



apache nginx 5900HX

AMD Ryzen 9 5900HX testing with a ASUS G513QY v1.0 (G513QY.311 BIOS) and ASUS AMD Cezanne 512MB on Ubuntu 21.10 via the Phoronix Test Suite.

Automated Executive Summary

2 had the most wins, coming in first place for 58% of the tests.

Based on the geometric mean of all complete results, the fastest (3) was 1.008x the speed of the slowest (2). 1 was 0.994x the speed of 3 and 2 was 0.998x the speed of 1.

Test Systems:

1

2

3

Processor: AMD Ryzen 9 5900HX @ 3.30GHz (8 Cores / 16 Threads), Motherboard: ASUS G513QY v1.0 (G513QY.311 BIOS), Chipset: AMD Renoir Root Complex, Memory: 16GB, Disk: 512GB SAMSUNG MZVLQ512HBLU-00B00, Graphics: ASUS AMD Cezanne 512MB (2600/1000MHz), Audio: AMD Device ab28, Monitor: LQ156M1JW25, Network: Realtek RTL8111/8168/8411 + MEDIATEK Device 7961

OS: Ubuntu 21.10, Kernel: 5.14.0-051400rc4daily20210804-generic (x86_64) 20210803, Desktop: GNOME Shell 40.2, Display Server: X Server + Wayland, OpenGL: 4.6 Mesa 21.1.4 (LLVM 12.0.1), Vulkan: 1.2.168, Compiler: GCC 10.3.0, File-System: ext4, Screen Resolution: 1920x1080

Kernel Notes: Transparent Huge Pages: madvise

Compiler Notes: --build=x86_64-linux-gnu --disable-vtable-verify --disable-werror --enable-bootstrap --enable-checking=release --enable-clocale=gnu --enable-default-pie --enable-gnu-unique-object --enable-languages=c,ada,c++,go,brig,d,fortran,objc,obj-c++,m2 --enable-libphobos-checking=release --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-link-mutex --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none=/build/gcc-10-lAajkb/gcc-10-10.3.0/debian/tmp-nvptx/usr,amdgc-n-amdhsa=/build/gcc-10-lAajkb/gcc-10-10.3.0/debian/tmp-gcn/usr,hsa --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-build-config=bootstrap-lto-lean --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: 0xa50000c

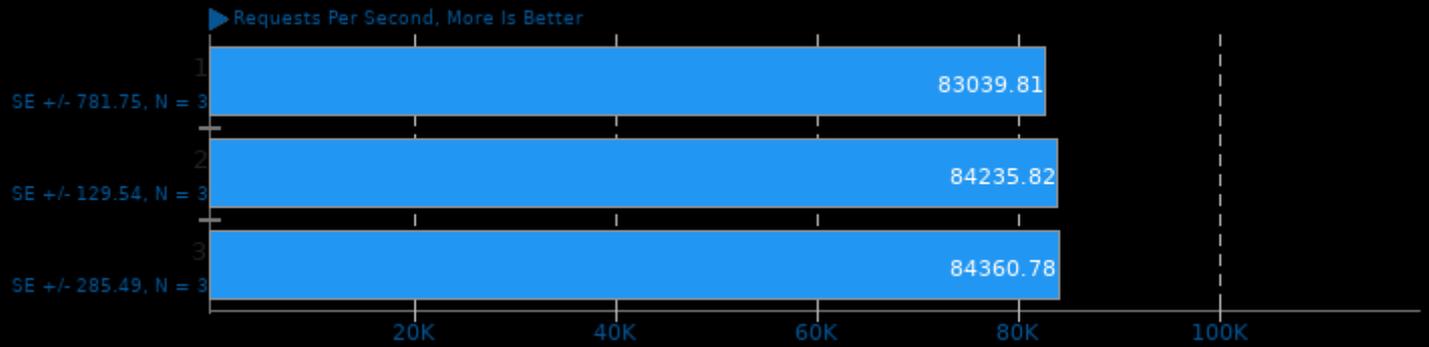
Security Notes: itlb_multihit: Not affected + l1tf: Not affected + mds: Not affected + meltdown: Not affected + spec_store_bypass: Mitigation of SSB disabled via prctl and seccomp + spectre_v1: Mitigation of usercopy/swaps barriers and __user pointer sanitization + spectre_v2: Mitigation of Full AMD retpoline IBPB: conditional IBRS_FW STIBP: always-on RSB filling + srbds: Not affected + tsx_async_abort: Not affected

	1	2	3
nginx - 1 (Reqs/sec)	83040	84236	84361
Normalized	98.43%	99.85%	100%
Standard Deviation	1.6%	0.3%	0.6%
nginx - 20 (Reqs/sec)	533821	538998	528271
Normalized	99.04%	100%	98.01%
Standard Deviation	0.2%	0.3%	0.3%
nginx - 100 (Reqs/sec)	528566	530168	521784
Normalized	99.7%	100%	98.42%
Standard Deviation	0.4%	0.4%	0.2%
nginx - 200 (Reqs/sec)	523275	527056	516856
Normalized	99.28%	100%	98.06%
Standard Deviation	0.1%	0.2%	0.2%
nginx - 500 (Reqs/sec)	512741	518750	508426
Normalized	98.84%	100%	98.01%
Standard Deviation	0.2%	0.2%	0.1%
nginx - 1000 (Reqs/sec)	495502	501035	491750
Normalized	98.9%	100%	98.15%
Standard Deviation	0.4%	0.2%	0.2%
Apache HTTP Server - 1 (Reqs/sec)	10039	9190	11380
Normalized	88.22%	80.75%	100%
Standard Deviation	21%	0.6%	18.2%
Apache HTTP Server - 20 (Reqs/sec)	68521	69505	70955
Normalized	96.57%	97.96%	100%
Standard Deviation	0.8%	0.8%	0.2%
Apache HTTP Server - 100 (Reqs/sec)	49261	49766	48804
Normalized	98.99%	100%	98.07%
Standard Deviation	0.8%	0.4%	0.4%
Apache HTTP Server - 200 (Reqs/sec)	37479	37374	37498
Normalized	99.95%	99.67%	100%

	Standard Deviation	0.6%	0.1%	0.8%
Apache HTTP Server - 500 (Reqs/sec)	38618	37483	37324	
	Normalized	100%	97.06%	96.65%
	Standard Deviation	0.2%	0.3%	0.3%
Apache HTTP Server - 1000 (Reqs/sec)	36997	37373	36652	
	Normalized	98.99%	100%	98.07%
	Standard Deviation	0.1%	0.8%	0.4%

nginx 1.21.1

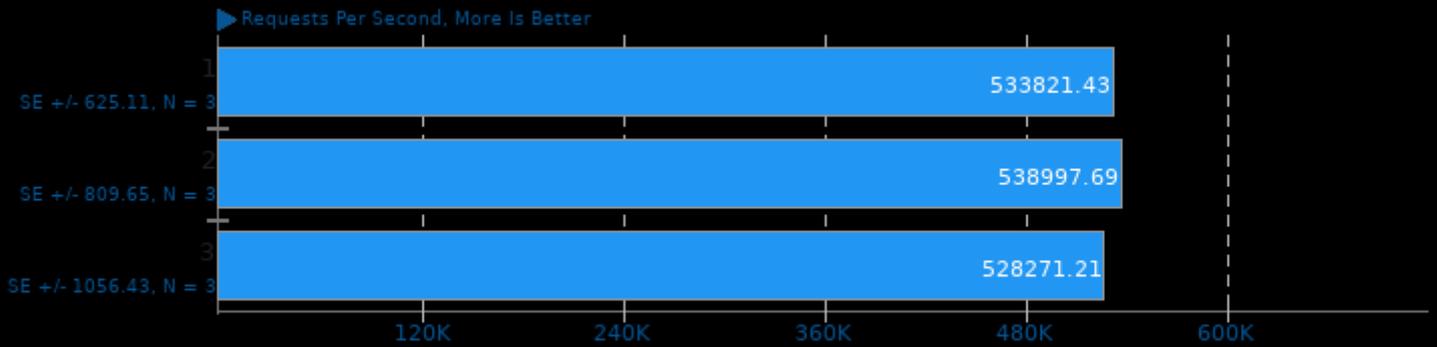
Concurrent Requests: 1



1. (CC) gcc options: -ldl -lpthread -lcrypt -lz -O3 -march=native

nginx 1.21.1

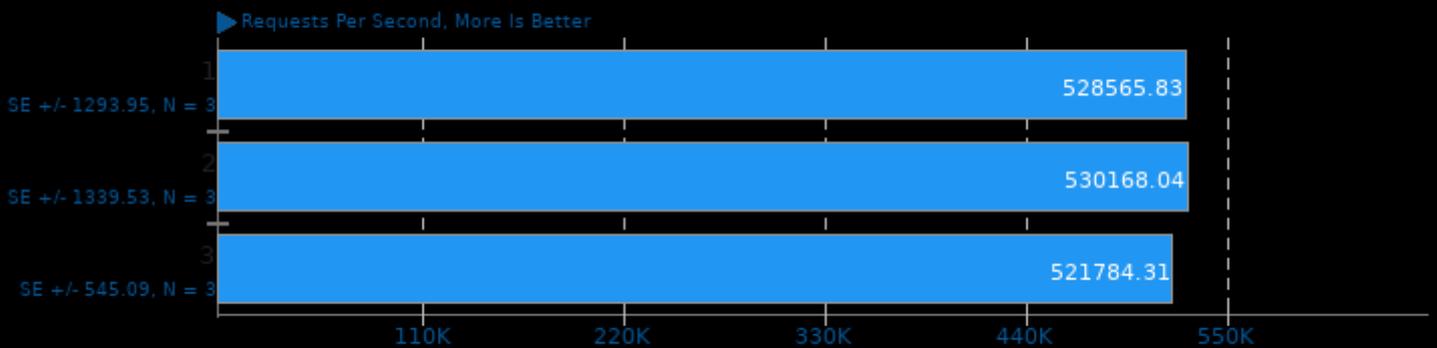
Concurrent Requests: 20



1. (CC) gcc options: -ldl -lpthread -lcrypt -lz -O3 -march=native

nginx 1.21.1

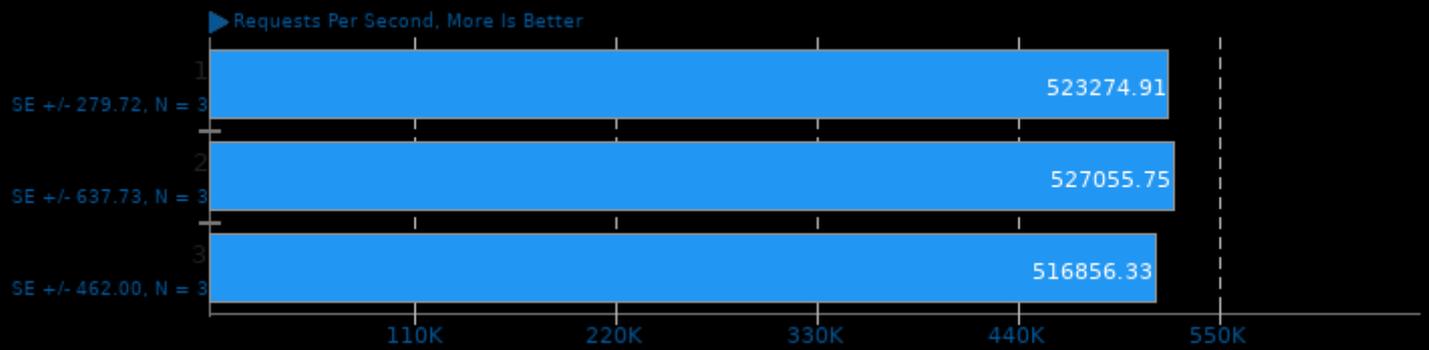
Concurrent Requests: 100



1. (CC) gcc options: -ldl -lpthread -lcrypt -lz -O3 -march=native

nginx 1.21.1

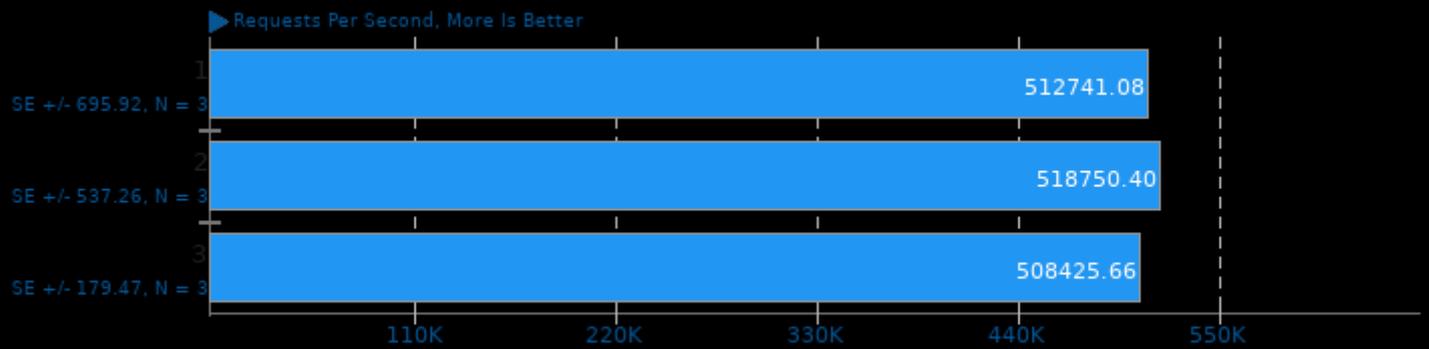
Concurrent Requests: 200



1. (CC) gcc options: -ldl -lpthread -lcrypt -lz -O3 -march=native

nginx 1.21.1

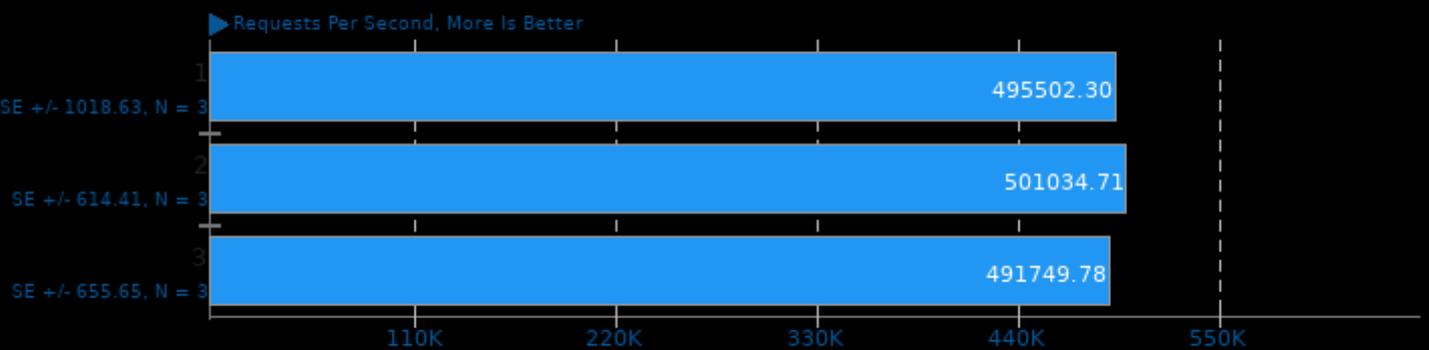
Concurrent Requests: 500



1. (CC) gcc options: -ldl -lpthread -lcrypt -lz -O3 -march=native

nginx 1.21.1

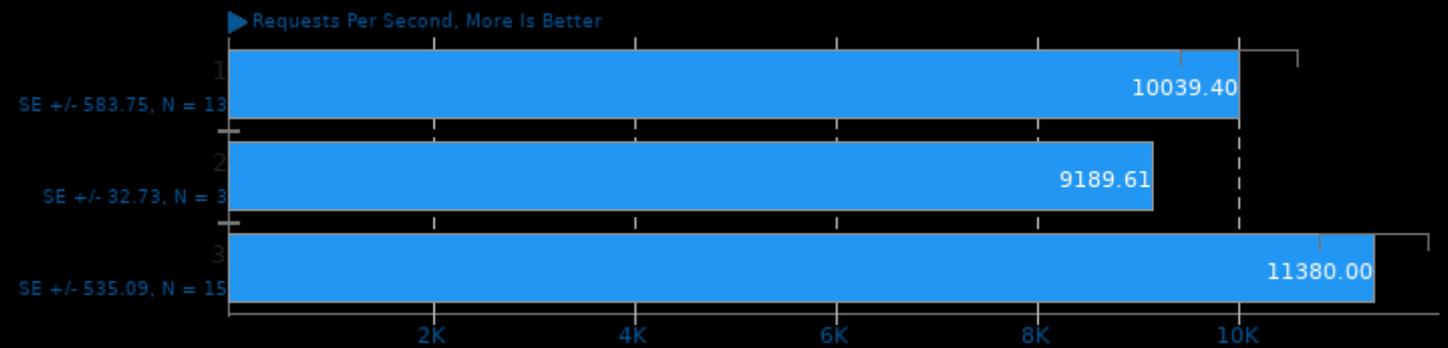
Concurrent Requests: 1000



1. (CC) gcc options: -ldl -lpthread -lcrypt -lz -O3 -march=native

Apache HTTP Server 2.4.48

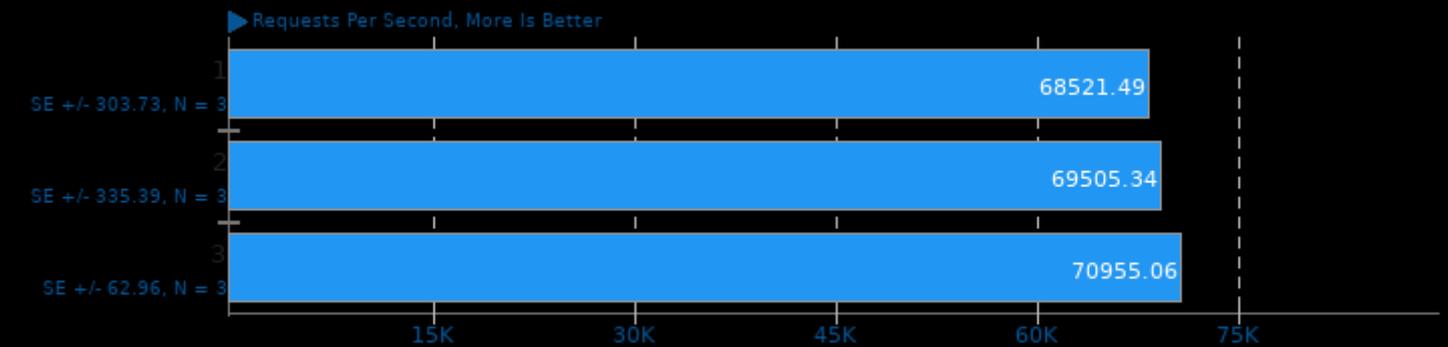
Concurrent Requests: 1



1. (CC) gcc options: -shared -fPIC -O2 -pthread

Apache HTTP Server 2.4.48

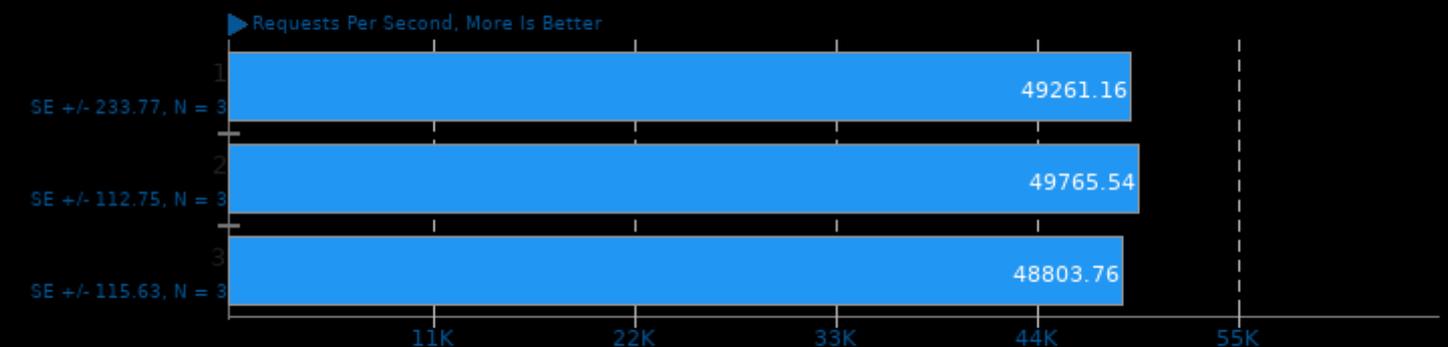
Concurrent Requests: 20



1. (CC) gcc options: -shared -fPIC -O2 -pthread

Apache HTTP Server 2.4.48

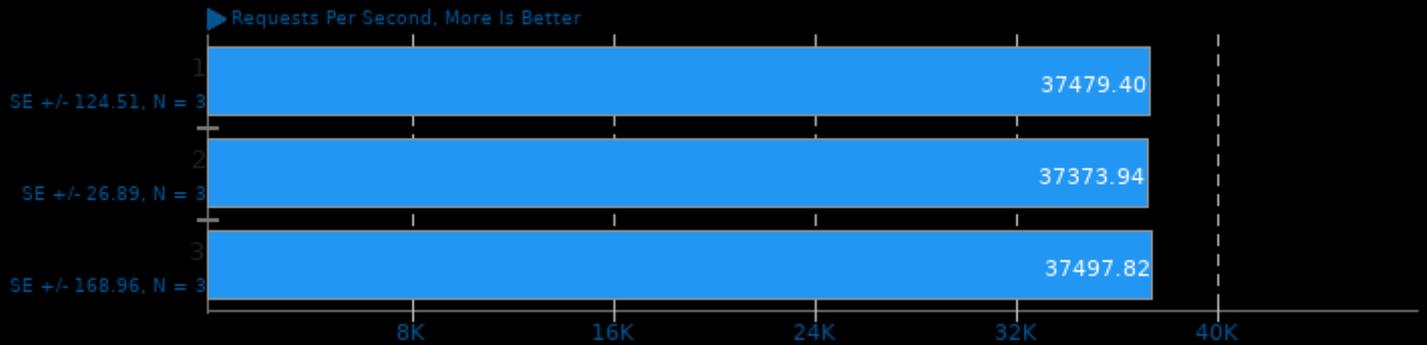
Concurrent Requests: 100



1. (CC) gcc options: -shared -fPIC -O2 -pthread

Apache HTTP Server 2.4.48

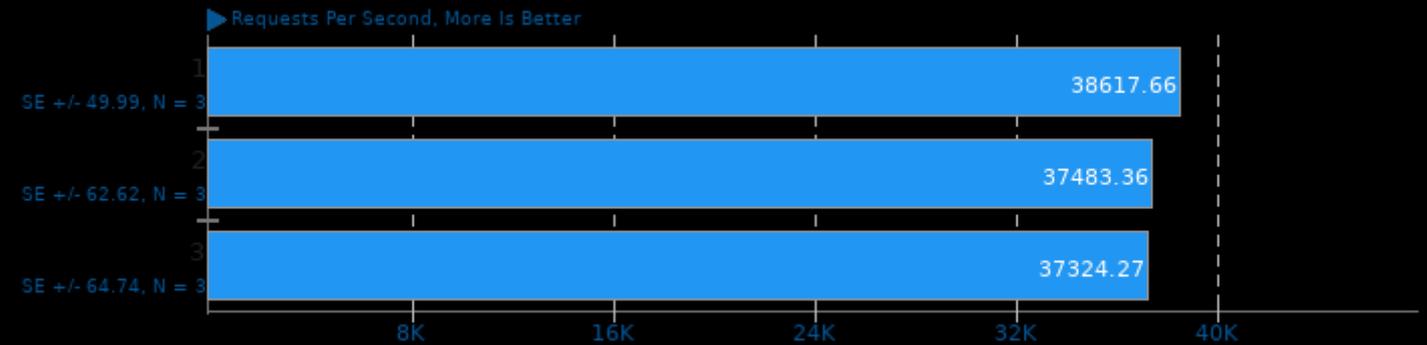
Concurrent Requests: 200



1. (CC) gcc options: -shared -fPIC -O2 -pthread

Apache HTTP Server 2.4.48

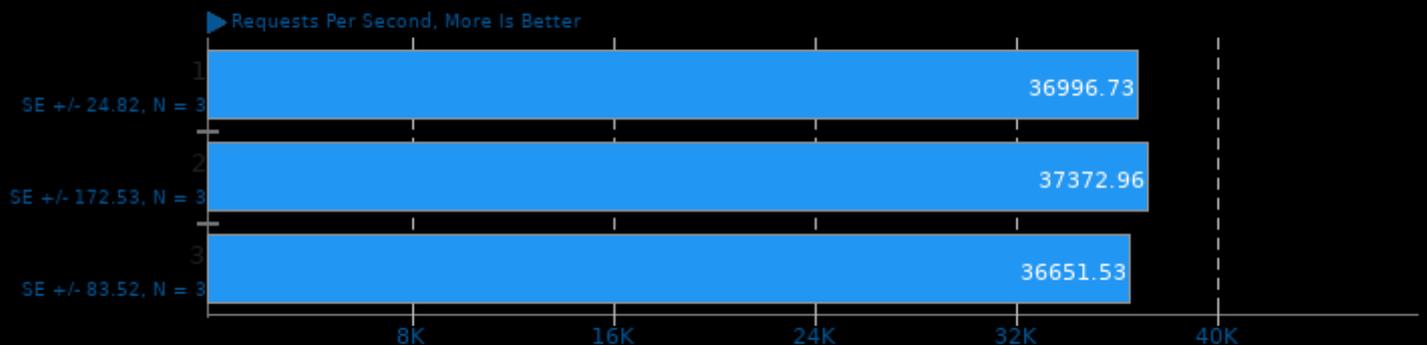
Concurrent Requests: 500



1. (CC) gcc options: -shared -fPIC -O2 -pthread

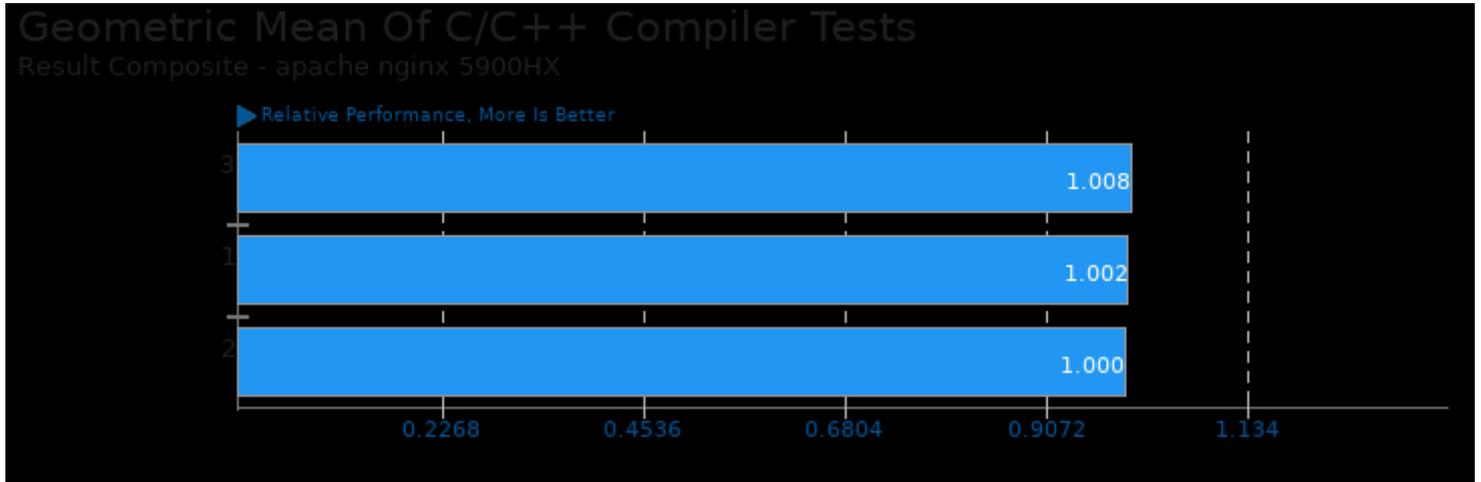
Apache HTTP Server 2.4.48

Concurrent Requests: 1000

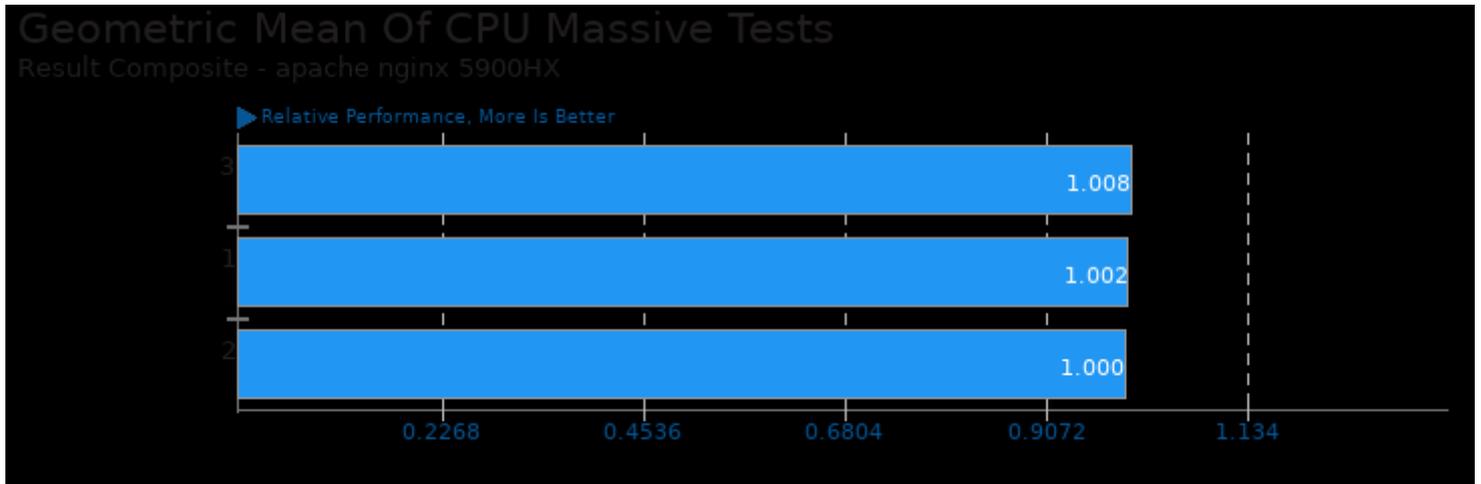


1. (CC) gcc options: -shared -fPIC -O2 -pthread

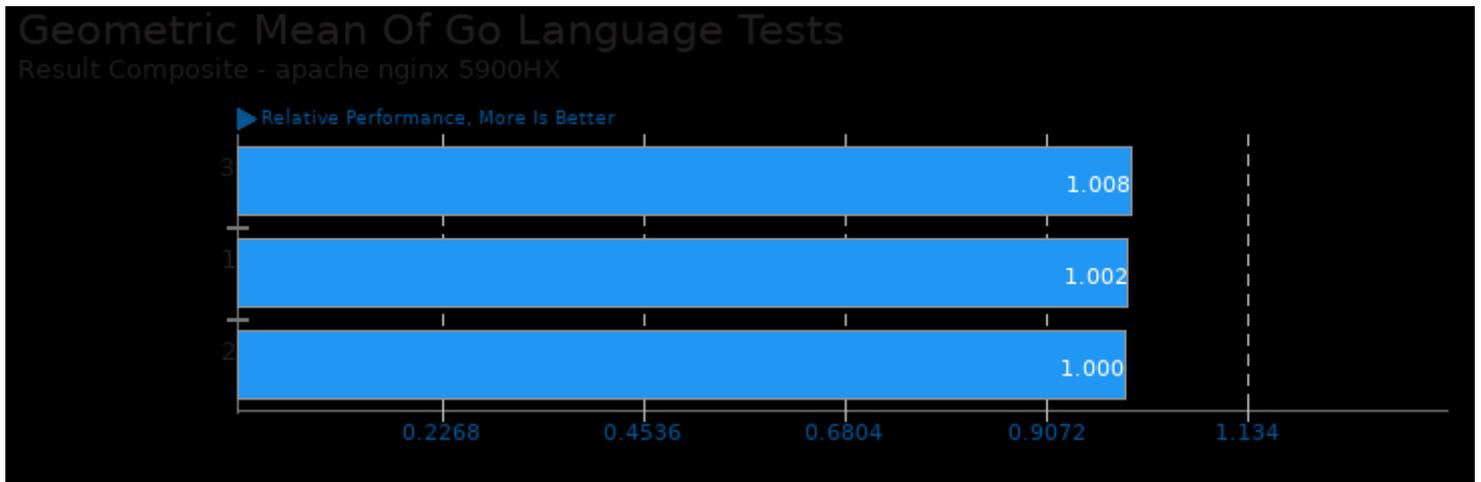
These geometric means are based upon test groupings / test suites for this result file.



Geometric mean based upon tests: pts/apache and pts/nginx



Geometric mean based upon tests: pts/apache and pts/nginx



Geometric mean based upon tests: pts/nginx and pts/apache



Geometric mean based upon tests: pts/apache and pts/nginx

This file was automatically generated via the Phoronix Test Suite benchmarking software on Thursday, 28 March 2024 07:52.