www.phoronix-test-suite.com

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-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 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-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

	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%

NVIDIA AMD Linux GPU Compute December 2018

SHOC Scalable HeterOgeneous Computing - OpenCL - T.R.B	451	433	530	595	1101	1119	**1134**	**250**	384	442
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 (GB/s)	**187**	188	209	317	330	328	**454**	205	203	221
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 - Luxball HDR	17288	16886	**13823**	21562	30091	29641	**42693**	23448	31179	32815
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%	0.3%	0.3%	0.1%	3%
clpeak - Kernel Latency (us)	3.74	3.66	3.72	3.80	3.52	**3.48**	3.57	5.68	**6.98**	6.94
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 (GIOPS)	1692	2082	2437	3321	8007	10059	**14385**	**1430**	1991	2491
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 HeterOgeneous Computing - OpenCL - T.R.B	3.35	3.53	3.66	3.13	**6.76**	6.59	5.08	**1.45**	1.94	1.98
Normalized	49.56%	52.22%	54.14%	46.3%	100%	97.49%	75.15%	21.45%	28.7%	29.29%
cl-mem - Copy (GB/s/Watt)	1.36	1.77	1.40	1.77	**3.88**	3.53	2.52	1.24	1.03	**1.02**
Normalized	35.05%	45.62%	36.08%	45.62%	100%	90.98%	64.95%	31.96%	26.55%	26.29%
Chaos Group V-RAY - CUDA GPU	90.43	86.30	**102.07**	66.41	66.02	72.07	**56.09**			
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 - Original ((NBody^2)/s)	**91.99**	112	111	186	244	304	**426**			
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 TensorFlow - ResNet-50, FP16 (Images/sec)	**160**	175	193	271	309	335	**449**			
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 AMD Linux GPU Compute December 2018

NVIDIA GPU Cloud TensorFlow - ResNet-50, FP16 (Images/sec/Watt)	**0.92**	1.14	1.02	1.22	1.90	1.86	**2.10**
Normalized	43.81%	54.29%	48.57%	58.1%	90.48%	88.57%	100%
NVIDIA GPU Cloud TensorFlow - ResNet-50, FP32 (Images/sec)	**125**	133	143	210		205	**285**
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 TensorFlow - ResNet-50, FP32 (Images/sec/Watt)	**0.73**	0.84	**0.73**	0.89		0.99	**1.15**
Normalized	63.48%	73.04%	63.48%	77.39%		86.09%	100%
NVIDIA GPU Cloud TensorFlow - AlexNet, FP16 (Images/sec)	**1563**	1683	1875	2674		3153	**4432**
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 TensorFlow - AlexNet, FP16 (Images/sec/Watt)	**9.86**	12.39	11.19	12.53		16.53	**24.95**
Normalized	39.52%	49.66%	44.85%	50.22%		66.25%	100%
NVIDIA GPU Cloud TensorFlow - AlexNet, FP32 (Images/sec)	**1516**	1591	1764	2539	2174	2419	**3344**
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 TensorFlow - AlexNet, FP32 (Images/sec/Watt)	**9.61**	10.99	9.90	12.19	11.72	11.44	**12.99**
Normalized	73.98%	84.6%	76.21%	93.84%	90.22%	88.07%	100%
NVIDIA GPU Cloud TensorFlow - Googlenet, FP16 (Images/sec)	**375**	413	458	629		738	**1015**
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 AMD Linux GPU Compute December 2018

NVIDIA GPU Cloud TensorFlow - Googlenet, FP16 (Images/sec/Watt)	**2.06**	2.53	2.13	2.45		3.21	**4.07**		
Normalized	50.61%	62.16%	52.33%	60.2%		78.87%	100%		
NVIDIA GPU Cloud TensorFlow - Inception v4, FP16 (Images/sec)	**44.77**	49.23	55.03	74.63		102.77	**135.90**		
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 TensorFlow - Inception v4, FP16 (Images/sec/Watt)	**0.28**	0.34	0.30	0.35		0.58	**0.74**		
Normalized	37.84%	45.95%	40.54%	47.3%		78.38%	100%		
NVIDIA GPU Cloud TensorFlow - VGG-16, FP16 (Images/sec)	**76.93**	84.43	93.13	131.37		153.33	**206.97**		
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 TensorFlow - VGG-16, FP16 (Images/sec/Watt)	**0.46**	0.53	0.48	0.54		0.69	**0.94**		
Normalized	48.94%	56.38%	51.06%	57.45%		73.4%	100%		
NVIDIA GPU Cloud TensorFlow - VGG-16, FP32 (Images/sec)	**71.57**	76.40	82.90	119.50		108.57	**152.40**		
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 TensorFlow - VGG-16, FP32 (Images/sec/Watt)	**0.42**	0.47	**0.42**	0.49		0.45	**0.56**		
Normalized	75%	83.93%	75%	87.5%		80.36%	100%		
SHOC Scalable HeterOgeneous Computing - Performance / Cost - OpenCL - FFT SP (GFLOPS/Dollar)	1.13	1.11	**1.05**	1.39	1.67	1.36	1.20	**2.26**	2.25
Normalized	50%	49.12%	46.46%	61.5%	73.89%	60.18%	53.1%	100%	99.56%

NVIDIA AMD Linux GPU Compute December 2018

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	61.41%	52.17%	52.72%	46.2%	100%	76.09%	51.63%	51.09%	50.54%
cl-mem - Performance / Cost - Copy	0.47	0.42	**0.38**	0.45	**0.55**	0.41	**0.38**	0.50	0.47
Normalized	85.45%	76.36%	69.09%	81.82%	100%	74.55%	69.09%	90.91%	85.45%
LuxMark - Performance / Cost - GPU - Luxball HDR (Score/Dollar)	43.33	37.61	**25.18**	30.85	50.24	37.14	35.61	**76.23**	69.08
Normalized	56.84%	49.34%	33.03%	40.47%	65.91%	48.72%	46.71%	100%	90.62%
clpeak - Performance / Cost - Kernel Latency (us x Dollar)	**1492**	1643	2042	2656	2108	2777	**4280**	2855	3297
Normalized	100%	90.81%	73.07%	56.18%	70.77%	53.74%	34.86%	52.27%	45.27%
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%

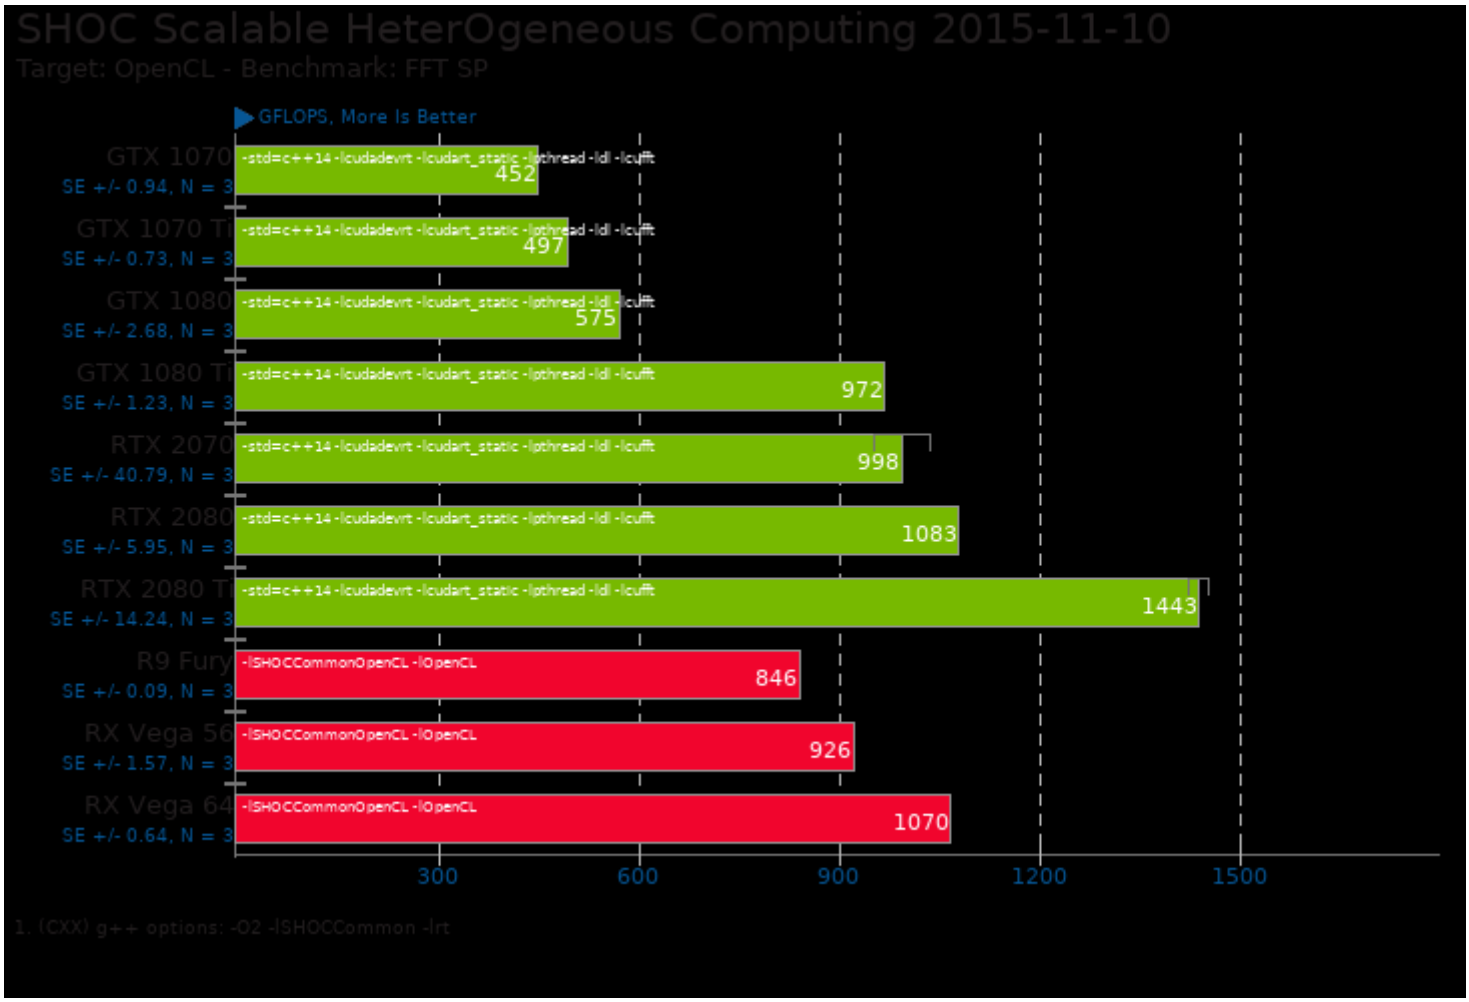
SHOC Scalable HeterOgeneous Computing 2015-11-10
Target: OpenCL - Benchmark: FFT SP
GFLOPS, More Is Better
GTX 1070
SE +/- 0.94, N = 3
452
GTX 1070 Ti
SE +/- 0.73, N = 3
497
GTX 1080
SE +/- 2.68, N = 3
575
GTX 1080 Ti
SE +/- 1.23, N = 3
972
RTX 2070
SE +/- 40.79, N = 3
998
RTX 2080
SE +/- 5.95, N = 3
1083
RTX 2080 Ti
SE +/- 14.24, N = 3
1443
R9 Fury
SE +/- 0.09, N = 3
846
RX Vega 56
SE +/- 1.57, N = 3
926
RX Vega 64
SE +/- 0.64, N = 3
1070
300
600
900
1200
1500
1. (CXX) g++ options: -O2 -lSHOCCommon -lrt

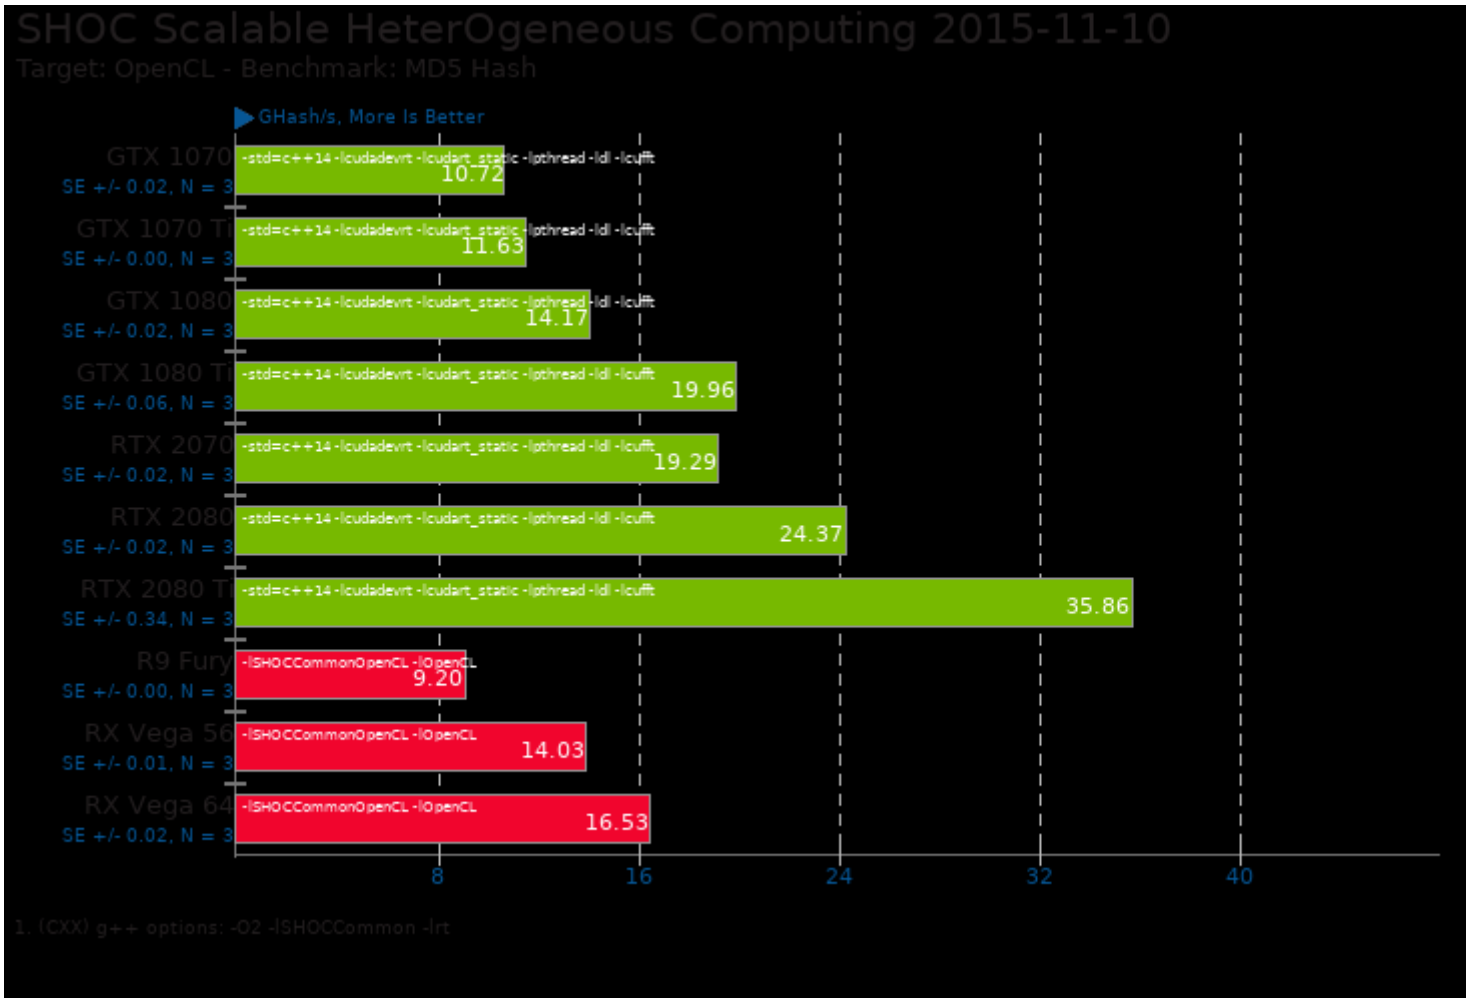
SHOC Scalable HeterOgeneous Computing 2015-11-10
Target: OpenCL - Benchmark: MD5 Hash
GHash/s, More Is Better
GTX 1070
SE +/- 0.02, N = 3
-std=c++14 -lcudadevrt -lcudart_static -lpthread -ldl -lcufft
10.72
GTX 1070 Ti
SE +/- 0.00, N = 3
11.63
GTX 1080
SE +/- 0.02, N = 3
14.17
GTX 1080 Ti
SE +/- 0.06, N = 3
19.96
RTX 2070
SE +/- 0.02, N = 3
19.29
RTX 2080
SE +/- 0.02, N = 3
24.37
RTX 2080 Ti
SE +/- 0.34, N = 3
35.86
R9 Fury
SE +/- 0.00, N = 3
-lSHOCCommonOpenCL -lOpenCL
9.20
RX Vega 56
SE +/- 0.01, N = 3
14.03
RX Vega 64
SE +/- 0.02, N = 3
16.53
8
16
24
32
40
1. (CXX) g++ options: -O2 -lSHOCCommon -lrt

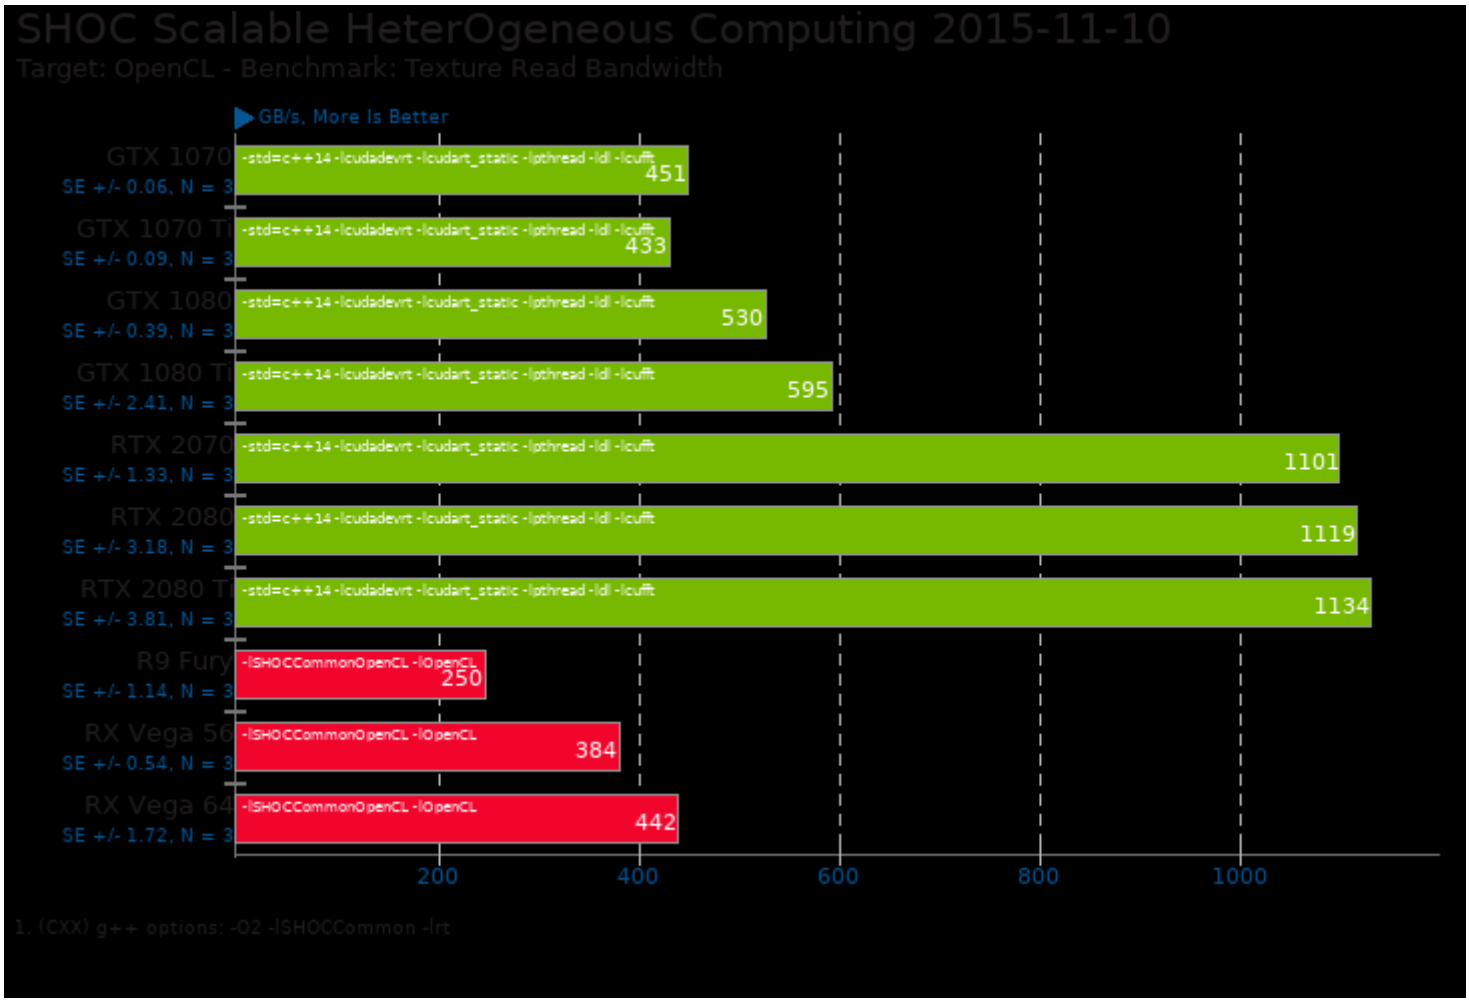
SHOC Scalable HeterOgeneous Computing 2015-11-10
Target: OpenCL - Benchmark: Texture Read Bandwidth
GB/s, More Is Better
GTX 1070
SE +/- 0.06, N = 3
-std=c++14 -lcudadevrt -lcudart_static -lpthread -ldl -lcufft
451
GTX 1070 Ti
SE +/- 0.09, N = 3
433
GTX 1080
SE +/- 0.39, N = 3
530
GTX 1080 Ti
SE +/- 2.41, N = 3
595
RTX 2070
SE +/- 1.33, N = 3
1101
RTX 2080
SE +/- 3.18, N = 3
1119
RTX 2080 Ti
SE +/- 3.81, N = 3
1134
R9 Fury
SE +/- 1.14, N = 3
-lSHOCCommonOpenCL -lOpenCL
250
RX Vega 56
SE +/- 0.54, N = 3
384
RX Vega 64
SE +/- 1.72, N = 3
442
200
400
600
800
1000
1. (CXX) g++ options: -O2 -lSHOCCommon -lrt

NVIDIA AMD Linux GPU Compute December 2018

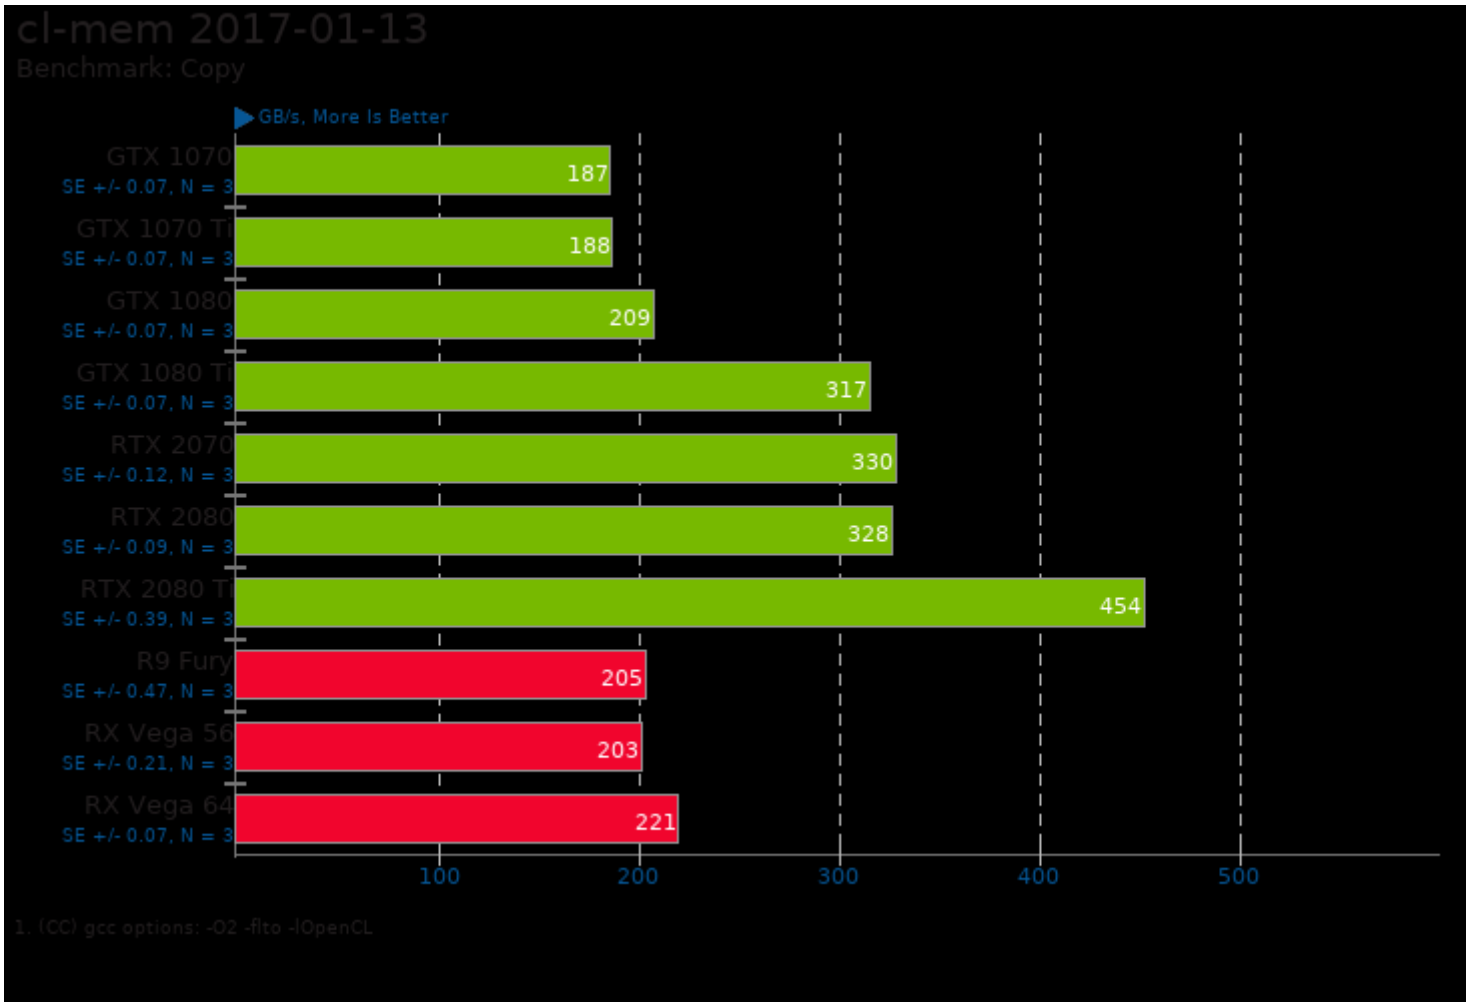

NVIDIA AMD Linux GPU Compute December 2018

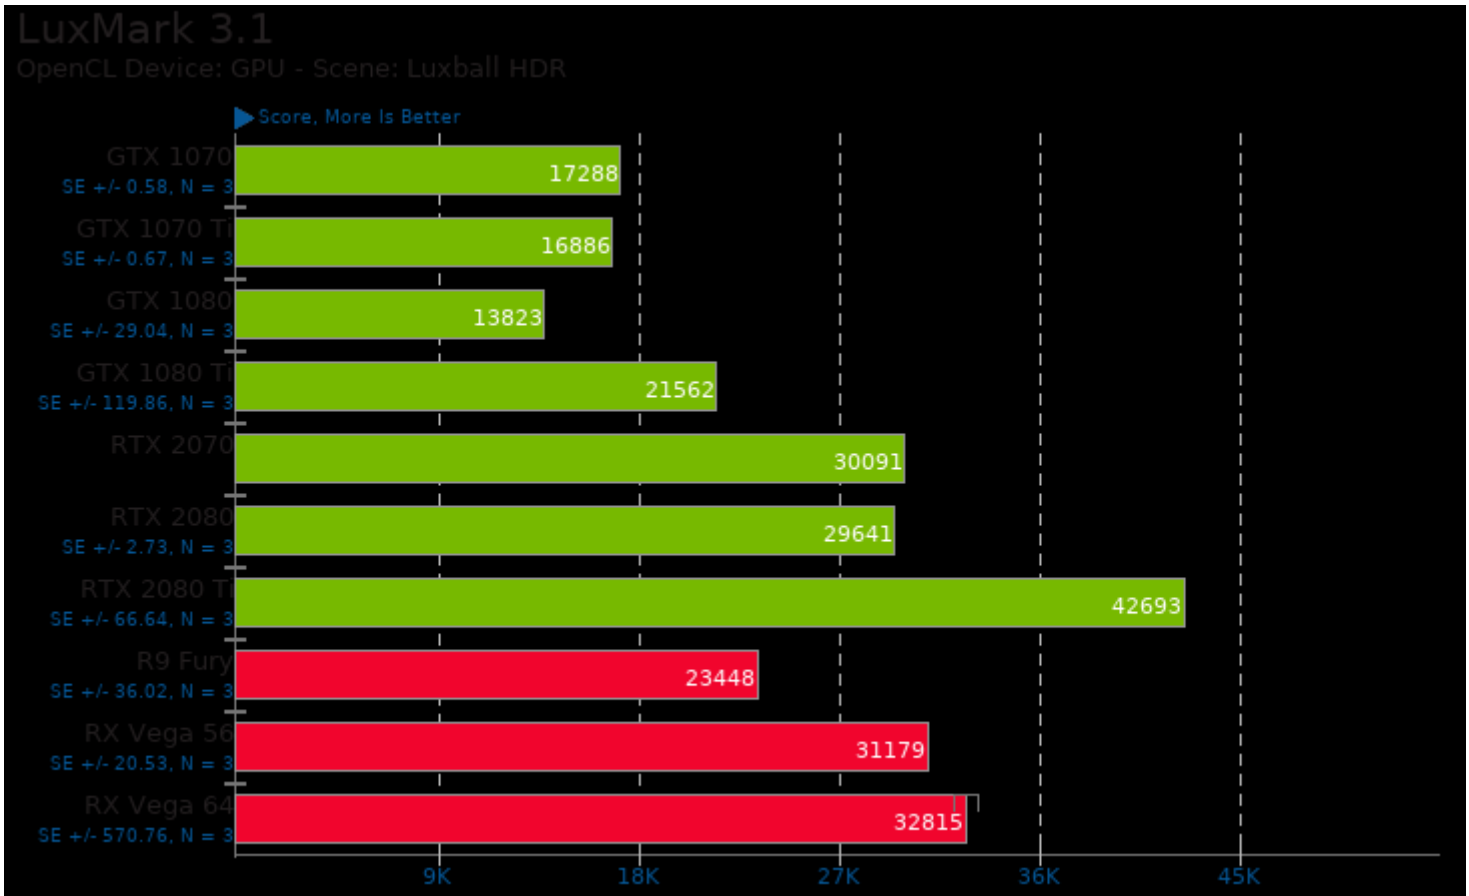

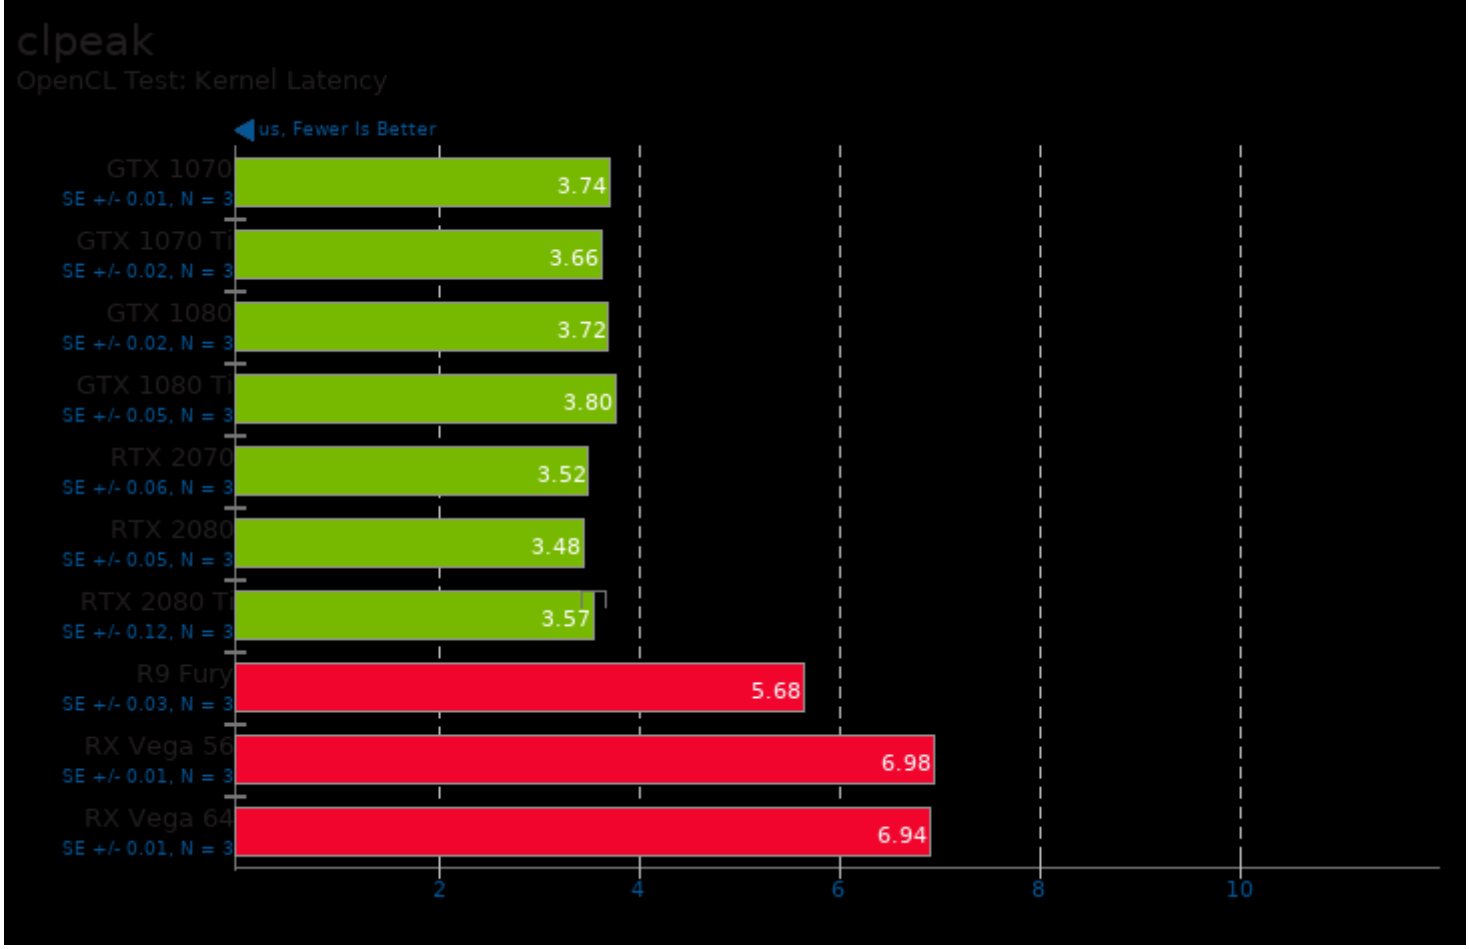

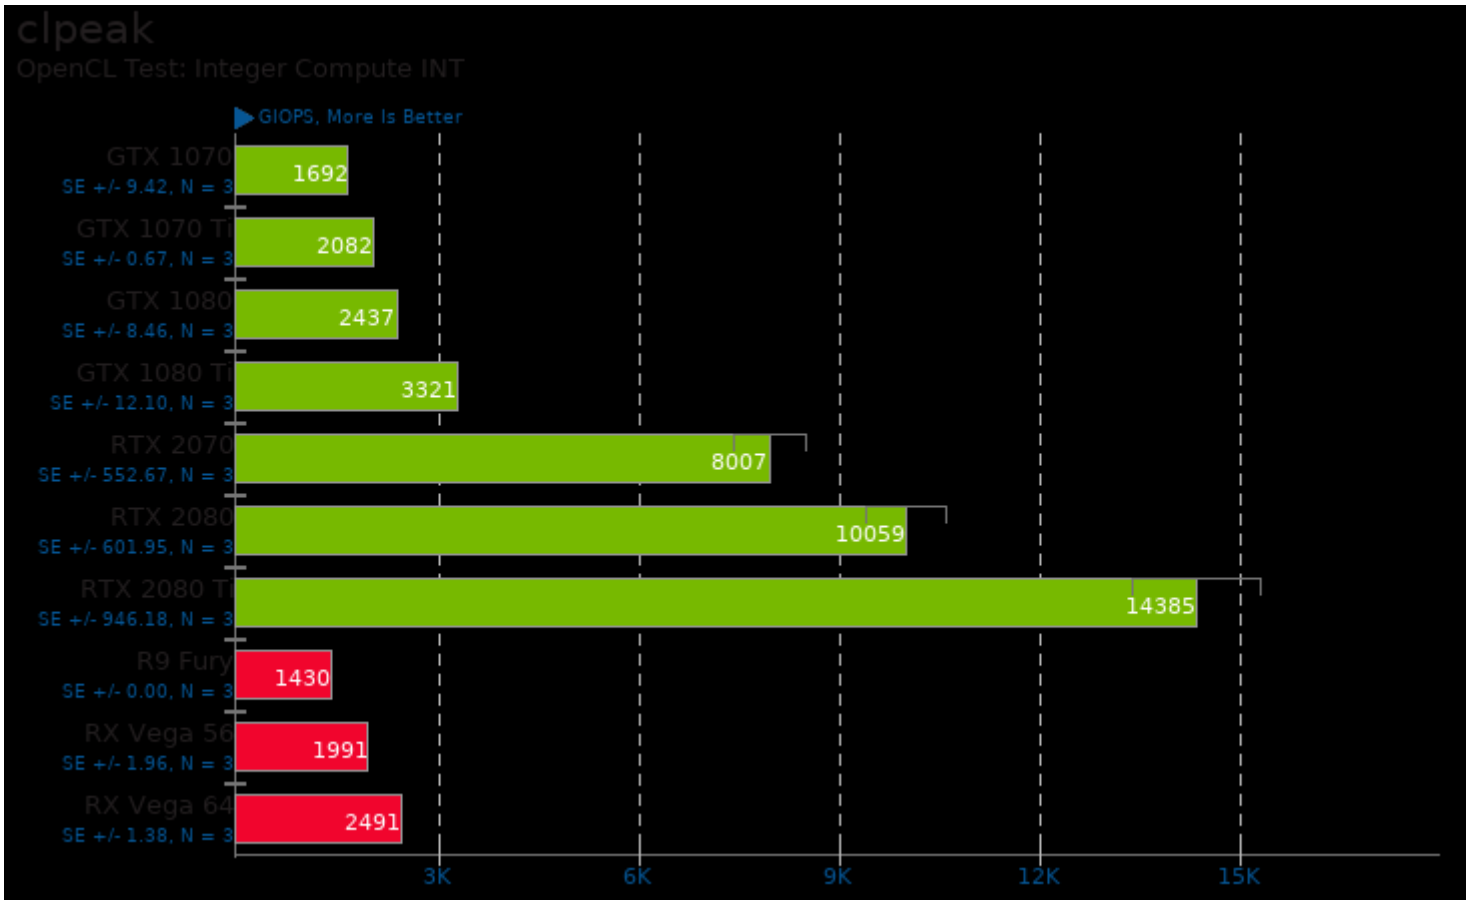
clpeak
OpenCL Test: Integer Compute INT
GIOPS, More Is Better
GTX 1070
SE +/- 9.42, N = 3
1692
GTX 1070 Ti
SE +/- 0.67, N = 3
2082
GTX 1080
SE +/- 8.46, N = 3
2437
GTX 1080 Ti
SE +/- 12.10, N = 3
3321
RTX 2070
SE +/- 552.67, N = 3
8007
RTX 2080
SE +/- 601.95, N = 3
10059
RTX 2080 Ti
SE +/- 946.18, N = 3
14385
R9 Fury
SE +/- 0.00, N = 3
1430
RX Vega 56
SE +/- 1.96, N = 3
1991
RX Vega 64
SE +/- 1.38, N = 3
2491
3K
6K
9K
12K
15K

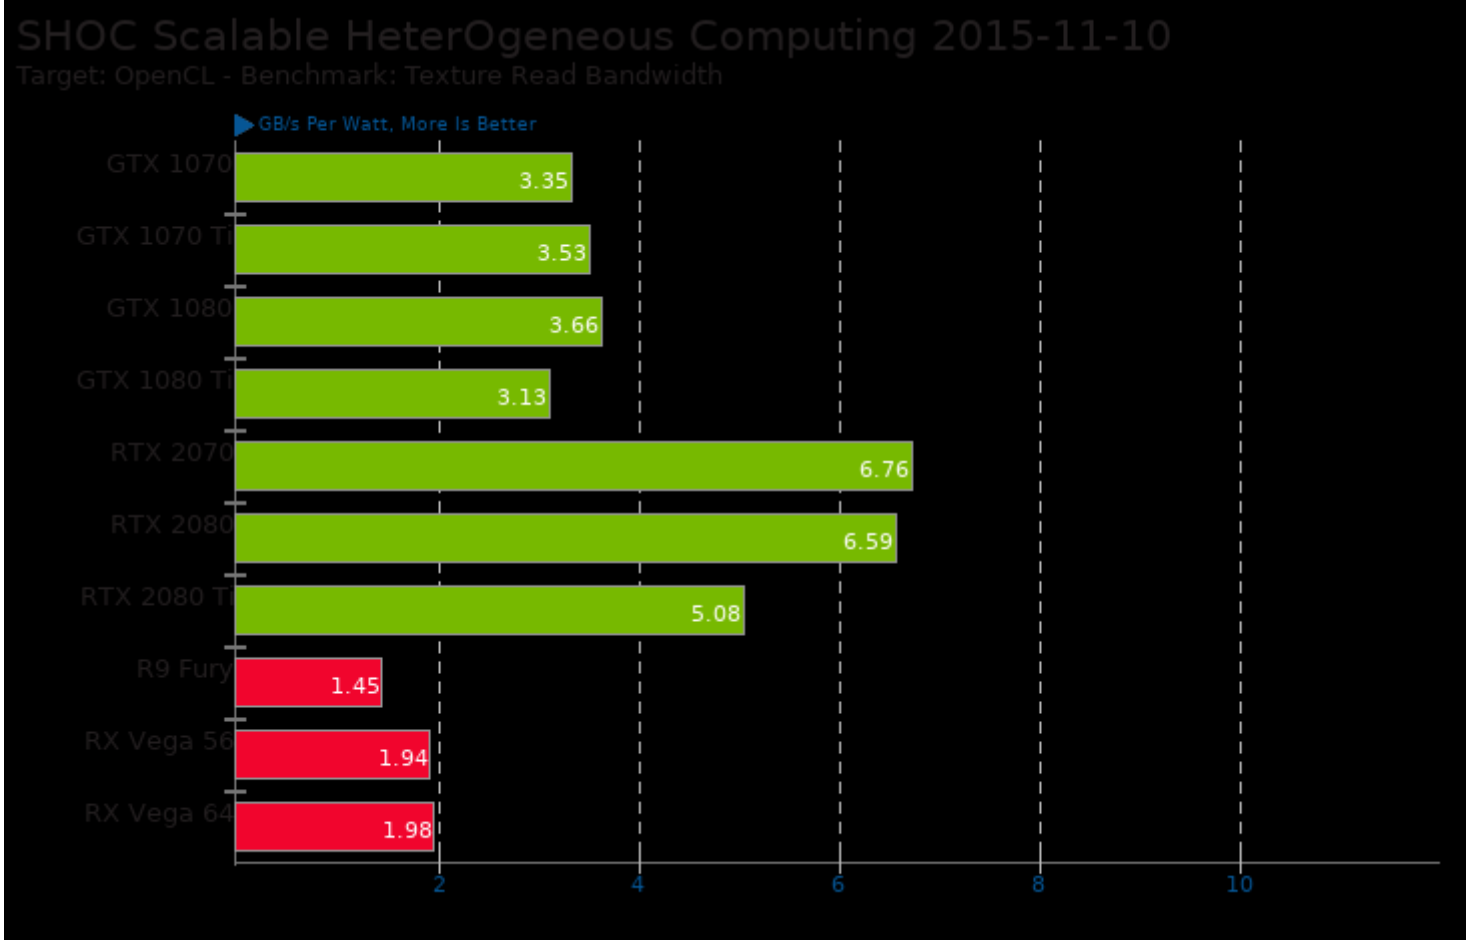
SHOC Scalable HeterOgeneous Computing 2015-11-10
Target: OpenCL - Benchmark: Texture Read Bandwidth
GB/s Per Watt, More Is Better
GTX 1070
3.35
GTX 1070 Ti
3.53
GTX 1080
3.66
GTX 1080 Ti
3.13
RTX 2070
6.76
RTX 2080
6.59
RTX 2080 Ti
5.08
R9 Fury
1.45
RX Vega 56
1.94
RX Vega 64
1.98
2
4
6
8
10

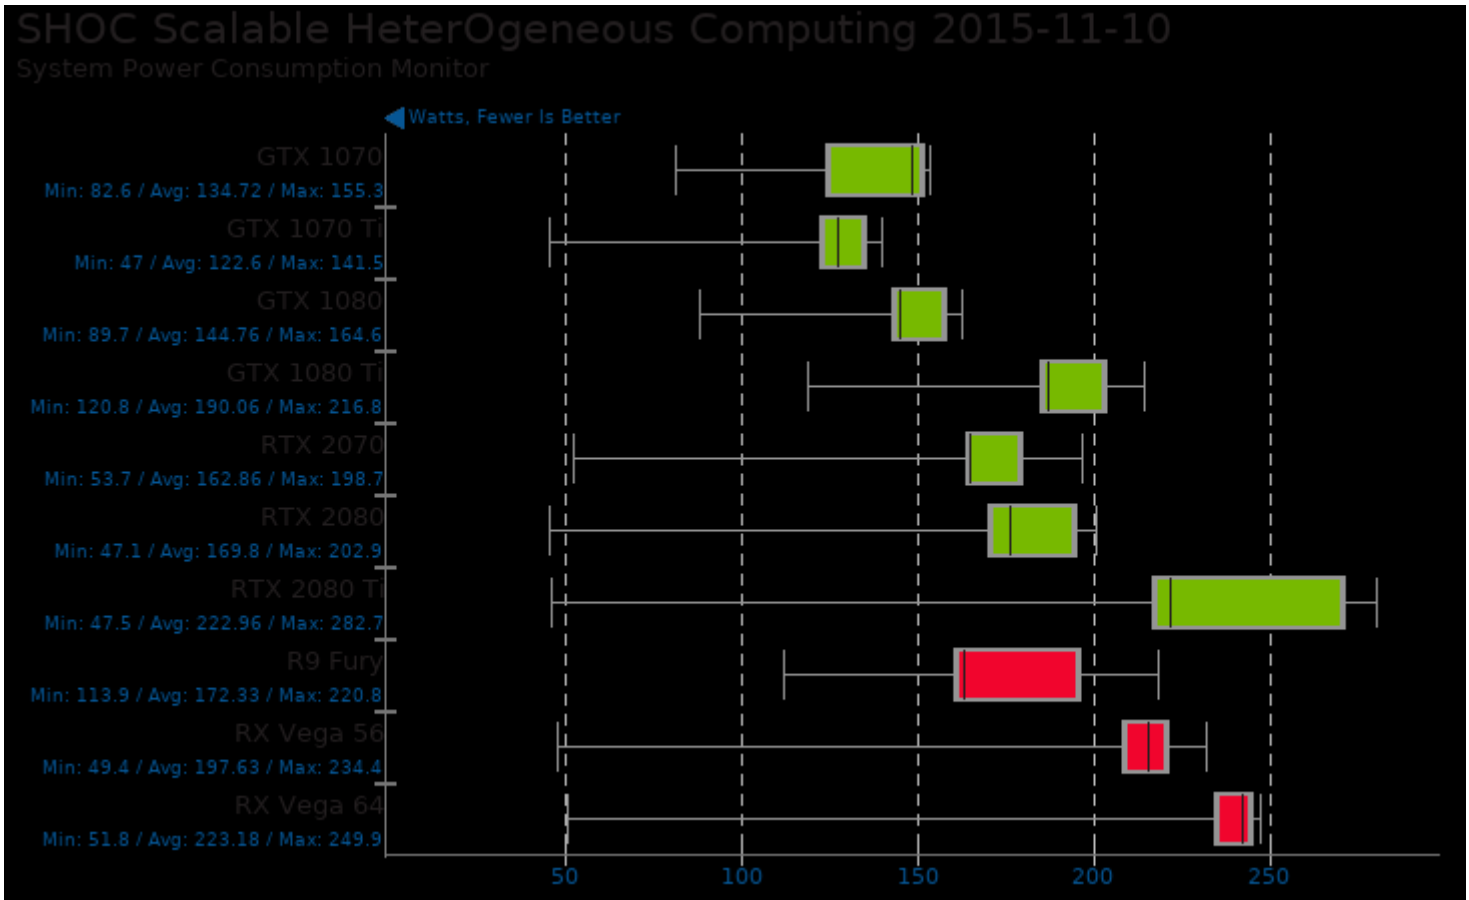
SHOC Scalable HeterOgeneous Computing 2015-11-10
System Power Consumption Monitor
Watts, Fewer Is Better
GTX 1070
Min: 82.6 / Avg: 134.72 / Max: 155.3
GTX 1070 Ti
Min: 47 / Avg: 122.6 / Max: 141.5
GTX 1080
Min: 89.7 / Avg: 144.76 / Max: 164.6
GTX 1080 Ti
Min: 120.8 / Avg: 190.06 / Max: 216.8
RTX 2070
Min: 53.7 / Avg: 162.86 / Max: 198.7
RTX 2080
Min: 47.1 / Avg: 169.8 / Max: 202.9
RTX 2080 Ti
Min: 47.5 / Avg: 222.96 / Max: 282.7
R9 Fury
Min: 113.9 / Avg: 172.33 / Max: 220.8
RX Vega 56
Min: 49.4 / Avg: 197.63 / Max: 234.4
RX Vega 64
Min: 51.8 / Avg: 223.18 / Max: 249.9
50
100
150
200
250

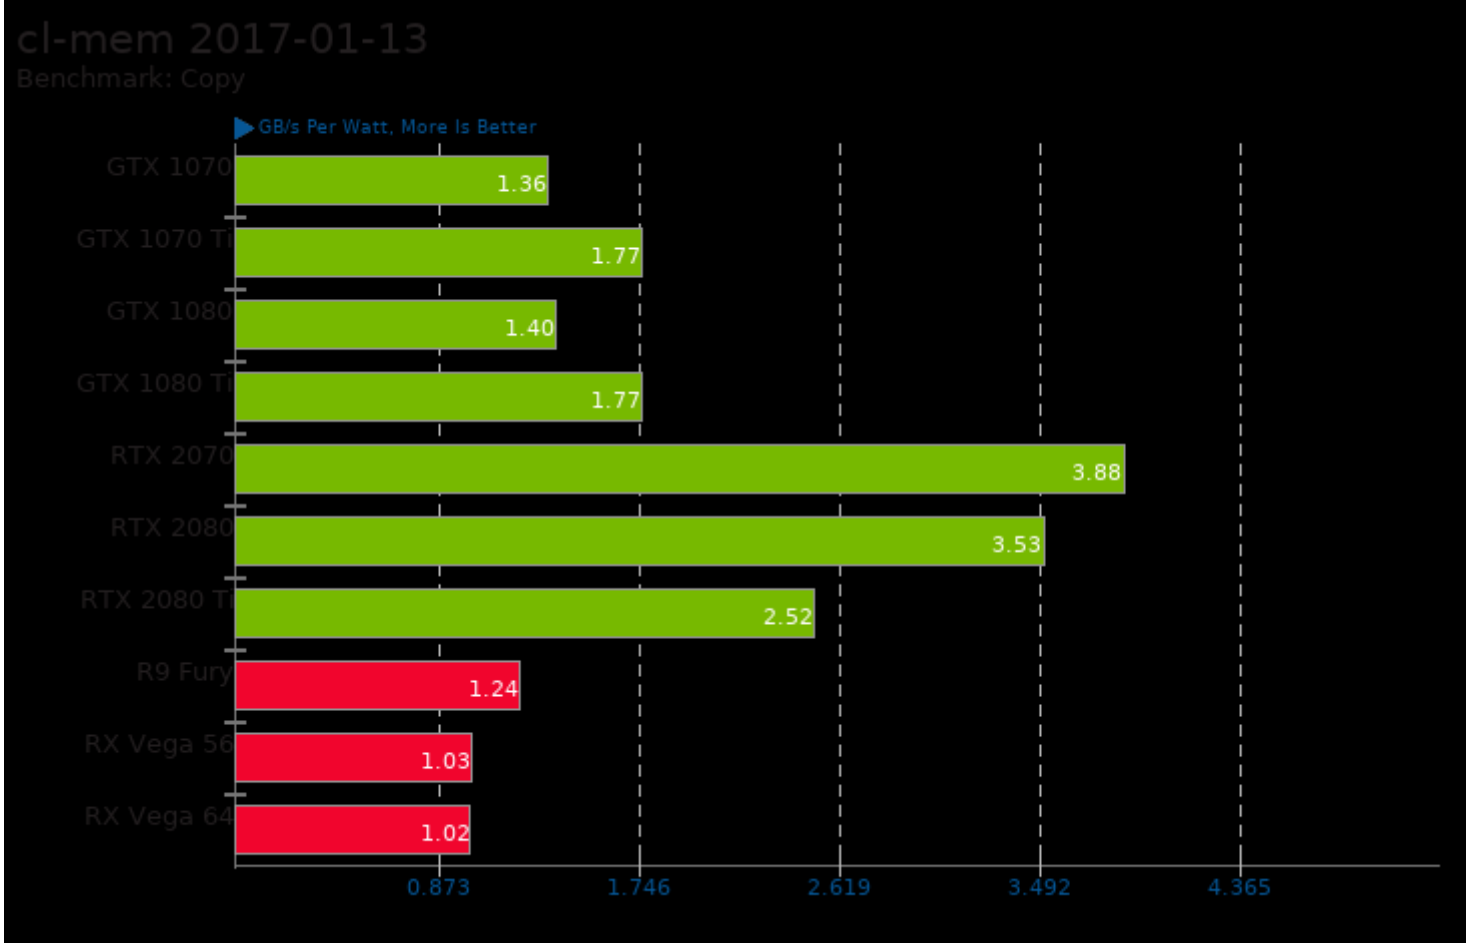
cl-mem 2017-01-13
Benchmark: Copy
GB/s Per Watt, More Is Better
GTX 1070
1.36
GTX 1070 Ti
1.77
GTX 1080
1.40
GTX 1080 Ti
1.77
RTX 2070
3.88
RTX 2080
3.53
RTX 2080 Ti
2.52
R9 Fury
1.24
RX Vega 56
1.03
RX Vega 64
1.02
0.873
1.746
2.619
3.492
4.365

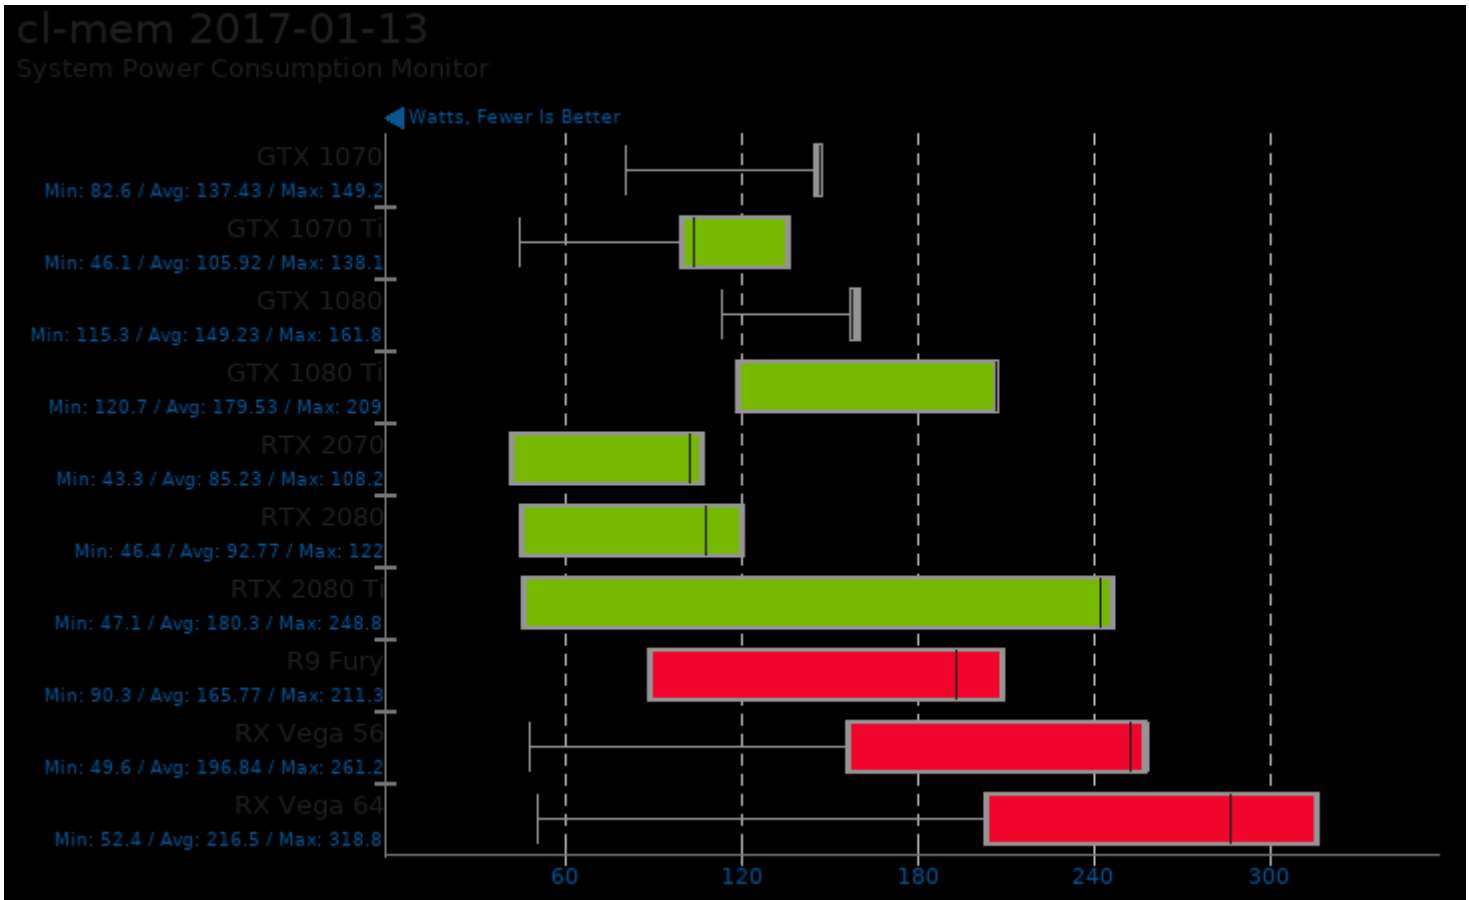
cl-mem 2017-01-13
System Power Consumption Monitor
Watts, Fewer Is Better
GTX 1070
Min: 82.6 / Avg: 137.43 / Max: 149.2
GTX 1070 Ti
Min: 46.1 / Avg: 105.92 / Max: 138.1
GTX 1080
Min: 115.3 / Avg: 149.23 / Max: 161.8
GTX 1080 Ti
Min: 120.7 / Avg: 179.53 / Max: 209
RTX 2070
Min: 43.3 / Avg: 85.23 / Max: 108.2
RTX 2080
Min: 46.4 / Avg: 92.77 / Max: 122
RTX 2080 Ti
Min: 47.1 / Avg: 180.3 / Max: 248.8
R9 Fury
Min: 90.3 / Avg: 165.77 / Max: 211.3
RX Vega 56
Min: 49.6 / Avg: 196.84 / Max: 261.2
RX Vega 64
Min: 52.4 / Avg: 216.5 / Max: 318.8
60
120
180
240
300

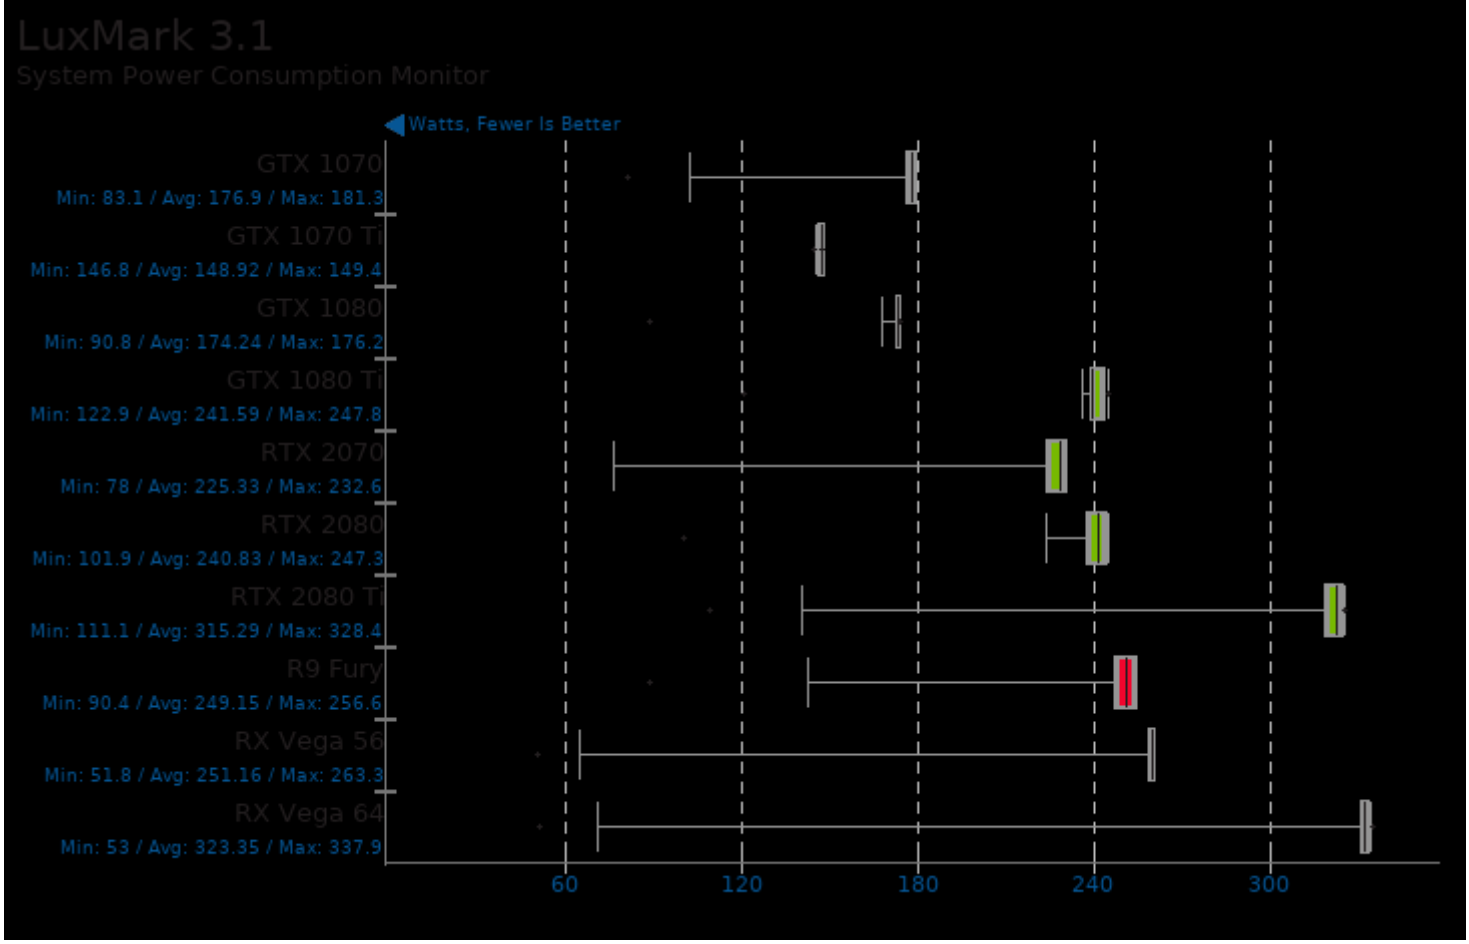
LuxMark 3.1
System Power Consumption Monitor
Watts, Fewer Is Better
GTX 1070
Min: 83.1 / Avg: 176.9 / Max: 181.3
GTX 1070 Ti
Min: 146.8 / Avg: 148.92 / Max: 149.4
GTX 1080
Min: 90.8 / Avg: 174.24 / Max: 176.2
GTX 1080 Ti
Min: 122.9 / Avg: 241.59 / Max: 247.8
RTX 2070
Min: 78 / Avg: 225.33 / Max: 232.6
RTX 2080
Min: 101.9 / Avg: 240.83 / Max: 247.3
RTX 2080 Ti
Min: 111.1 / Avg: 315.29 / Max: 328.4
R9 Fury
Min: 90.4 / Avg: 249.15 / Max: 256.6
RX Vega 56
Min: 51.8 / Avg: 251.16 / Max: 263.3
RX Vega 64
Min: 53 / Avg: 323.35 / Max: 337.9
60
120
180
240
300

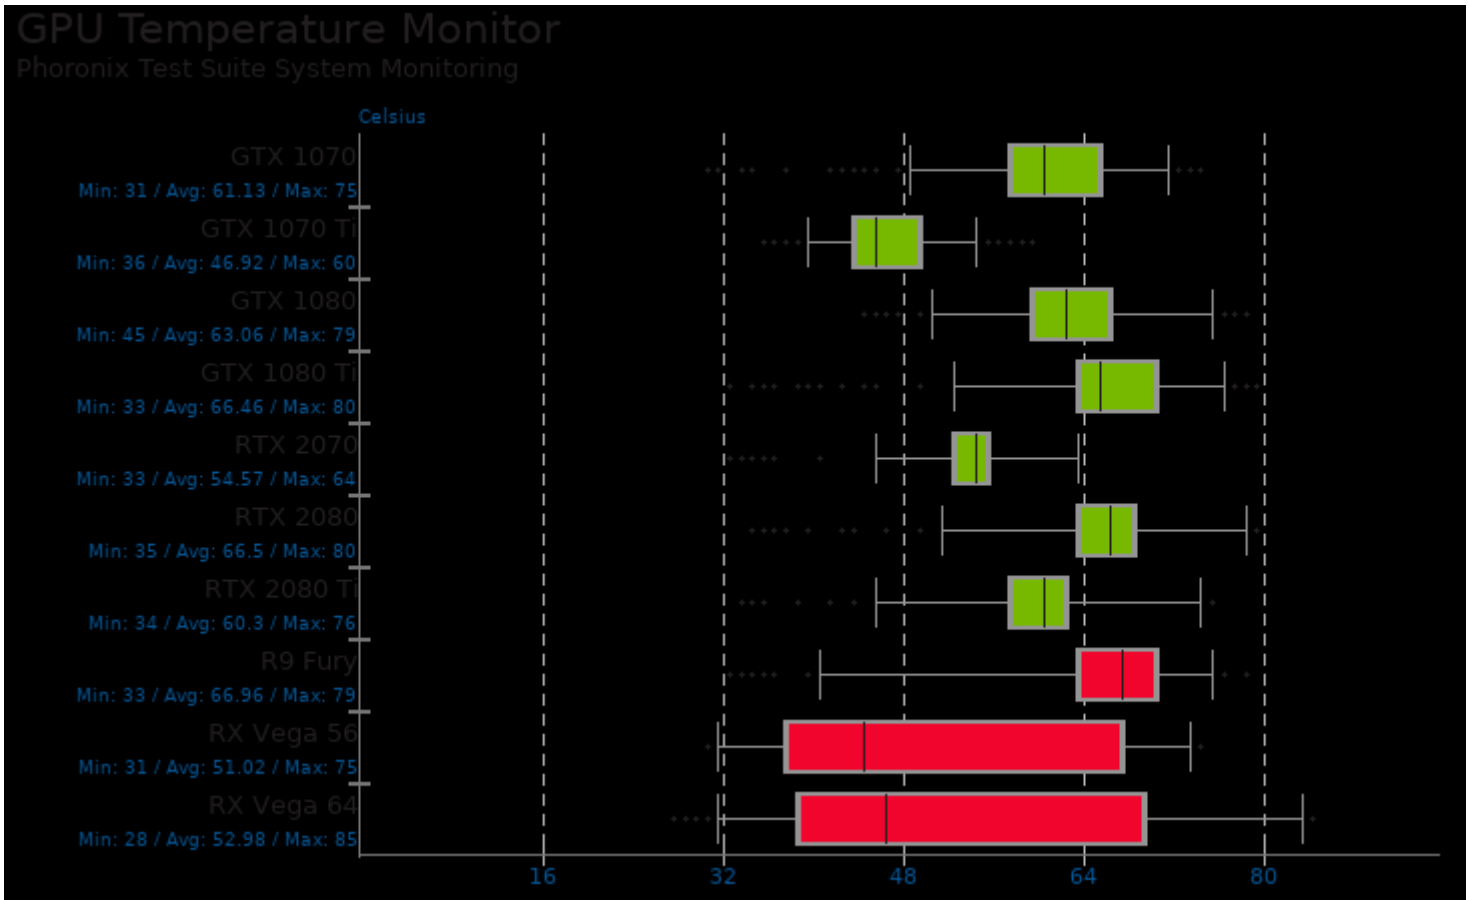
GPU Temperature Monitor
Phoronix Test Suite System Monitoring
Celsius
GTX 1070
Min: 31 / Avg: 61.13 / Max: 75
GTX 1070 Ti
Min: 36 / Avg: 46.92 / Max: 60
GTX 1080
Min: 45 / Avg: 63.06 / Max: 79
GTX 1080 Ti
Min: 33 / Avg: 66.46 / Max: 80
RTX 2070
Min: 33 / Avg: 54.57 / Max: 64
RTX 2080
Min: 35 / Avg: 66.5 / Max: 80
RTX 2080 Ti
Min: 34 / Avg: 60.3 / Max: 76
R9 Fury
Min: 33 / Avg: 66.96 / Max: 79
RX Vega 56
Min: 31 / Avg: 51.02 / Max: 75
RX Vega 64
Min: 28 / Avg: 52.98 / Max: 85
16
32
48
64
80

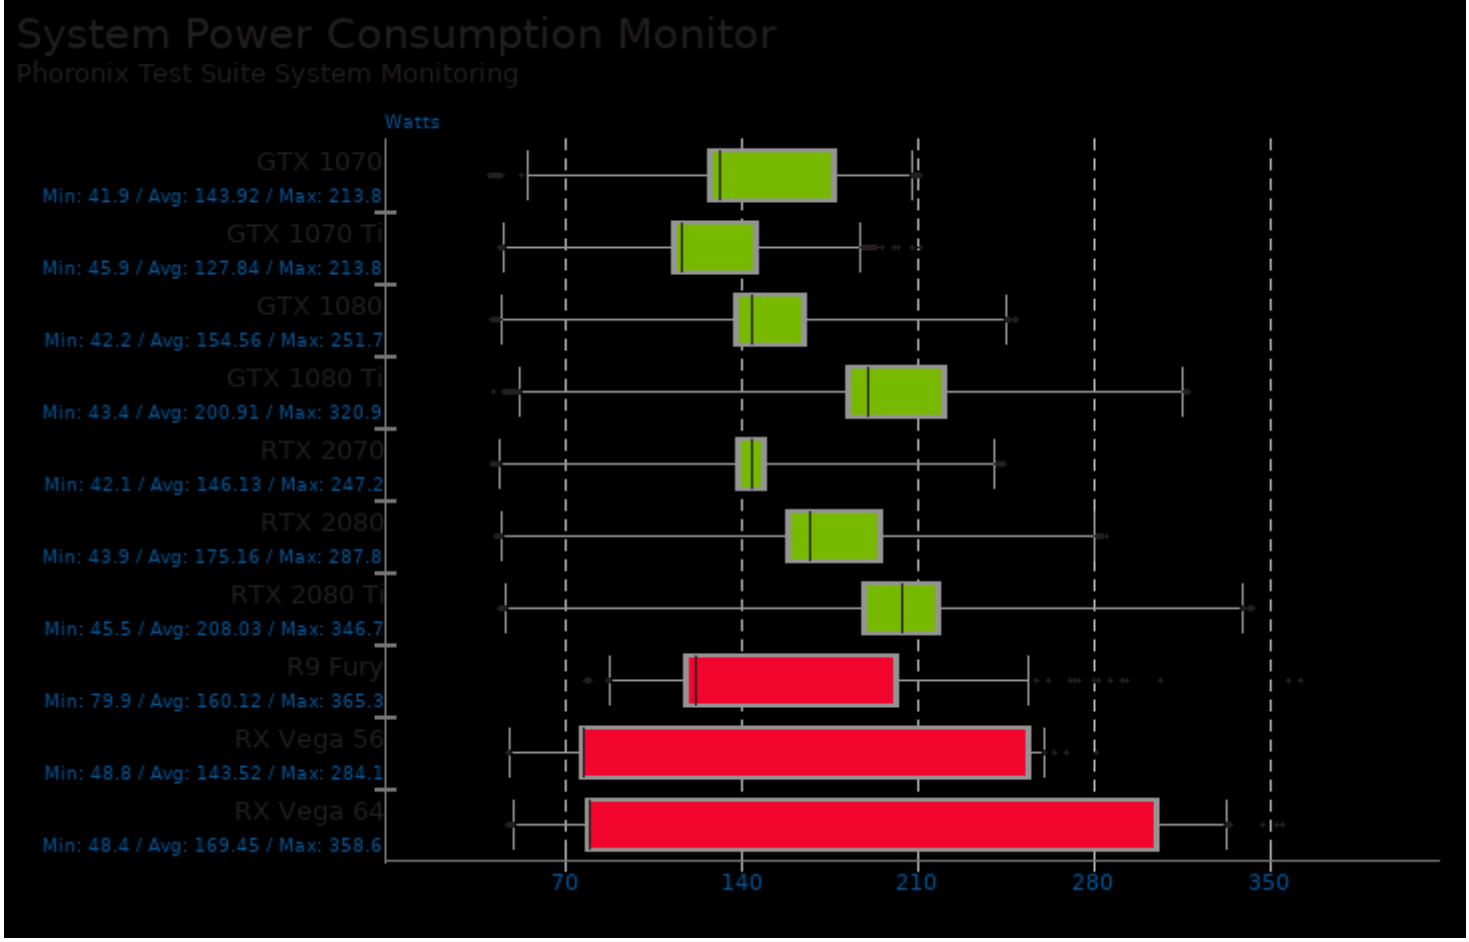
System Power Consumption Monitor
Phoronix Test Suite System Monitoring
Watts
GTX 1070
Min: 41.9 / Avg: 143.92 / Max: 213.8
GTX 1070 Ti
Min: 45.9 / Avg: 127.84 / Max: 213.8
GTX 1080
Min: 42.2 / Avg: 154.56 / Max: 251.7
GTX 1080 Ti
Min: 43.4 / Avg: 200.91 / Max: 320.9
RTX 2070
Min: 42.1 / Avg: 146.13 / Max: 247.2
RTX 2080
Min: 43.9 / Avg: 175.16 / Max: 287.8
RTX 2080 Ti
Min: 45.5 / Avg: 208.03 / Max: 346.7
R9 Fury
Min: 79.9 / Avg: 160.12 / Max: 365.3
RX Vega 56
Min: 48.8 / Avg: 143.52 / Max: 284.1
RX Vega 64
Min: 48.4 / Avg: 169.45 / Max: 358.6
70
140
210
280
350

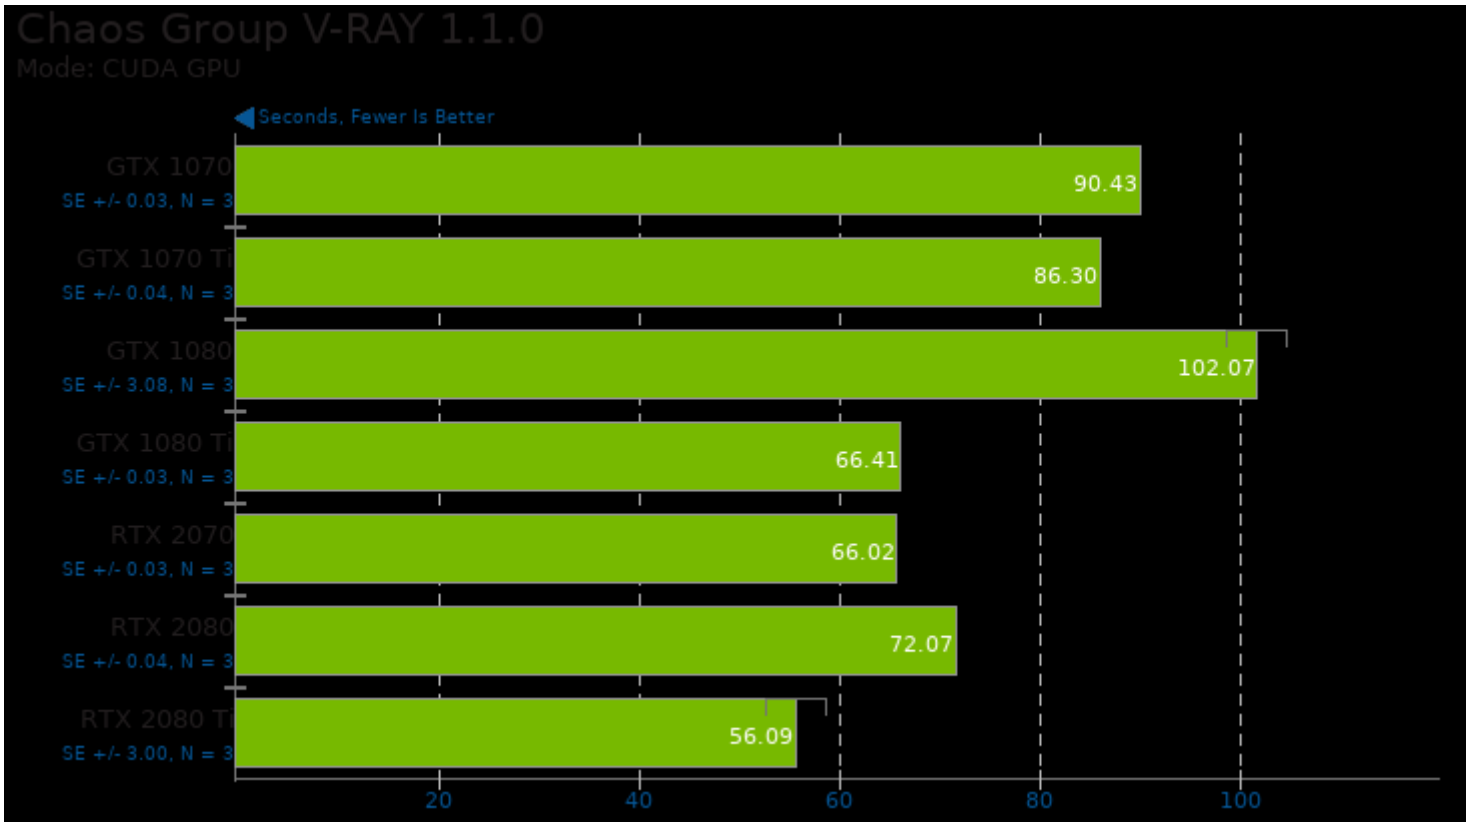
Chaos Group V-RAY 1.1.0
Mode: CUDA GPU
Seconds, Fewer Is Better
GTX 1070
SE +/- 0.03, N = 3
90.43
GTX 1070 Ti
SE +/- 0.04, N = 3
86.30
GTX 1080
SE +/- 3.08, N = 3
102.07
GTX 1080 Ti
SE +/- 0.03, N = 3
66.41
RTX 2070
SE +/- 0.03, N = 3
66.02
RTX 2080
SE +/- 0.04, N = 3
72.07
RTX 2080 Ti
SE +/- 3.00, N = 3
56.09
20
40
60
80
100

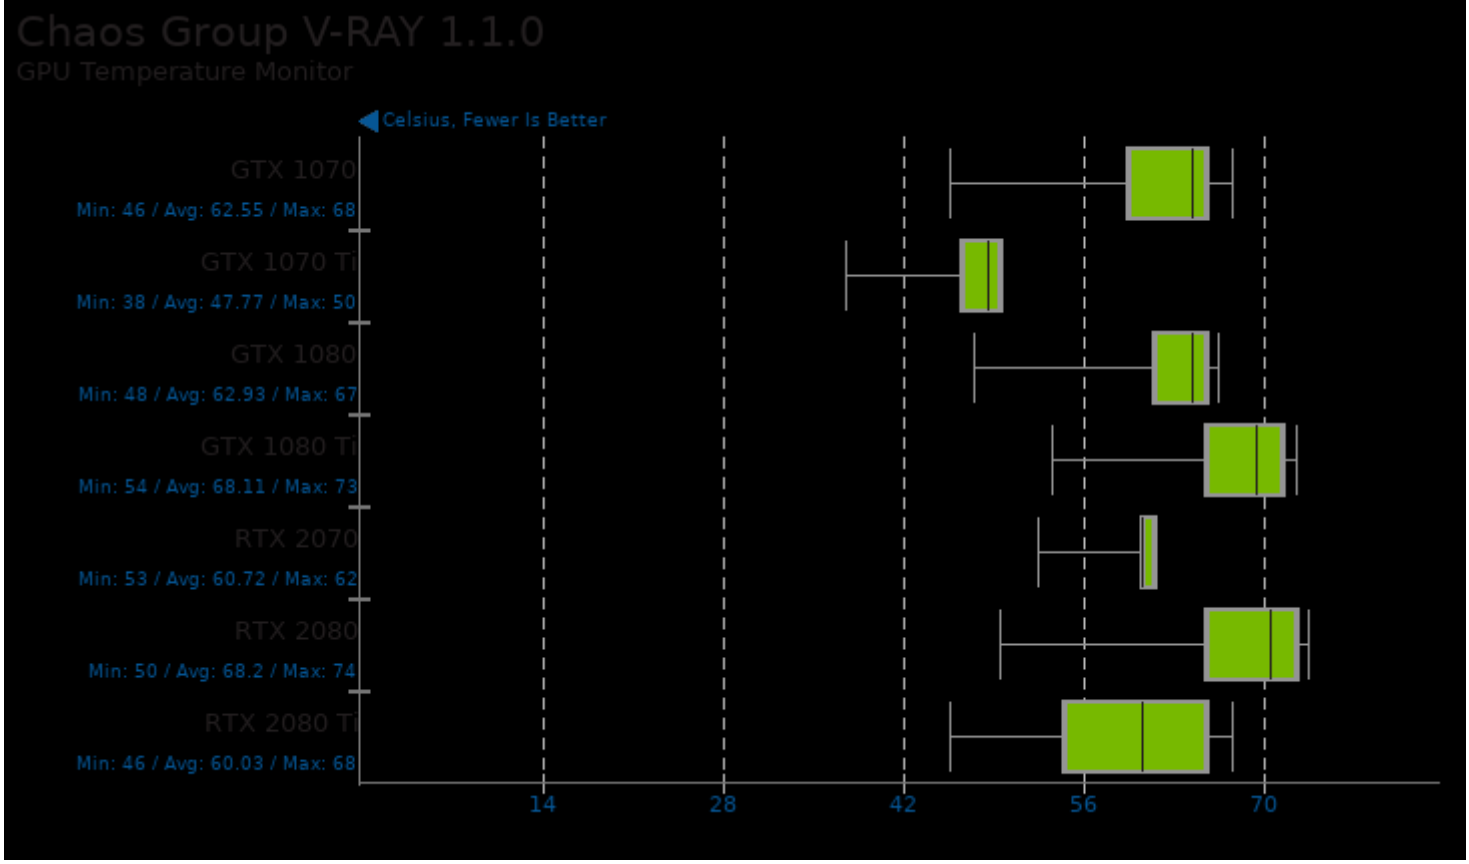
Chaos Group V-RAY 1.1.0
GPU Temperature Monitor
Celsius, Fewer Is Better
GTX 1070
Min: 46 / Avg: 62.55 / Max: 68
GTX 1070 Ti
Min: 38 / Avg: 47.77 / Max: 50
GTX 1080
Min: 48 / Avg: 62.93 / Max: 67
GTX 1080 Ti
Min: 54 / Avg: 68.11 / Max: 73
RTX 2070
Min: 53 / Avg: 60.72 / Max: 62
RTX 2080
Min: 50 / Avg: 68.2 / Max: 74
RTX 2080 Ti
Min: 46 / Avg: 60.03 / Max: 68
14
28
42
56
70

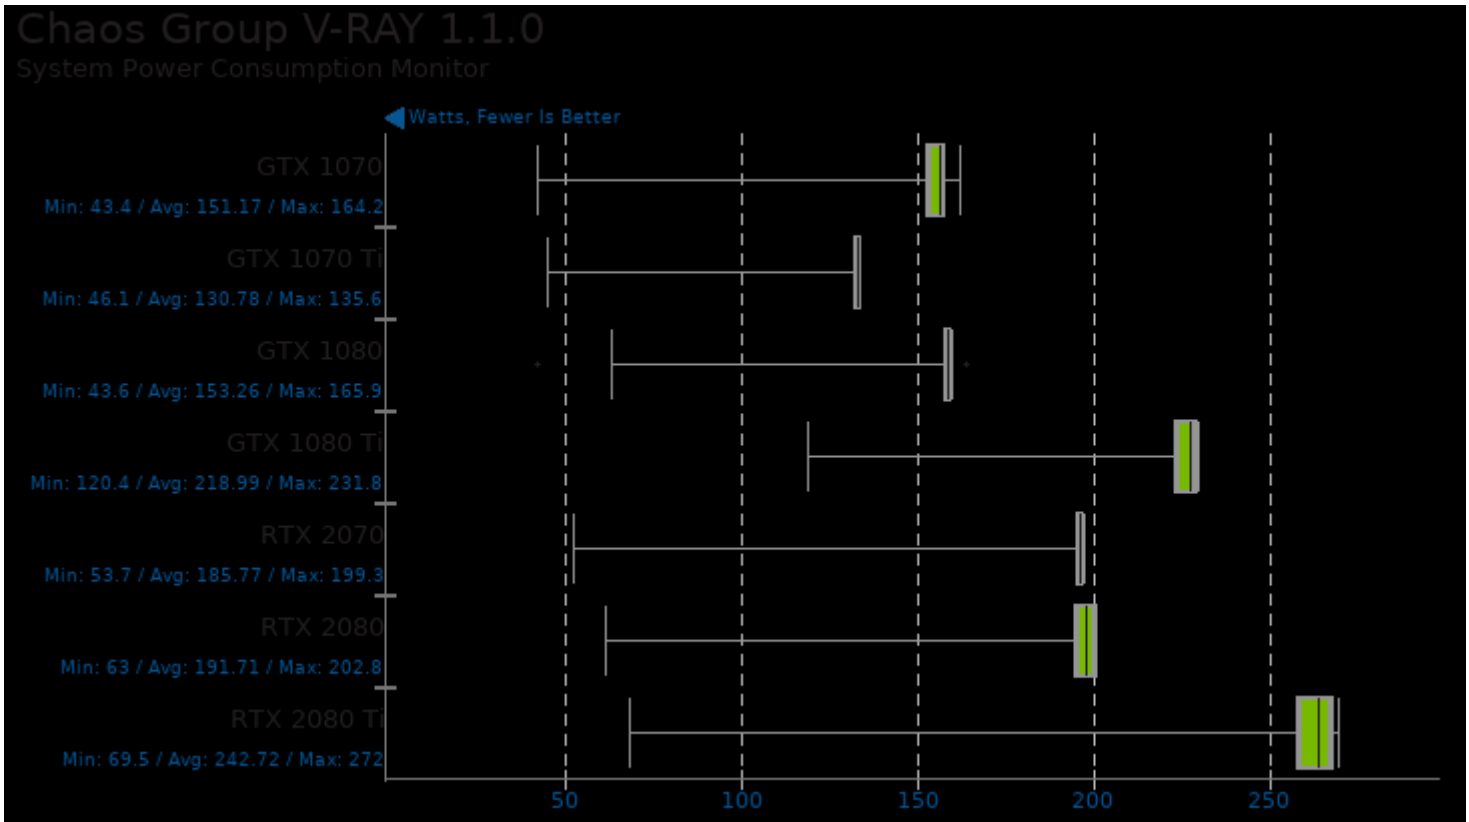
Chaos Group V-RAY 1.1.0
System Power Consumption Monitor
Watts, Fewer Is Better
GTX 1070
Min: 43.4 / Avg: 151.17 / Max: 164.2
GTX 1070 Ti
Min: 46.1 / Avg: 130.78 / Max: 135.6
GTX 1080
Min: 43.6 / Avg: 153.26 / Max: 165.9
GTX 1080 Ti
Min: 120.4 / Avg: 218.99 / Max: 231.8
RTX 2070
Min: 53.7 / Avg: 185.77 / Max: 199.3
RTX 2080
Min: 63 / Avg: 191.71 / Max: 202.8
RTX 2080 Ti
Min: 69.5 / Avg: 242.72 / Max: 272
50
100
150
200
250

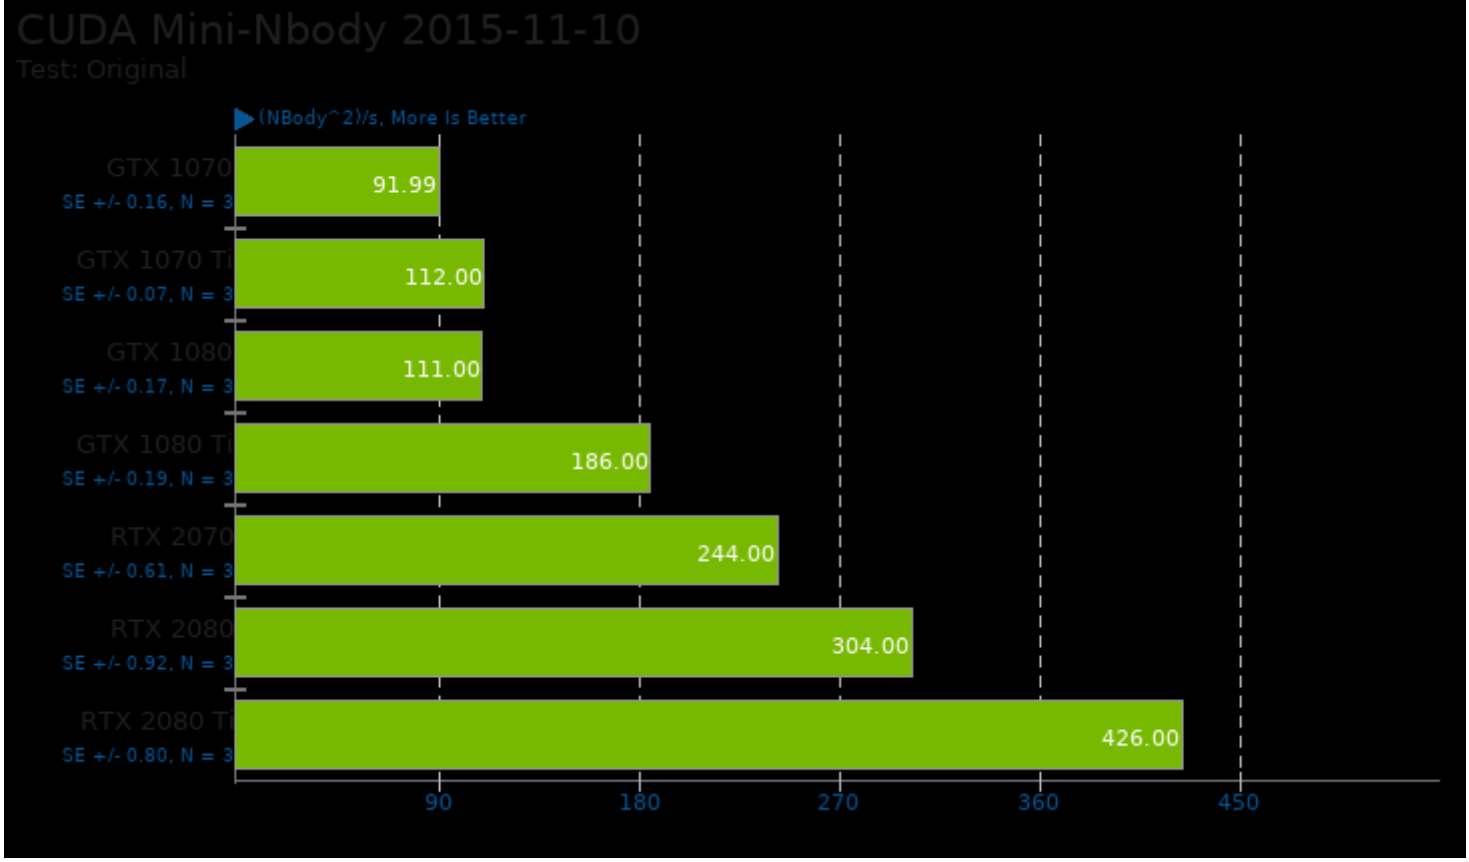
CUDA Mini-Nbody 2015-11-10
Test: Original
(NBody^2)/s, More Is Better
GTX 1070
SE +/- 0.16, N = 3
91.99
GTX 1070 Ti
SE +/- 0.07, N = 3
112.00
GTX 1080
SE +/- 0.17, N = 3
111.00
GTX 1080 Ti
SE +/- 0.19, N = 3
186.00
RTX 2070
SE +/- 0.61, N = 3
244.00
RTX 2080
SE +/- 0.92, N = 3
304.00
RTX 2080 Ti
SE +/- 0.80, N = 3
426.00
90
180
270
360
450

NVIDIA AMD Linux GPU Compute December 2018

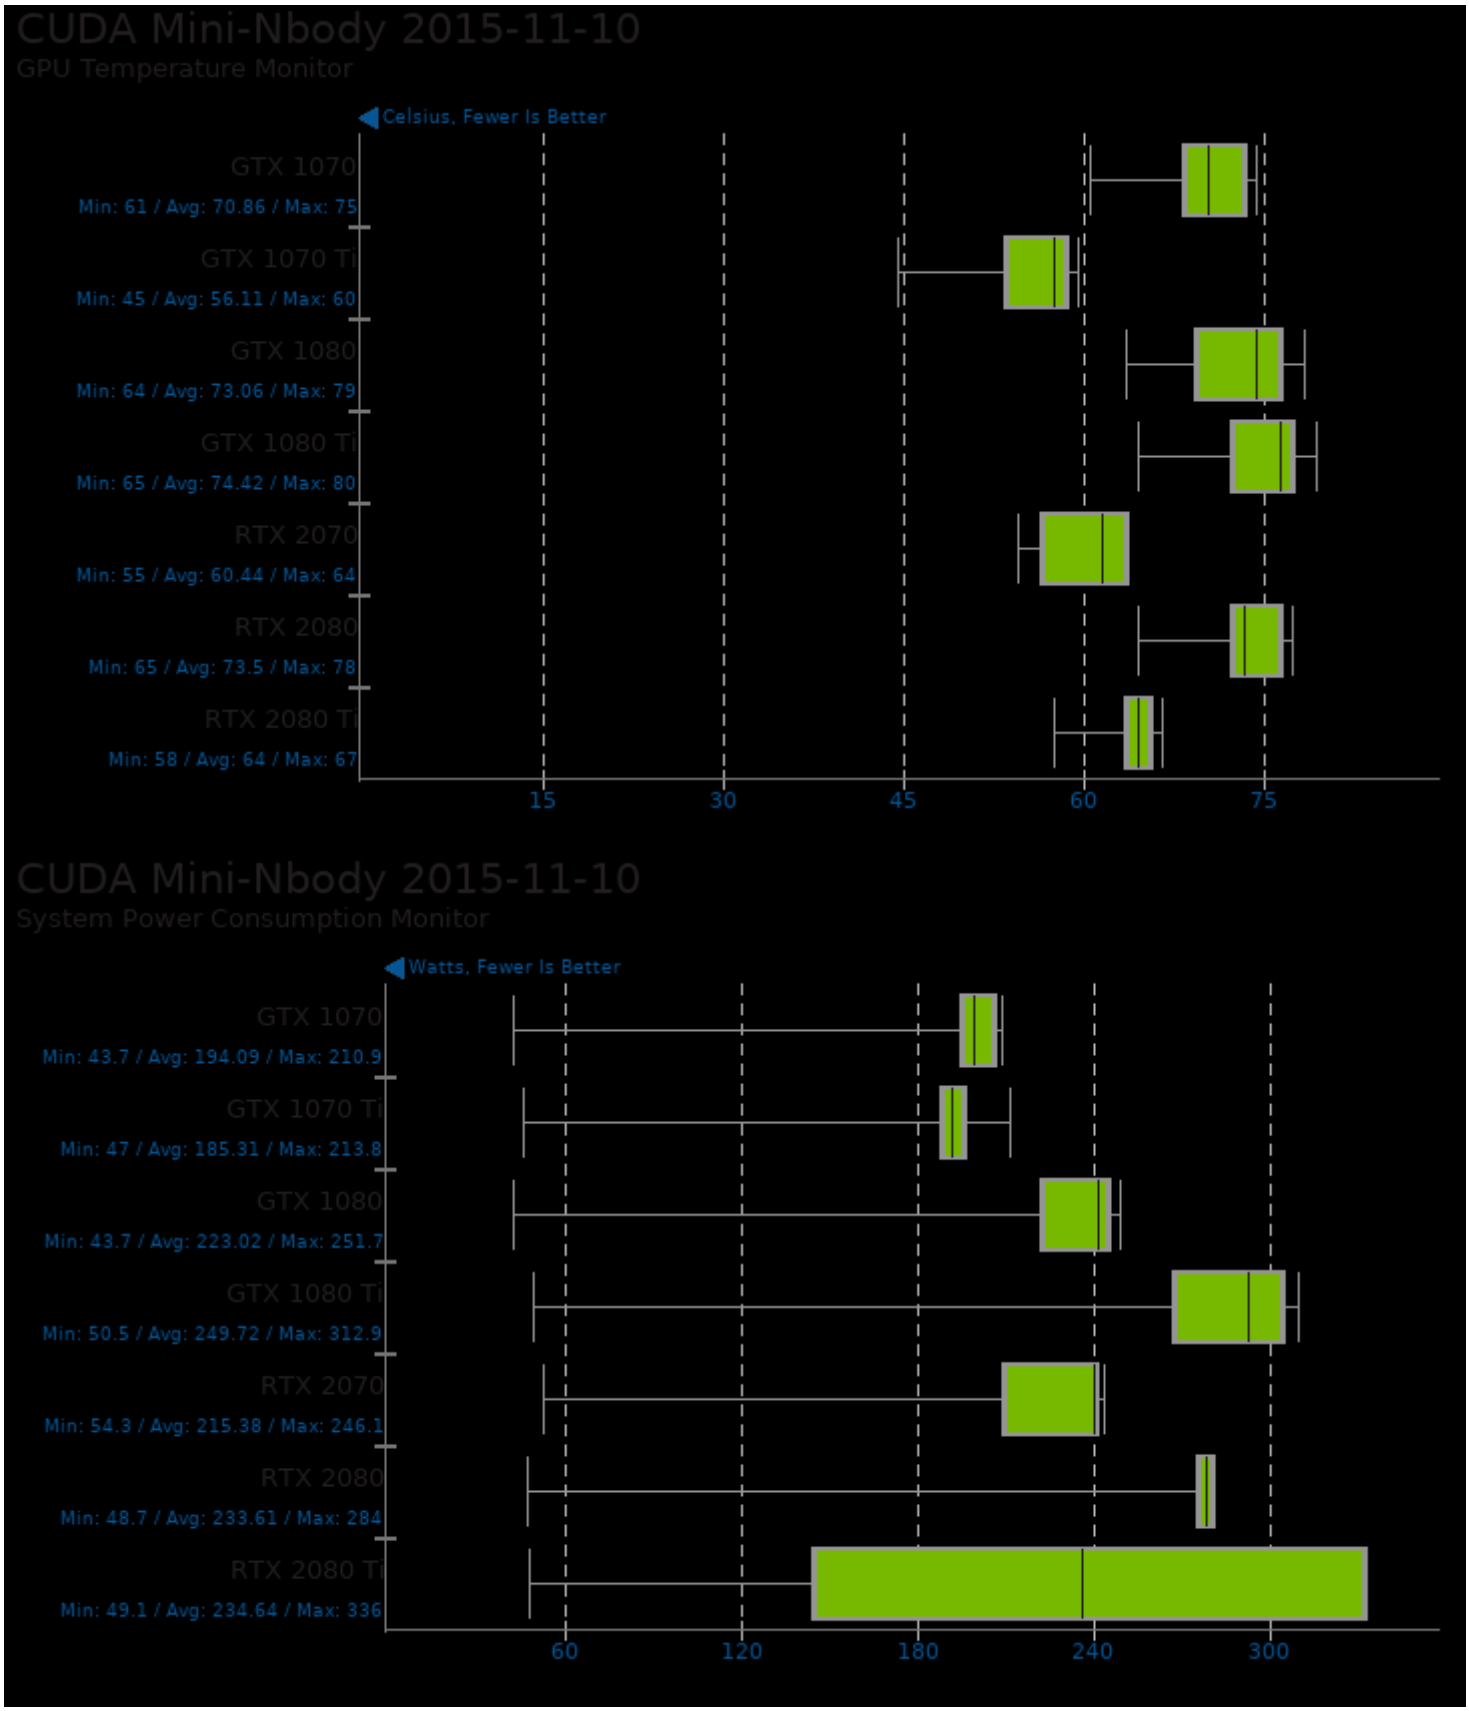

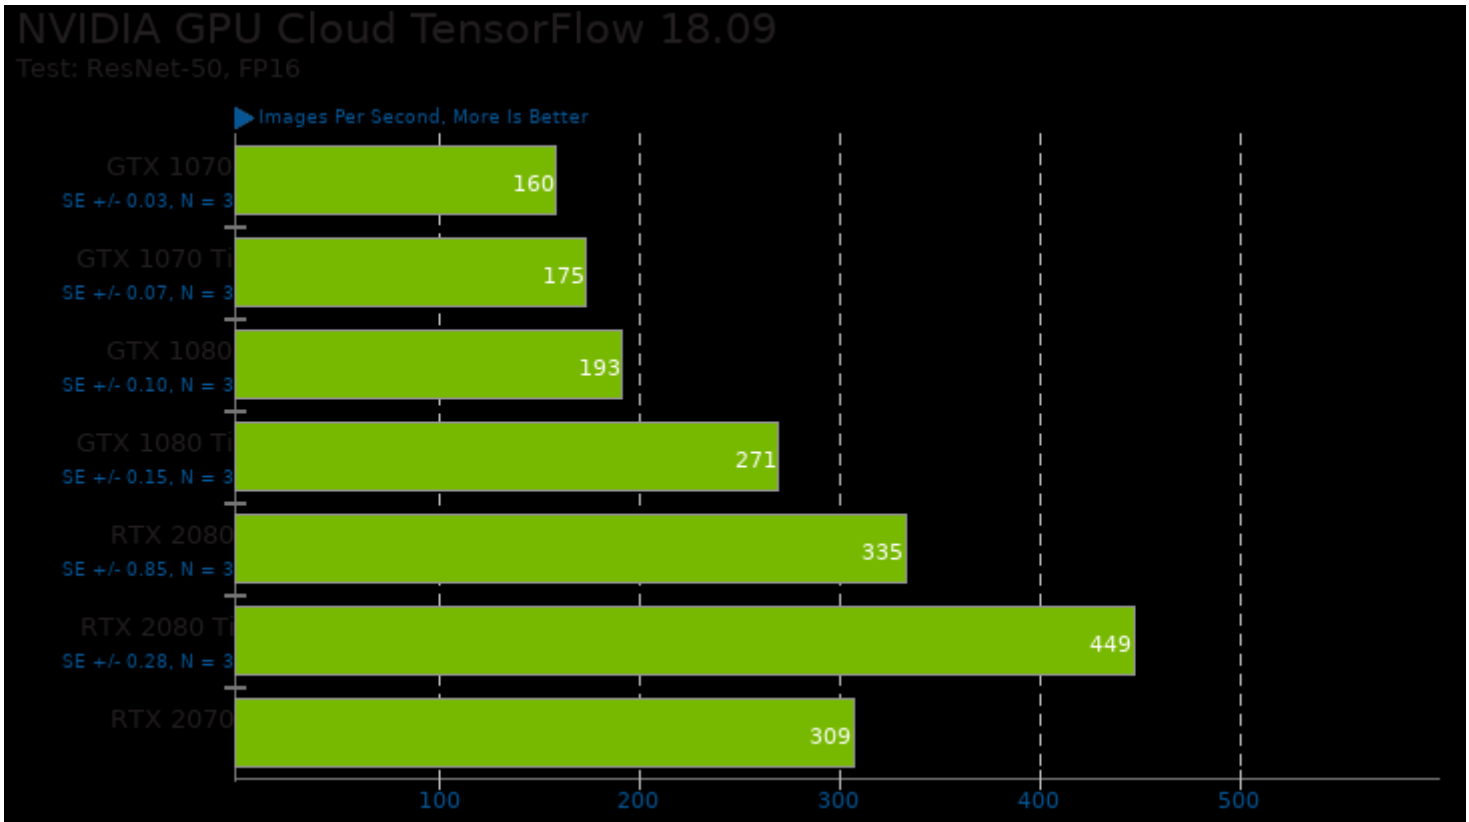
NVIDIA GPU Cloud TensorFlow 18.09
Test: ResNet-50, FP16
Images Per Second, More Is Better
GTX 1070
SE +/- 0.03, N = 3
160
GTX 1070 Ti
SE +/- 0.07, N = 3
175
GTX 1080
SE +/- 0.10, N = 3
193
GTX 1080 Ti
SE +/- 0.15, N = 3
271
RTX 2080
SE +/- 0.85, N = 3
335
RTX 2080 Ti
SE +/- 0.28, N = 3
449
RTX 2070
309
100
200
300
400
500

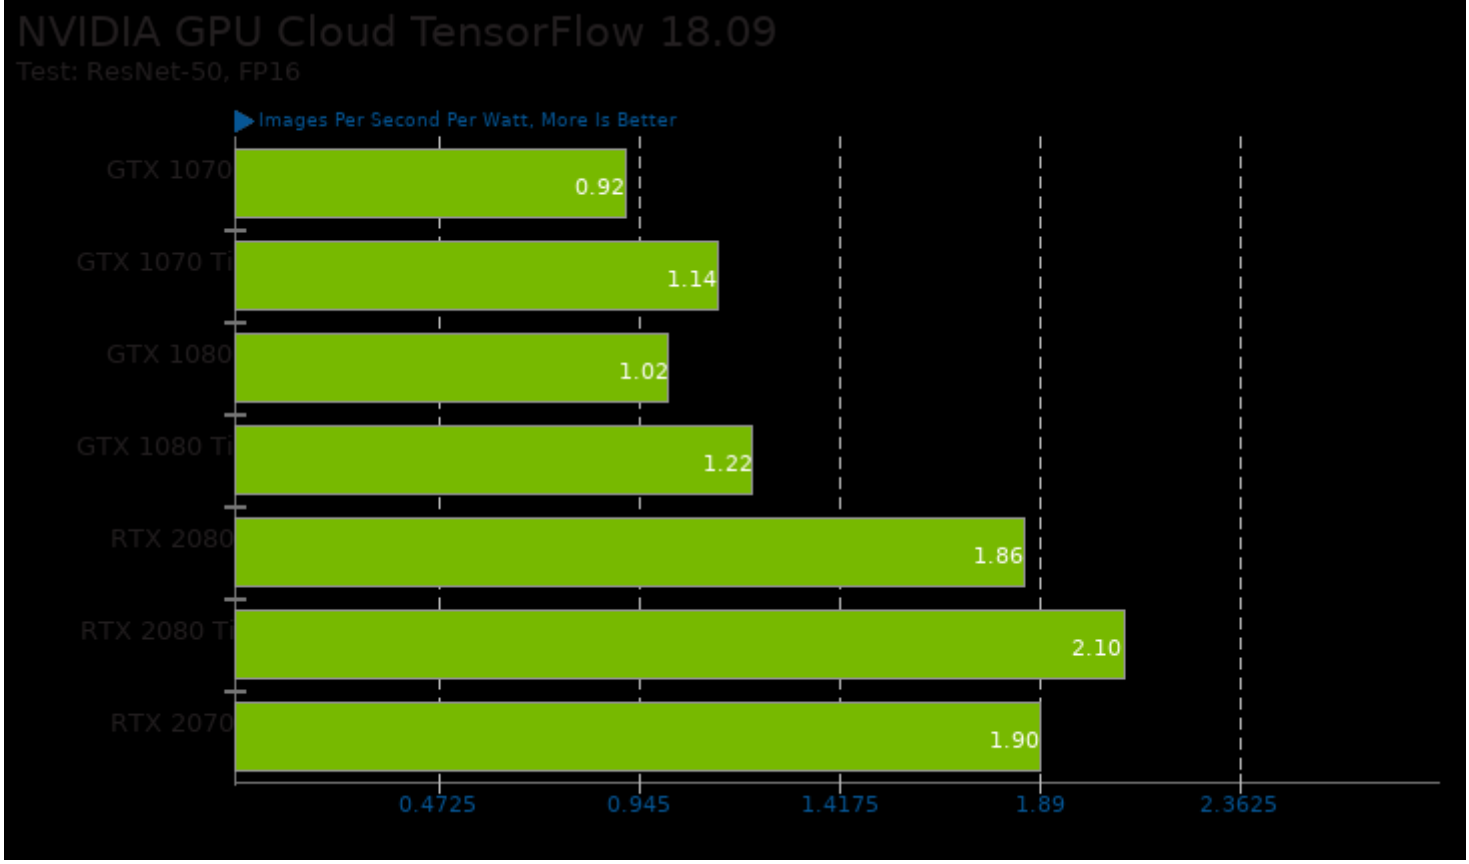
NVIDIA GPU Cloud TensorFlow 18.09
Test: ResNet-50, FP16
Images Per Second Per Watt, More Is Better
GTX 1070
0.92
GTX 1070 Ti
1.14
GTX 1080
1.02
GTX 1080 Ti
1.22
RTX 2080
1.86
RTX 2080 Ti
2.10
RTX 2070
1.90
0.4725
0.945
1.4175
1.89
2.3625

NVIDIA AMD Linux GPU Compute December 2018

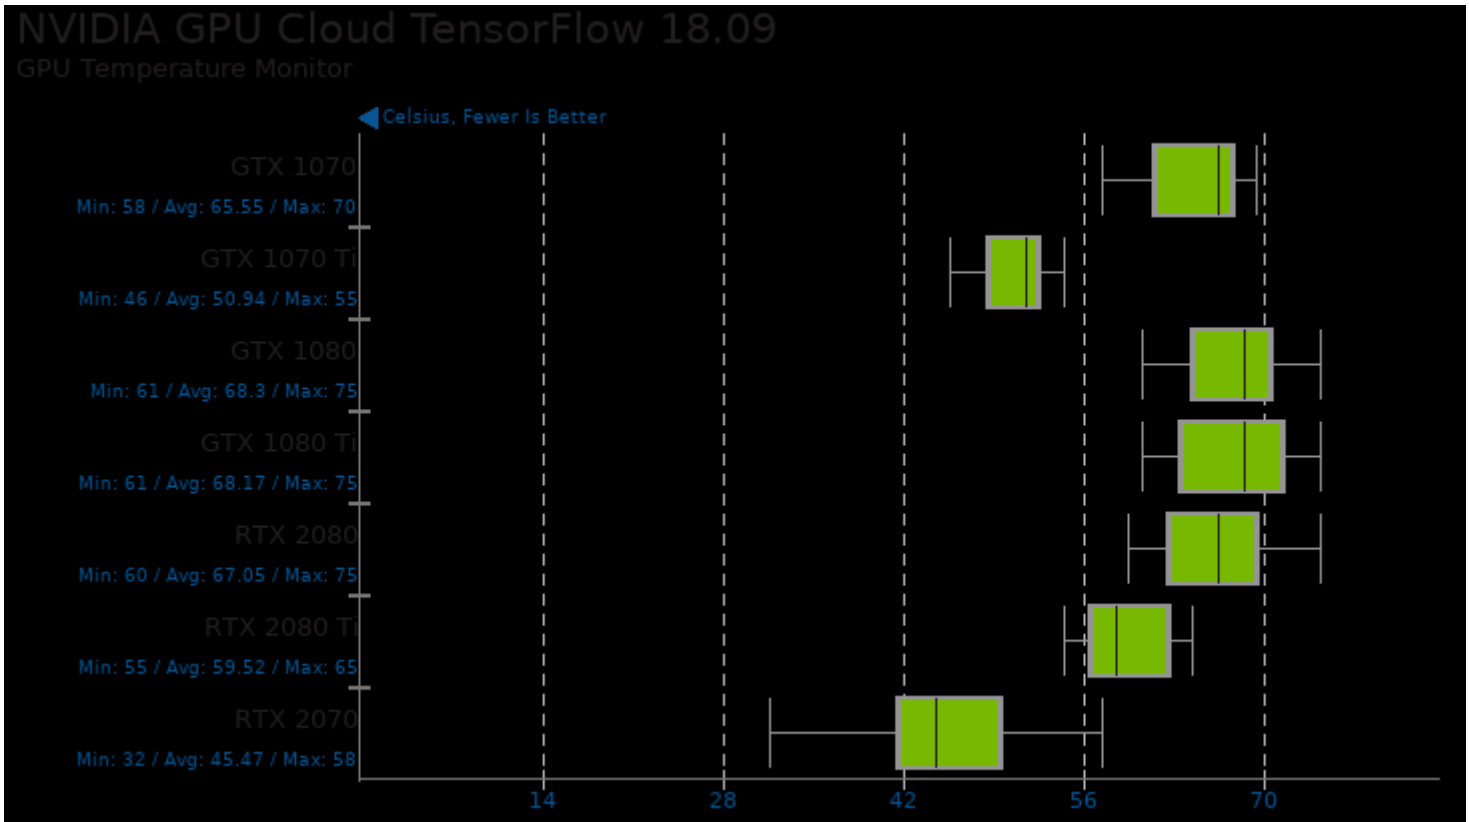

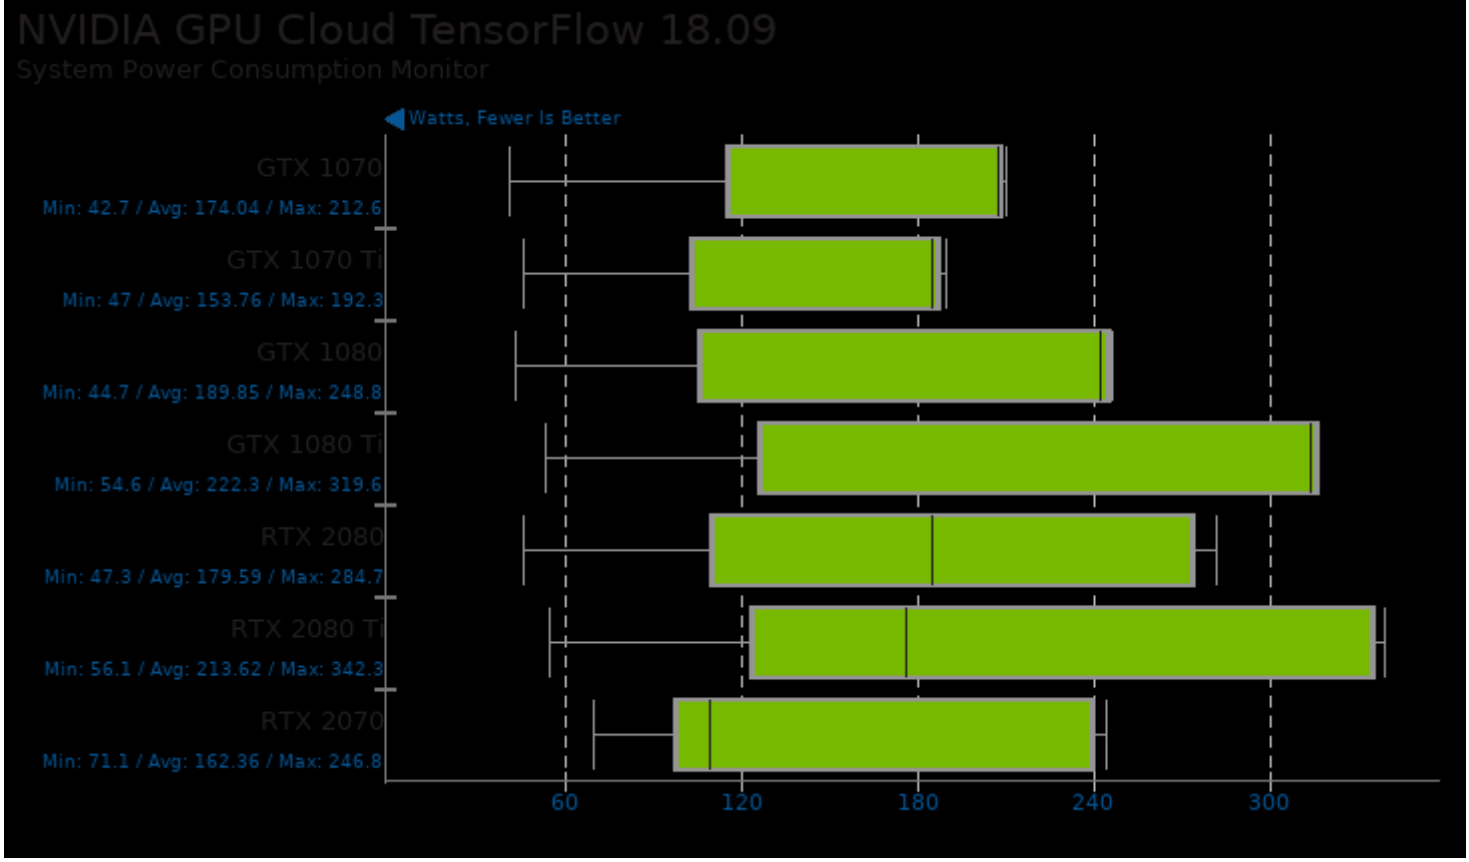

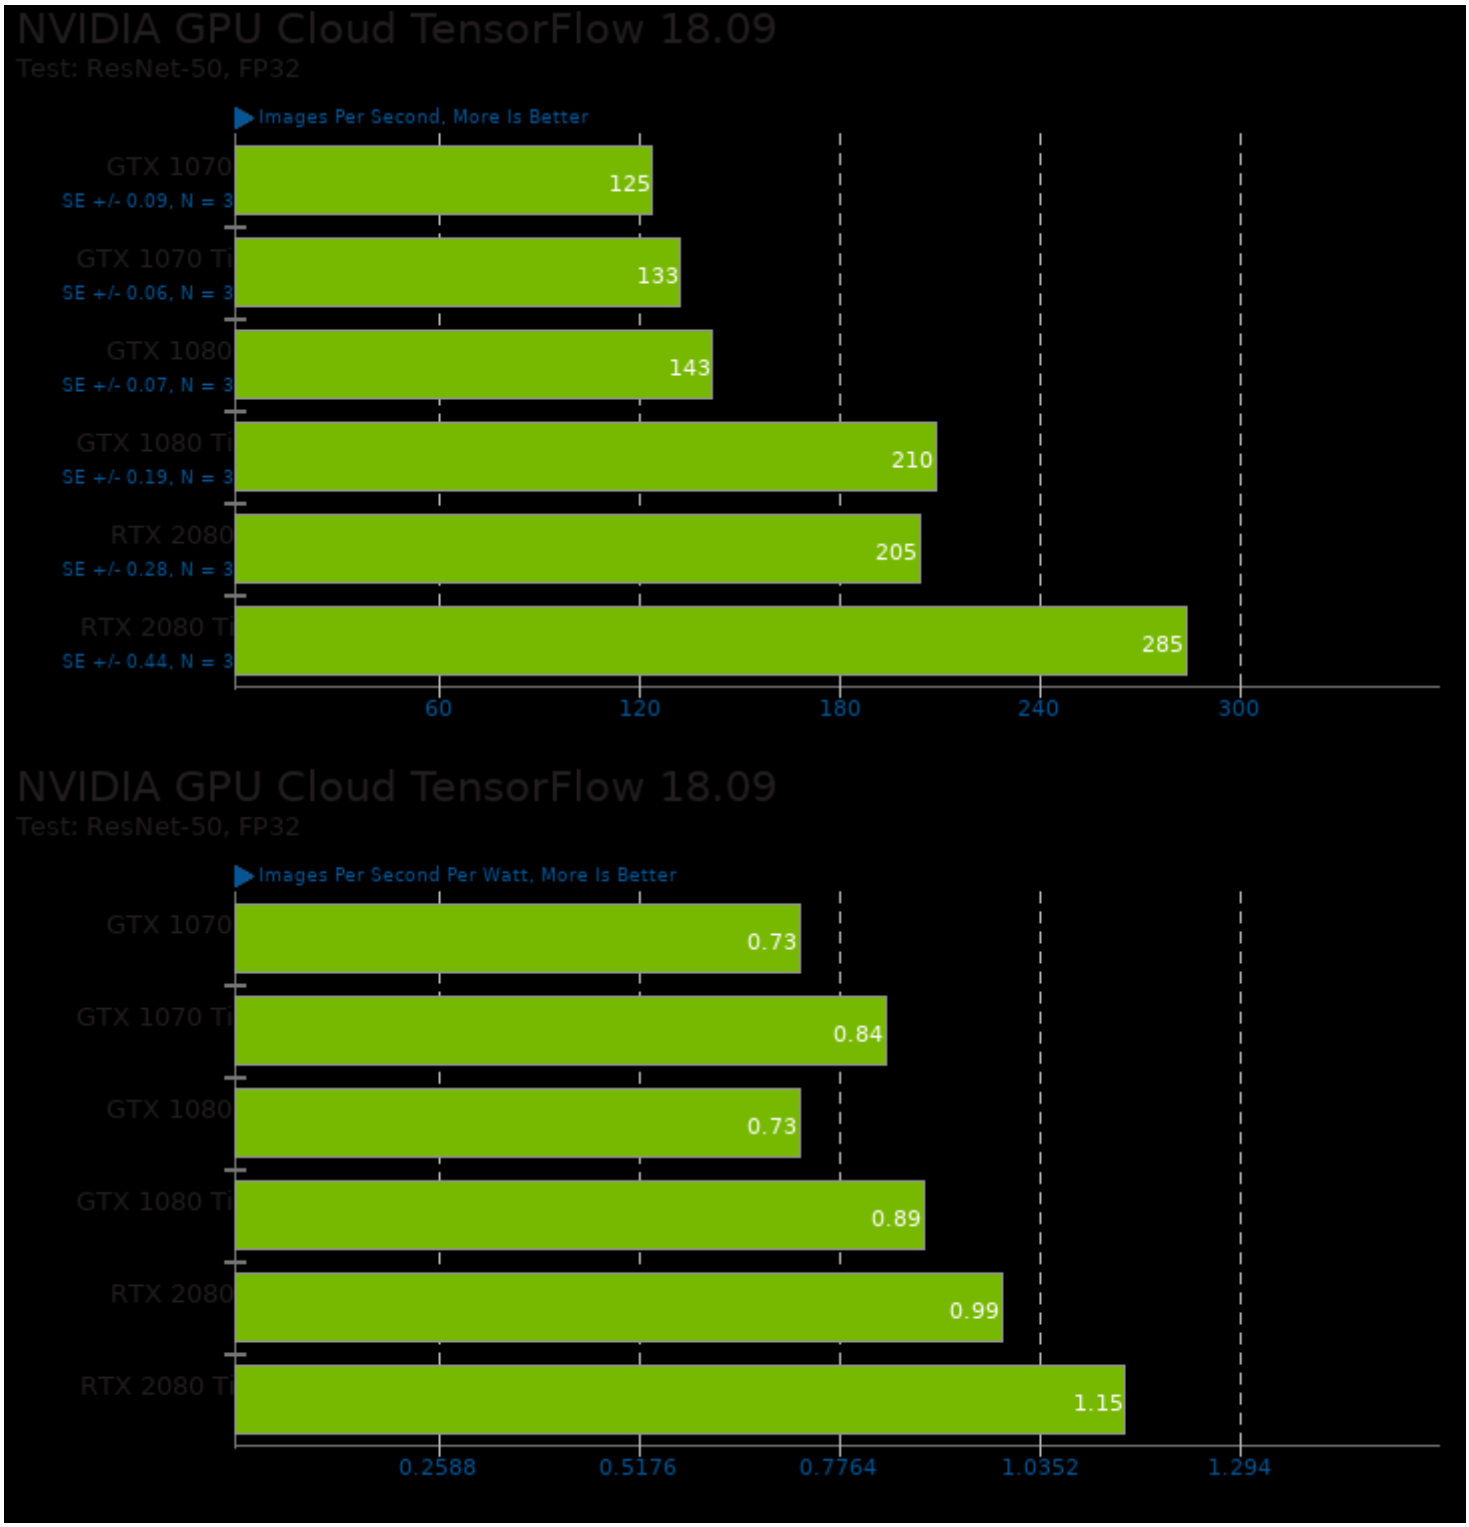
NVIDIA GPU Cloud TensorFlow 18.09
Test: ResNet-50, FP32
Images Per Second, More Is Better
GTX 1070
SE +/- 0.09, N = 3
125
GTX 1070 Ti
SE +/- 0.06, N = 3
133
GTX 1080
SE +/- 0.07, N = 3
143
GTX 1080 Ti
SE +/- 0.19, N = 3
210
RTX 2080
SE +/- 0.28, N = 3
205
RTX 2080 Ti
SE +/- 0.44, N = 3
285
60
120
180
240
300
NVIDIA GPU Cloud TensorFlow 18.09
Test: ResNet-50, FP32
Images Per Second Per Watt, More Is Better
GTX 1070
0.73
GTX 1070 Ti
0.84
GTX 1080
0.73
GTX 1080 Ti
0.89
RTX 2080
0.99
RTX 2080 Ti
1.15
0.2588
0.5176
0.7764
1.0352
1.294

NVIDIA AMD Linux GPU Compute December 2018

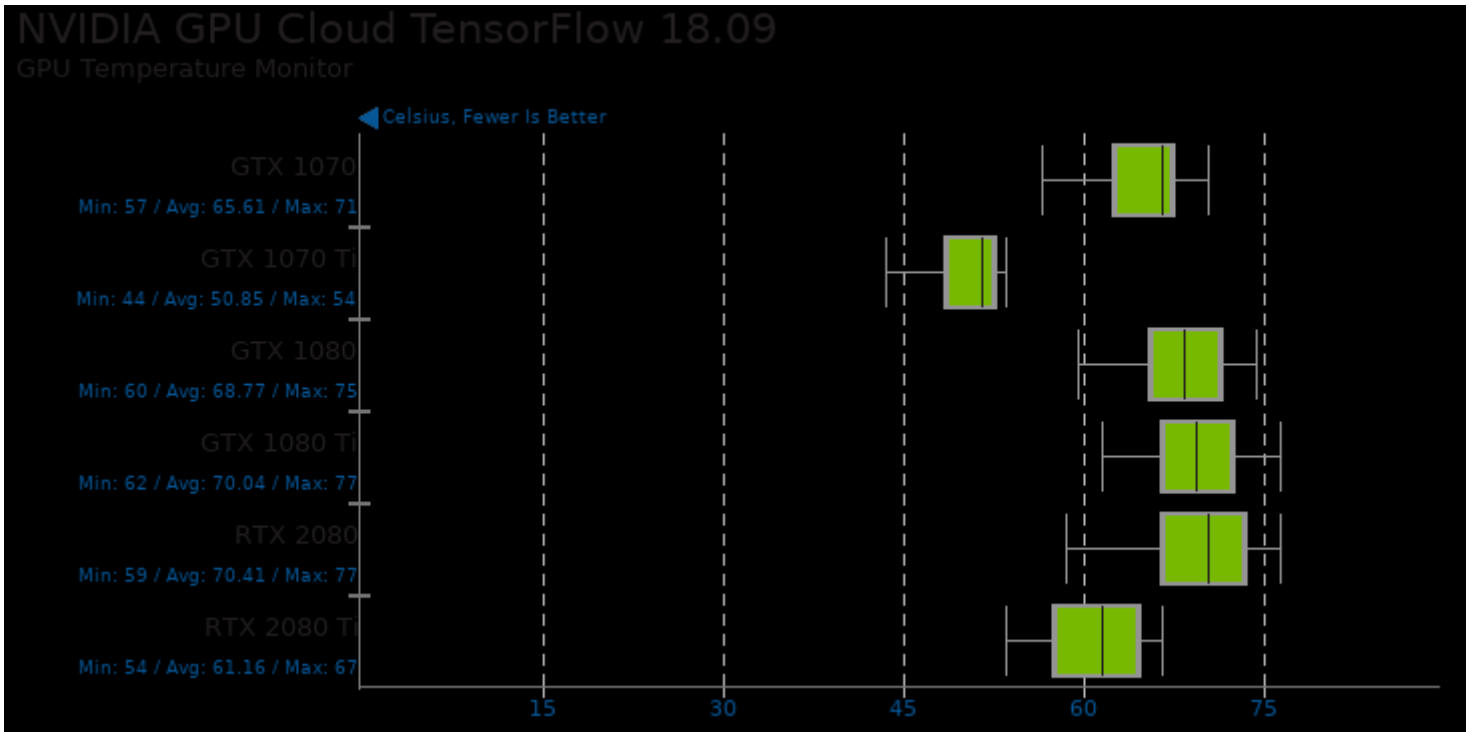

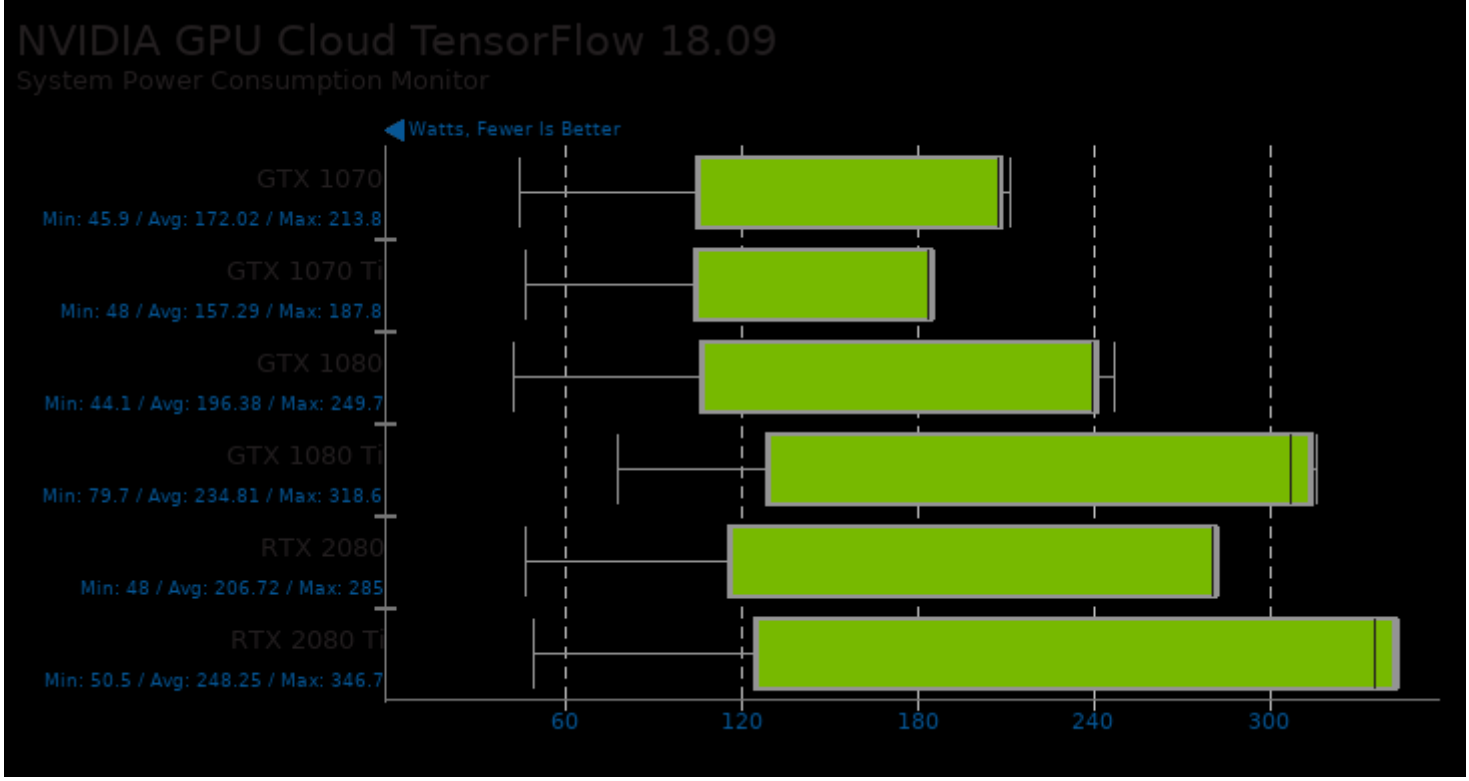

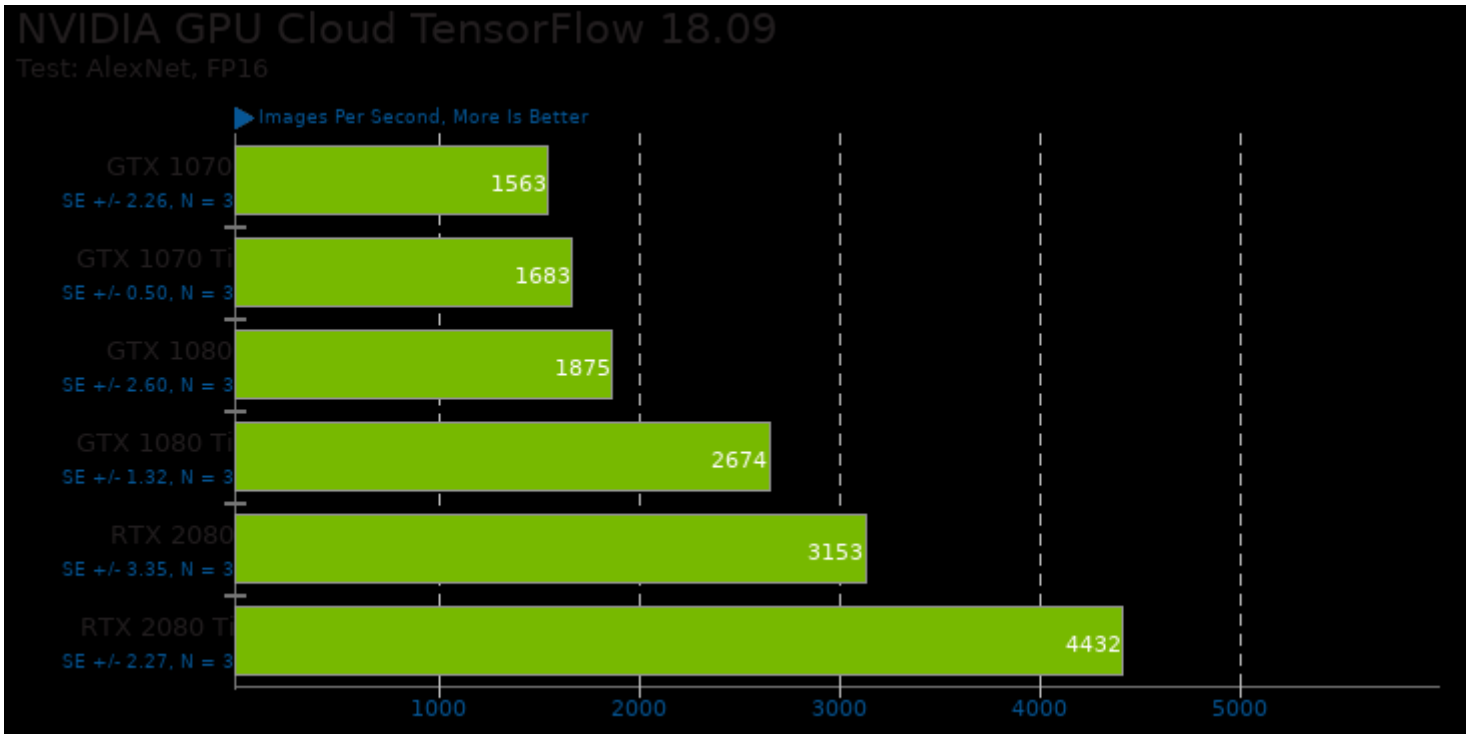
NVIDIA GPU Cloud TensorFlow 18.09
Test: AlexNet, FP16
Images Per Second, More Is Better
GTX 1070
SE +/- 2.26, N = 3
1563
GTX 1070 Ti
SE +/- 0.50, N = 3
1683
GTX 1080
SE +/- 2.60, N = 3
1875
GTX 1080 Ti
SE +/- 1.32, N = 3
2674
RTX 2080
SE +/- 3.35, N = 3
3153
RTX 2080 Ti
SE +/- 2.27, N = 3
4432
1000
2000
3000
4000
5000

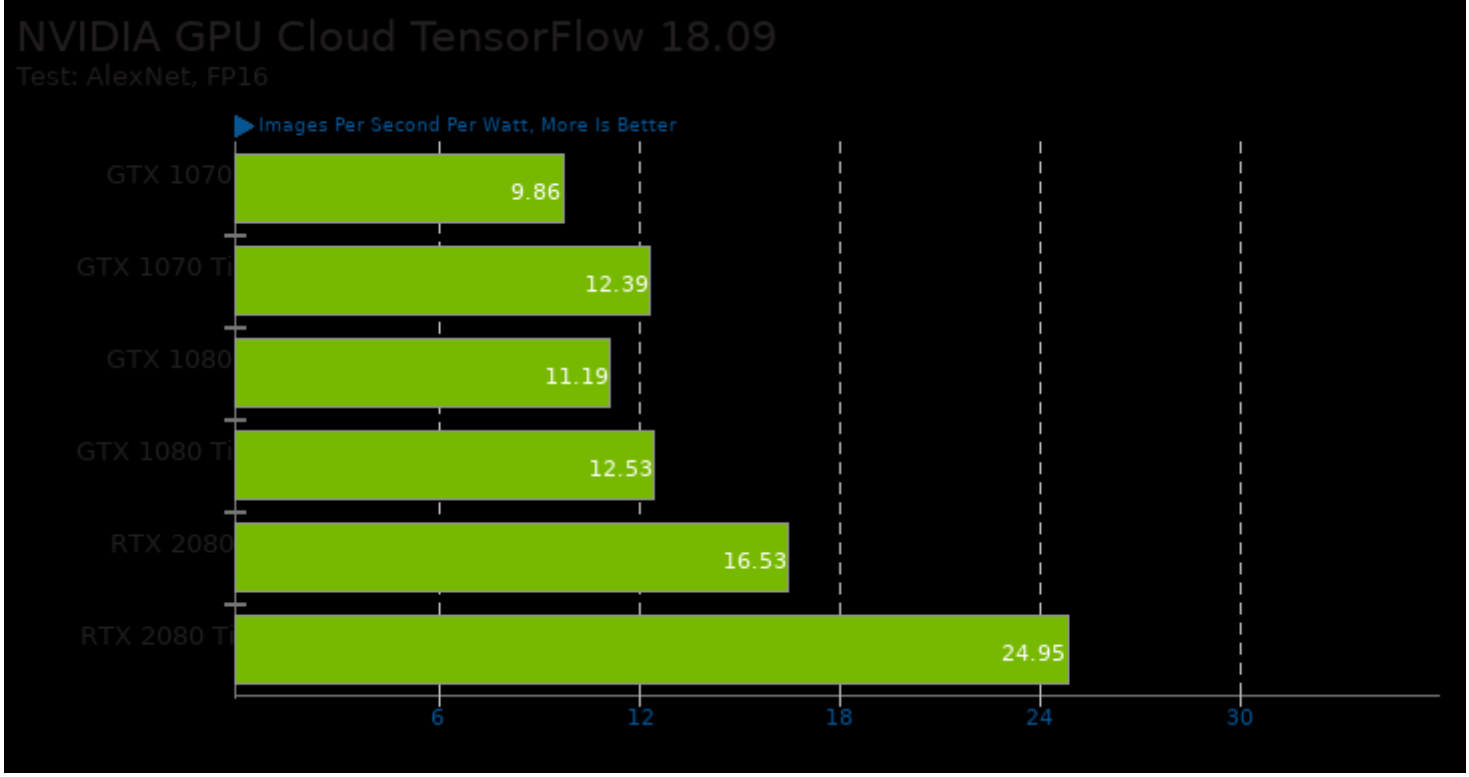
NVIDIA GPU Cloud TensorFlow 18.09
Test: AlexNet, FP16
Images Per Second Per Watt, More Is Better
GTX 1070
9.86
GTX 1070 Ti
12.39
GTX 1080
11.19
GTX 1080 Ti
12.53
RTX 2080
16.53
RTX 2080 Ti
24.95
6
12
18
24
30

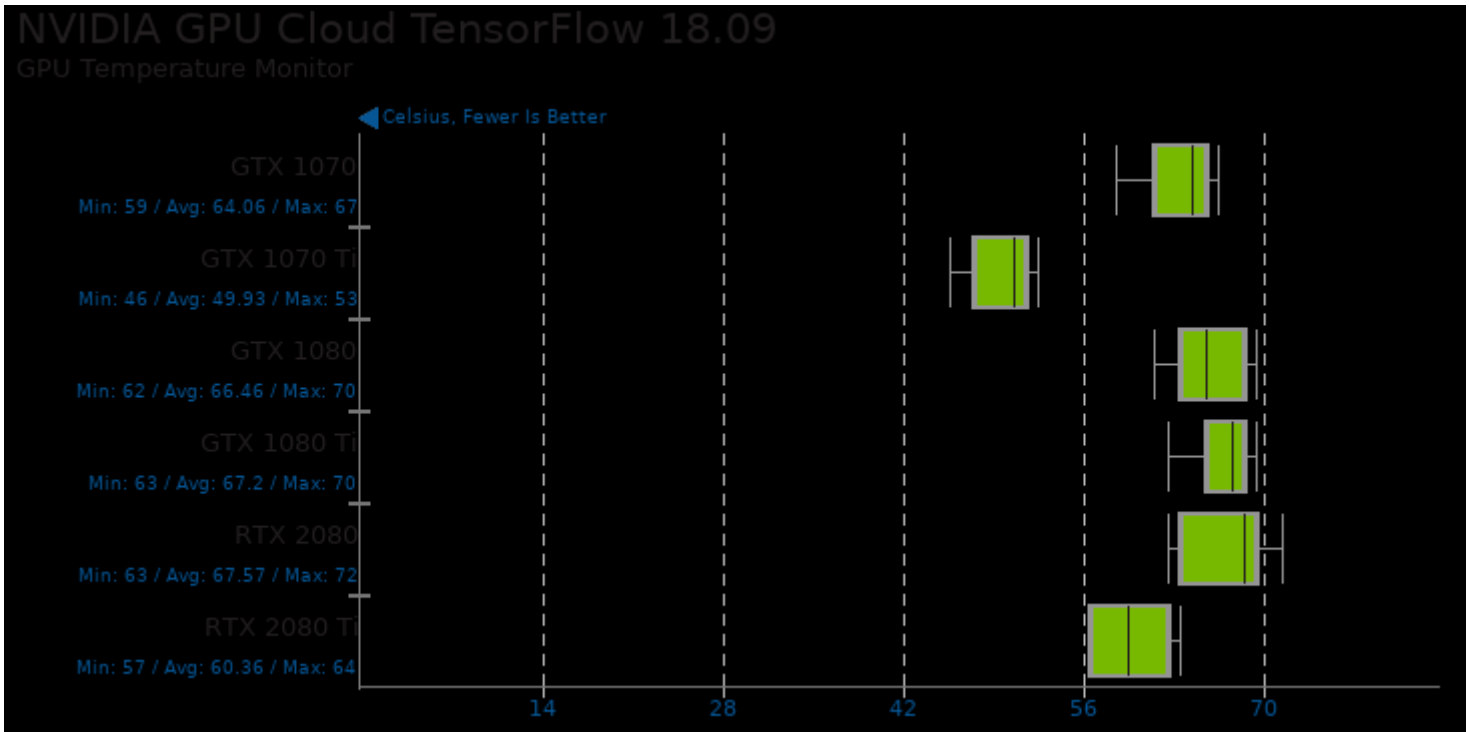
NVIDIA GPU Cloud TensorFlow 18.09
GPU Temperature Monitor
Celsius, Fewer Is Better
GTX 1070
Min: 59 / Avg: 64.06 / Max: 67
GTX 1070 Ti
Min: 46 / Avg: 49.93 / Max: 53
GTX 1080
Min: 62 / Avg: 66.46 / Max: 70
GTX 1080 Ti
Min: 63 / Avg: 67.2 / Max: 70
RTX 2080
Min: 63 / Avg: 67.57 / Max: 72
RTX 2080 Ti
Min: 57 / Avg: 60.36 / Max: 64
14
28
42
56
70

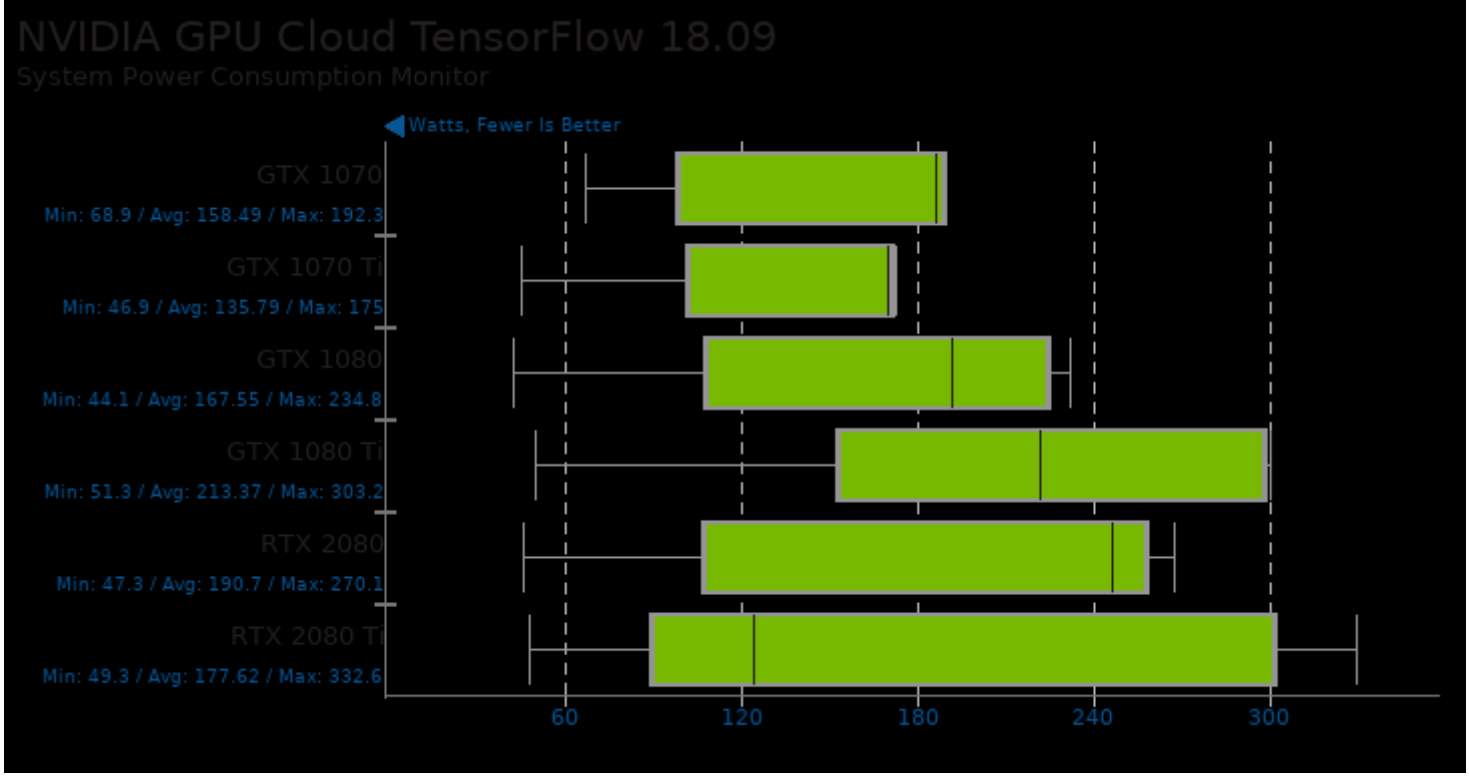
NVIDIA GPU Cloud TensorFlow 18.09
System Power Consumption Monitor
Watts, Fewer Is Better
GTX 1070
Min: 68.9 / Avg: 158.49 / Max: 192.3
GTX 1070 Ti
Min: 46.9 / Avg: 135.79 / Max: 175
GTX 1080
Min: 44.1 / Avg: 167.55 / Max: 234.8
GTX 1080 Ti
Min: 51.3 / Avg: 213.37 / Max: 303.2
RTX 2080
Min: 47.3 / Avg: 190.7 / Max: 270.1
RTX 2080 Ti
Min: 49.3 / Avg: 177.62 / Max: 332.6
60
120
180
240
300

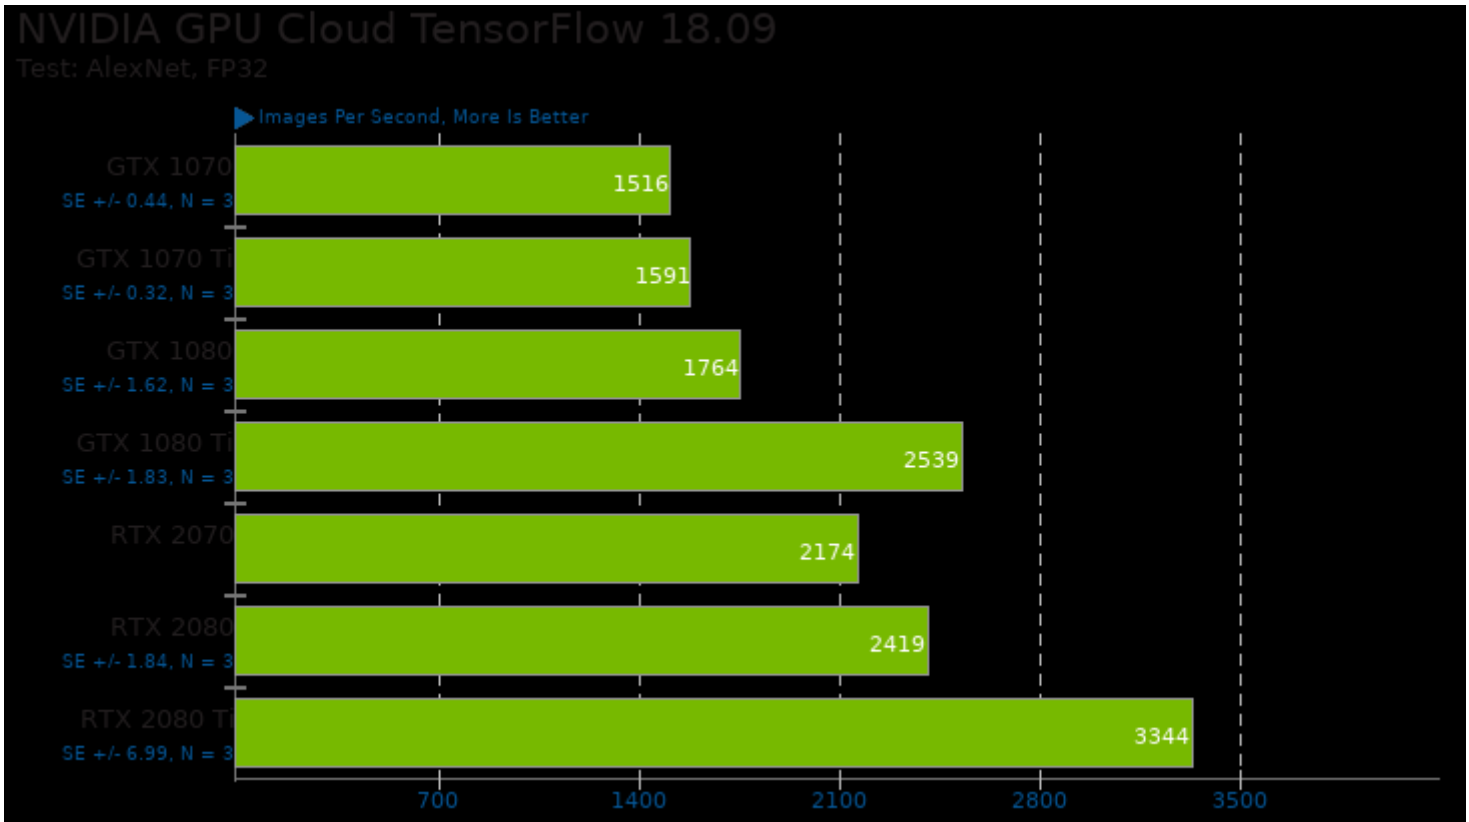
NVIDIA GPU Cloud TensorFlow 18.09
Test: AlexNet, FP32
Images Per Second, More Is Better
GTX 1070
SE +/- 0.44, N = 3
1516
GTX 1070 Ti
SE +/- 0.32, N = 3
1591
GTX 1080
SE +/- 1.62, N = 3
1764
GTX 1080 Ti
SE +/- 1.83, N = 3
2539
RTX 2070
2174
RTX 2080
SE +/- 1.84, N = 3
2419
RTX 2080 Ti
SE +/- 6.99, N = 3
3344
700
1400
2100
2800
3500

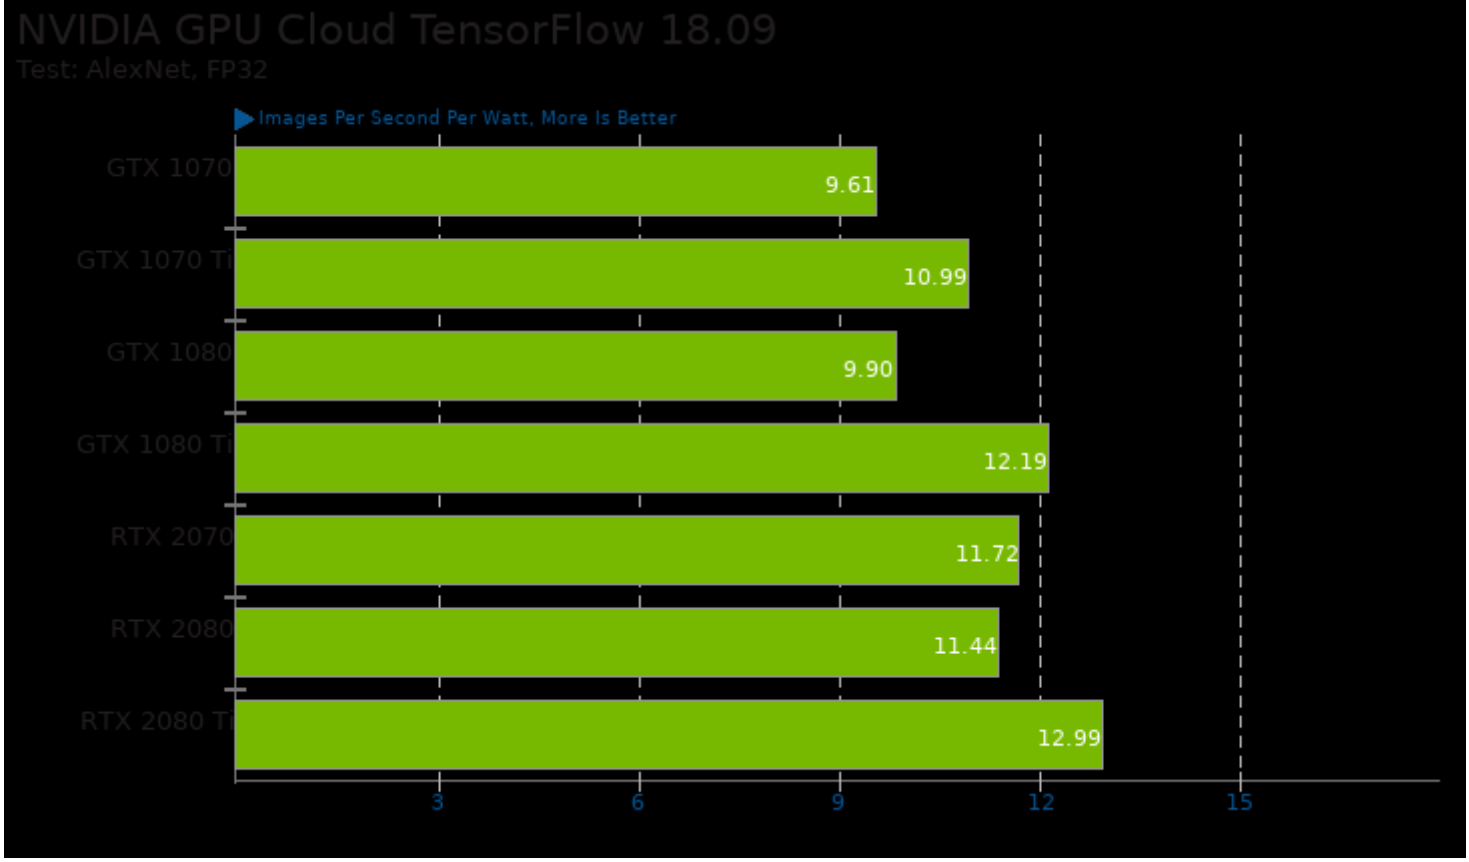
NVIDIA GPU Cloud TensorFlow 18.09
Test: AlexNet, FP32
Images Per Second Per Watt, More Is Better
GTX 1070
9.61
GTX 1070 Ti
10.99
GTX 1080
9.90
GTX 1080 Ti
12.19
RTX 2070
11.72
RTX 2080
11.44
RTX 2080 Ti
12.99
3
6
9
12
15

NVIDIA AMD Linux GPU Compute December 2018

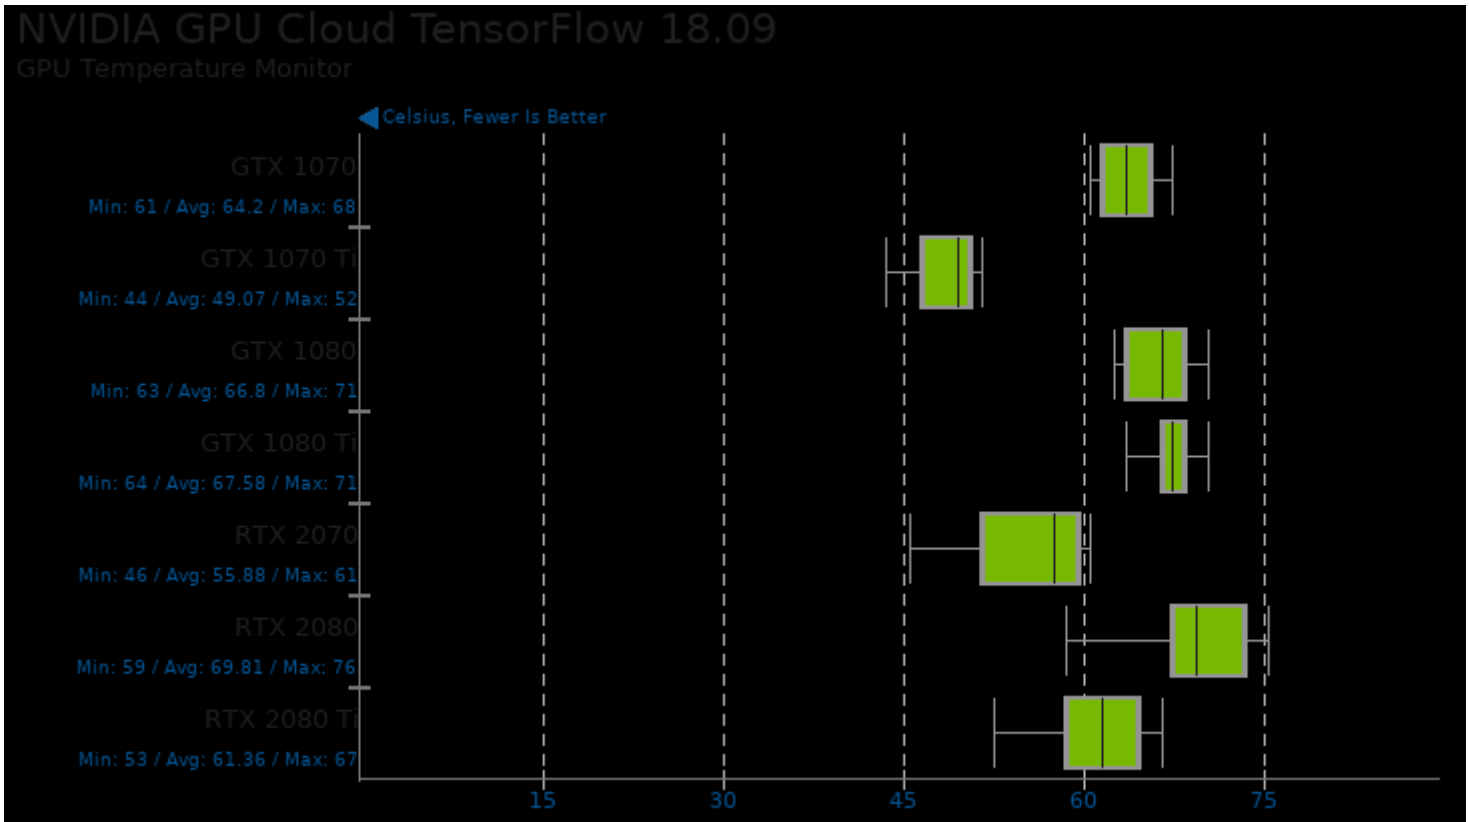

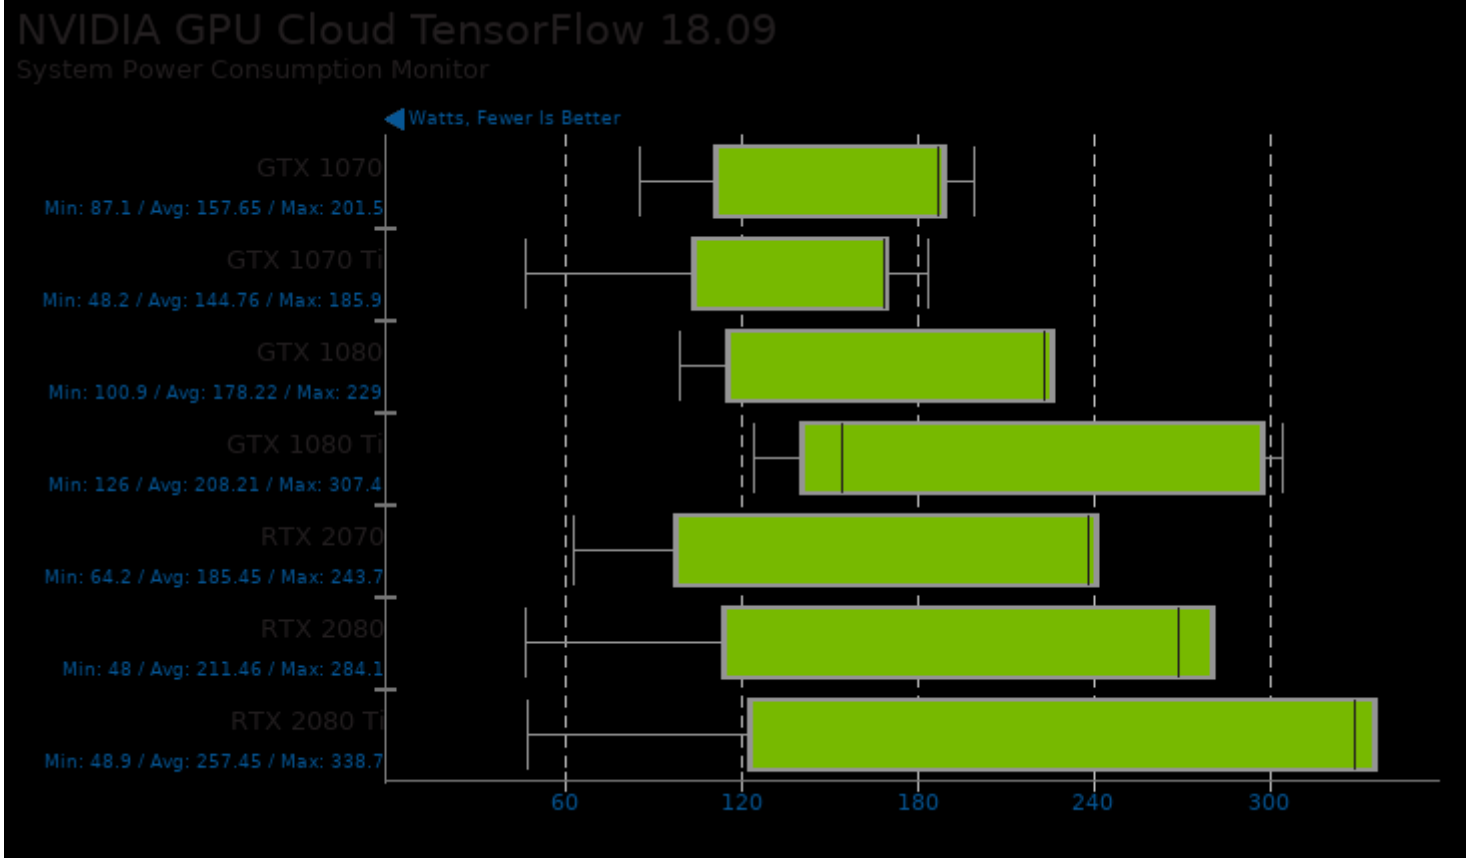

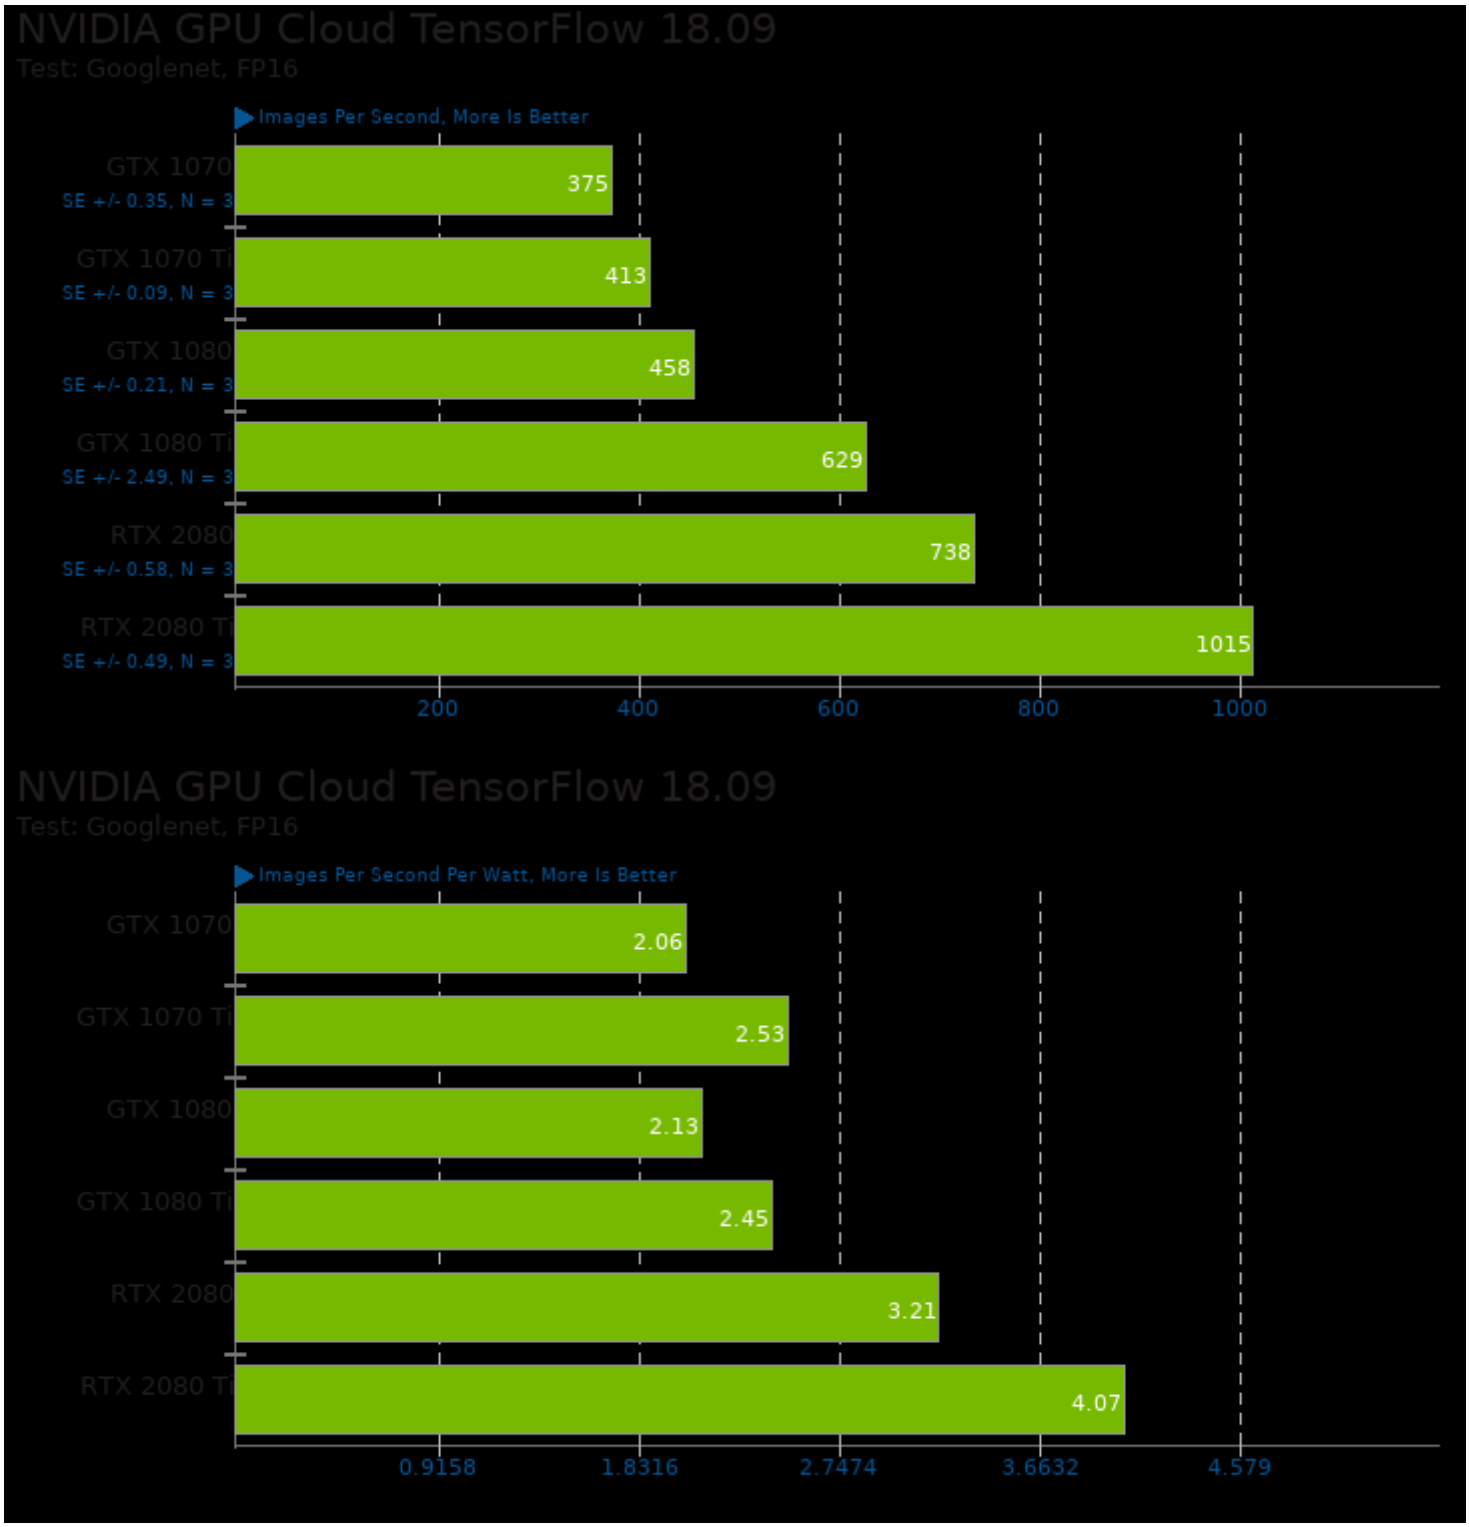
NVIDIA GPU Cloud TensorFlow 18.09
Test: Googlenet, FP16
Images Per Second, More Is Better
GTX 1070
SE +/- 0.35, N = 3
375
GTX 1070 Ti
SE +/- 0.09, N = 3
413
GTX 1080
SE +/- 0.21, N = 3
458
GTX 1080 Ti
SE +/- 2.49, N = 3
629
RTX 2080
SE +/- 0.58, N = 3
738
RTX 2080 Ti
SE +/- 0.49, N = 3
1015
200
400
600
800
1000
NVIDIA GPU Cloud TensorFlow 18.09
Test: Googlenet, FP16
Images Per Second Per Watt, More Is Better
GTX 1070
2.06
GTX 1070 Ti
2.53
GTX 1080
2.13
GTX 1080 Ti
2.45
RTX 2080
3.21
RTX 2080 Ti
4.07
0.9158
1.8316
2.7474
3.6632
4.579

NVIDIA AMD Linux GPU Compute December 2018

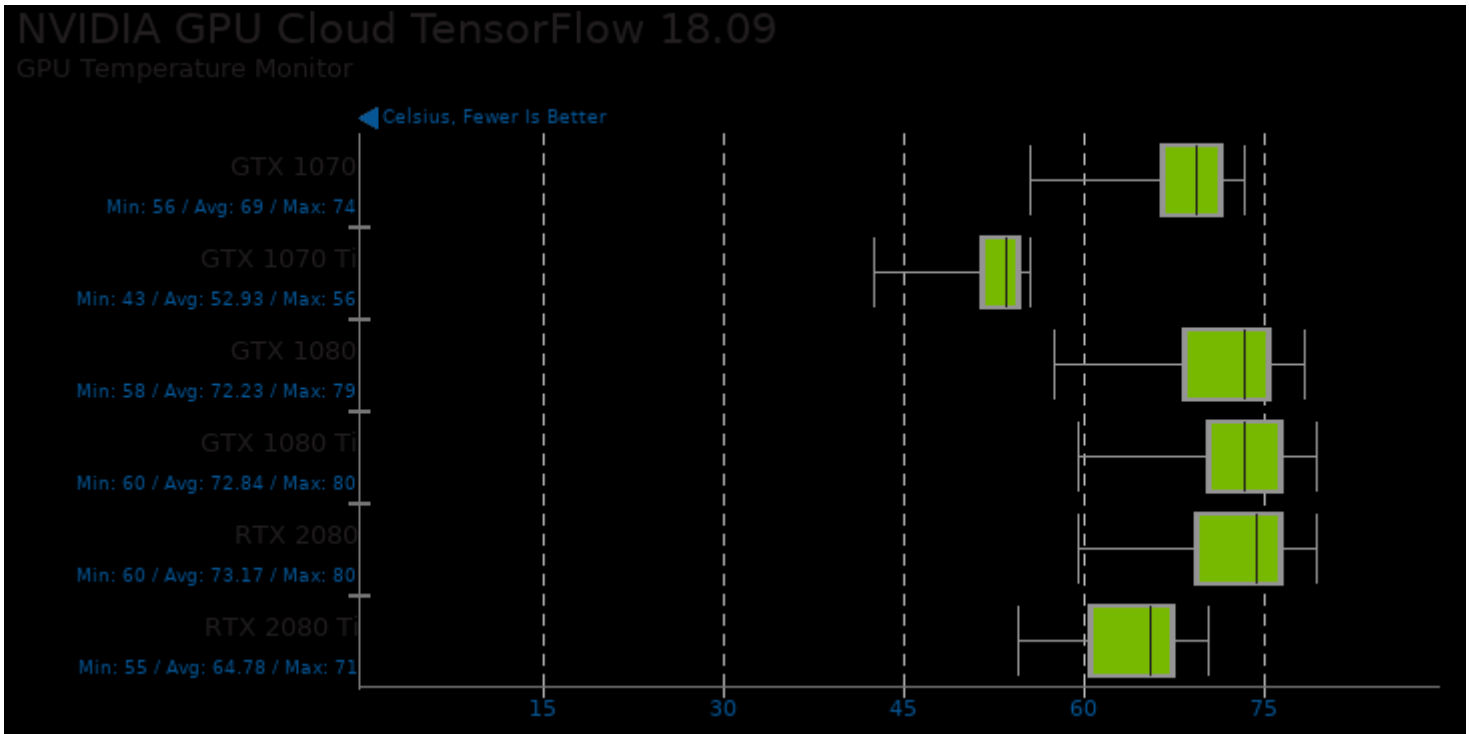

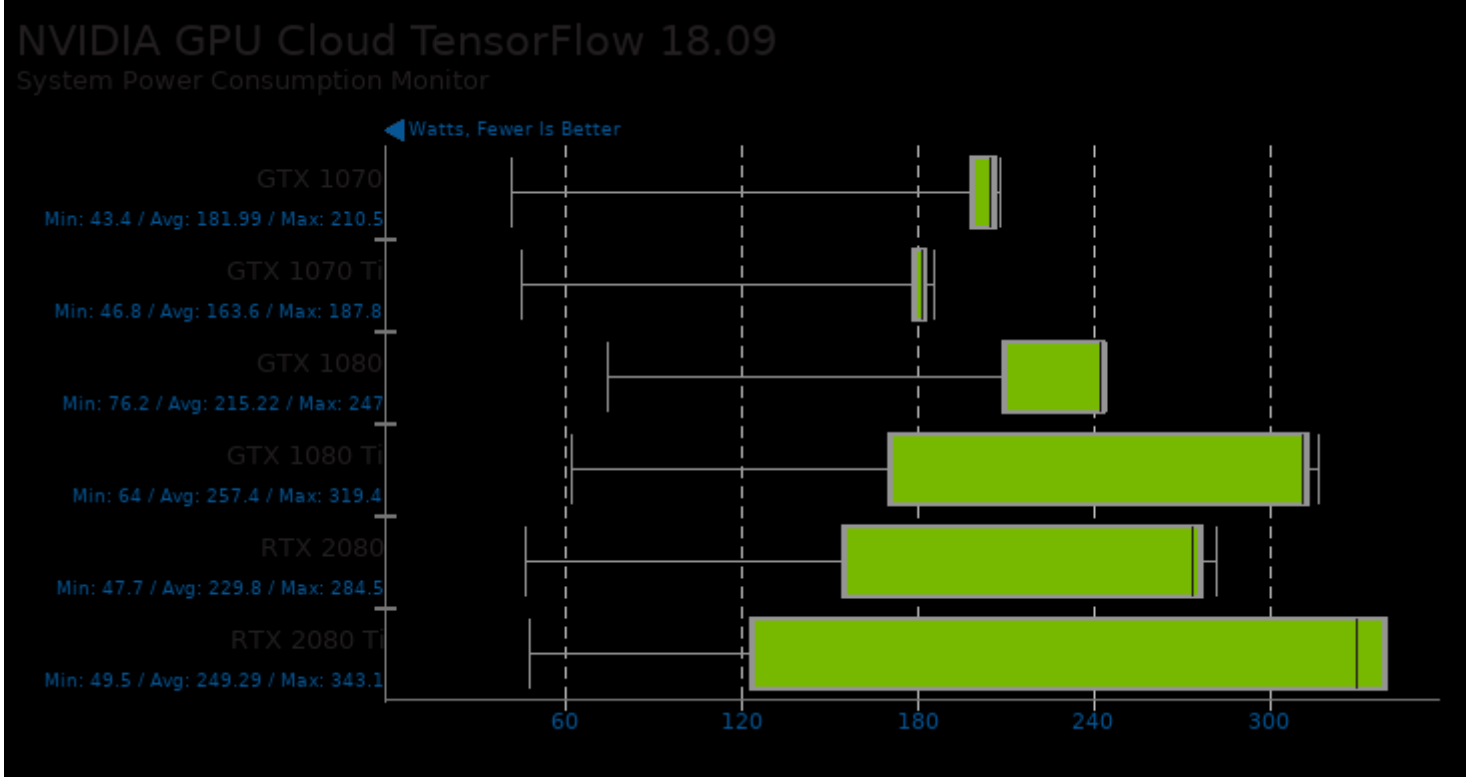

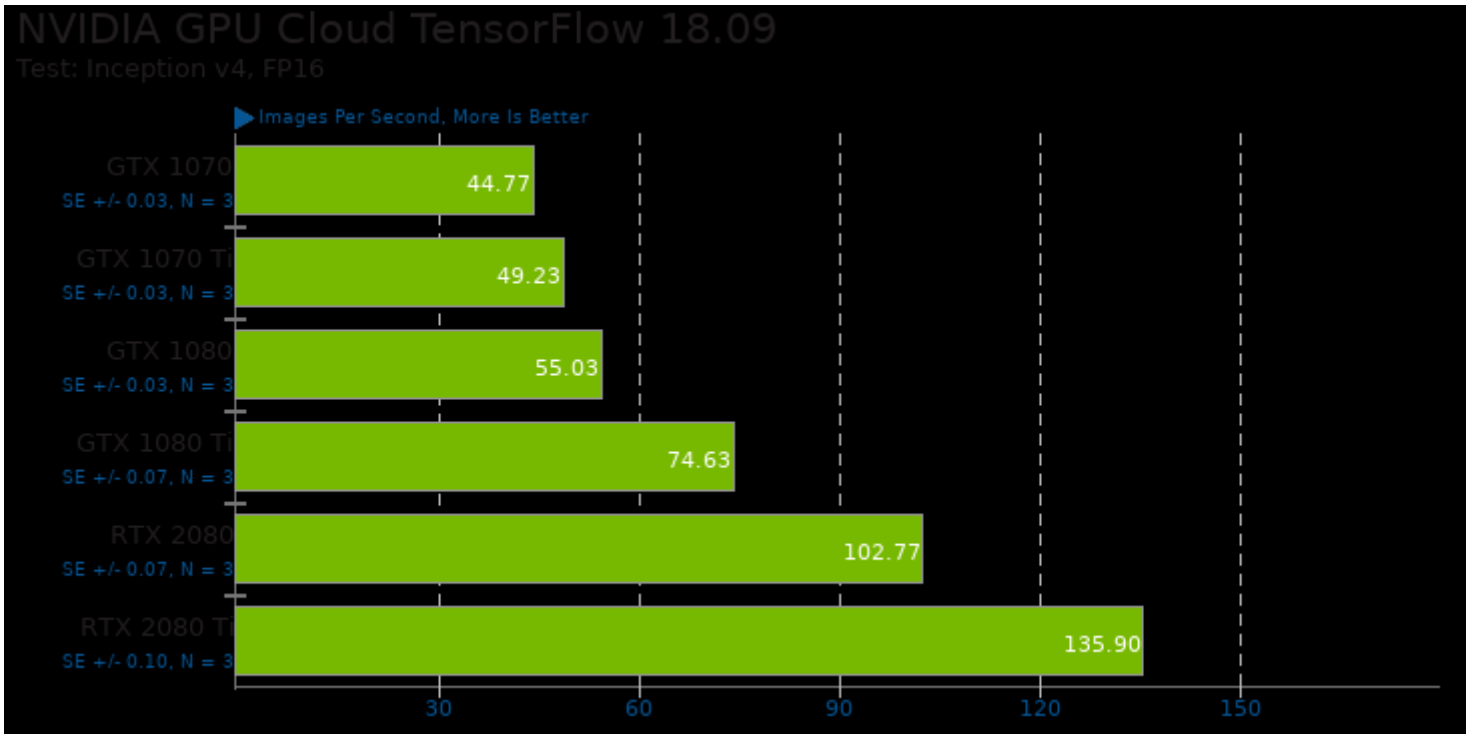
NVIDIA GPU Cloud TensorFlow 18.09
Test: Inception v4, FP16
Images Per Second, More Is Better
GTX 1070
SE +/- 0.03, N = 3
44.77
GTX 1070 Ti
SE +/- 0.03, N = 3
49.23
GTX 1080
SE +/- 0.03, N = 3
55.03
GTX 1080 Ti
SE +/- 0.07, N = 3
74.63
RTX 2080
SE +/- 0.07, N = 3
102.77
RTX 2080 Ti
SE +/- 0.10, N = 3
135.90
30
60
90
120
150

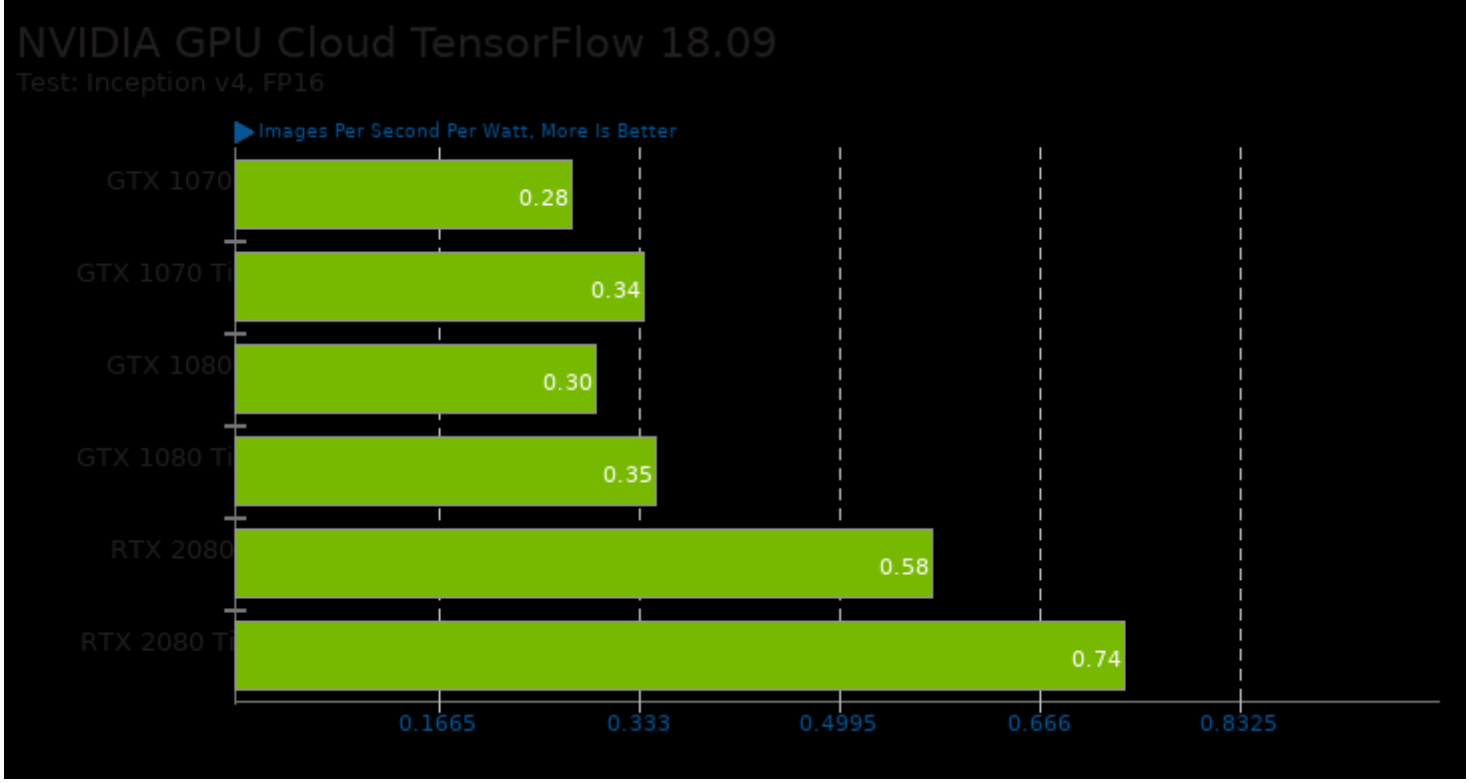
NVIDIA GPU Cloud TensorFlow 18.09
Test: Inception v4, FP16
Images Per Second Per Watt, More Is Better
GTX 1070
0.28
GTX 1070 Ti
0.34
GTX 1080
0.30
GTX 1080 Ti
0.35
RTX 2080
0.58
RTX 2080 Ti
0.74
0.1665
0.333
0.4995
0.666
0.8325

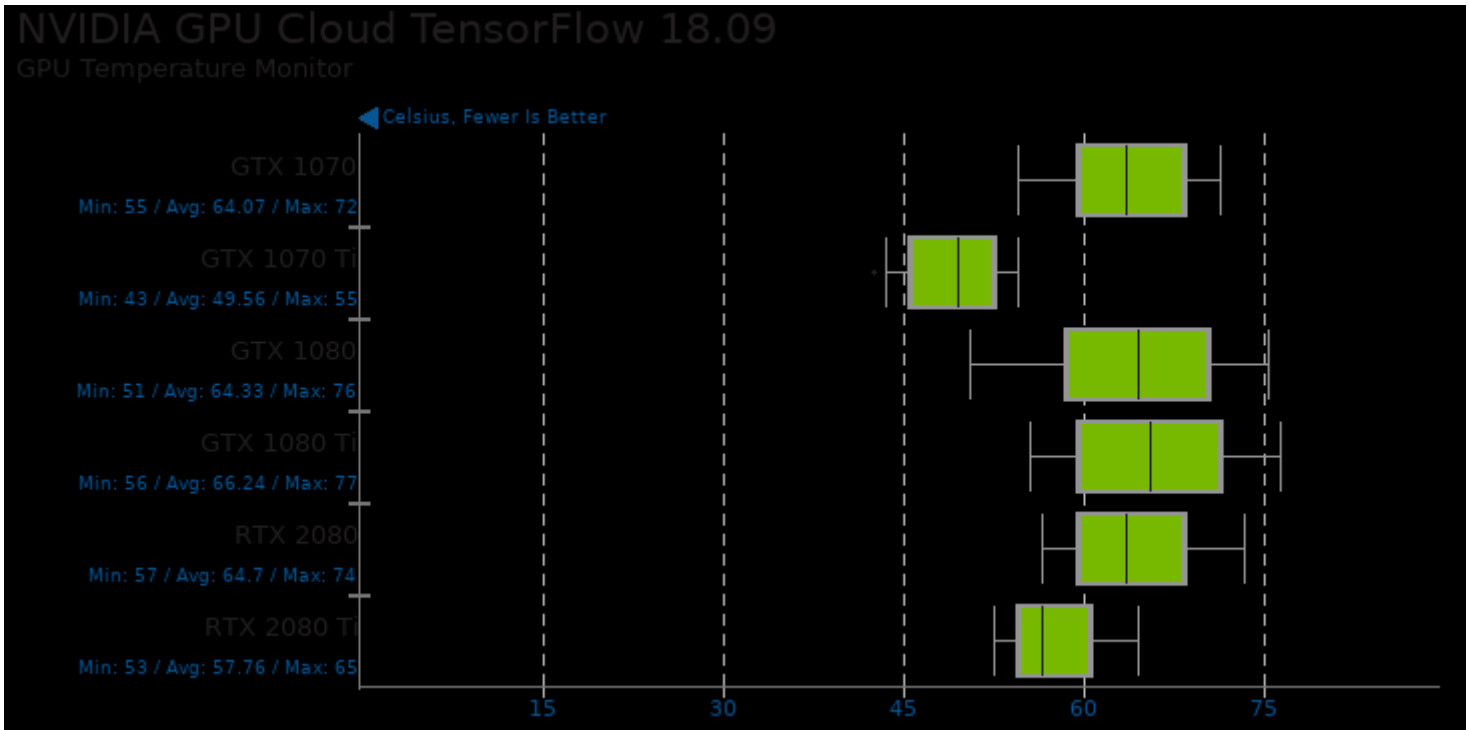
NVIDIA GPU Cloud TensorFlow 18.09
GPU Temperature Monitor
Celsius, Fewer Is Better
GTX 1070
Min: 55 / Avg: 64.07 / Max: 72
GTX 1070 Ti
Min: 43 / Avg: 49.56 / Max: 55
GTX 1080
Min: 51 / Avg: 64.33 / Max: 76
GTX 1080 Ti
Min: 56 / Avg: 66.24 / Max: 77
RTX 2080
Min: 57 / Avg: 64.7 / Max: 74
RTX 2080 Ti
Min: 53 / Avg: 57.76 / Max: 65
15
30
45
60
75

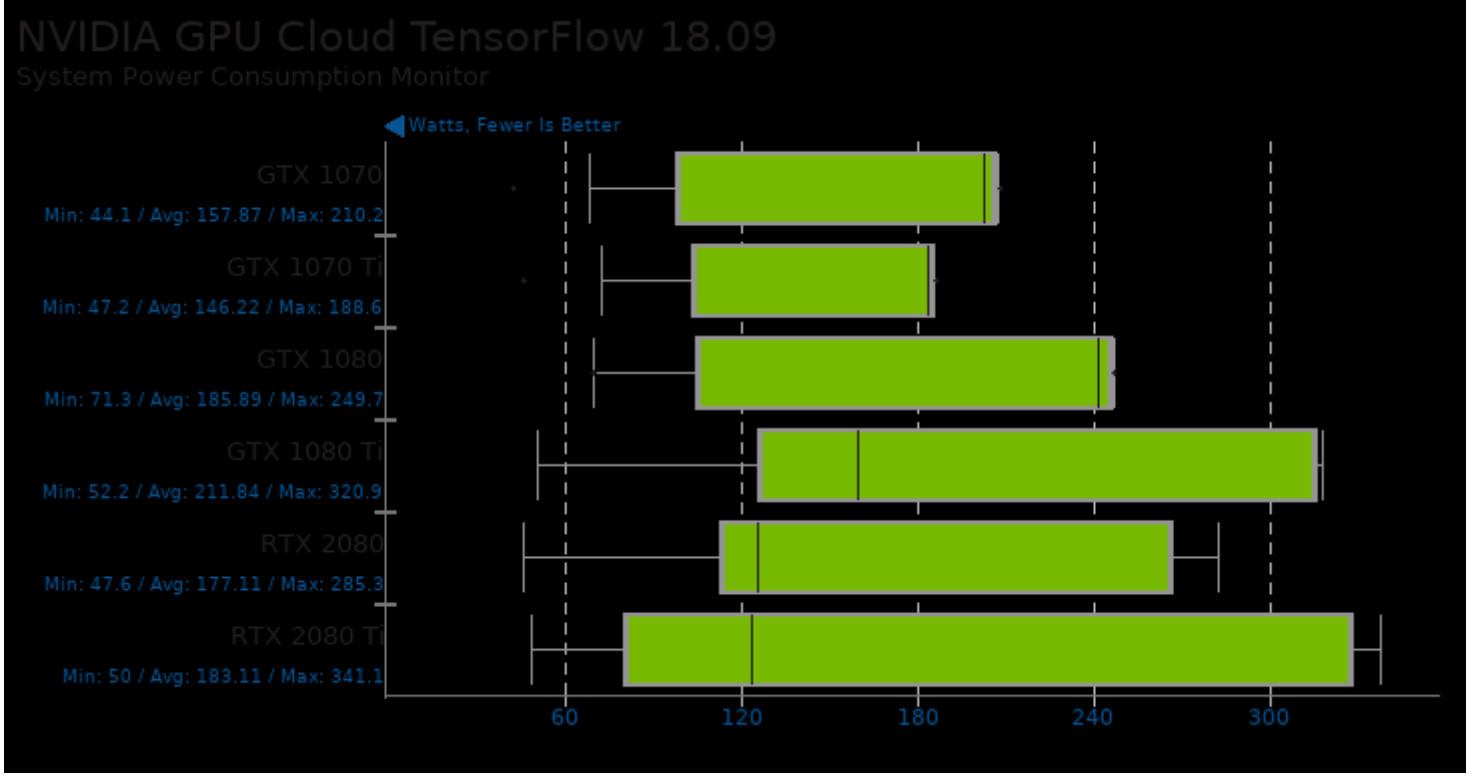
NVIDIA GPU Cloud TensorFlow 18.09
System Power Consumption Monitor
Watts, Fewer Is Better
GTX 1070
Min: 44.1 / Avg: 157.87 / Max: 210.2
GTX 1070 Ti
Min: 47.2 / Avg: 146.22 / Max: 188.6
GTX 1080
Min: 71.3 / Avg: 185.89 / Max: 249.7
GTX 1080 Ti
Min: 52.2 / Avg: 211.84 / Max: 320.9
RTX 2080
Min: 47.6 / Avg: 177.11 / Max: 285.3
RTX 2080 Ti
Min: 50 / Avg: 183.11 / Max: 341.1
60
120
180
240
300

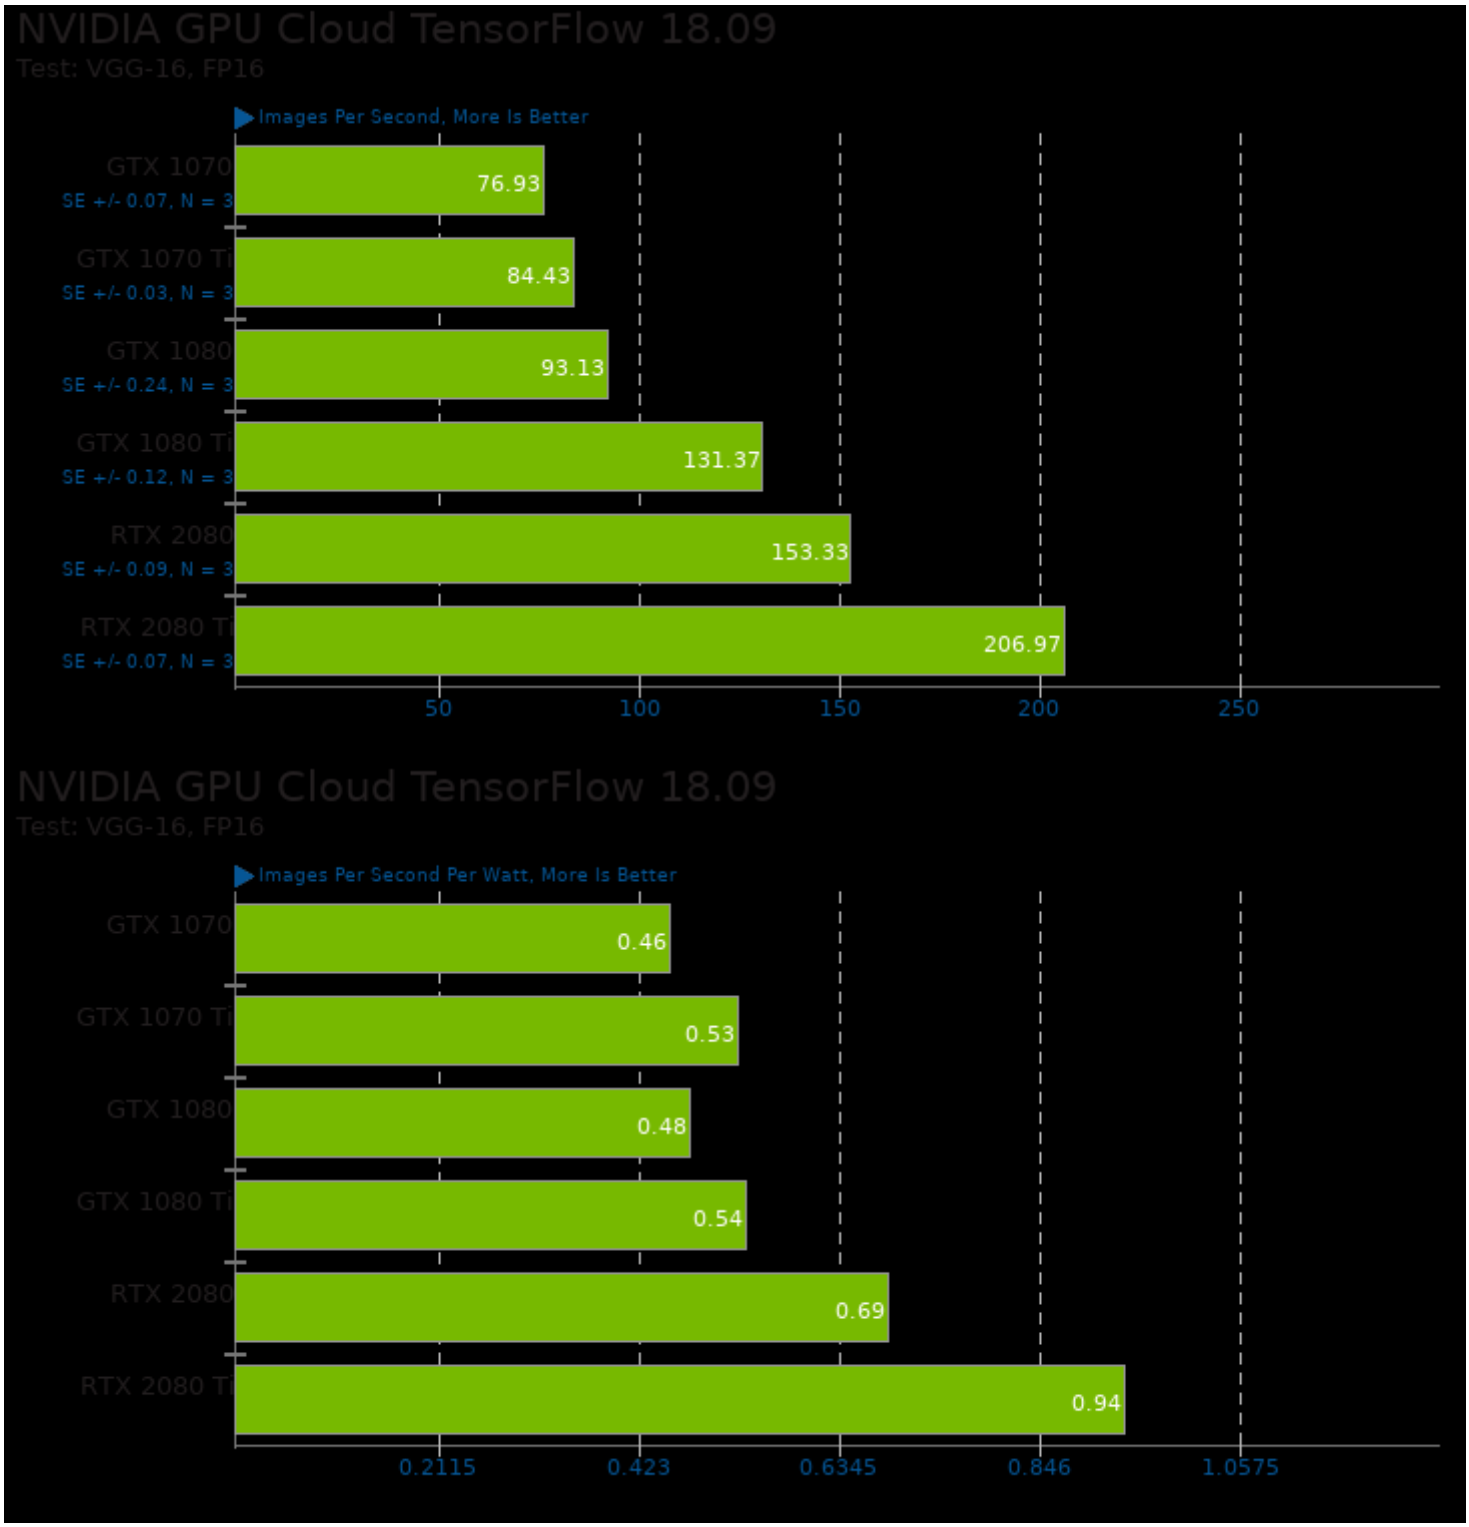
NVIDIA GPU Cloud TensorFlow 18.09
Test: VGG-16, FP16
Images Per Second, More Is Better
GTX 1070
SE +/- 0.07, N = 3
76.93
GTX 1070 Ti
SE +/- 0.03, N = 3
84.43
GTX 1080
SE +/- 0.24, N = 3
93.13
GTX 1080 Ti
SE +/- 0.12, N = 3
131.37
RTX 2080
SE +/- 0.09, N = 3
153.33
RTX 2080 Ti
SE +/- 0.07, N = 3
206.97
50
100
150
200
250
NVIDIA GPU Cloud TensorFlow 18.09
Test: VGG-16, FP16
Images Per Second Per Watt, More Is Better
GTX 1070
0.46
GTX 1070 Ti
0.53
GTX 1080
0.48
GTX 1080 Ti
0.54
RTX 2080
0.69
RTX 2080 Ti
0.94
0.2115
0.423
0.6345
0.846
1.0575

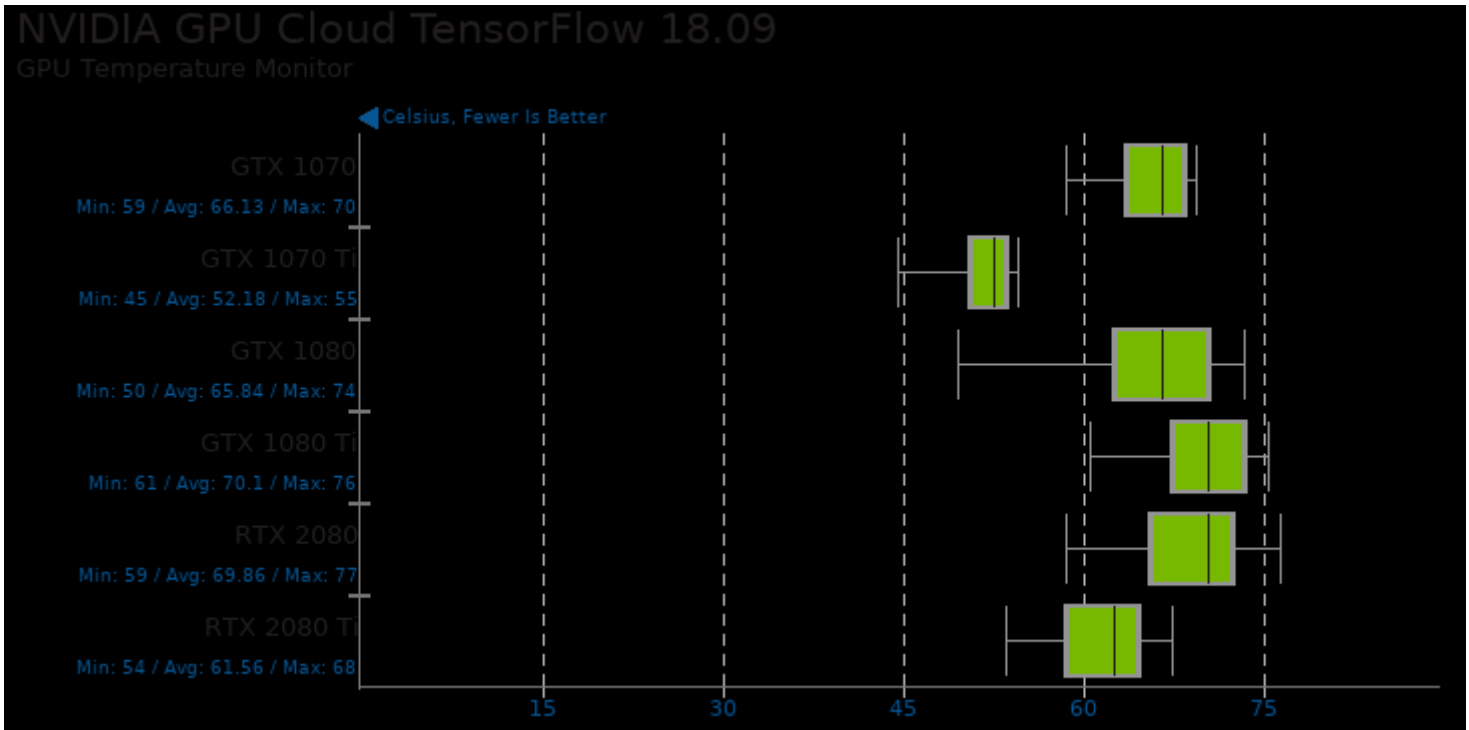
NVIDIA GPU Cloud TensorFlow 18.09
GPU Temperature Monitor
Celsius, Fewer Is Better
GTX 1070
Min: 59 / Avg: 66.13 / Max: 70
GTX 1070 Ti
Min: 45 / Avg: 52.18 / Max: 55
GTX 1080
Min: 50 / Avg: 65.84 / Max: 74
GTX 1080 Ti
Min: 61 / Avg: 70.1 / Max: 76
RTX 2080
Min: 59 / Avg: 69.86 / Max: 77
RTX 2080 Ti
Min: 54 / Avg: 61.56 / Max: 68
15
30
45
60
75

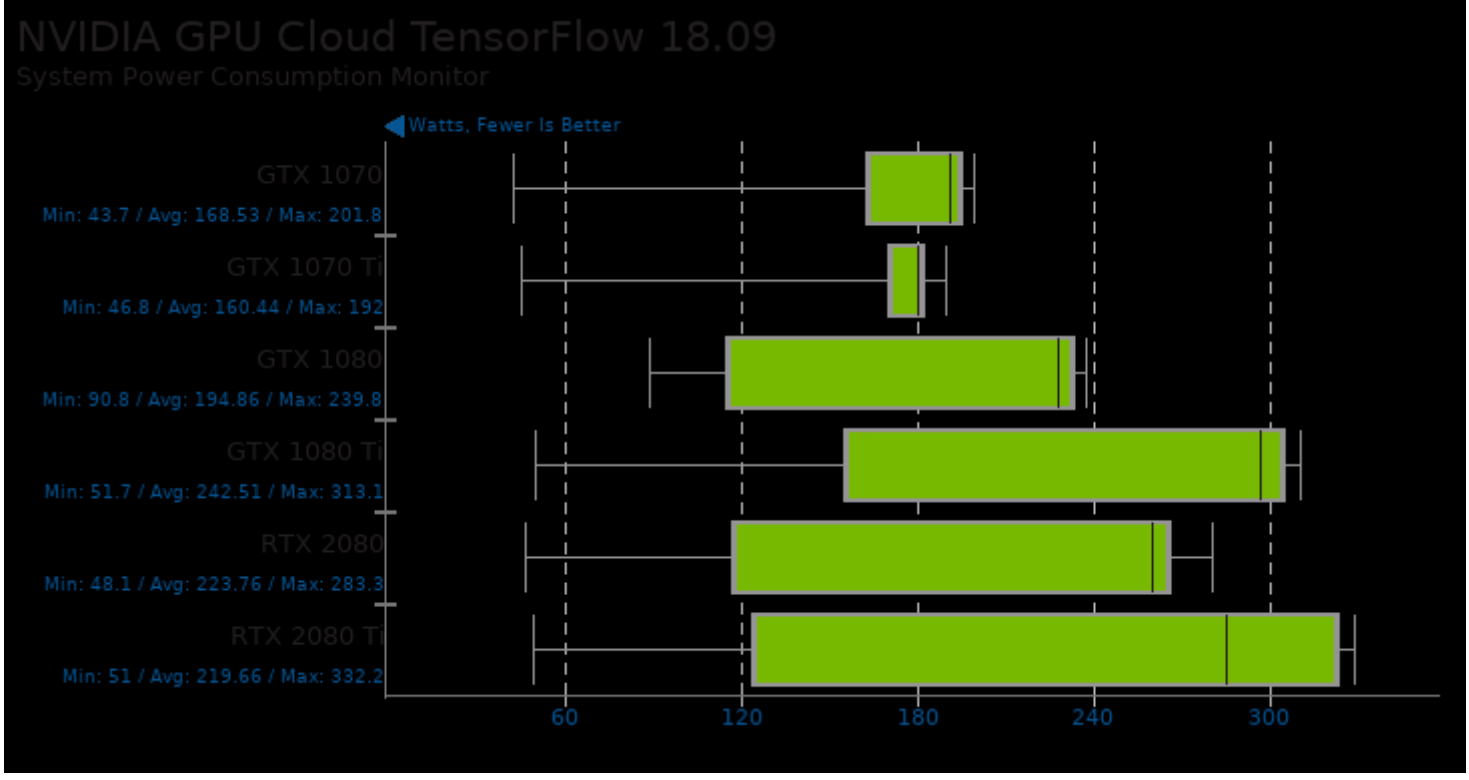
NVIDIA GPU Cloud TensorFlow 18.09
System Power Consumption Monitor
Watts, Fewer Is Better
GTX 1070
Min: 43.7 / Avg: 168.53 / Max: 201.8
GTX 1070 Ti
Min: 46.8 / Avg: 160.44 / Max: 192
GTX 1080
Min: 90.8 / Avg: 194.86 / Max: 239.8
GTX 1080 Ti
Min: 51.7 / Avg: 242.51 / Max: 313.1
RTX 2080
Min: 48.1 / Avg: 223.76 / Max: 283.3
RTX 2080 Ti
Min: 51 / Avg: 219.66 / Max: 332.2
60
120
180
240
300

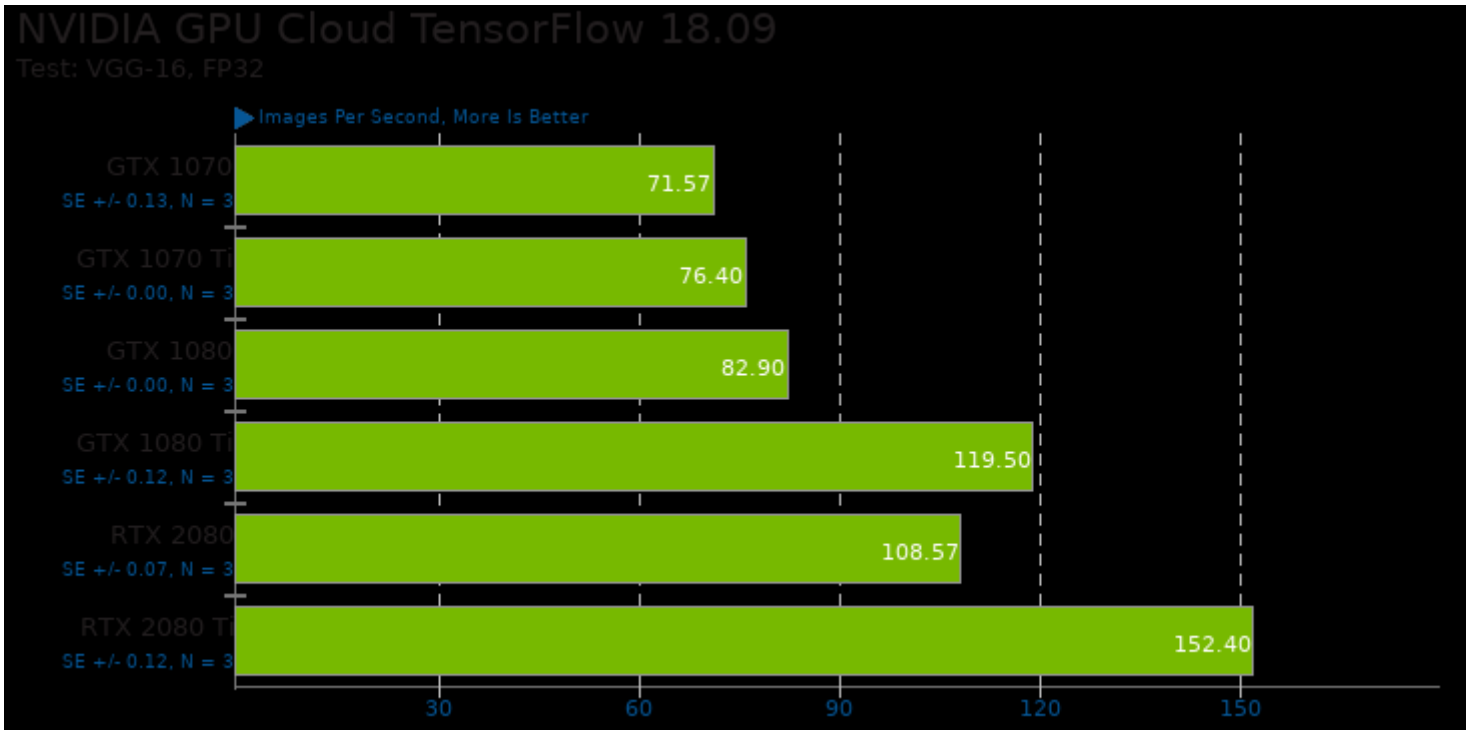
NVIDIA GPU Cloud TensorFlow 18.09
Test: VGG-16, FP32
Images Per Second, More Is Better
GTX 1070
SE +/- 0.13, N = 3
71.57
GTX 1070 Ti
SE +/- 0.00, N = 3
76.40
GTX 1080
SE +/- 0.00, N = 3
82.90
GTX 1080 Ti
SE +/- 0.12, N = 3
119.50
RTX 2080
SE +/- 0.07, N = 3
108.57
RTX 2080 Ti
SE +/- 0.12, N = 3
152.40
30
60
90
120
150

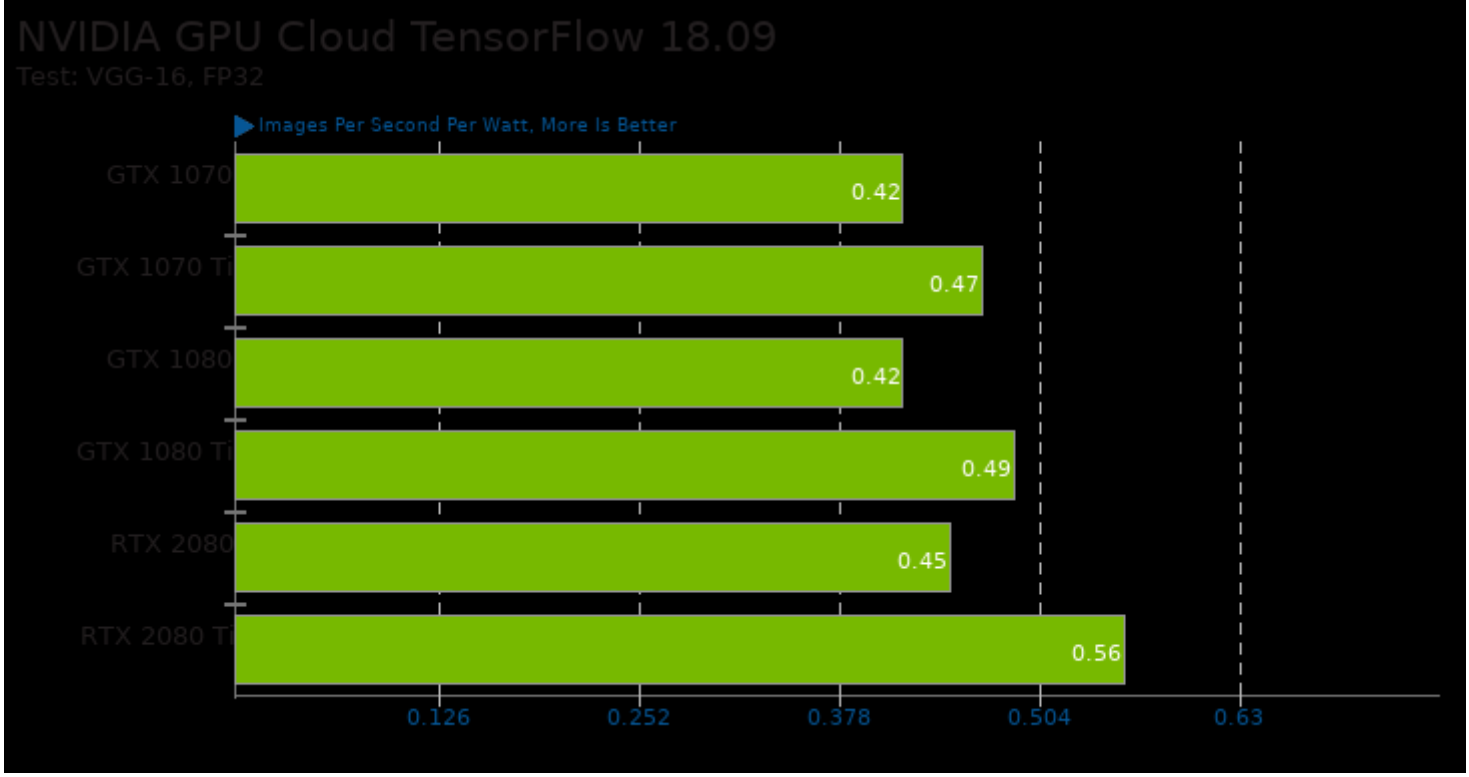
NVIDIA GPU Cloud TensorFlow 18.09
Test: VGG-16, FP32
Images Per Second Per Watt, More Is Better
GTX 1070
0.42
GTX 1070 Ti
0.47
GTX 1080
0.42
GTX 1080 Ti
0.49
RTX 2080
0.45
RTX 2080 Ti
0.56
0.126
0.252
0.378
0.504
0.63

NVIDIA AMD Linux GPU Compute December 2018

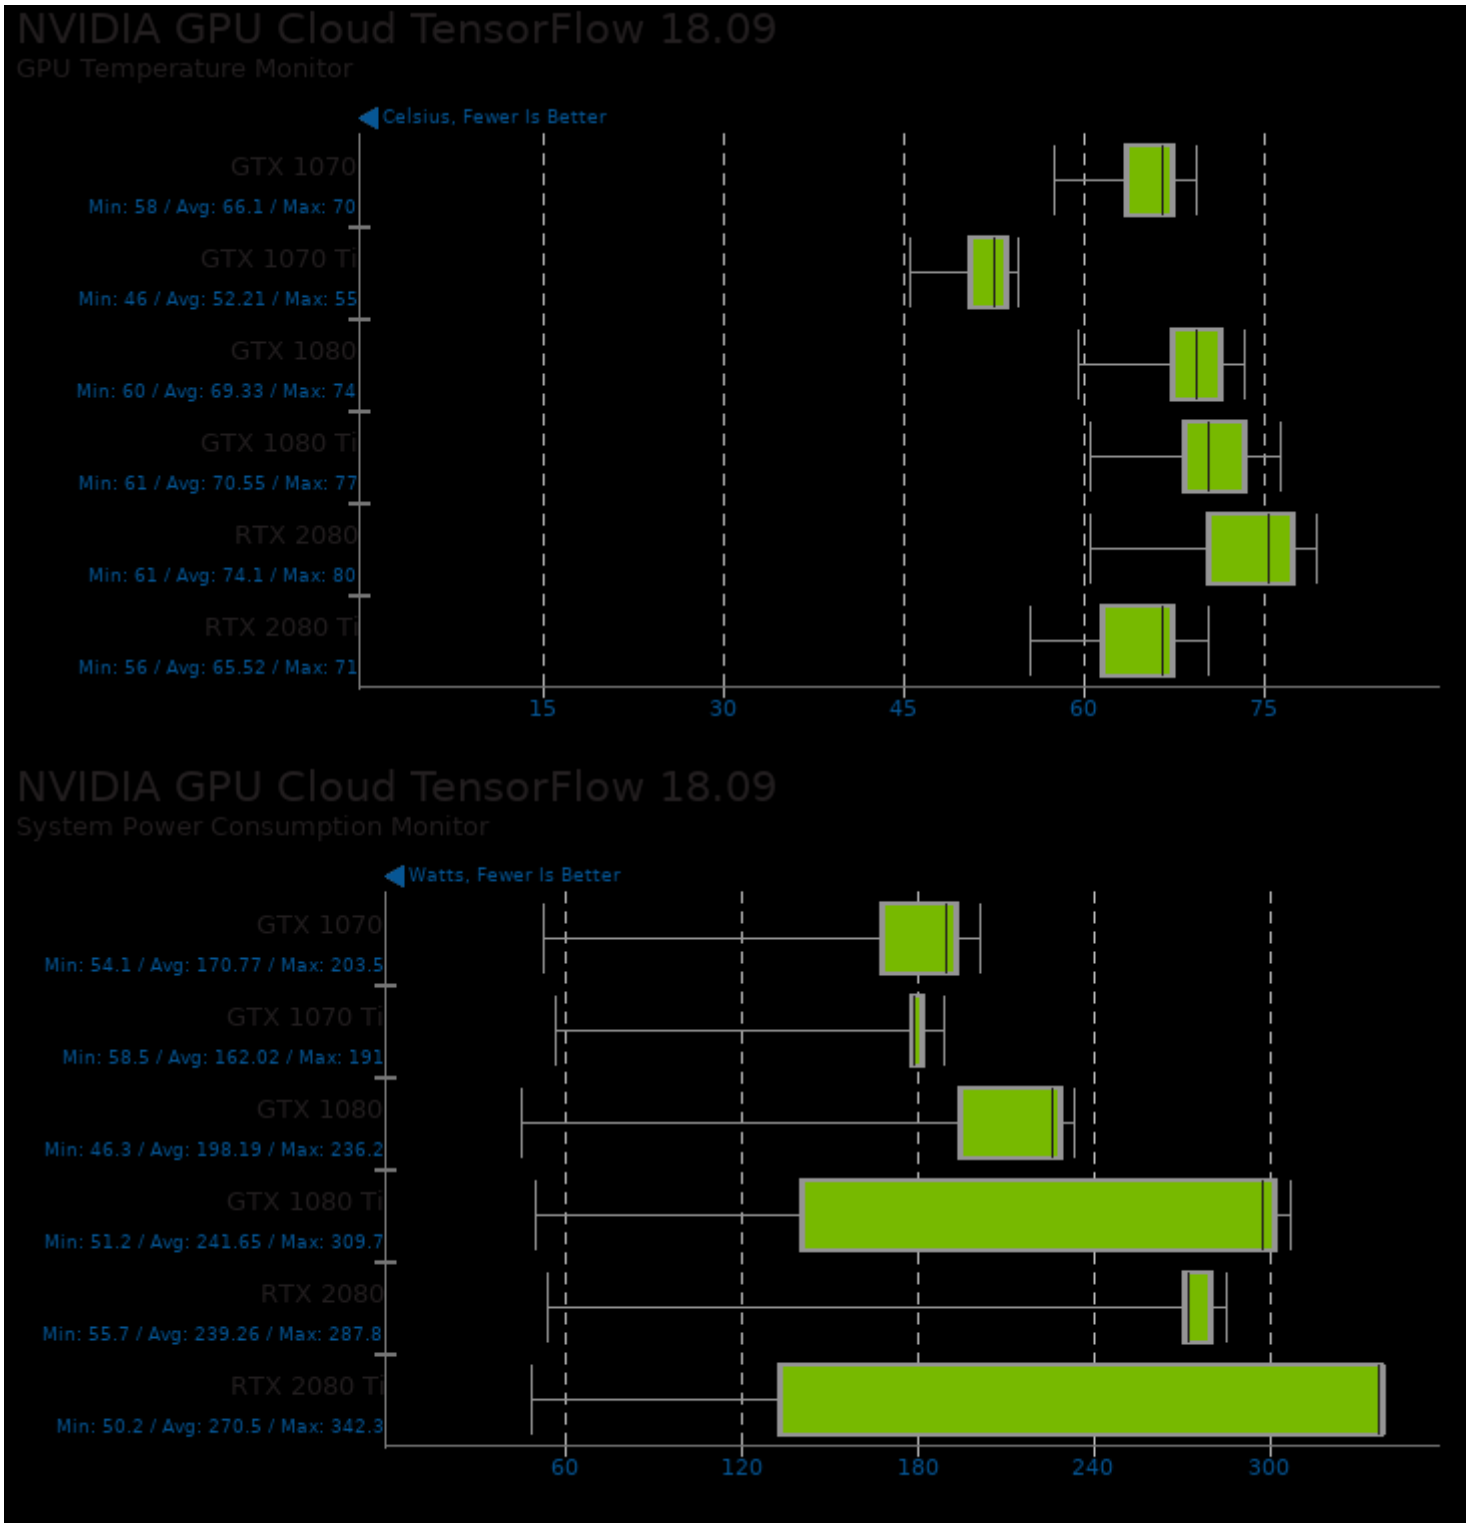

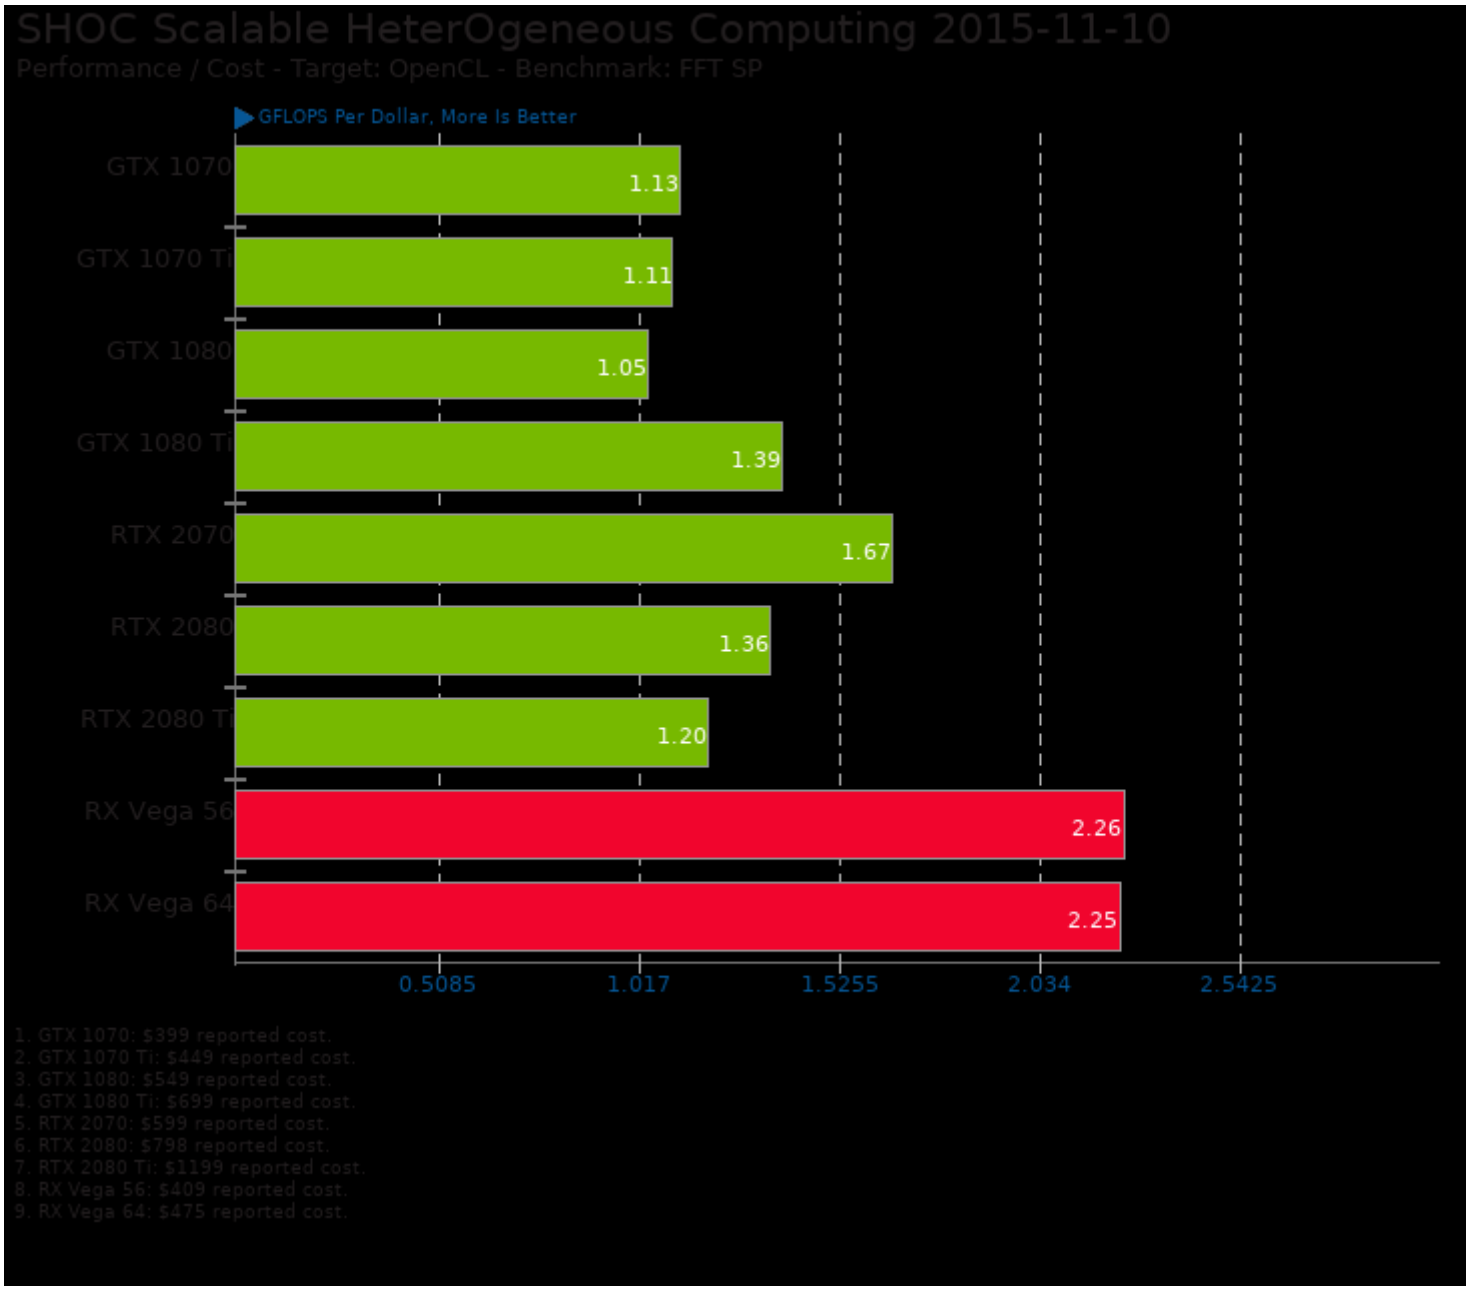
SHOC Scalable HeterOgeneous Computing 2015-11-10
Performance / Cost - Target: OpenCL - Benchmark: FFT SP
GFLOPS Per Dollar, More Is Better
GTX 1070
1.13
GTX 1070 Ti
1.11
GTX 1080
1.05
GTX 1080 Ti
1.39
RTX 2070
1.67
RTX 2080
1.36
RTX 2080 Ti
1.20
RX Vega 56
2.26
RX Vega 64
2.25
0.5085
1.017
1.5255
2.034
2.5425
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.

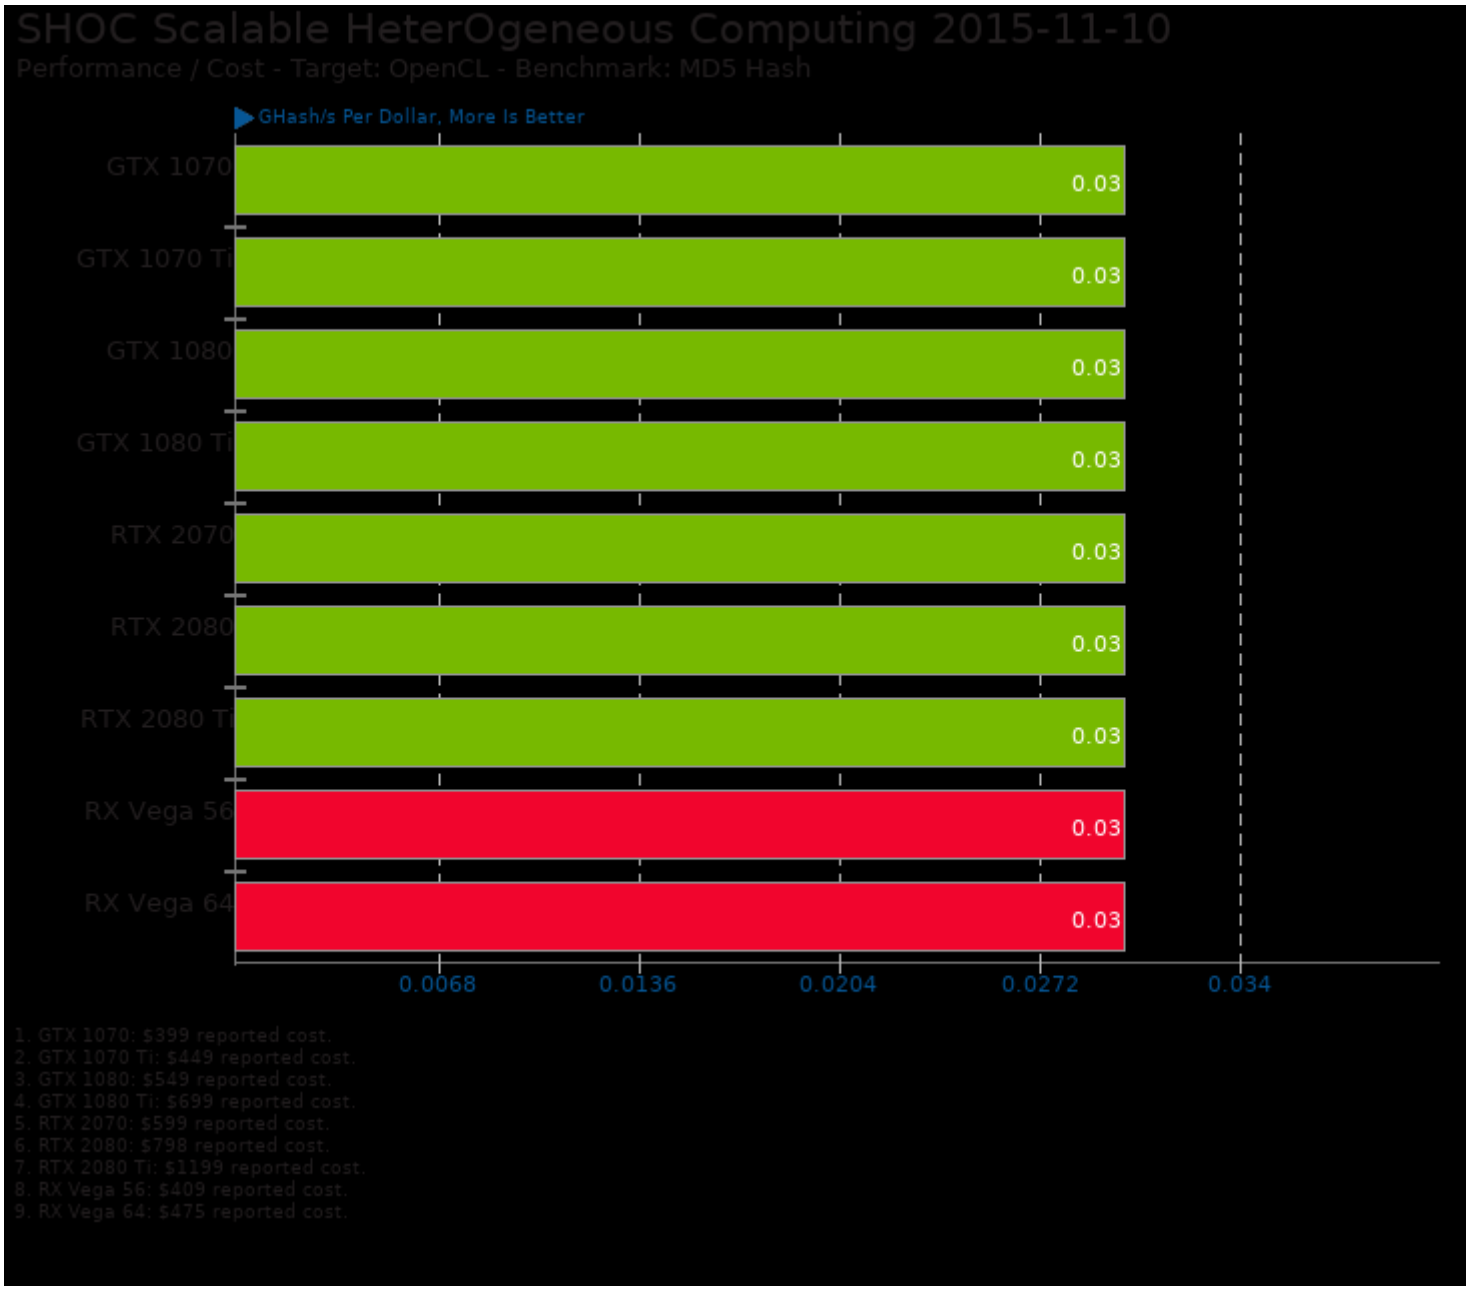
SHOC Scalable HeterOgeneous Computing 2015-11-10
Performance / Cost - Target: OpenCL - Benchmark: MD5 Hash
GHash/s Per Dollar, More Is Better
GTX 1070
GTX 1070 Ti
GTX 1080
GTX 1080 Ti
RTX 2070
RTX 2080
RTX 2080 Ti
RX Vega 56
RX Vega 64
0.03
0.0068
0.0136
0.0204
0.0272
0.034
1. GTX 1070: $399 reported 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.

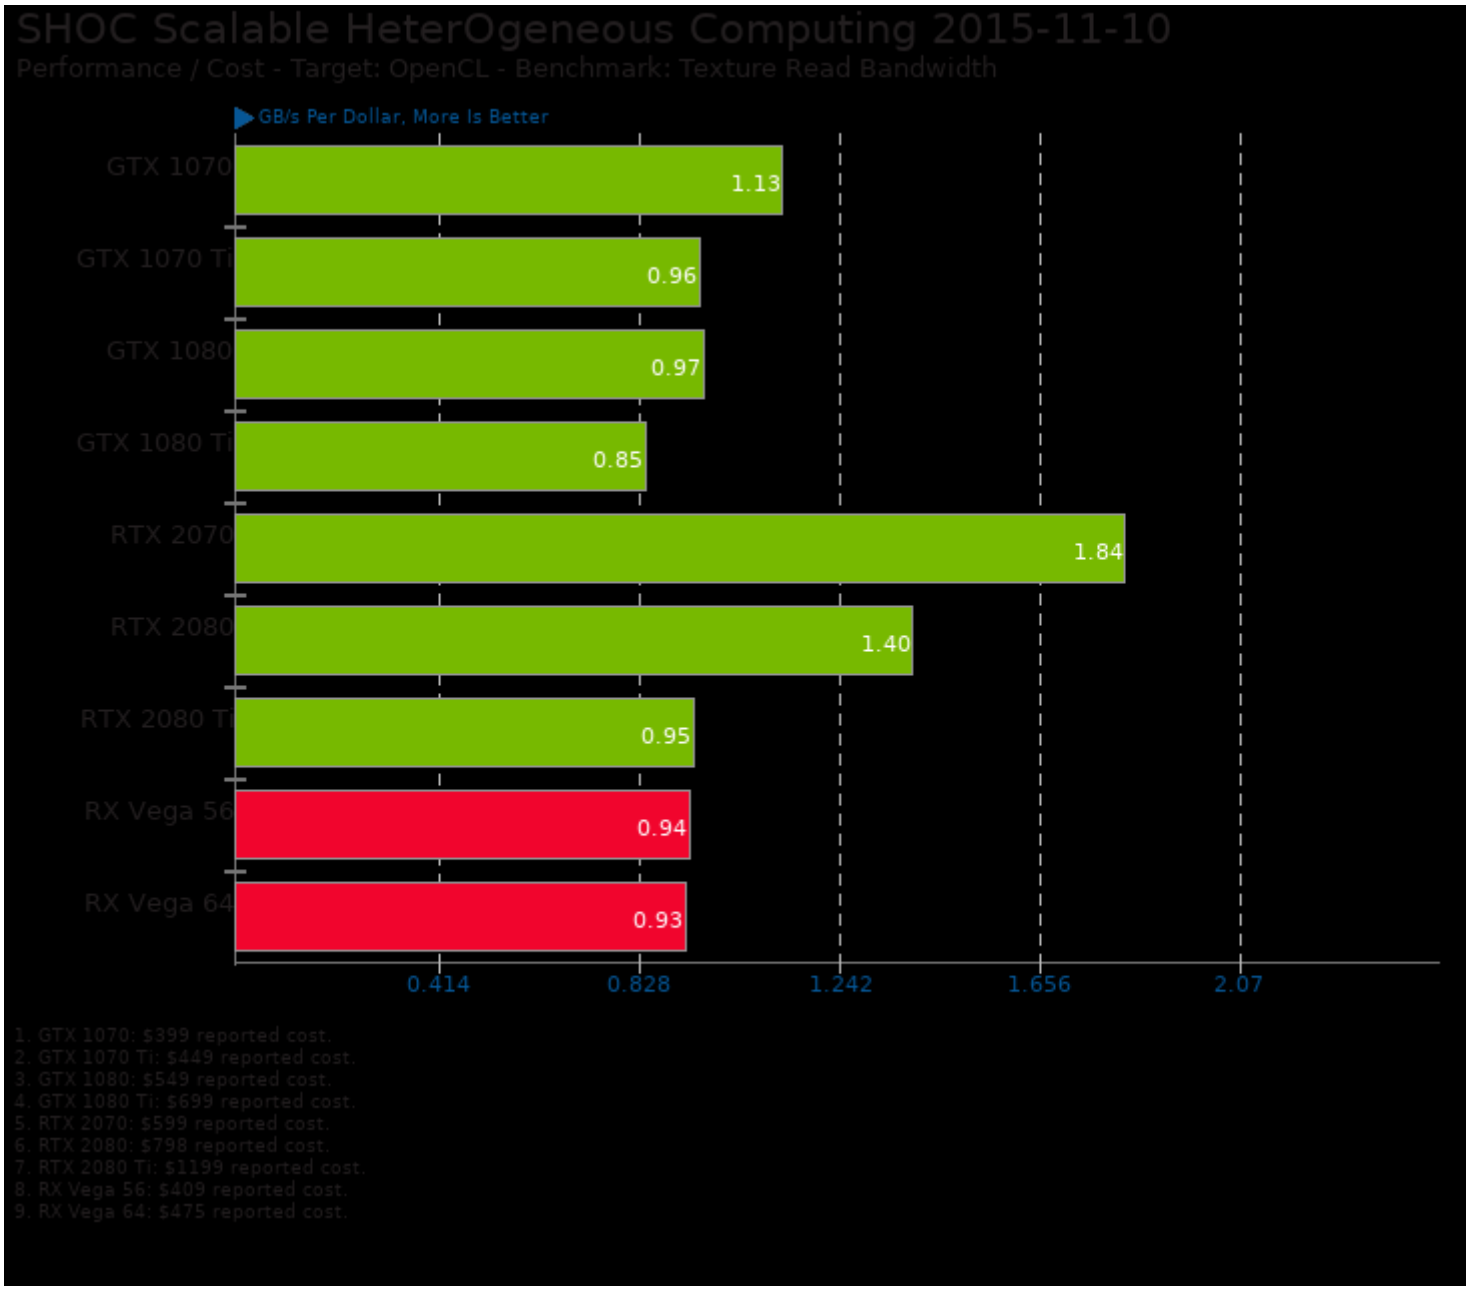
SHOC Scalable HeterOgeneous Computing 2015-11-10
Performance / Cost - Target: OpenCL - Benchmark: Texture Read Bandwidth
GB/s Per Dollar, More Is Better
GTX 1070
1.13
GTX 1070 Ti
0.96
GTX 1080
0.97
GTX 1080 Ti
0.85
RTX 2070
1.84
RTX 2080
1.40
RTX 2080 Ti
0.95
RX Vega 56
0.94
RX Vega 64
0.93
0.414
0.828
1.242
1.656
2.07
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.

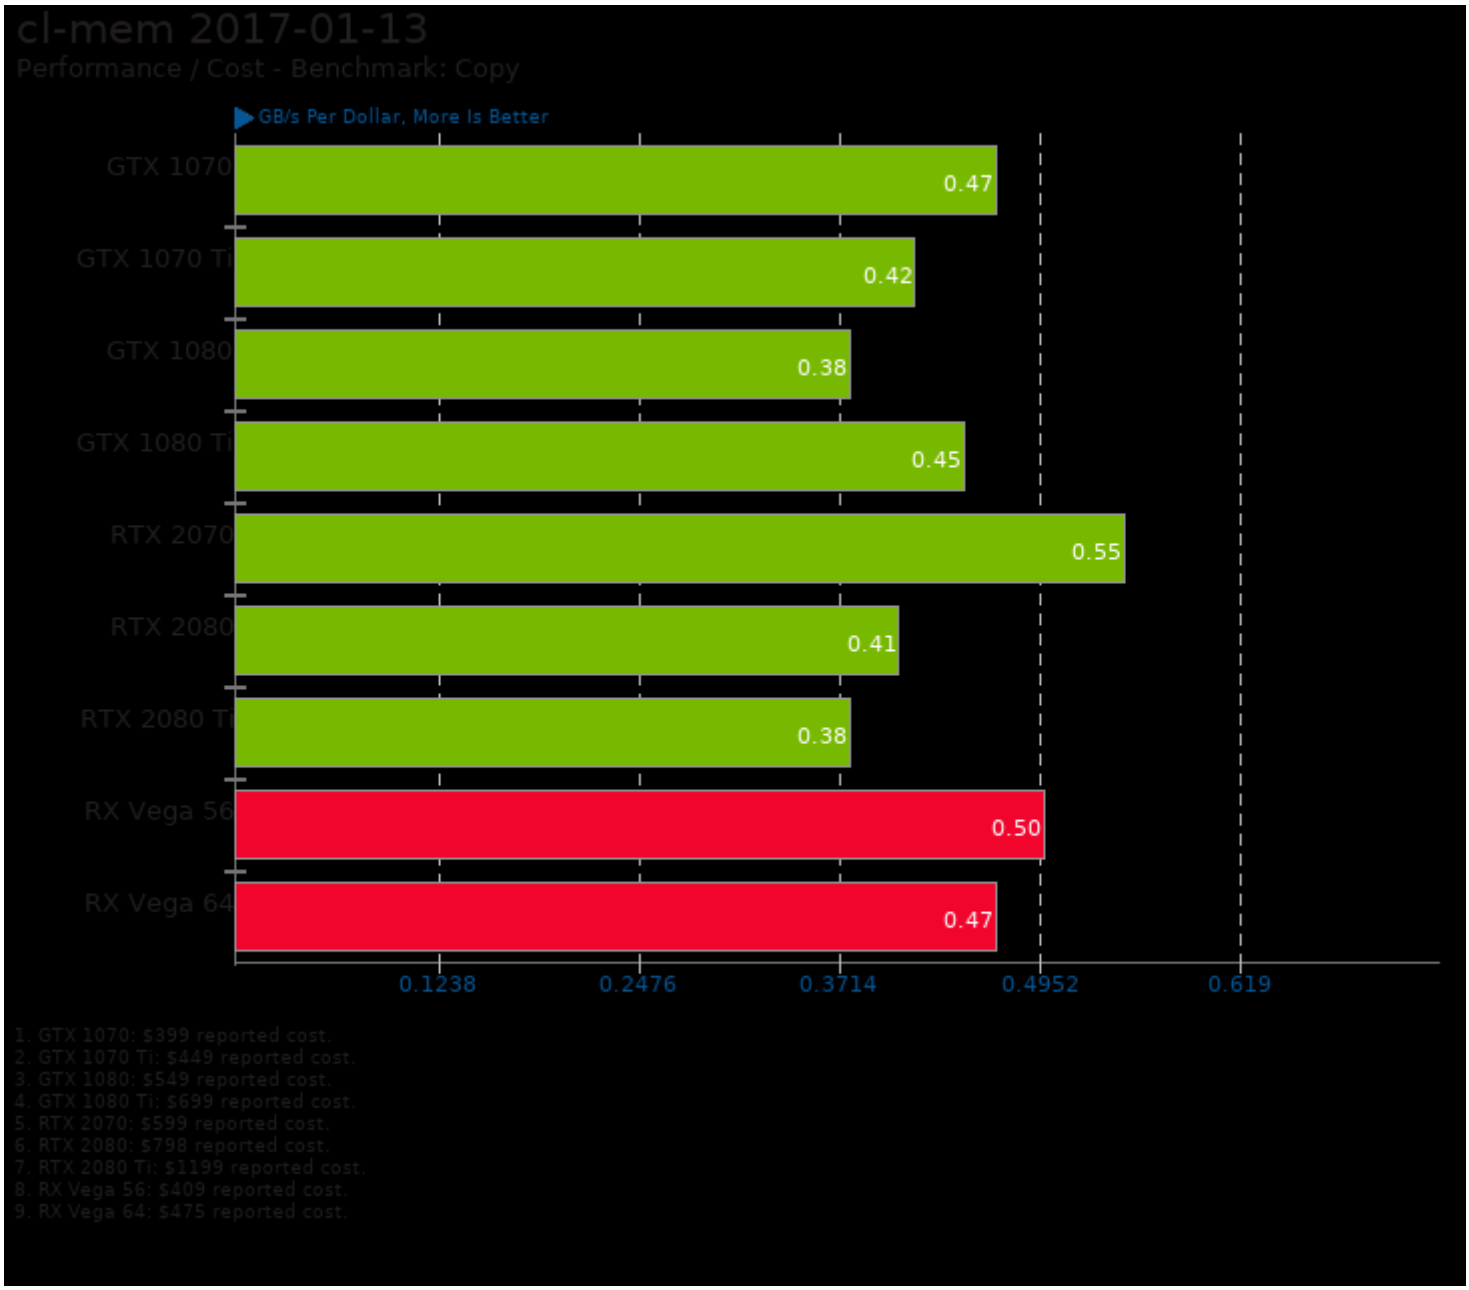
cl-mem 2017-01-13
Performance / Cost - Benchmark: Copy
GB/s Per Dollar, More Is Better
GTX 1070
0.47
GTX 1070 Ti
0.42
GTX 1080
0.38
GTX 1080 Ti
0.45
RTX 2070
0.55
RTX 2080
0.41
RTX 2080 Ti
0.38
RX Vega 56
0.50
RX Vega 64
0.47
0.1238
0.2476
0.3714
0.4952
0.619
1. GTX 1070: $399 reported cost.
2. GTX 1070 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.

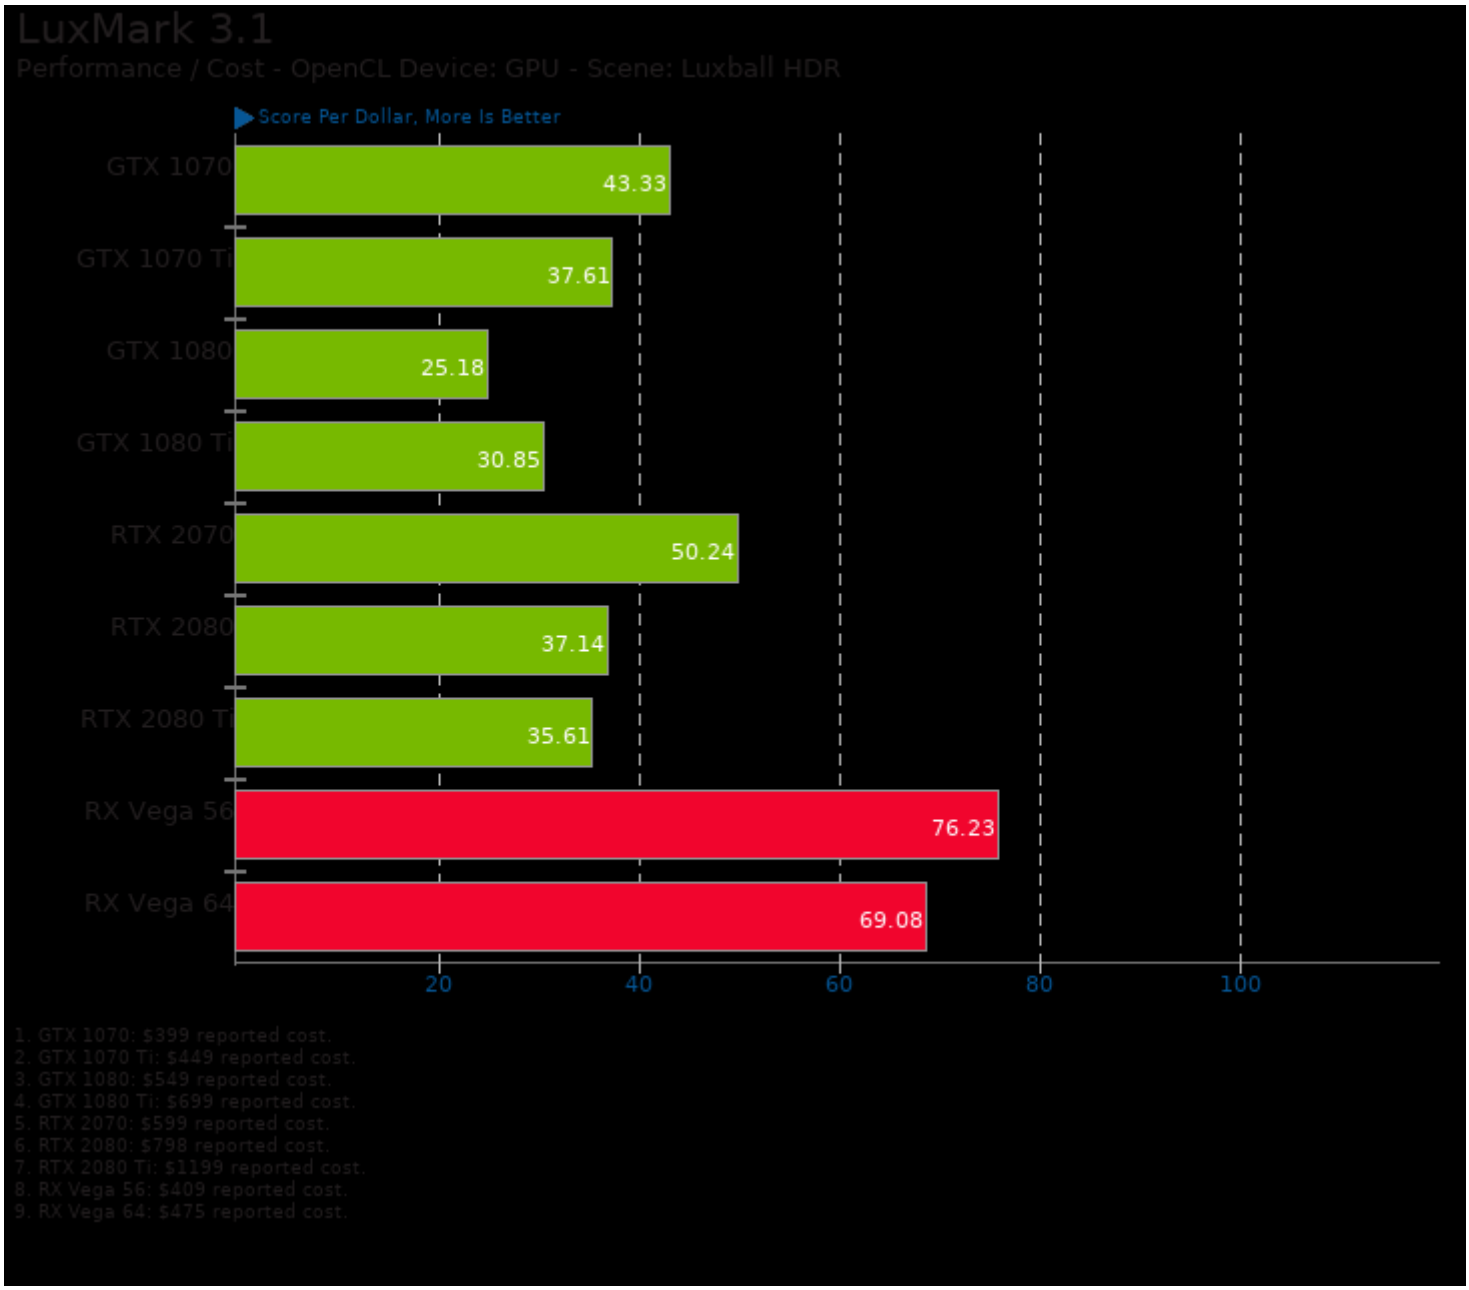
LuxMark 3.1
Performance / Cost - OpenCL Device: GPU - Scene: Luxball HDR
Score Per Dollar, More Is Better
GTX 1070
43.33
GTX 1070 Ti
37.61
GTX 1080
25.18
GTX 1080 Ti
30.85
RTX 2070
50.24
RTX 2080
37.14
RTX 2080 Ti
35.61
RX Vega 56
76.23
RX Vega 64
69.08
20
40
60
80
100
1. GTX 1070: $399 reported cost.
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.

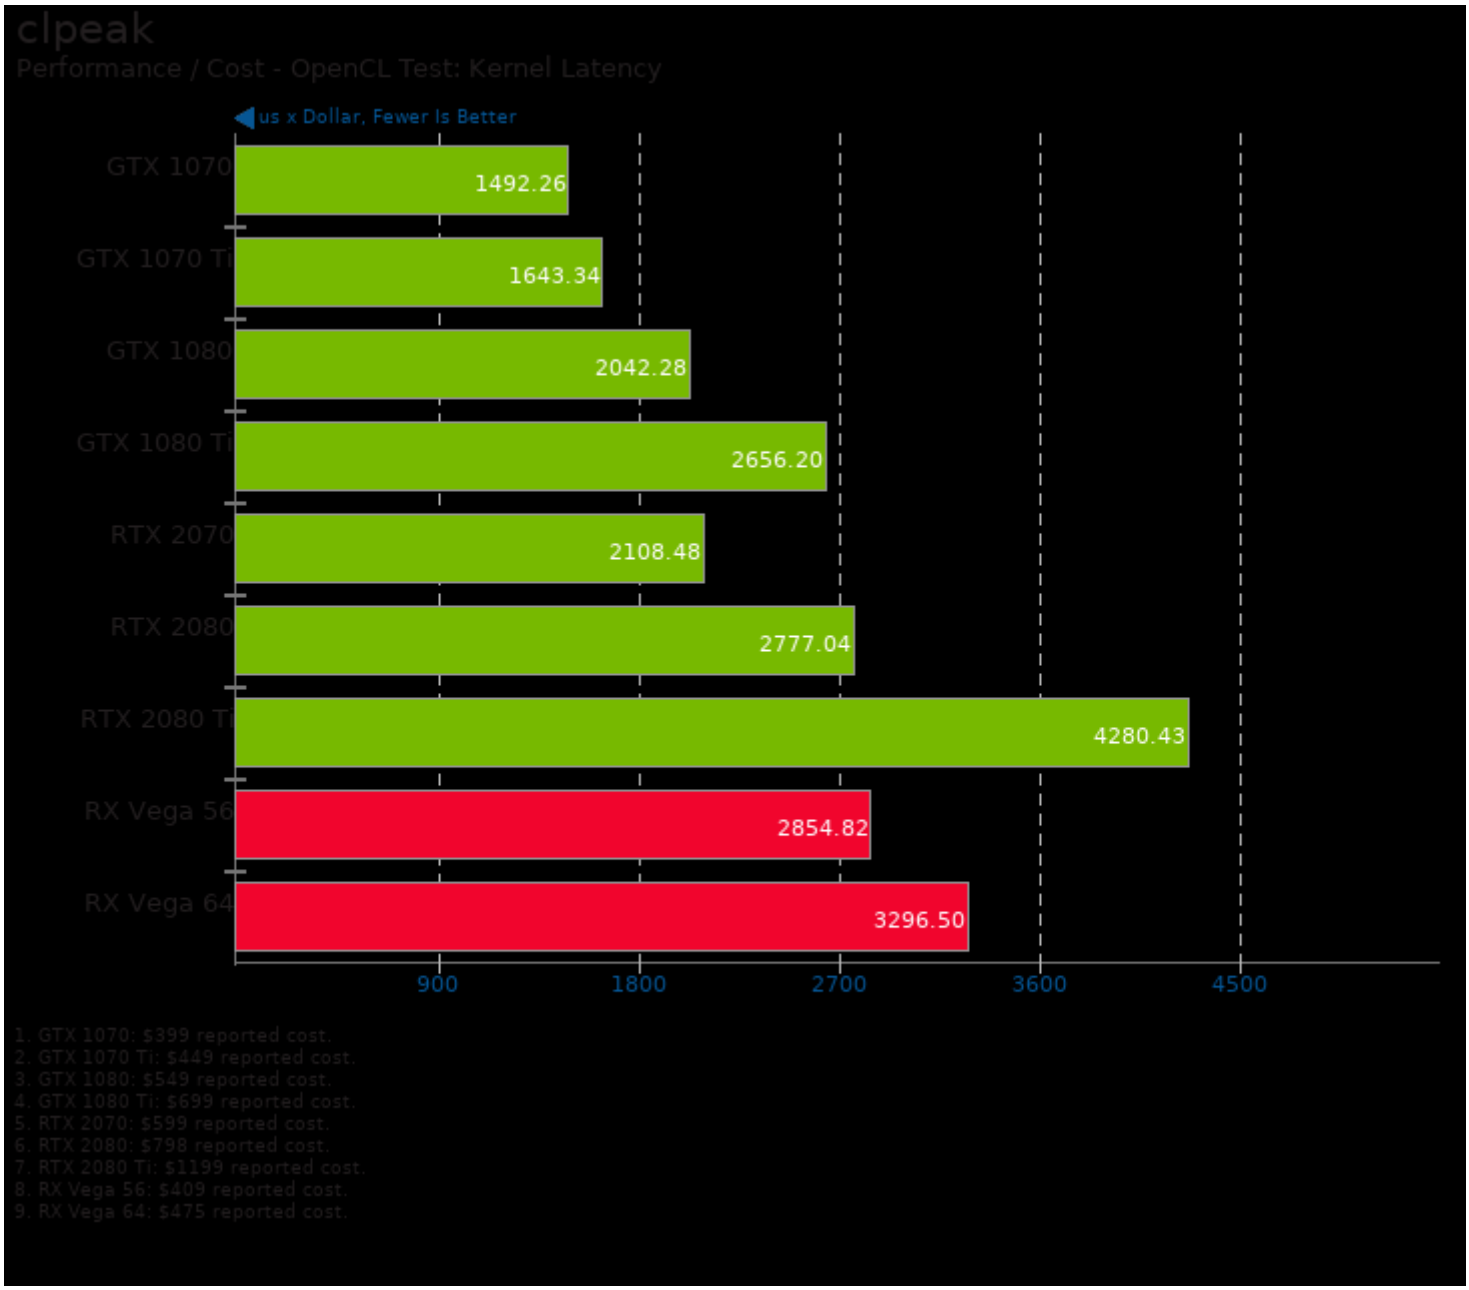
clpeak
Performance / Cost - OpenCL Test: Kernel Latency
us x Dollar, Fewer Is Better
GTX 1070
1492.26
GTX 1070 Ti
1643.34
GTX 1080
2042.28
GTX 1080 Ti
2656.20
RTX 2070
2108.48
RTX 2080
2777.04
RTX 2080 Ti
4280.43
RX Vega 56
2854.82
RX Vega 64
3296.50
900
1800
2700
3600
4500
1. GTX 1070: $399 reported 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.

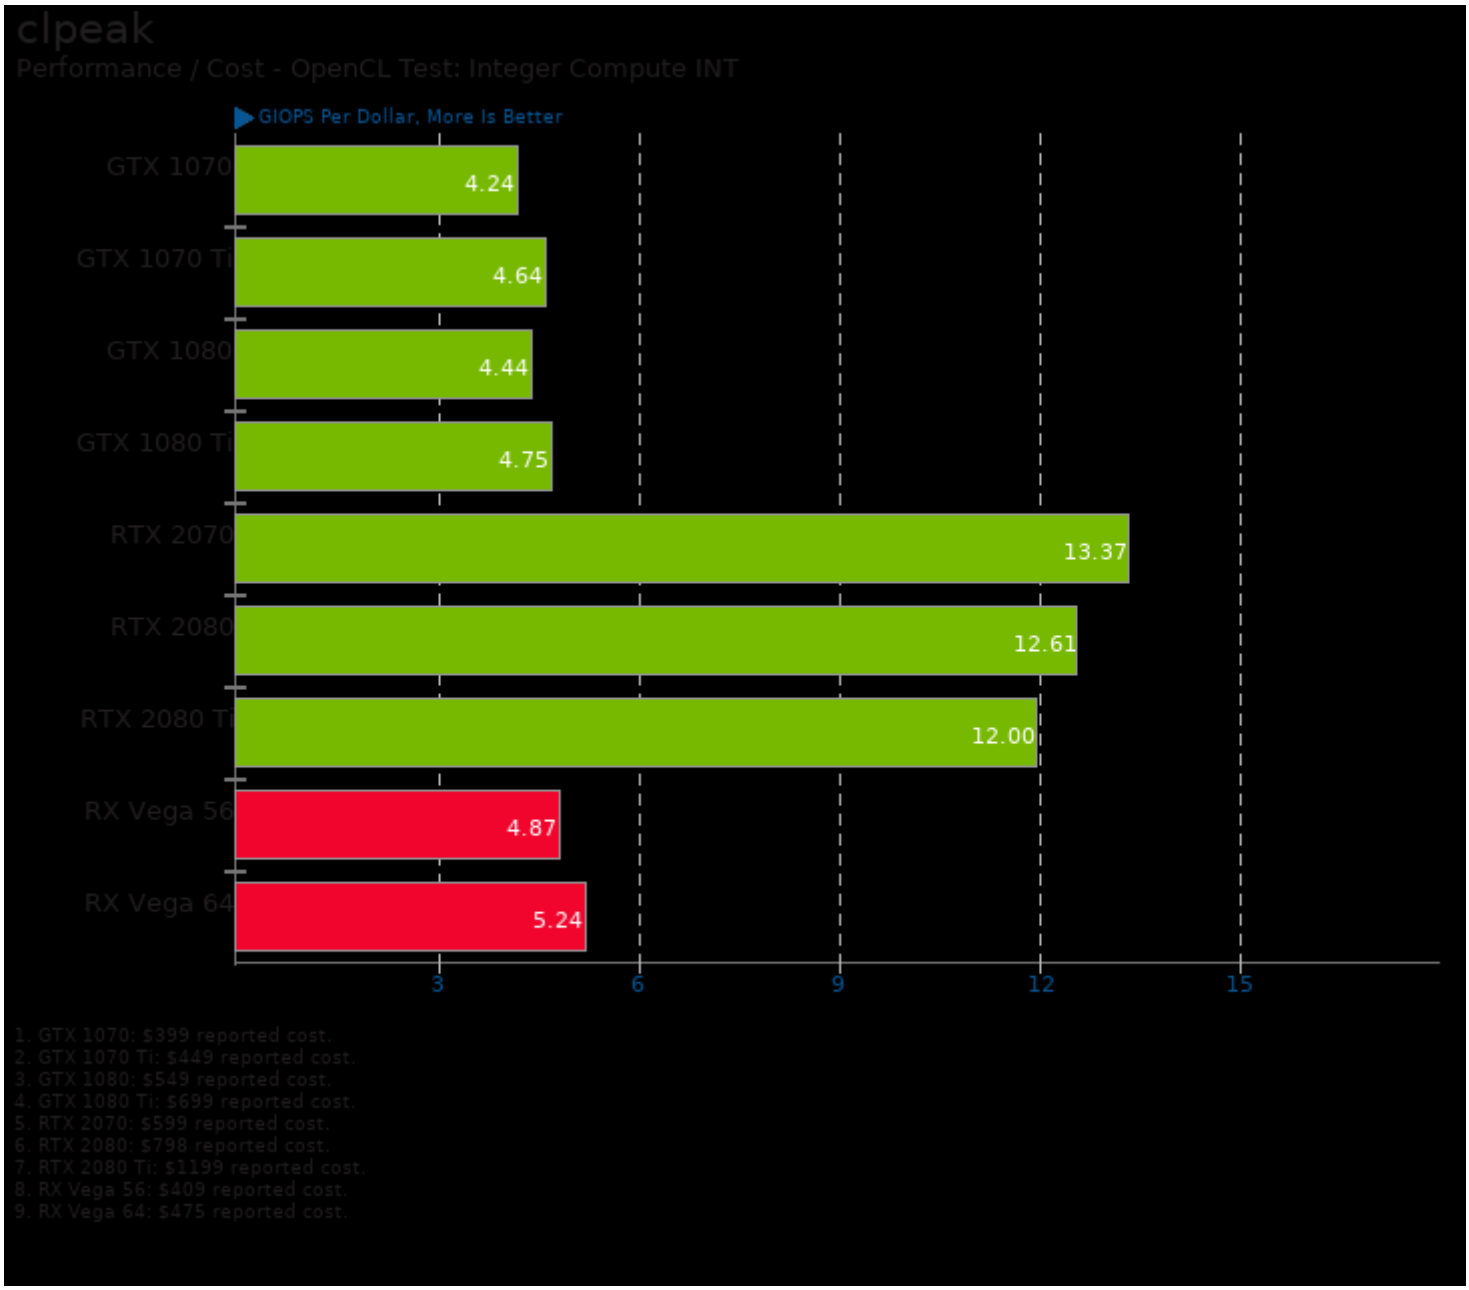
clpeak
Performance / Cost - OpenCL Test: Integer Compute INT
GIOPS Per Dollar, More Is Better
GTX 1070
4.24
GTX 1070 Ti
4.64
GTX 1080
4.44
GTX 1080 Ti
4.75
RTX 2070
13.37
RTX 2080
12.61
RTX 2080 Ti
12.00
RX Vega 56
4.87
RX Vega 64
5.24
3
6
9
12
15
1. GTX 1070: $399 reported cost.
2. GTX 1070 Ti: $449 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.

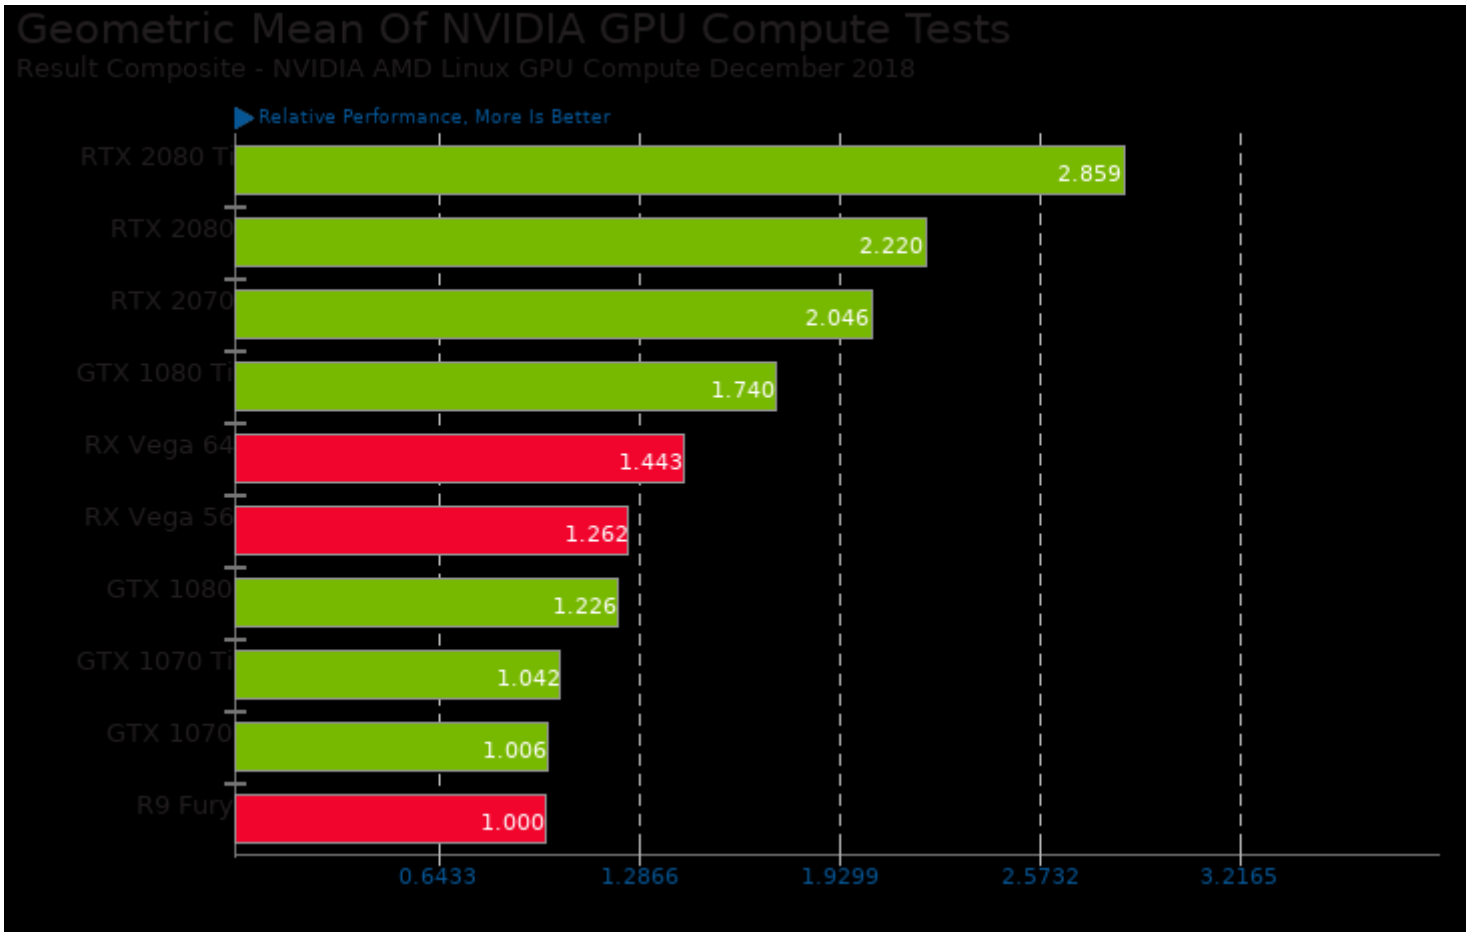

Geometric mean based upon tests: pts/cl-mem, pts/shoc and pts/v-ray

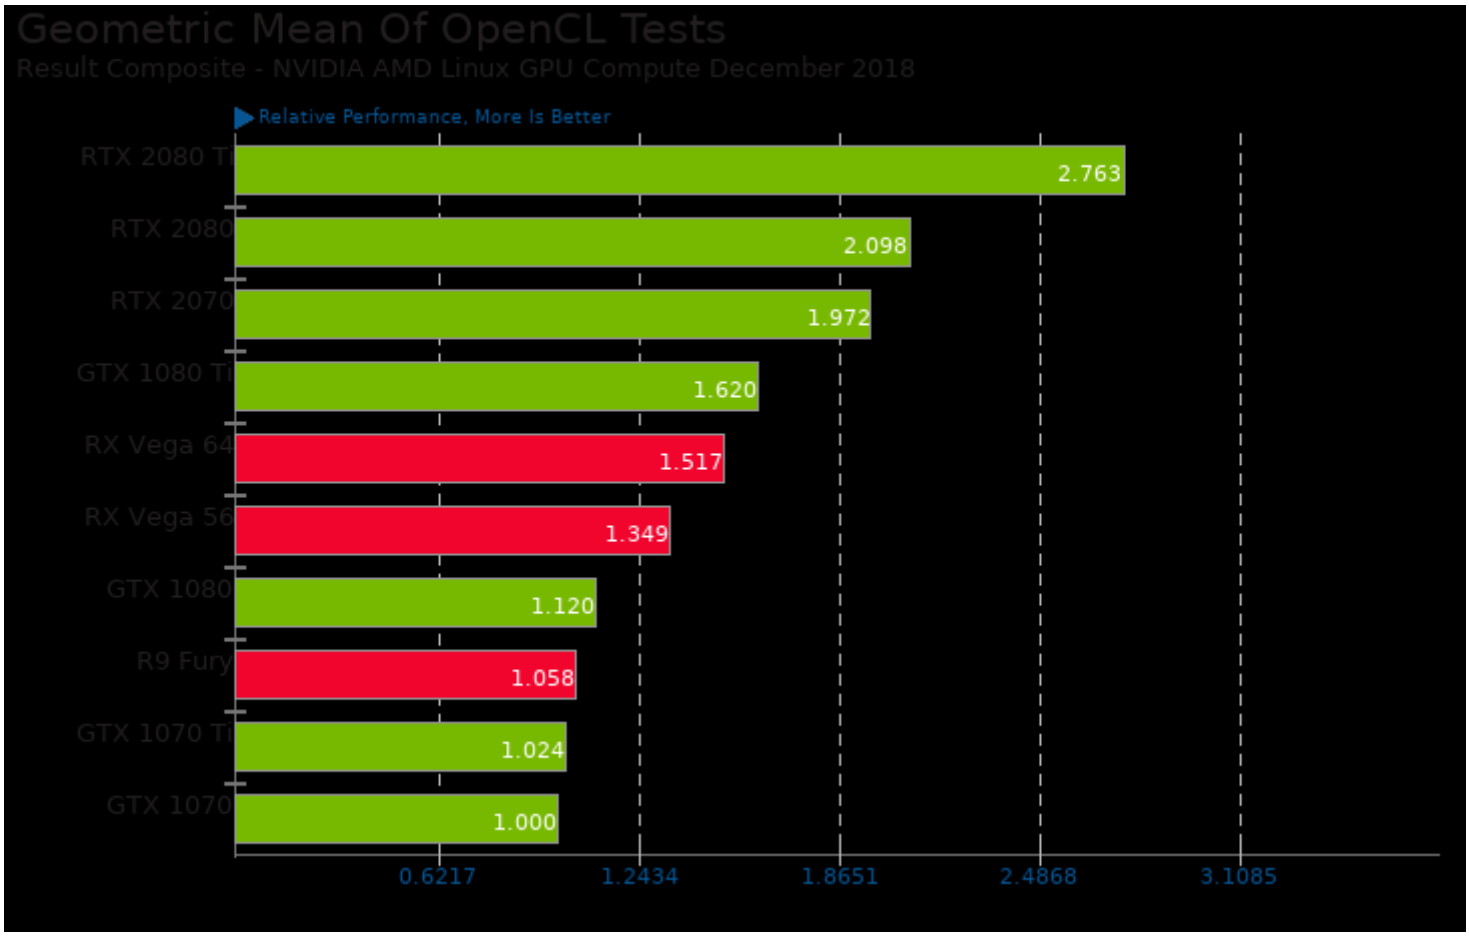

Geometric mean based upon tests: pts/luxmark, pts/shoc and pts/cl-mem

This file was automatically generated via the Phoronix Test Suite benchmarking software on Thursday, 28 March 2024 15:33.