www.phoronix-test-suite.com

newa

AMD Ryzen 7 3800XT 8-Core testing with a MSI X370 XPOWER GAMING TITANIUM (MS-7A31) v1.0 (1.MS BIOS) and Sapphire AMD Radeon HD 4650 on Debian 12 via the Phoronix Test Suite.

Automated Executive Summary

b had the most wins, coming in first place for 54% of the tests.

Based on the geometric mean of all complete results, the fastest (a) was 1.002x the speed of the slowest (b).

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

Timed Mesa Compilation (Time To Compile) at 1.025x
TensorFlow (Device: CPU - Batch Size: 32 - Model: AlexNet) at 1.02x
Blender (Blend File: BMW27 - Compute: CPU-Only) at 1.017x
TensorFlow (Device: CPU - Batch Size: 16 - Model: ResNet-50) at 1.008x
TensorFlow (Device: CPU - Batch Size: 32 - Model: GoogLeNet) at 1.007x
TensorFlow (Device: CPU - Batch Size: 1 - Model: VGG-16) at 1.006x
TensorFlow (Device: CPU - Batch Size: 16 - Model: GoogLeNet) at 1.004x
Blender (Blend File: Junkshop - Compute: CPU-Only) at 1.004x
TensorFlow (Device: CPU - Batch Size: 1 - Model: GoogLeNet) at 1.003x
TensorFlow (Device: CPU - Batch Size: 64 - Model: ResNet-50) at 1.003x.

Test Systems:

a

b

Processor: AMD Ryzen 7 3800XT 8-Core @ 3.90GHz (8 Cores / 16 Threads), Motherboard: MSI X370 XPOWER GAMING TITANIUM (MS-7A31) v1.0 (1.MS BIOS), Chipset: AMD Starship/Matisse, Memory: 2 x 8GB DDR4-3200MT/s CL16-18-18 D4-3200, Disk: 128GB INTEL SSDPEKKW128G7, Graphics: Sapphire AMD Radeon HD 4650, Audio: AMD RV710/730, Network: Intel I211

OS: Debian 12, Kernel: 6.1.0-18-amd64 (x86_64), Display Server: X Server 1.20.11, Compiler: GCC 12.2.0, File-System: ext4, Screen Resolution: 1024x768

Kernel Notes: Transparent Huge Pages: always
Compiler Notes: --build=x86_64-linux-gnu --disable-vtable-verify --disable-werror --enable-cet --enable-checking=release --enable-clocale=gnu --enable-default-pie --enable-gnu-unique-object --enable-languages=c,ada,c++,go,d,fortran,objc,obj-c++,m2 --enable-libphobos-checking=release --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-defaulted --enable-offload-targets=nvptx-none=/build/gcc-12-bTRWOB/gcc-12-12.2.0/debian/tmp-nvptx/usr,amdgcn-amdhsa=/build/gcc-12-bTRWOB/gcc-12-12.2.0/debian/tmp-gcn/usr --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-default-libstdcxx-abi=new --with-gcc-major-version-only --with-multilib-list=m32,m64,mx32 --with-target-system-zlib=auto --with-tune=generic --without-cuda-driver -v
Processor Notes: Scaling Governor: acpi-cpufreq schedutil (Boost: Enabled) - CPU Microcode: 0x8701021
Python Notes: Python 3.11.2
Security Notes: gather_data_sampling: Not affected + itlb_multihit: Not affected + l1tf: Not affected + mds: Not affected + meltdown: Not affected + mmio_stale_data: Not affected + retbleed: Mitigation of untrained return thunk; SMT enabled with STIBP protection + spec_rstack_overflow: Mitigation of safe RET + spec_store_bypass: Mitigation of SSB disabled via prctl + spectre_v1: Mitigation of usercopy/swapgs barriers and __user pointer sanitization + spectre_v2: Mitigation of Retpolines IBPB: conditional STIBP: always-on RSB filling PBRSB-eIBRS: Not affected + srbds: Not affected + tsx_async_abort: Not affected

	a	b
Timed Mesa Compilation - Time To Compile (sec)	36.759	37.68
Normalized	100%	97.56%
TensorFlow - CPU - 1 - VGG-16 (images/sec)	1.62	1.61
Normalized	100%	99.38%
TensorFlow - CPU - 1 - AlexNet (images/sec)	4.86	4.85
Normalized	100%	99.79%
TensorFlow - CPU - 16 - VGG-16 (images/sec)	4.12	4.13
Normalized	99.76%	100%
TensorFlow - CPU - 32 - VGG-16 (images/sec)	4.46	4.45
Normalized	100%	99.78%
TensorFlow - CPU - 16 - AlexNet (images/sec)	37.6	37.72
Normalized	99.68%	100%
TensorFlow - CPU - 32 - AlexNet (images/sec)	56.76	55.62
Normalized	100%	97.99%
TensorFlow - CPU - 64 - AlexNet (images/sec)	73.2	73.38
Normalized	99.75%	100%
TensorFlow - CPU - 1 - GoogLeNet (images/sec)	14.72	14.77
Normalized	99.66%	100%
TensorFlow - CPU - 1 - ResNet-50 (images/sec)	6.32	6.33
Normalized	99.84%	100%
TensorFlow - CPU - 256 - AlexNet (images/sec)	94.1	93.99
Normalized	100%	99.88%

newa

TensorFlow - CPU - 512 - AlexNet (images/sec)	96.96	97.01
Normalized	99.95%	100%
TensorFlow - CPU - 16 - GoogLeNet (images/sec)	29.31	29.2
Normalized	100%	99.62%
TensorFlow - CPU - 16 - ResNet-50 (images/sec)	9.1	9.17
Normalized	99.24%	100%
TensorFlow - CPU - 32 - GoogLeNet (images/sec)	29.17	29.36
Normalized	99.35%	100%
TensorFlow - CPU - 32 - ResNet-50 (images/sec)	8.96	8.96
TensorFlow - CPU - 64 - GoogLeNet (images/sec)	28.9	28.92
Normalized	99.93%	100%
TensorFlow - CPU - 64 - ResNet-50 (images/sec)	8.89	8.92
Normalized	99.66%	100%
TensorFlow - CPU - 256 - GoogLeNet (images/sec)	27.08	27.15
Normalized	99.74%	100%
Blender - BMW27 - CPU-Only (sec)	151.3	153.91
Normalized	100%	98.3%
Blender - Junkshop - CPU-Only (sec)	224.11	223.28
Normalized	99.63%	100%
Blender - Fishy Cat - CPU-Only (sec)	190.55	190.8
Normalized	100%	99.87%

newa

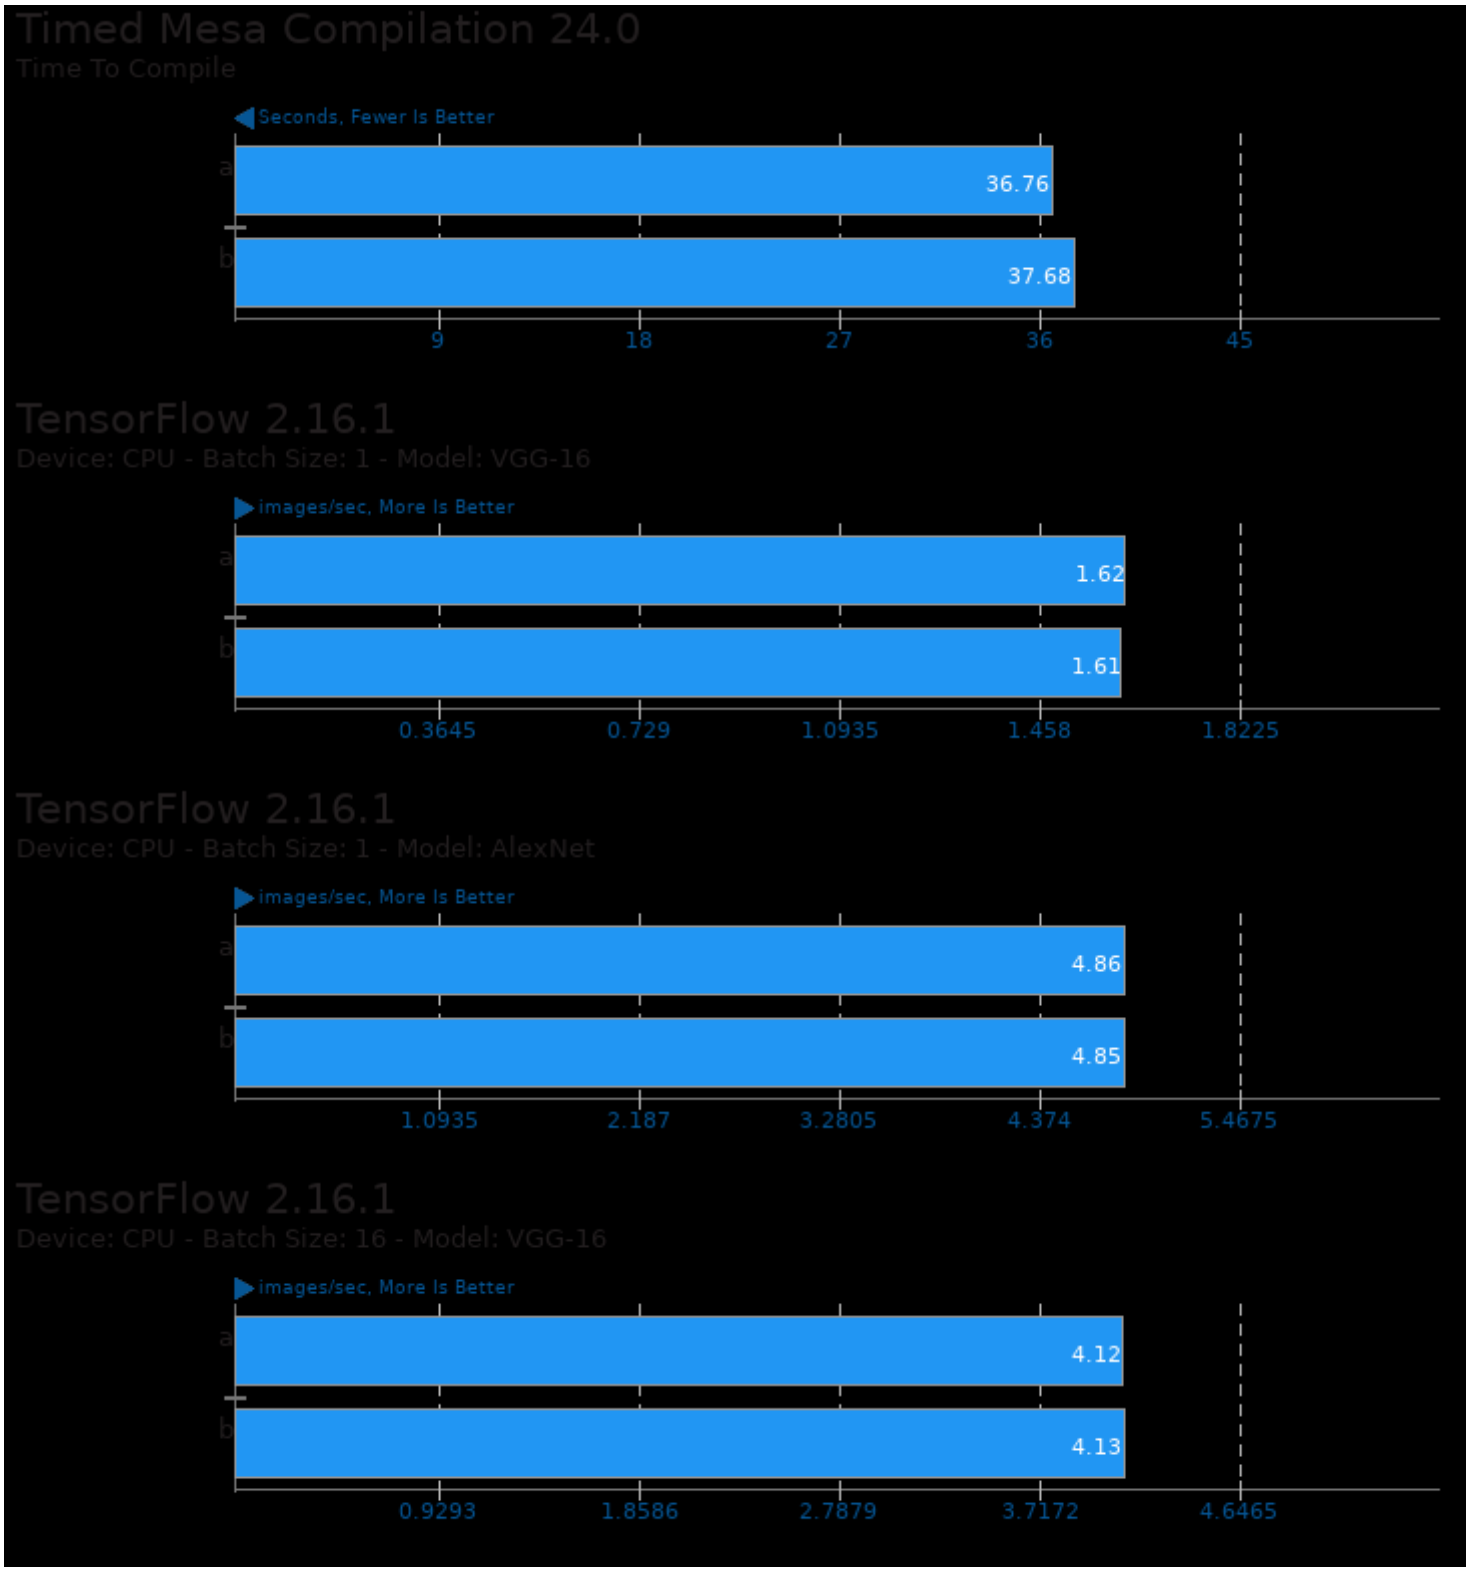
Timed Mesa Compilation 24.0
Time To Compile
Seconds, Fewer Is Better
a
36.76
b
37.68
9
18
27
36
45
TensorFlow 2.16.1
Device: CPU - Batch Size: 1 - Model: VGG-16
images/sec, More Is Better
a
1.62
b
1.61
0.3645
0.729
1.0935
1.458
1.8225
TensorFlow 2.16.1
Device: CPU - Batch Size: 1 - Model: AlexNet
images/sec, More Is Better
a
4.86
b
4.85
1.0935
2.187
3.2805
4.374
5.4675
TensorFlow 2.16.1
Device: CPU - Batch Size: 16 - Model: VGG-16
images/sec, More Is Better
a
4.12
b
4.13
0.9293
1.8586
2.7879
3.7172
4.6465

newa

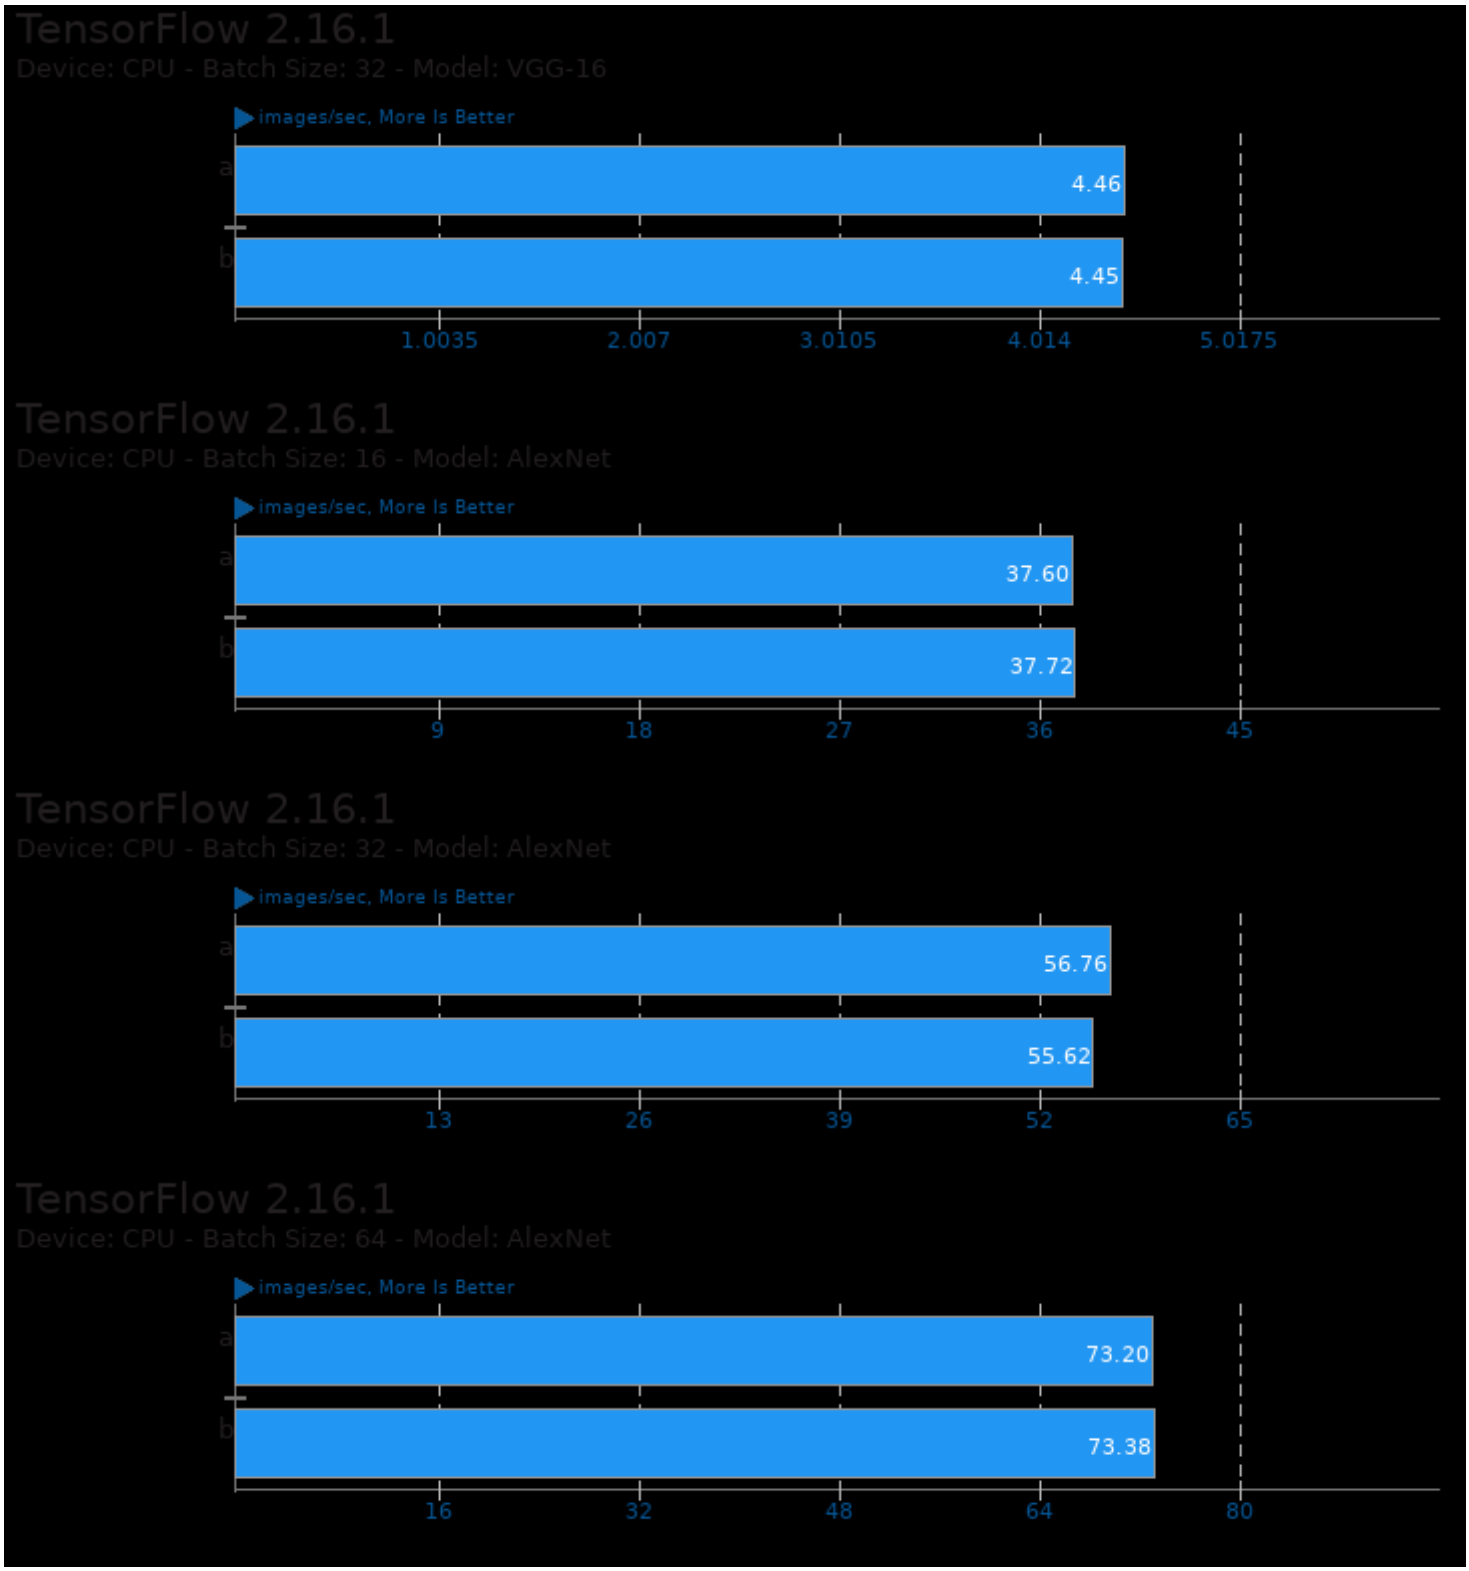
TensorFlow 2.16.1
Device: CPU - Batch Size: 32 - Model: VGG-16
images/sec, More Is Better
a
b
4.46
4.45
1.0035
2.007
3.0105
4.014
5.0175
TensorFlow 2.16.1
Device: CPU - Batch Size: 16 - Model: AlexNet
images/sec, More Is Better
a
b
37.60
37.72
9
18
27
36
45
TensorFlow 2.16.1
Device: CPU - Batch Size: 32 - Model: AlexNet
images/sec, More Is Better
a
b
56.76
55.62
13
26
39
52
65
TensorFlow 2.16.1
Device: CPU - Batch Size: 64 - Model: AlexNet
images/sec, More Is Better
a
b
73.20
73.38
16
32
48
64
80

newa

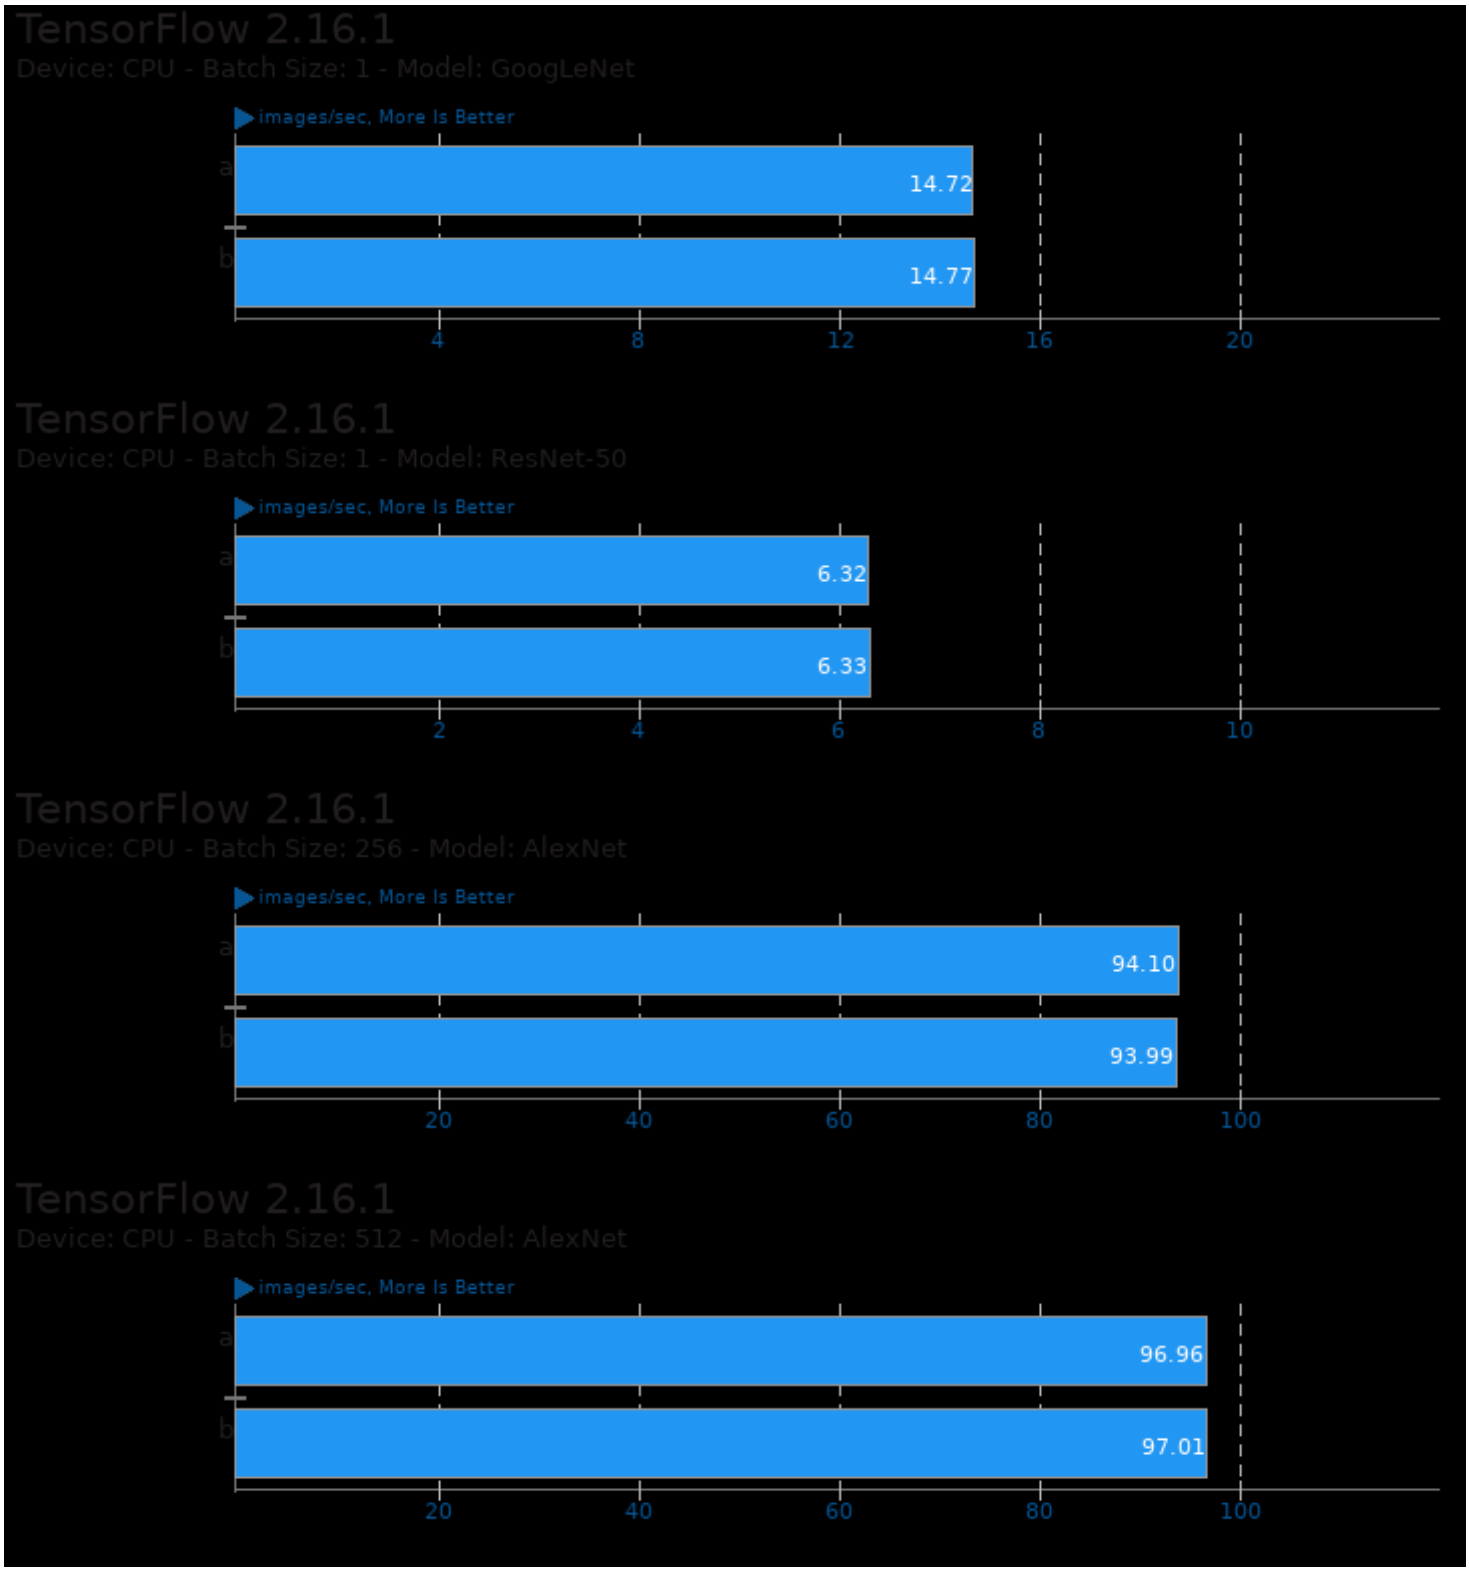
TensorFlow 2.16.1
Device: CPU - Batch Size: 1 - Model: GoogLeNet
images/sec, More Is Better
a
14.72
b
14.77
4
8
12
16
20
TensorFlow 2.16.1
Device: CPU - Batch Size: 1 - Model: ResNet-50
images/sec, More Is Better
a
6.32
b
6.33
2
4
6
8
10
TensorFlow 2.16.1
Device: CPU - Batch Size: 256 - Model: AlexNet
images/sec, More Is Better
a
94.10
b
93.99
20
40
60
80
100
TensorFlow 2.16.1
Device: CPU - Batch Size: 512 - Model: AlexNet
images/sec, More Is Better
a
96.96
b
97.01
20
40
60
80
100

newa

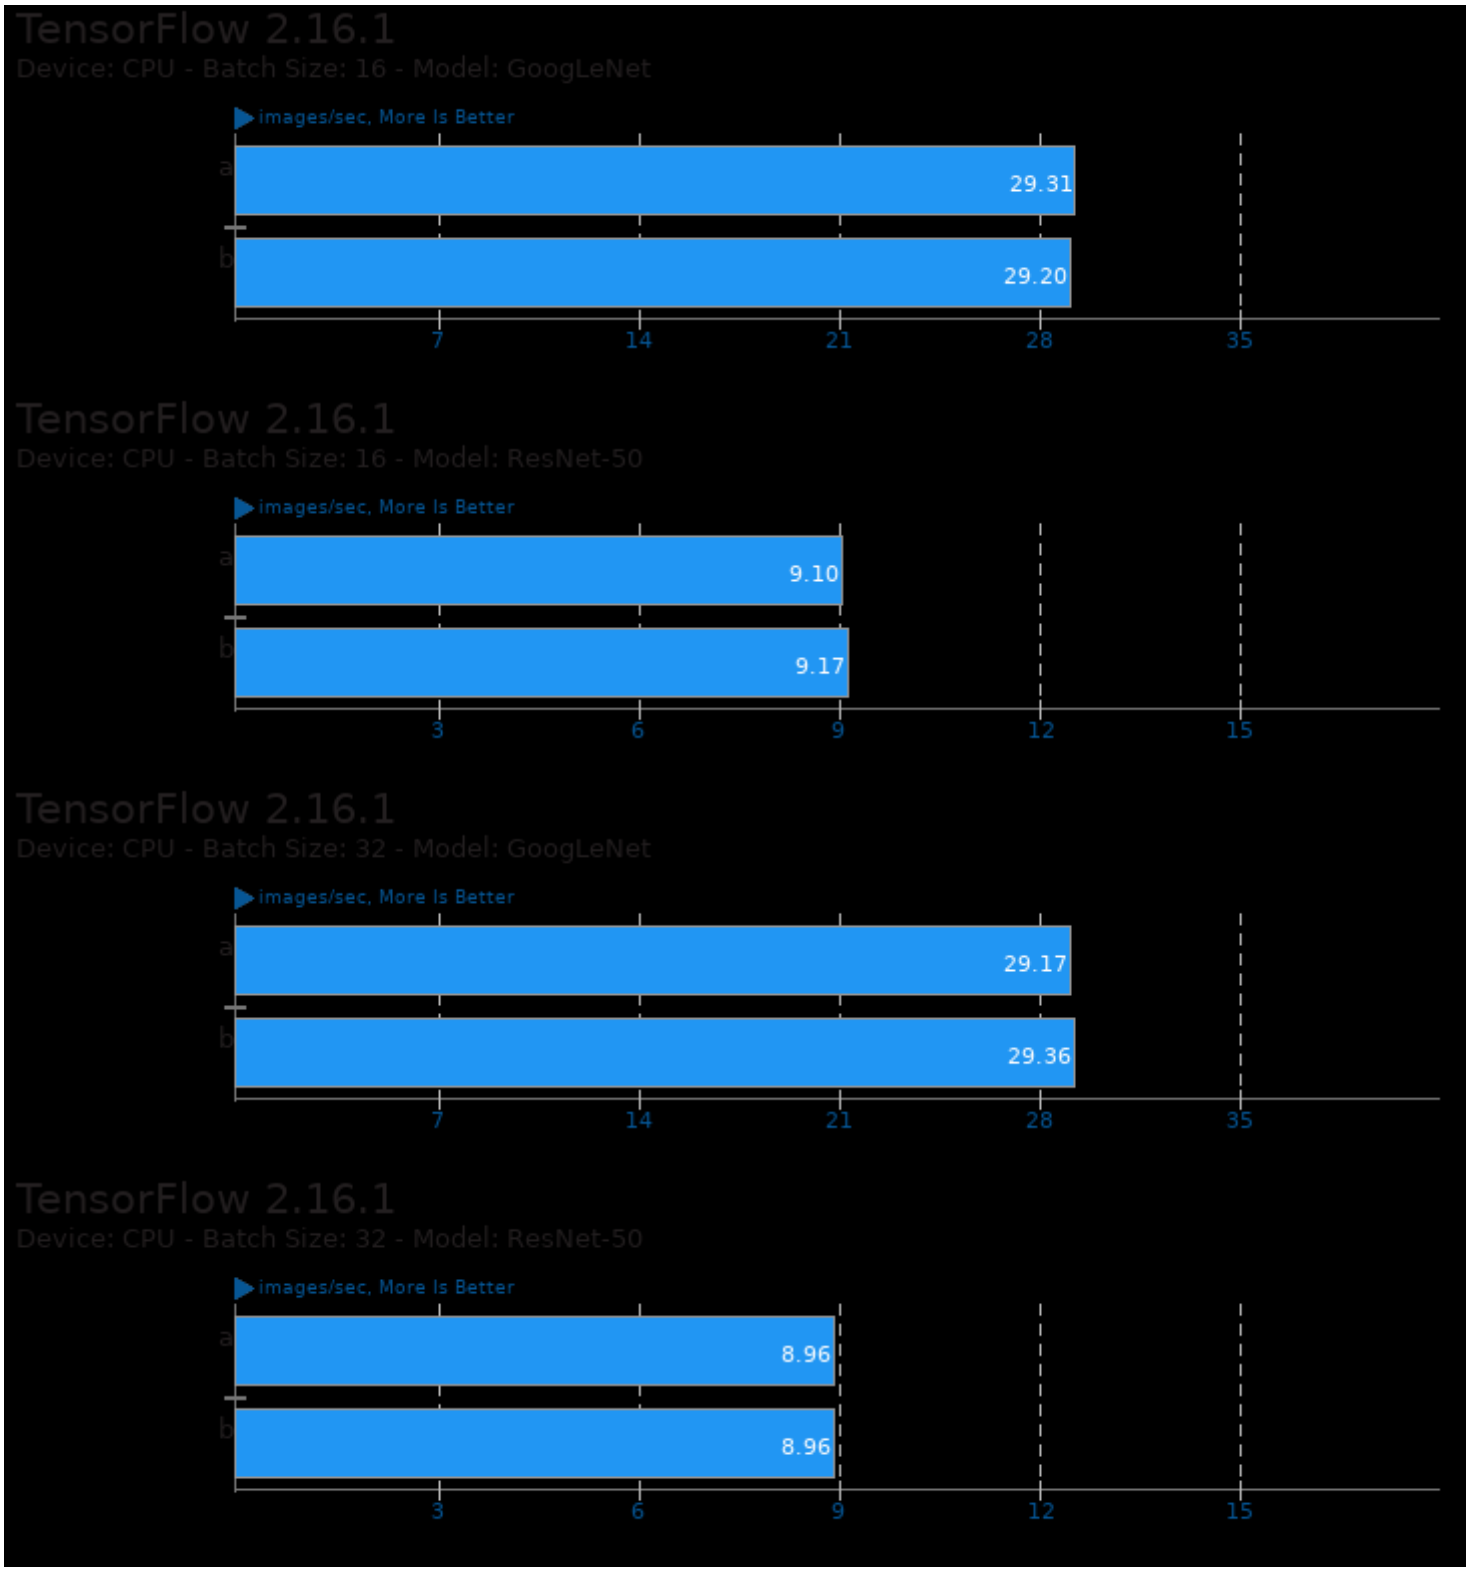
TensorFlow 2.16.1
Device: CPU - Batch Size: 16 - Model: GoogLeNet
images/sec, More Is Better
a
29.31
b
29.20
7
14
21
28
35
TensorFlow 2.16.1
Device: CPU - Batch Size: 16 - Model: ResNet-50
images/sec, More Is Better
a
9.10
b
9.17
3
6
9
12
15
TensorFlow 2.16.1
Device: CPU - Batch Size: 32 - Model: GoogLeNet
images/sec, More Is Better
a
29.17
b
29.36
7
14
21
28
35
TensorFlow 2.16.1
Device: CPU - Batch Size: 32 - Model: ResNet-50
images/sec, More Is Better
a
8.96
b
8.96
3
6
9
12
15

newa

TensorFlow 2.16.1

Device: CPU - Batch Size: 64 - Model: GoogLeNet

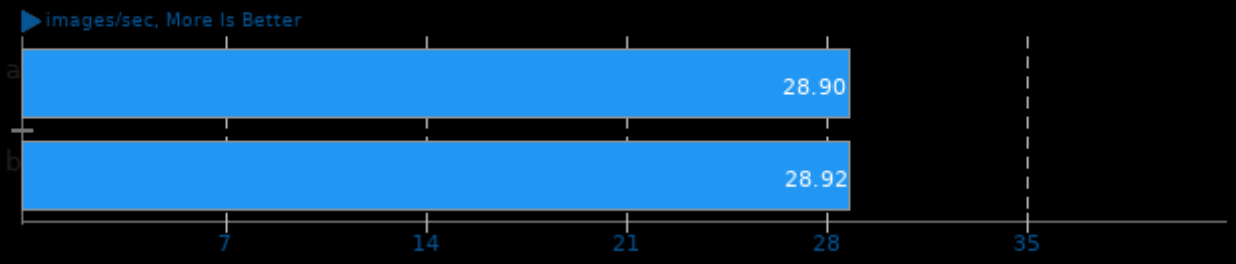

TensorFlow 2.16.1

Device: CPU - Batch Size: 64 - Model: ResNet-50

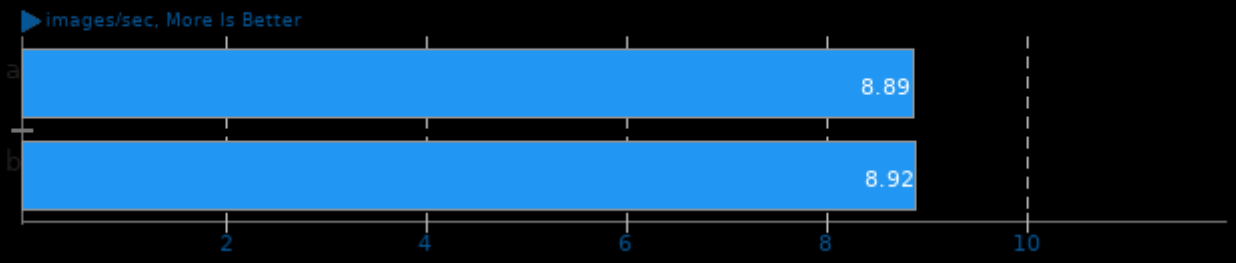

TensorFlow 2.16.1

Device: CPU - Batch Size: 256 - Model: GoogLeNet

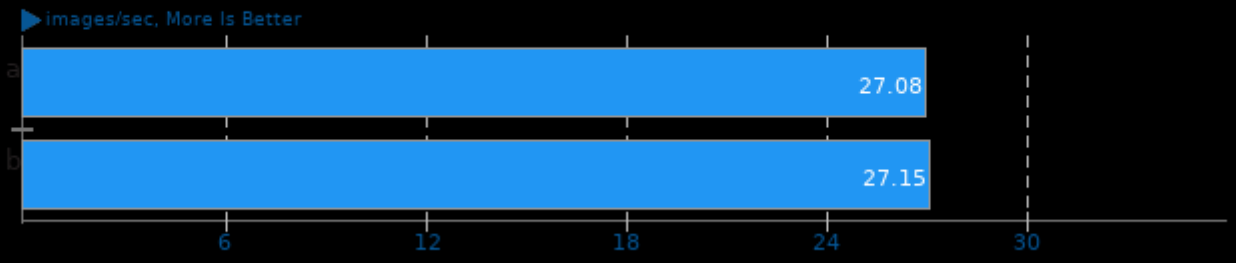

Blender 4.1

Blend File: BMW27 - Compute: CPU-Only

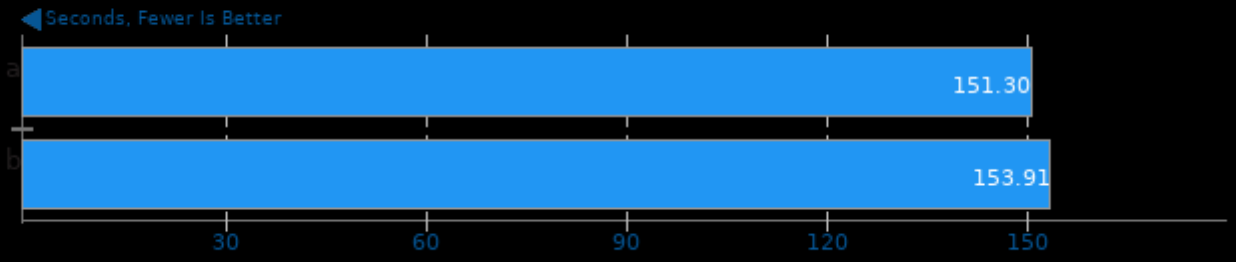

newa

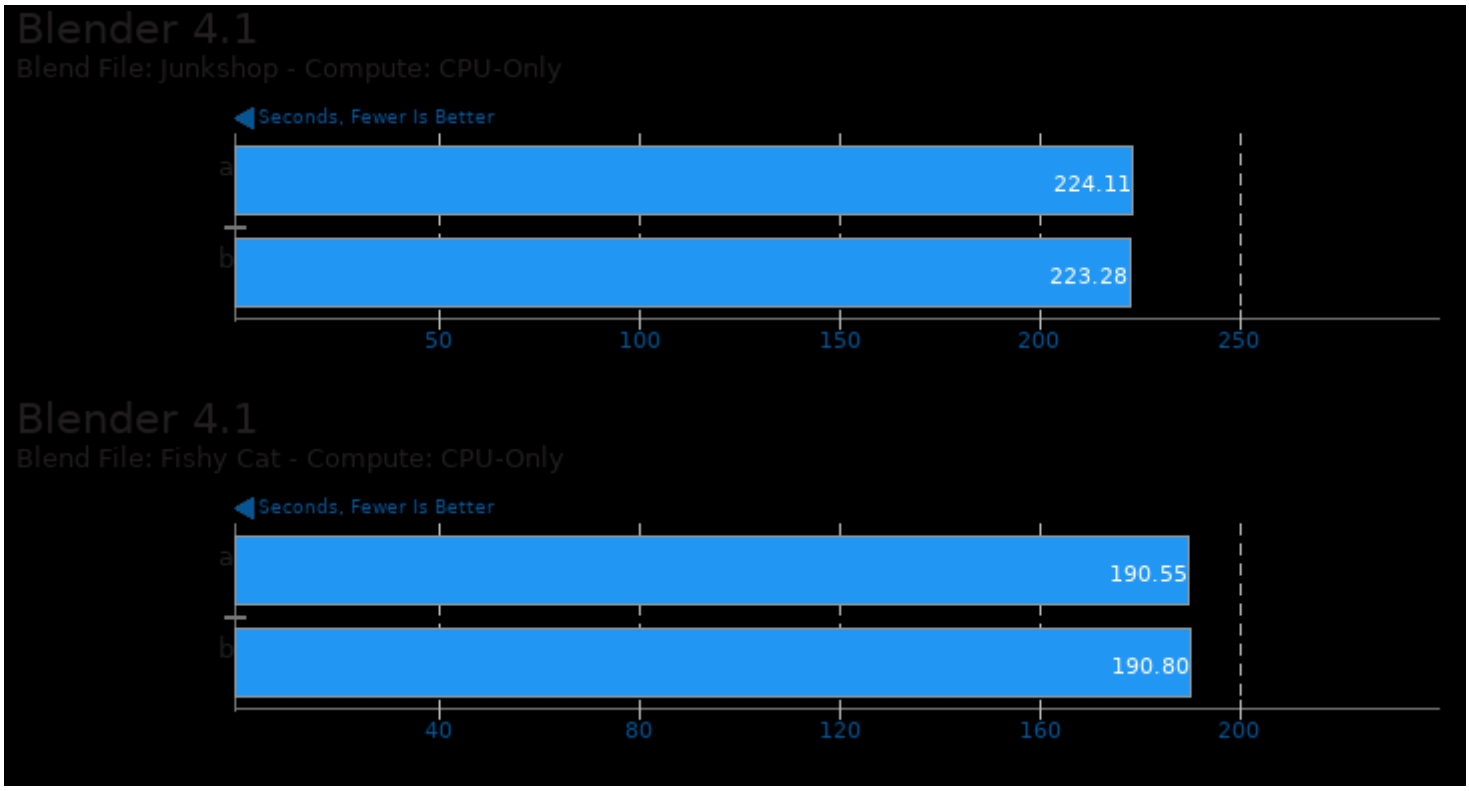
Blender 4.1
Blend File: Junkshop - Compute: CPU-Only
Seconds, Fewer Is Better
a
224.11
b
223.28
50
100
150
200
250
Blender 4.1
Blend File: Fishy Cat - Compute: CPU-Only
Seconds, Fewer Is Better
a
190.55
b
190.80
40
80
120
160
200

newa

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

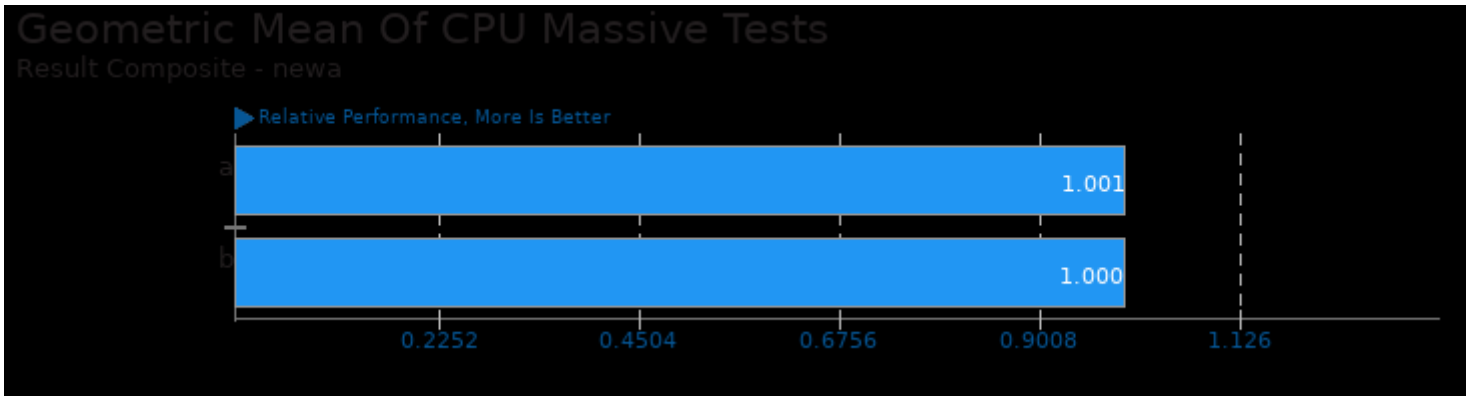

Geometric mean based upon tests: pts/tensorflow and pts/blender

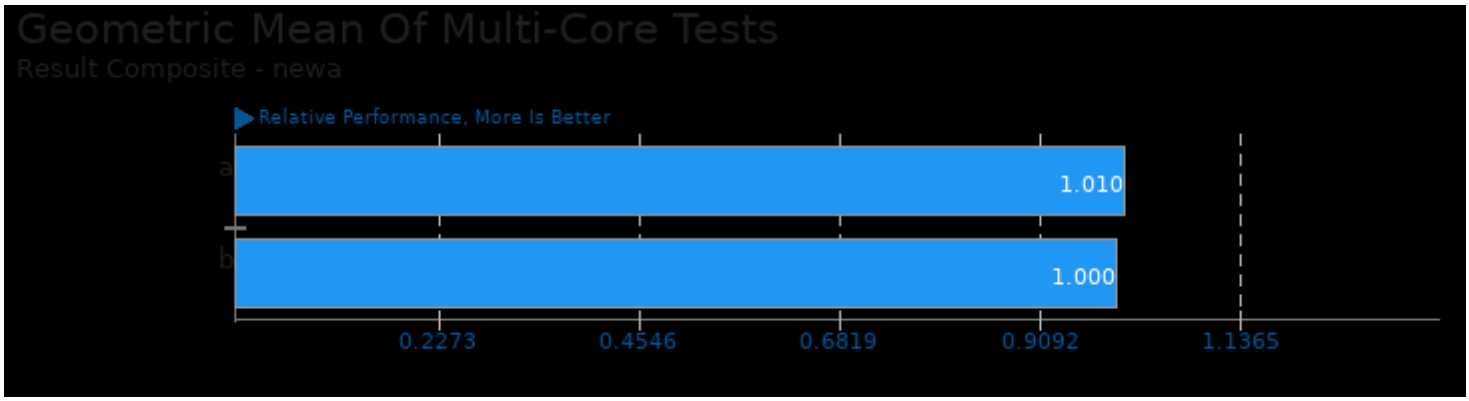

Geometric mean based upon tests: pts/blender and pts/build-mesa

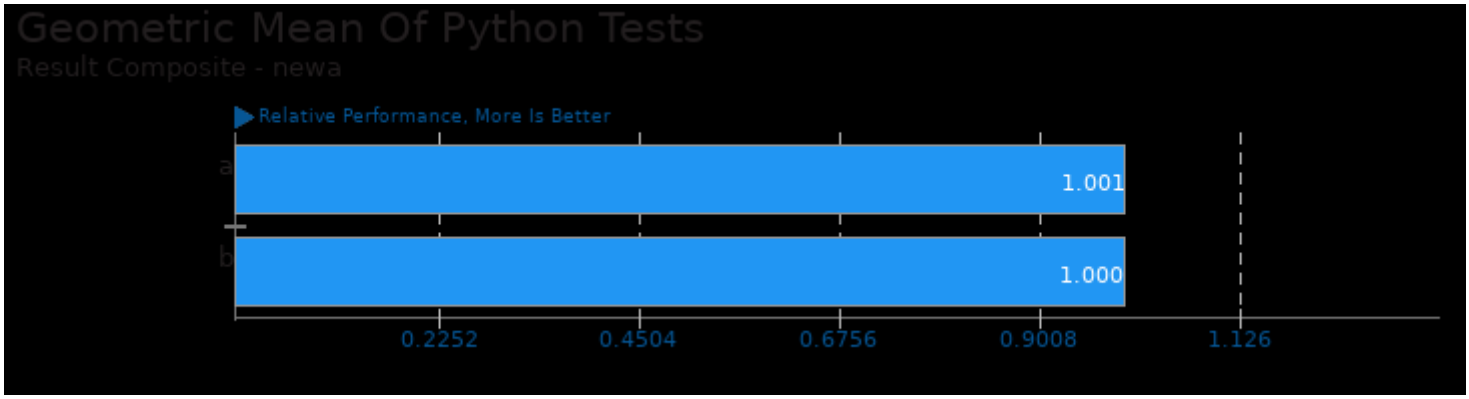

Geometric mean based upon tests: pts/build-mesa and pts/tensorflow

This file was automatically generated via the Phoronix Test Suite benchmarking software on Tuesday, 10 December 2024 22:37.