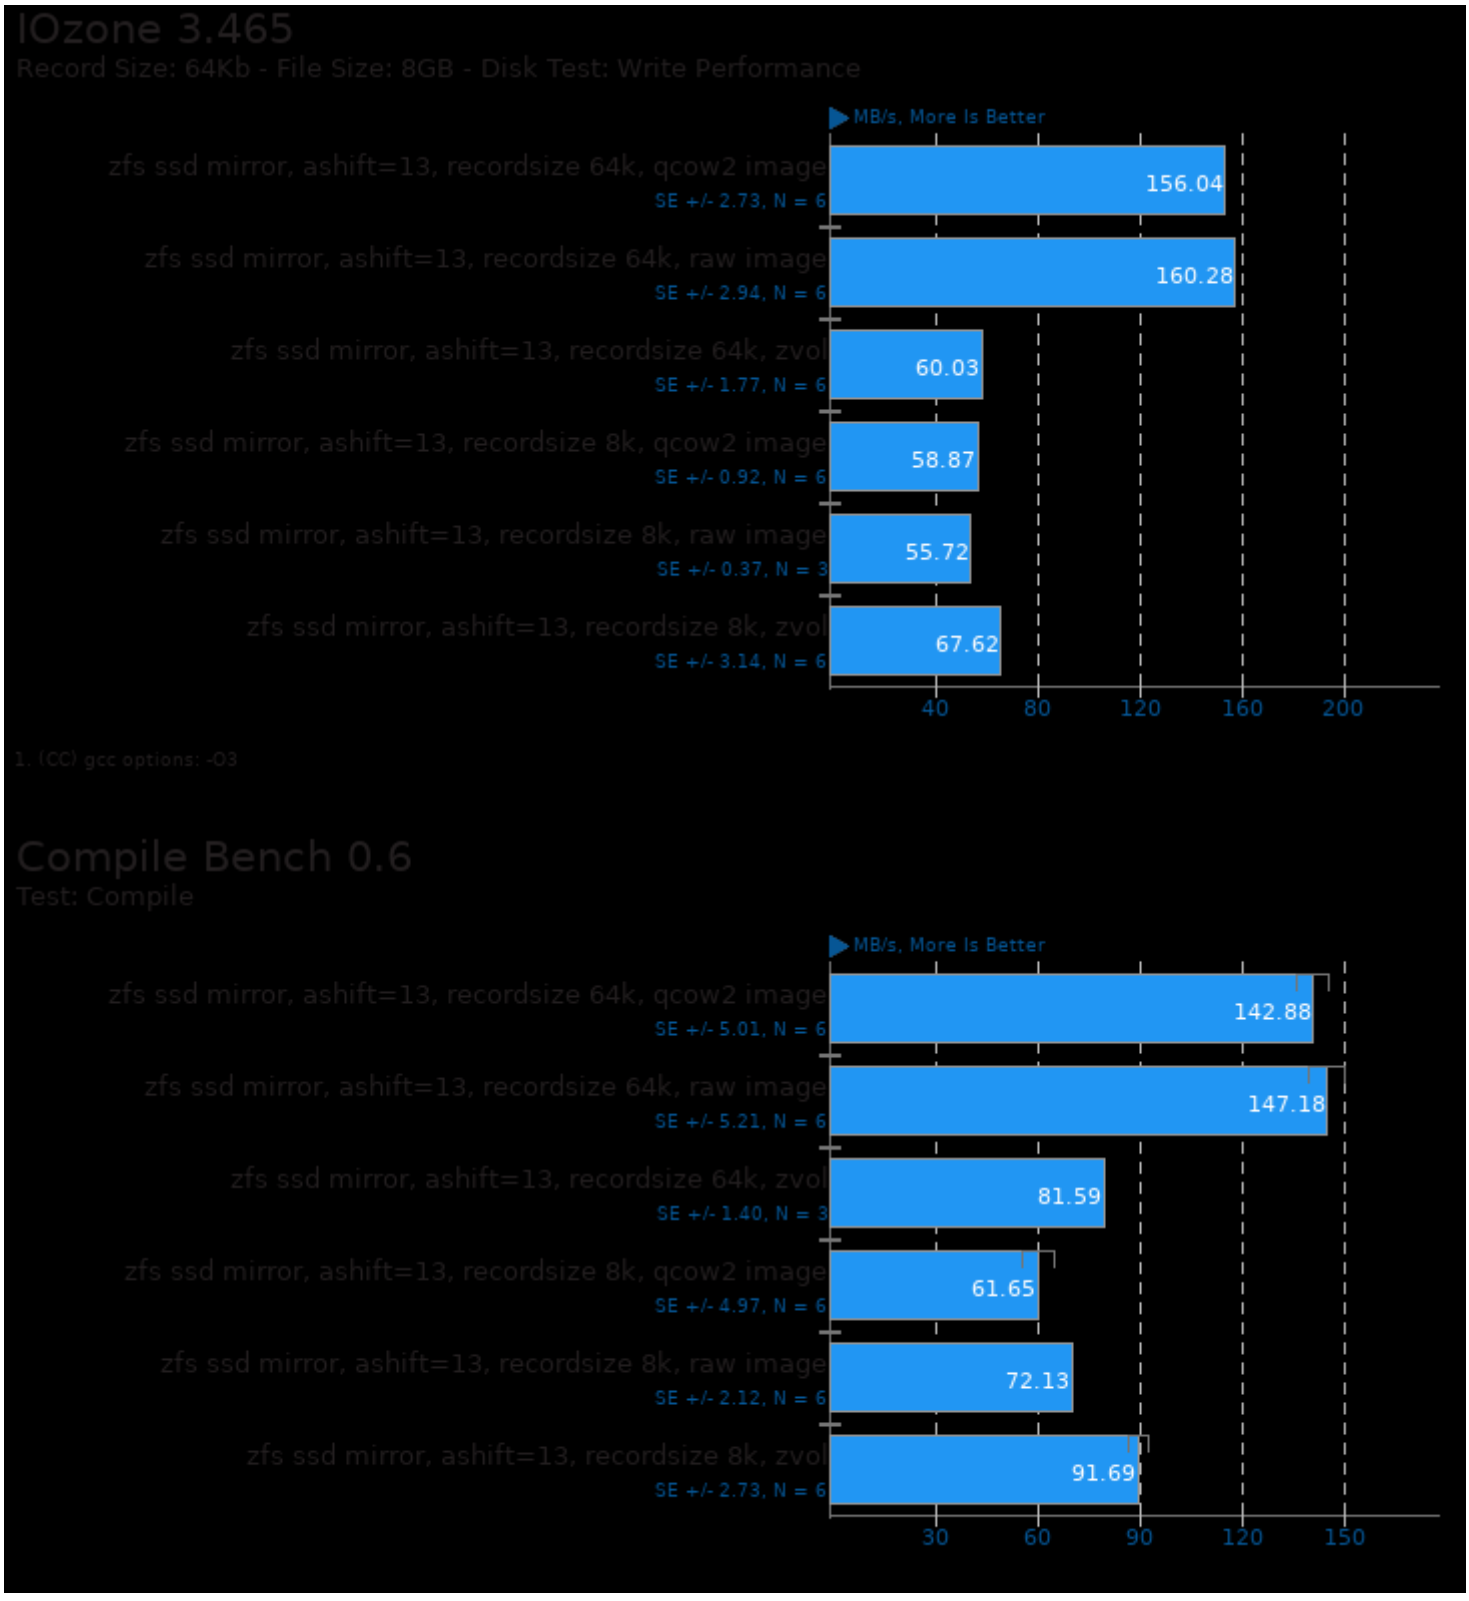
IOzone 3.465
Record Size: 64Kb - File Size: 8GB - Disk Test: Write Performance
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 2.73, N = 6
156.04
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 2.94, N = 6
160.28
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 1.77, N = 6
60.03
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 0.92, N = 6
58.87
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 0.37, N = 3
55.72
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 3.14, N = 6
67.62
40
80
120
160
200
1. (CC) gcc options: -O3
Compile Bench 0.6
Test: Compile
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 5.01, N = 6
142.88
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 5.21, N = 6
147.18
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 1.40, N = 3
81.59
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 4.97, N = 6
61.65
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 2.12, N = 6
72.13
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 2.73, N = 6
91.69
30
60
90
120
150

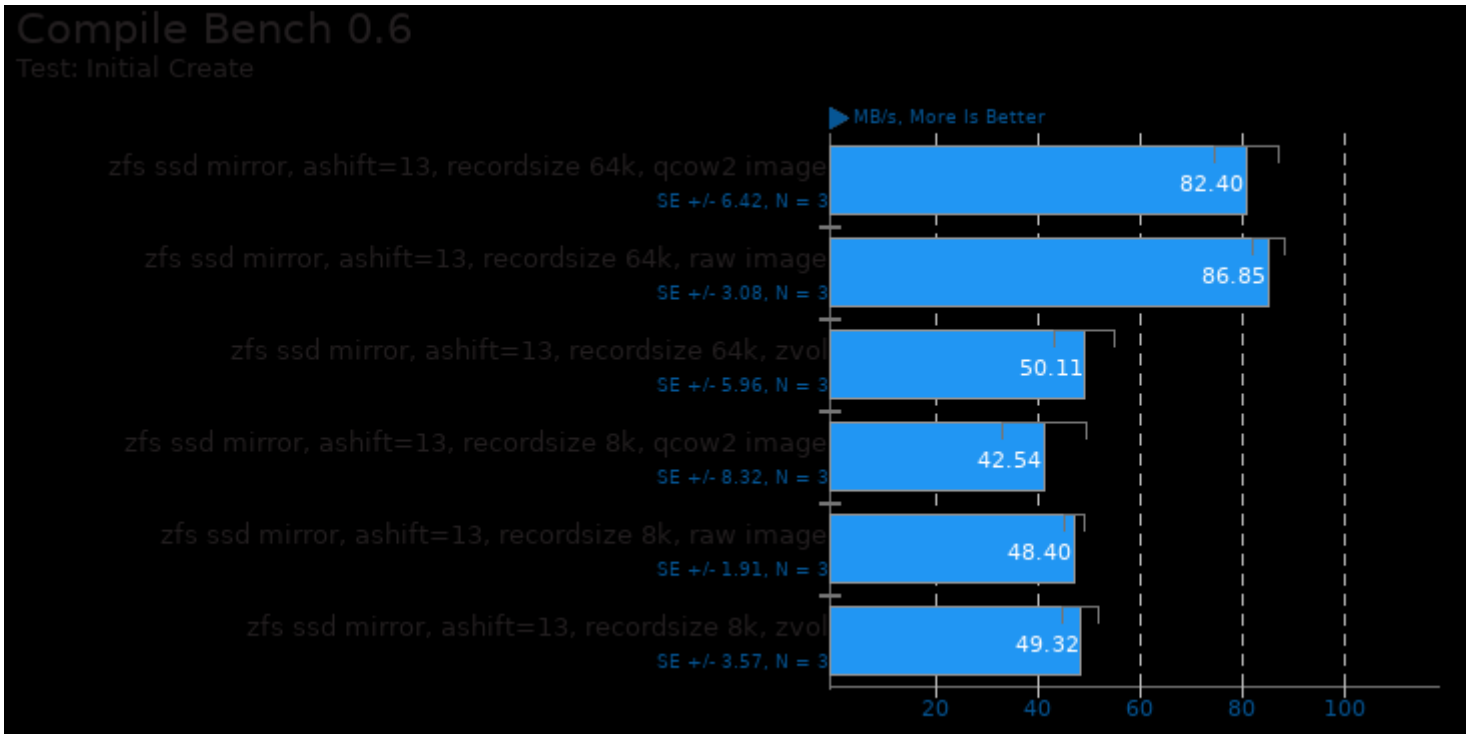
Compile Bench 0.6
Test: Initial Create
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 6.42, N = 3
82.40
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 3.08, N = 3
86.85
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 5.96, N = 3
50.11
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 8.32, N = 3
42.54
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 1.91, N = 3
48.40
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 3.57, N = 3
49.32
20
40
60
80
100

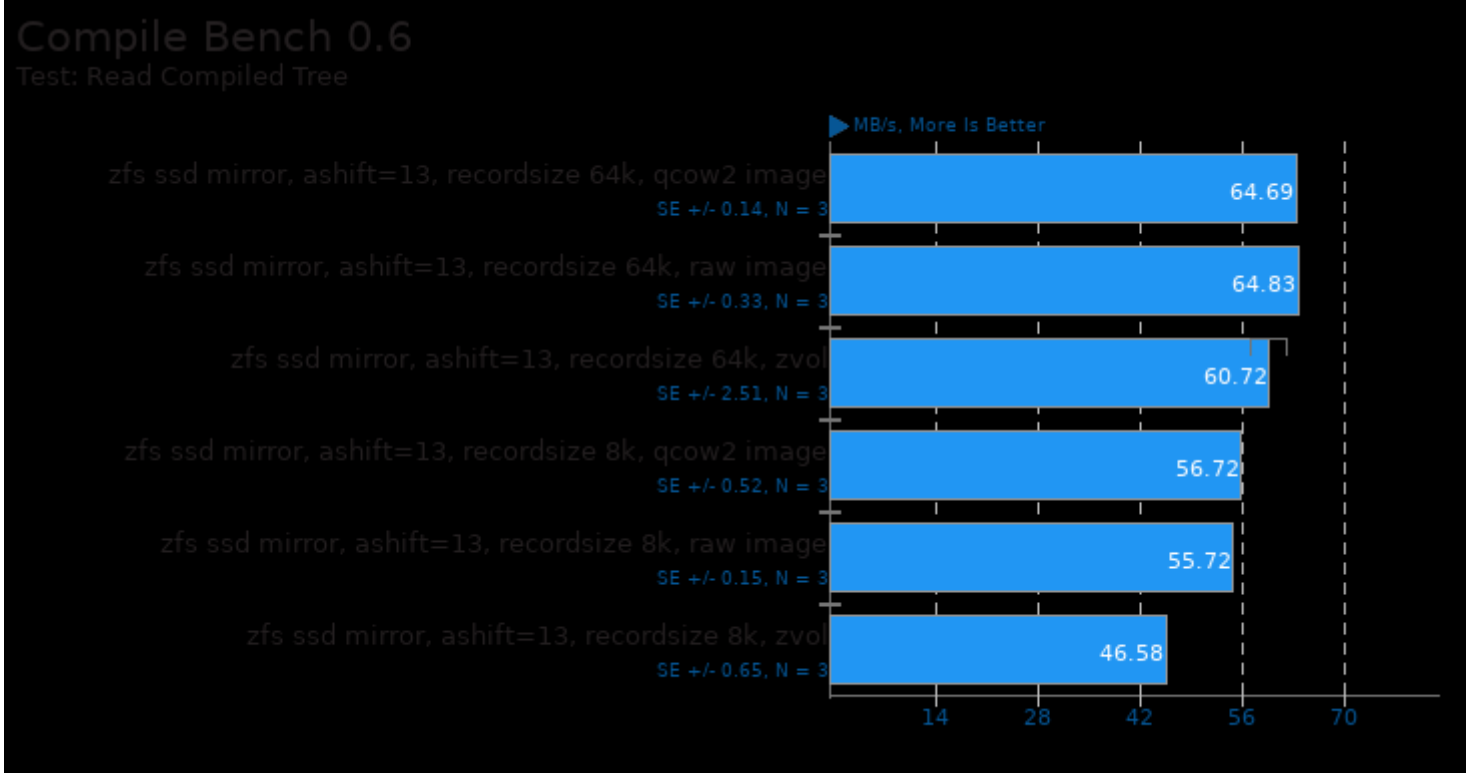
Compile Bench 0.6
Test: Read Compiled Tree
MB/s, More Is Better
zfs ssd mirror, ashift=13, recordsize 64k, qcow2 image
SE +/- 0.14, N = 3
64.69
zfs ssd mirror, ashift=13, recordsize 64k, raw image
SE +/- 0.33, N = 3
64.83
zfs ssd mirror, ashift=13, recordsize 64k, zvol
SE +/- 2.51, N = 3
60.72
zfs ssd mirror, ashift=13, recordsize 8k, qcow2 image
SE +/- 0.52, N = 3
56.72
zfs ssd mirror, ashift=13, recordsize 8k, raw image
SE +/- 0.15, N = 3
55.72
zfs ssd mirror, ashift=13, recordsize 8k, zvol
SE +/- 0.65, N = 3
46.58
14
28
42
56
70

These geometric means are based upon test groupings / test suites for this result file.

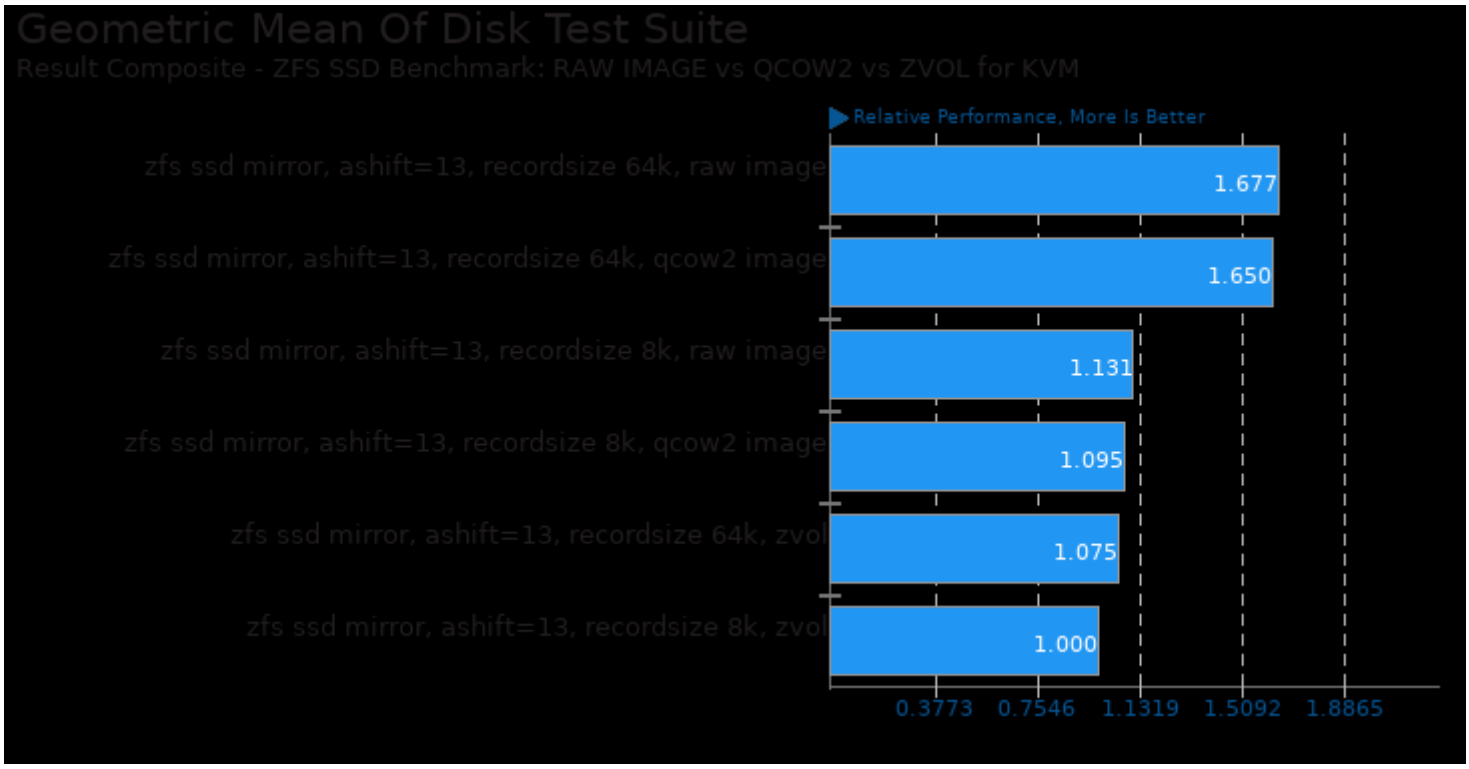

Geometric mean based upon tests: pts/compilebench, pts/iozone, pts/dbench and pts/fio

This file was automatically generated via the Phoronix Test Suite benchmarking software on Thursday, 28 March 2024 21:23.