



NVIDIA AMD Linux GPU Compute December 2018

NVIDIA and AMD GPU Linux compute benchmarks December 2018 by Michael Larabel for a future article..

Automated Executive Summary

RTX 2080 Ti had the most wins, coming in first place for 94% of the tests.

Based on the geometric mean of all complete results, the fastest (RTX 2080 Ti) was 2.95x the speed of the slowest (R9 Fury). RTX 2080 was 0.783x the speed of RTX 2080 Ti, RTX 2070 was 0.924x the speed of RTX 2080, GTX 1080 Ti was 0.758x the speed of RTX 2070, RX Vega 64 was 0.84x the speed of GTX 1080 Ti, RX Vega 56 was 0.89x the speed of RX Vega 64, GTX 1080 was 0.986x the speed of RX Vega 56, GTX 1070 Ti was 0.919x the speed of GTX 1080, GTX 1070 was 0.952x the speed of GTX 1070 Ti, R9 Fury was 0.957x the speed of GTX 1070.

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

NVIDIA GPU Cloud TensorFlow (System Power Consumption Monitor) at 8.016x
NVIDIA GPU Cloud TensorFlow (System Power Consumption Monitor) at 7.906x
NVIDIA GPU Cloud TensorFlow (System Power Consumption Monitor) at 7.862x
NVIDIA GPU Cloud TensorFlow (System Power Consumption Monitor) at 7.735x
CUDA Mini-Nbody (System Power Consumption Monitor) at 7.689x
NVIDIA GPU Cloud TensorFlow (System Power Consumption Monitor) at 7.602x
NVIDIA GPU Cloud TensorFlow (System Power Consumption Monitor) at 7.542x

NVIDIA GPU Cloud TensorFlow (System Power Consumption Monitor) at 7.393x
cl-mem (System Power Consumption Monitor) at 7.363x
NVIDIA GPU Cloud TensorFlow (System Power Consumption Monitor) at 7.056x.

Test Systems:

GTX 1070

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0602 BIOS), Chipset: Intel Cannon Lake PCH Shared SRAM, Memory: 16384MB, Disk: 2000GB SABRENT + Samsung SSD 970 EVO 250GB, Graphics: NVIDIA GeForce GTX 1070 8GB (1506/4006MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel Connection

OS: Ubuntu 18.04, Kernel: 4.19.5-041905-generic (x86_64), Desktop: GNOME Shell 3.28.3, Display Server: X Server 1.19.6, Display Driver: NVIDIA 415.22, OpenGL: 4.6.0, OpenCL: OpenCL 1.2 CUDA 10.0.132, Vulkan: 1.1.84, Compiler: GCC 7.3.0 + CUDA 10.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86_64-linux-gnu --disable-vtable-verify --disable-werror --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel_pstate performance

OpenCL Notes: GPU Compute Cores: 1920

Security Notes: __user pointer sanitization + Full generic retpoline IBPB IBRS_FW + SSB disabled via prctl and seccomp

GTX 1070 Ti

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0602 BIOS), Chipset: Intel Cannon Lake PCH Shared SRAM, Memory: 16384MB, Disk: 2000GB SABRENT + Samsung SSD 970 EVO 250GB, Graphics: Zotac NVIDIA GeForce GTX 1070 Ti 8GB (1607/4006MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel Connection

OS: Ubuntu 18.04, Kernel: 4.19.5-041905-generic (x86_64), Desktop: GNOME Shell 3.28.3, Display Server: X Server 1.19.6, Display Driver: NVIDIA 415.22, OpenGL: 4.6.0, OpenCL: OpenCL 1.2 CUDA 10.0.132, Vulkan: 1.1.84, Compiler: GCC 7.3.0 + CUDA 10.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86_64-linux-gnu --disable-vtable-verify --disable-werror --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v

Processor Notes: Scaling Governor: intel_pstate performance

OpenCL Notes: GPU Compute Cores: 2432

Security Notes: __user pointer sanitization + Full generic retpoline IBPB IBRS_FW + SSB disabled via prctl and seccomp

GTX 1080

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0602 BIOS), Chipset: Intel Cannon Lake PCH Shared SRAM, Memory: 16384MB, Disk: 2000GB SABRENT + Samsung SSD 970 EVO 250GB, Graphics: NVIDIA GeForce GTX 1080 8GB (1607/5005MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel Connection

OS: Ubuntu 18.04, Kernel: 4.19.5-041905-generic (x86_64), Desktop: GNOME Shell 3.28.3, Display Server: X Server 1.19.6, Display Driver: NVIDIA 415.22, OpenGL: 4.6.0, OpenCL: OpenCL 1.2 CUDA 10.0.132, Vulkan: 1.1.84, Compiler: GCC 7.3.0 + CUDA 10.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86_64-linux-gnu --disable-vtable-verify --disable-werror --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v
Processor Notes: Scaling Governor: intel_pstate performance
OpenCL Notes: GPU Compute Cores: 2560
Security Notes: __user pointer sanitization + Full generic retpoline IBPB IBRS_FW + SSB disabled via prctl and seccomp

GTX 1080 Ti

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0602 BIOS), Chipset: Intel Cannon Lake PCH Shared SRAM, Memory: 16384MB, Disk: 2000GB SABRENT + Samsung SSD 970 EVO 250GB, Graphics: NVIDIA GeForce GTX 1080 Ti 11GB (1480/5508MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel Connection

OS: Ubuntu 18.04, Kernel: 4.19.5-041905-generic (x86_64), Desktop: GNOME Shell 3.28.3, Display Server: X Server 1.19.6, Display Driver: NVIDIA 415.22, OpenGL: 4.6.0, OpenCL: OpenCL 1.2 CUDA 10.0.132, Vulkan: 1.1.84, Compiler: GCC 7.3.0 + CUDA 10.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86_64-linux-gnu --disable-vtable-verify --disable-werror --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v
Processor Notes: Scaling Governor: intel_pstate performance
OpenCL Notes: GPU Compute Cores: 3584
Security Notes: __user pointer sanitization + Full generic retpoline IBPB IBRS_FW + SSB disabled via prctl and seccomp

RTX 2070

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0602 BIOS), Chipset: Intel Cannon Lake PCH Shared SRAM, Memory: 16384MB, Disk: 2000GB SABRENT + Samsung SSD 970 EVO 250GB, Graphics: eVGA NVIDIA GeForce RTX 2070 8GB (1410/7000MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel Connection

OS: Ubuntu 18.04, Kernel: 4.19.5-041905-generic (x86_64), Desktop: GNOME Shell 3.28.3, Display Server: X Server 1.19.6, Display Driver: NVIDIA 415.22, OpenGL: 4.6.0, OpenCL: OpenCL 1.2 CUDA 10.0.132, Vulkan: 1.1.84, Compiler: GCC 7.3.0 + CUDA 10.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86_64-linux-gnu --disable-vtable-verify --disable-werror --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v
Processor Notes: Scaling Governor: intel_pstate performance
OpenCL Notes: GPU Compute Cores: 2304
Security Notes: __user pointer sanitization + Full generic retpoline IBPB IBRS_FW + SSB disabled via prctl and seccomp

RTX 2080

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0602 BIOS), Chipset: Intel Cannon Lake PCH Shared SRAM, Memory: 16384MB, Disk: 2000GB SABRENT + Samsung SSD 970 EVO 250GB, Graphics: Zotac NVIDIA GeForce RTX 2080 8GB (1515/7000MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel Connection

OS: Ubuntu 18.04, Kernel: 4.19.5-041905-generic (x86_64), Desktop: GNOME Shell 3.28.3, Display Server: X Server 1.19.6, Display Driver: NVIDIA 415.22, OpenGL: 4.6.0, OpenCL: OpenCL 1.2 CUDA 10.0.132, Vulkan: 1.1.84, Compiler: GCC 7.3.0 + CUDA 10.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86_64-linux-gnu --disable-vtable-verify --disable-werror --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v
Processor Notes: Scaling Governor: intel_pstate performance
OpenCL Notes: GPU Compute Cores: 2944
Security Notes: __user pointer sanitization + Full generic retpoline IBPB IBRS_FW + SSB disabled via prctl and seccomp

RTX 2080 Ti

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0602 BIOS), Chipset: Intel Cannon Lake PCH Shared SRAM, Memory: 16384MB, Disk: 2000GB SABRENT + Samsung SSD 970 EVO 250GB, Graphics: NVIDIA GeForce RTX 2080 Ti 11GB (1350/7000MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel Connection

OS: Ubuntu 18.04, Kernel: 4.19.5-041905-generic (x86_64), Desktop: GNOME Shell 3.28.3, Display Server: X Server 1.19.6, Display Driver: NVIDIA 415.22, OpenGL: 4.6.0, OpenCL: OpenCL 1.2 CUDA 10.0.132, Vulkan: 1.1.84, Compiler: GCC 7.3.0 + CUDA 10.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86_64-linux-gnu --disable-vtable-verify --disable-werror --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v
Processor Notes: Scaling Governor: intel_pstate performance
OpenCL Notes: GPU Compute Cores: 4352
Security Notes: __user pointer sanitization + Full generic retpoline IBPB IBRS_FW + SSB disabled via prctl and seccomp

R9 Fury

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0602 BIOS), Chipset: Intel Cannon Lake PCH Shared SRAM, Memory: 16384MB, Disk: 2000GB SABRENT + Samsung SSD 970 EVO 250GB, Graphics: Sapphire AMD Radeon R9 FURY / NANO 4GB (1000/500MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel Connection

OS: Ubuntu 18.04, Kernel: 4.19.5-041905-generic (x86_64), Desktop: GNOME Shell 3.28.3, Display Server: X Server 1.19.6, OpenGL: 4.5 Mesa 19.0.0-devel padoka PPA (LLVM 8.0.0), OpenCL: OpenCL 2.1 AMD-APP (2679.0), Vulkan: 1.1.70, Compiler: GCC 7.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86_64-linux-gnu --disable-vtable-verify --disable-werror --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++ --enable-libmpx --enable-libstdcxx-debug --enable-libstdcxx-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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 --with-tune=generic --without-cuda-driver -v
Processor Notes: Scaling Governor: intel_pstate performance
Security Notes: __user pointer sanitization + Full generic retpoline IBPB IBRS_FW + SSB disabled via prctl and seccomp

RX Vega 56

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0602 BIOS), Chipset: Intel Cannon Lake PCH Shared SRAM, Memory: 16384MB, Disk: 2000GB SABRENT + Samsung SSD 970 EVO 250GB, Graphics: AMD Radeon RX Vega 8GB (1590/800MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel Connection

OS: Ubuntu 18.04, Kernel: 4.19.5-041905-generic (x86_64), Desktop: GNOME Shell 3.28.3, Display Server: X Server 1.19.6, OpenGL: 4.5 Mesa 19.0.0-devel padoka PPA (LLVM 8.0.0), OpenCL: OpenCL 2.1 AMD-APP (2679.0), Vulkan: 1.1.70, Compiler: GCC 7.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86_64-linux-gnu --disable-vtable-verify --disable-werror --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++ --enable-libmpx --enable-libstdc++-debug --enable-libstdc++-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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-libstdc++-abi=new --with-gcc-major-version-only --with-multilib-list=m32,m64,mx32 --with-target-system-zlib --with-tune=generic --without-cuda-driver -v
Processor Notes: Scaling Governor: intel_pstate performance
Security Notes: __user pointer sanitization + Full generic retpoline IBPB IBRS_FW + SSB disabled via prctl and seccomp

RX Vega 64

Processor: Intel Core i9-9900K @ 5.00GHz (8 Cores / 16 Threads), Motherboard: ASUS PRIME Z390-A (0602 BIOS), Chipset: Intel Cannon Lake PCH Shared SRAM, Memory: 16384MB, Disk: 2000GB SABRENT + Samsung SSD 970 EVO 250GB, Graphics: AMD Radeon RX Vega 8GB (1630/945MHz), Audio: Realtek ALC1220, Monitor: Acer B286HK, Network: Intel Connection

OS: Ubuntu 18.04, Kernel: 4.19.5-041905-generic (x86_64), Desktop: GNOME Shell 3.28.3, Display Server: X Server 1.19.6, OpenGL: 4.5 Mesa 19.0.0-devel padoka PPA (LLVM 8.0.0), OpenCL: OpenCL 2.1 AMD-APP (2679.0), Vulkan: 1.1.70, Compiler: GCC 7.3.0, File-System: ext4, Screen Resolution: 3840x2160

Compiler Notes: --build=x86_64-linux-gnu --disable-vtable-verify --disable-werror --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++ --enable-libmpx --enable-libstdc++-debug --enable-libstdc++-time=yes --enable-multiarch --enable-multilib --enable-nls --enable-objc-gc=auto --enable-offload-targets=nvptx-none --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-libstdc++-abi=new --with-gcc-major-version-only --with-multilib-list=m32,m64,mx32 --with-target-system-zlib --with-tune=generic --without-cuda-driver -v
Processor Notes: Scaling Governor: intel_pstate performance
Security Notes: __user pointer sanitization + Full generic retpoline IBPB IBRS_FW + SSB disabled via prctl and seccomp


	GTX 1070	GTX 1070 Ti	GTX 1080	GTX 1080 Ti	RTX 2070	RTX 2080	RTX 2080 Ti	R9 Fury	RX Vega 56	RX Vega 64
SHOC Scalable Heterogeneous Computing - OpenCL - FFT SP	452	497	575	972	998	1083	1443	846	926	1070
Normalized	31.32%	34.44%	39.85%	67.36%	69.16%	75.05%	100%	58.63%	64.17%	74.15%
Standard Deviation	0.4%	0.3%	0.8%	0.2%	7.1%	1%	1.7%	0%	0.3%	0.1%
SHOC Scalable Heterogeneous Computing - OpenCL - MD5 Hash (GHash/s)	10.72	11.63	14.17	19.96	19.29	24.37	35.86	9.20	14.03	16.53
Normalized	29.89%	32.43%	39.51%	55.66%	53.79%	67.96%	100%	25.66%	39.12%	46.1%
Standard Deviation	0.3%	0%	0.2%	0.5%	0.2%	0.2%	1.6%	0%	0.1%	0.2%

SHOC Scalable	451	433	530	595	1101	1119	1134	250	384	442
Heterogeneous										
Computing -										
OpenCL - T.R.B										
Normalized	39.77%	38.18%	46.74%	52.47%	97.09%	98.68%	100%	22.05%	33.86%	38.98%
Standard Deviation	0%	0%	0.1%	0.7%	0.2%	0.5%	0.6%	0.8%	0.2%	0.7%
cl-mem - Copy	187	188	209	317	330	328	454	205	203	221
(GB/s)										
Normalized	41.19%	41.41%	46.04%	69.82%	72.69%	72.25%	100%	45.15%	44.71%	48.68%
Standard Deviation	0.1%	0.1%	0.1%	0%	0.1%	0%	0.1%	0.4%	0.2%	0.1%
LuxMark - GPU -	17288	16886	13823	21562	30091	29641	42693	23448	31179	32815
Luxball HDR										
Normalized	40.49%	39.55%	32.38%	50.5%	70.48%	69.43%	100%	54.92%	73.03%	76.86%
Standard Deviation	0%	0%	0.4%	1%	0%	0%	0.3%	0.3%	0.1%	3%
clpeak - Kernel	3.74	3.66	3.72	3.80	3.52	3.48	3.57	5.68	6.98	6.94
Latency (us)										
Normalized	93.05%	95.08%	93.55%	91.58%	98.86%	100%	97.48%	61.27%	49.86%	50.14%
Standard Deviation	0.7%	1%	1%	2.1%	2.9%	2.4%	5.8%	1.1%	0.2%	0.3%
clpeak - I.C.I	1692	2082	2437	3321	8007	10059	14385	1430	1991	2491
(GIOPS)										
Normalized	11.76%	14.47%	16.94%	23.09%	55.66%	69.93%	100%	9.94%	13.84%	17.32%
Standard Deviation	1%	0.1%	0.6%	0.6%	12%	10.4%	11.4%	0%	0.2%	0.1%
SHOC Scalable	3.35	3.53	3.66	3.13	6.76	6.59	5.08	1.45	1.94	1.98
Heterogeneous										
Computing -										
OpenCL - T.R.B										
Normalized	49.56%	52.22%	54.14%	46.3%	100%	97.49%	75.15%	21.45%	28.7%	29.29%
cl-mem - Copy	1.36	1.77	1.40	1.77	3.88	3.53	2.52	1.24	1.03	1.02
(GB/s/Watt)										
Normalized	35.05%	45.62%	36.08%	45.62%	100%	90.98%	64.95%	31.96%	26.55%	26.29%
Chaos Group	90.43	86.30	102.07	66.41	66.02	72.07	56.09			
V-RAY - CUDA GPU										
Normalized	62.03%	64.99%	54.95%	84.46%	84.96%	77.83%	100%			
Standard Deviation	0.1%	0.1%	5.2%	0.1%	0.1%	0.1%	9.3%			
CUDA Mini-Nbody -	91.99	112	111	186	244	304	426			
Original										
((NBody^2)/s)										
Normalized	21.59%	26.29%	26.06%	43.66%	57.28%	71.36%	100%			
Standard Deviation	0.3%	0.1%	0.3%	0.2%	0.4%	0.5%	0.3%			
NVIDIA GPU Cloud	160	175	193	271	309	335	449			
TensorFlow -										
ResNet-50, FP16										
(Images/sec)										
Normalized	35.63%	38.98%	42.98%	60.36%	68.82%	74.61%	100%			
Standard Deviation	0%	0.1%	0.1%	0.1%		0.4%	0.1%			

NVIDIA GPU Cloud	0.92	1.14	1.02	1.22	1.90	1.86	2.10
TensorFlow -							
ResNet-50, FP16							
(Images/sec/Watt)							
Normalized	43.81%	54.29%	48.57%	58.1%	90.48%	88.57%	100%
NVIDIA GPU Cloud	125	133	143	210		205	285
TensorFlow -							
ResNet-50, FP32							
(Images/sec)							
Normalized	43.86%	46.67%	50.18%	73.68%		71.93%	100%
Standard Deviation	0.1%	0.1%	0.1%	0.2%		0.2%	0.3%
NVIDIA GPU Cloud	0.73	0.84	0.73	0.89		0.99	1.15
TensorFlow -							
ResNet-50, FP32							
(Images/sec/Watt)							
Normalized	63.48%	73.04%	63.48%	77.39%		86.09%	100%
NVIDIA GPU Cloud	1563	1683	1875	2674		3153	4432
TensorFlow -							
AlexNet, FP16							
(Images/sec)							
Normalized	35.27%	37.97%	42.31%	60.33%		71.14%	100%
Standard Deviation	0.2%	0.1%	0.2%	0.1%		0.2%	0.1%
NVIDIA GPU Cloud	9.86	12.39	11.19	12.53		16.53	24.95
TensorFlow -							
AlexNet, FP16							
(Images/sec/Watt)							
Normalized	39.52%	49.66%	44.85%	50.22%		66.25%	100%
NVIDIA GPU Cloud	1516	1591	1764	2539	2174	2419	3344
TensorFlow -							
AlexNet, FP32							
(Images/sec)							
Normalized	45.33%	47.58%	52.75%	75.93%	65.01%	72.34%	100%
Standard Deviation	0%	0%	0.2%	0.1%		0.1%	0.4%
NVIDIA GPU Cloud	9.61	10.99	9.90	12.19	11.72	11.44	12.99
TensorFlow -							
AlexNet, FP32							
(Images/sec/Watt)							
Normalized	73.98%	84.6%	76.21%	93.84%	90.22%	88.07%	100%
NVIDIA GPU Cloud	375	413	458	629		738	1015
TensorFlow -							
Googlenet, FP16							
(Images/sec)							
Normalized	36.95%	40.69%	45.12%	61.97%		72.71%	100%
Standard Deviation	0.2%	0%	0.1%	0.7%		0.1%	0.1%


NVIDIA GPU Cloud	2.06	2.53	2.13	2.45	3.21	4.07	
TensorFlow -							
GoogLeNet, FP16							
(Images/sec/Watt)							
Normalized	50.61%	62.16%	52.33%	60.2%	78.87%	100%	
NVIDIA GPU Cloud	44.77	49.23	55.03	74.63	102.77	135.90	
TensorFlow -							
Inception v4, FP16							
(Images/sec)							
Normalized	32.94%	36.23%	40.49%	54.92%	75.62%	100%	
Standard Deviation	0.1%	0.1%	0.1%	0.2%	0.1%	0.1%	
NVIDIA GPU Cloud	0.28	0.34	0.30	0.35	0.58	0.74	
TensorFlow -							
Inception v4, FP16							
(Images/sec/Watt)							
Normalized	37.84%	45.95%	40.54%	47.3%	78.38%	100%	
NVIDIA GPU Cloud	76.93	84.43	93.13	131.37	153.33	206.97	
TensorFlow -							
VGG-16, FP16							
(Images/sec)							
Normalized	37.17%	40.79%	45%	63.47%	74.08%	100%	
Standard Deviation	0.2%	0.1%	0.4%	0.2%	0.1%	0.1%	
NVIDIA GPU Cloud	0.46	0.53	0.48	0.54	0.69	0.94	
TensorFlow -							
VGG-16, FP16							
(Images/sec/Watt)							
Normalized	48.94%	56.38%	51.06%	57.45%	73.4%	100%	
NVIDIA GPU Cloud	71.57	76.40	82.90	119.50	108.57	152.40	
TensorFlow -							
VGG-16, FP32							
(Images/sec)							
Normalized	46.96%	50.13%	54.4%	78.41%	71.24%	100%	
Standard Deviation	0.3%	0%	0%	0.2%	0.1%	0.1%	
NVIDIA GPU Cloud	0.42	0.47	0.42	0.49	0.45	0.56	
TensorFlow -							
VGG-16, FP32							
(Images/sec/Watt)							
Normalized	75%	83.93%	75%	87.5%	80.36%	100%	
SHOC Scalable	1.13	1.11	1.05	1.39	1.67	1.36	2.26
Heterogeneous							
Computing -							
Performance / Cost							
- OpenCL - FFT SP							
(GFLOPS/Dollar)							
Normalized	50%	49.12%	46.46%	61.5%	73.89%	60.18%	53.1%
							99.56%

SHOC Scalable Heterogeneous Computing - Performance / Cost - OpenCL - MD5 Hash	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03
SHOC Scalable Heterogeneous Computing - Performance / Cost - OpenCL - T.R.B (GB/s/Dollar)	1.13	0.96	0.97	0.85	1.84	1.40	0.95	0.94	0.93
Normalized cl-mem - Performance / Cost - Copy	61.41%	52.17%	52.72%	46.2%	100%	76.09%	51.63%	51.09%	50.54%
Normalized LuxMark - Performance / Cost - GPU - Luxball HDR (Score/Dollar)	0.47	0.42	0.38	0.45	0.55	0.41	0.38	0.50	0.47
Normalized clpeak - Performance / Cost - Kernel Latency (us x Dollar)	85.45%	76.36%	69.09%	81.82%	100%	74.55%	69.09%	90.91%	85.45%
Normalized clpeak - Performance / Cost - I.C.I	43.33	37.61	25.18	30.85	50.24	37.14	35.61	76.23	69.08
Normalized clpeak - Performance / Cost - I.C.I	56.84%	49.34%	33.03%	40.47%	65.91%	48.72%	46.71%	100%	90.62%
Normalized clpeak - Performance / Cost - I.C.I	100%	90.81%	73.07%	56.18%	70.77%	53.74%	34.86%	52.27%	45.27%
Normalized clpeak - Performance / Cost - I.C.I	4.24	4.64	4.44	4.75	13.37	12.61	12.00	4.87	5.24
Normalized	31.71%	34.7%	33.21%	35.53%	100%	94.32%	89.75%	36.42%	39.19%





SHOC Scalable HeterOgeneous Computing 2015-11-10

Target: OpenCL - Benchmark: MD5 Hash



1. (CXX) g++ options: -O2 -lSHOCCCommon -lirt






LuxMark 3.1

OpenCL Device: GPU - Scene: Luxball HDR




clpeak

OpenCL Test: Kernel Latency



clpeak

OpenCL Test: Integer Compute INT




SHOC Scalable Heterogeneous Computing 2015-11-10

Target: OpenCL - Benchmark: Texture Read Bandwidth




SHOC Scalable Heterogeneous Computing 2015-11-10

System Power Consumption Monitor




cl-mem 2017-01-13

Benchmark: Copy




cl-mem 2017-01-13

System Power Consumption Monitor




LuxMark 3.1

System Power Consumption Monitor




GPU Temperature Monitor

Phoronix Test Suite System Monitoring




System Power Consumption Monitor

Phoronix Test Suite System Monitoring




Chaos Group V-RAY 1.1.0

Mode: CUDA GPU




Chaos Group V-RAY 1.1.0

GPU Temperature Monitor




Chaos Group V-RAY 1.1.0

System Power Consumption Monitor



CUDA Mini-Nbody 2015-11-10


Test: Original



CUDA Mini-Nbody 2015-11-10

GPU Temperature Monitor


◀ Celsius, Fewer Is Better



CUDA Mini-Nbody 2015-11-10

System Power Consumption Monitor


◀ Watts, Fewer Is Better



NVIDIA GPU Cloud TensorFlow 18.09

Test: ResNet-50, FP16


Images Per Second, More Is Better



NVIDIA GPU Cloud TensorFlow 18.09


Test: ResNet-50, FP16

Images Per Second Per Watt, More Is Better




NVIDIA GPU Cloud TensorFlow 18.09

GPU Temperature Monitor



NVIDIA GPU Cloud TensorFlow 18.09


System Power Consumption Monitor



NVIDIA GPU Cloud TensorFlow 18.09

Test: ResNet-50, FP32


Images Per Second, More Is Better



NVIDIA GPU Cloud TensorFlow 18.09

Test: ResNet-50, FP32

Images Per Second Per Watt, More Is Better



NVIDIA GPU Cloud TensorFlow 18.09

GPU Temperature Monitor

◀ Celsius, Fewer Is Better

GTX 1070

Min: 57 / Avg: 65.61 / Max: 71

GTX 1070 Ti

Min: 44 / Avg: 50.85 / Max: 54

GTX 1080

Min: 60 / Avg: 68.77 / Max: 75

GTX 1080 Ti

Min: 62 / Avg: 70.04 / Max: 77

RTX 2080

Min: 59 / Avg: 70.41 / Max: 77

RTX 2080 Ti

Min: 54 / Avg: 61.16 / Max: 67

15

30

45

60

75

NVIDIA GPU Cloud TensorFlow 18.09

System Power Consumption Monitor

◀ Watts, Fewer Is Better

GTX 1070

Min: 45.9 / Avg: 172.02 / Max: 213.8

GTX 1070 Ti

Min: 48 / Avg: 157.29 / Max: 187.8

GTX 1080

Min: 44.1 / Avg: 196.38 / Max: 249.7

GTX 1080 Ti

Min: 79.7 / Avg: 234.81 / Max: 318.6

RTX 2080

Min: 48 / Avg: 206.72 / Max: 285

RTX 2080 Ti

Min: 50.5 / Avg: 248.25 / Max: 346.7

60

120


180

240

300


NVIDIA GPU Cloud TensorFlow 18.09

Test: AlexNet, FP16



NVIDIA GPU Cloud TensorFlow 18.09

Test: AlexNet, FP16



NVIDIA GPU Cloud TensorFlow 18.09

GPU Temperature Monitor

◀ Celsius, Fewer Is Better

GTX 1070

Min: 59 / Avg: 64.06 / Max: 67

GTX 1070 Ti

Min: 46 / Avg: 49.93 / Max: 53

GTX 1080

Min: 62 / Avg: 66.46 / Max: 70

GTX 1080 Ti

Min: 63 / Avg: 67.2 / Max: 70

RTX 2080

Min: 63 / Avg: 67.57 / Max: 72

RTX 2080 Ti

Min: 57 / Avg: 60.36 / Max: 64

14

28

42

56

70

NVIDIA GPU Cloud TensorFlow 18.09

System Power Consumption Monitor

◀ Watts, Fewer Is Better

GTX 1070

Min: 68.9 / Avg: 158.49 / Max: 192.3

GTX 1070 Ti

Min: 46.9 / Avg: 135.79 / Max: 175

GTX 1080

Min: 44.1 / Avg: 167.55 / Max: 234.8

GTX 1080 Ti

Min: 51.3 / Avg: 213.37 / Max: 303.2

RTX 2080

Min: 47.3 / Avg: 190.7 / Max: 270.1

RTX 2080 Ti

Min: 49.3 / Avg: 177.62 / Max: 332.6

60

120

180


240

300

NVIDIA GPU Cloud TensorFlow 18.09

Test: AlexNet, FP32


Images Per Second, More Is Better



NVIDIA GPU Cloud TensorFlow 18.09


Test: AlexNet, FP32

Images Per Second Per Watt, More Is Better




NVIDIA GPU Cloud TensorFlow 18.09

GPU Temperature Monitor




NVIDIA GPU Cloud TensorFlow 18.09

System Power Consumption Monitor




NVIDIA GPU Cloud TensorFlow 18.09

Test: Googlenet, FP16




NVIDIA GPU Cloud TensorFlow 18.09

Test: Googlenet, FP16




NVIDIA GPU Cloud TensorFlow 18.09

GPU Temperature Monitor




NVIDIA GPU Cloud TensorFlow 18.09

System Power Consumption Monitor




NVIDIA GPU Cloud TensorFlow 18.09

Test: Inception v4, FP16



NVIDIA GPU Cloud TensorFlow 18.09

Test: Inception v4, FP16



NVIDIA GPU Cloud TensorFlow 18.09

GPU Temperature Monitor

◀ Celsius, Fewer Is Better

GTX 1070

Min: 55 / Avg: 64.07 / Max: 72

GTX 1070 Ti

Min: 43 / Avg: 49.56 / Max: 55

GTX 1080

Min: 51 / Avg: 64.33 / Max: 76

GTX 1080 Ti

Min: 56 / Avg: 66.24 / Max: 77

RTX 2080

Min: 57 / Avg: 64.7 / Max: 74

RTX 2080 Ti

Min: 53 / Avg: 57.76 / Max: 65

15

30

45

60

75

NVIDIA GPU Cloud TensorFlow 18.09

System Power Consumption Monitor

◀ Watts, Fewer Is Better

GTX 1070

Min: 44.1 / Avg: 157.87 / Max: 210.2

GTX 1070 Ti

Min: 47.2 / Avg: 146.22 / Max: 188.6

GTX 1080

Min: 71.3 / Avg: 185.89 / Max: 249.7

GTX 1080 Ti

Min: 52.2 / Avg: 211.84 / Max: 320.9

RTX 2080

Min: 47.6 / Avg: 177.11 / Max: 285.3

RTX 2080 Ti

Min: 50 / Avg: 183.11 / Max: 341.1

60

120


180

240

300


NVIDIA GPU Cloud TensorFlow 18.09

Test: VGG-16, FP16



NVIDIA GPU Cloud TensorFlow 18.09

Test: VGG-16, FP16



NVIDIA GPU Cloud TensorFlow 18.09

GPU Temperature Monitor

◀ Celsius, Fewer Is Better

GTX 1070

Min: 59 / Avg: 66.13 / Max: 70

GTX 1070 Ti

Min: 45 / Avg: 52.18 / Max: 55

GTX 1080

Min: 50 / Avg: 65.84 / Max: 74

GTX 1080 Ti

Min: 61 / Avg: 70.1 / Max: 76

RTX 2080

Min: 59 / Avg: 69.86 / Max: 77

RTX 2080 Ti

Min: 54 / Avg: 61.56 / Max: 68

15

30

45

60

75

NVIDIA GPU Cloud TensorFlow 18.09

System Power Consumption Monitor

◀ Watts, Fewer Is Better

GTX 1070

Min: 43.7 / Avg: 168.53 / Max: 201.8

GTX 1070 Ti

Min: 46.8 / Avg: 160.44 / Max: 192

GTX 1080

Min: 90.8 / Avg: 194.86 / Max: 239.8

GTX 1080 Ti

Min: 51.7 / Avg: 242.51 / Max: 313.1

RTX 2080

Min: 48.1 / Avg: 223.76 / Max: 283.3

RTX 2080 Ti

Min: 51 / Avg: 219.66 / Max: 332.2

60

120

180


240

300

NVIDIA GPU Cloud TensorFlow 18.09

Test: VGG-16, FP32


Images Per Second, More Is Better



NVIDIA GPU Cloud TensorFlow 18.09

Test: VGG-16, FP32


Images Per Second Per Watt, More Is Better



NVIDIA GPU Cloud TensorFlow 18.09

GPU Temperature Monitor


◀ Celsius, Fewer Is Better



NVIDIA GPU Cloud TensorFlow 18.09


System Power Consumption Monitor

◀ Watts, Fewer Is Better



SHOC Scalable HeterOgeneous Computing 2015-11-10

Performance / Cost - Target: OpenCL - Benchmark: FFT SP



1. GTX 1070: \$399 reported cost.
2. GTX 1070 Ti: \$449 reported cost.
3. GTX 1080: \$549 reported cost.
4. GTX 1080 Ti: \$699 reported cost.
5. RTX 2070: \$599 reported cost.
6. RTX 2080: \$798 reported cost.
7. RTX 2080 Ti: \$1199 reported cost.
8. RX Vega 56: \$409 reported cost.
9. RX Vega 64: \$475 reported cost.

SHOC Scalable Heterogeneous Computing 2015-11-10

Performance / Cost - Target: OpenCL - Benchmark: MD5 Hash

► GHash/s Per Dollar, More Is Better




1. GTX 1070: \$399 reported cost.
2. GTX 1070 Ti: \$449 reported cost.
3. GTX 1080: \$549 reported cost.
4. GTX 1080 Ti: \$699 reported cost.
5. RTX 2070: \$599 reported cost.
6. RTX 2080: \$798 reported cost.
7. RTX 2080 Ti: \$1199 reported cost.
8. RX Vega 56: \$409 reported cost.
9. RX Vega 64: \$475 reported cost.

SHOC Scalable HeterOgeneous Computing 2015-11-10

Performance / Cost - Target: OpenCL - Benchmark: Texture Read Bandwidth

► GB/s Per Dollar, More Is Better




1. GTX 1070: \$399 reported cost.
2. GTX 1070 Ti: \$449 reported cost.
3. GTX 1080: \$549 reported cost.
4. GTX 1080 Ti: \$699 reported cost.
5. RTX 2070: \$599 reported cost.
6. RTX 2080: \$798 reported cost.
7. RTX 2080 Ti: \$1199 reported cost.
8. RX Vega 56: \$409 reported cost.
9. RX Vega 64: \$475 reported cost.

cl-mem 2017-01-13

Performance / Cost - Benchmark: Copy

► GB/s Per Dollar, More Is Better




1. GTX 1070: \$399 reported cost.
2. GTX 1070 Ti: \$449 reported cost.
3. GTX 1080: \$549 reported cost.
4. GTX 1080 Ti: \$699 reported cost.
5. RTX 2070: \$599 reported cost.
6. RTX 2080: \$798 reported cost.
7. RTX 2080 Ti: \$1199 reported cost.
8. RX Vega 56: \$409 reported cost.
9. RX Vega 64: \$475 reported cost.


LuxMark 3.1


Performance / Cost - OpenCL Device: GPU - Scene: Luxball HDR

Score Per Dollar, More Is Better




1. GTX 1070: \$399 reported cost.
2. GTX 1070 Ti: \$449 reported cost.
3. GTX 1080: \$549 reported cost.
4. GTX 1080 Ti: \$699 reported cost.
5. RTX 2070: \$599 reported cost.
6. RTX 2080: \$798 reported cost.
7. RTX 2080 Ti: \$1199 reported cost.
8. RX Vega 56: \$409 reported cost.
9. RX Vega 64: \$475 reported cost.






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



Geometric mean based upon tests: pts/cl-mem, pts/shoc and pts/v-ray



Geometric mean based upon tests: pts/luxmark, pts/shoc and pts/cl-mem

This file was automatically generated via the Phoronix Test Suite benchmarking software on Sunday, 22 December 2024 06:33.