



SPEChpc™ 2021 Tiny Result

Copyright 2021-2023 Standard Performance Evaluation Corporation

GIGA-BYTE TECHNOLOGY CO., LTD
(Test Sponsor: NVIDIA Corporation)


SPEChpc 2021_tny_base = 19.8

SPEChpc 2021_tny_peak = 23.9

GIGA-BYTE G242-P31 (Ampere Altra Q80-33, Tesla A100-PCIE-40GB)

hpc2021 License: 019
Test Sponsor: NVIDIA Corporation
Tested by: NVIDIA Corporation

Test Date: Sep-2021
Hardware Availability: Jun-2021
Software Availability: Sep-2021



Results Table

Benchmark	Base								Peak									
	Model	Ranks	Thrds/Rnk	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Model	Ranks	Thrds/Rnk	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
505.lbm_t	ACC	2	1	54.7	41.1	54.7	41.2			ACC	2	1	54.7	41.1	54.7	41.2		
513.soma_t	ACC	2	1	80.8	45.8	81.0	45.7			ACC	2	1	77.2	48.0	77.3	47.9		
518.tealeaf_t	ACC	2	1	180	9.15	180	9.16			ACC	2	1	152	10.8	152	10.9		
519.clvleaf_t	ACC	2	1	66.5	24.8	67.7	24.4			ACC	2	1	62.7	26.3	62.4	26.4		
521.miniswp_t	ACC	2	1	109	14.6	110	14.6			ACC	2	1	88.5	18.1	90.1	17.8		
528.pot3d_t	ACC	2	1	99.5	21.4	99.6	21.3			ACC	2	1	99.3	21.4	99.8	21.3		
532.sph_exa_t	ACC	2	1	289	6.74	289	6.76			ACC	16	1	95.8	20.3	96.4	20.2		
534.hpgmgfv_t	ACC	2	1	124	9.44	122	9.62			ACC	2	1	110	10.7	112	10.5		
535.weather_t	ACC	2	1	58.1	55.5	56.3	57.3			ACC	2	1	56.1	57.4	56.2	57.3		

SPEChpc 2021_tny_base = **19.8**

SPEChpc 2021_tny_peak = **23.9**

Results appear in the order in which they were run. Bold underlined text indicates a median measurement.



SPEChpc™ 2021 Tiny Result

Copyright 2021-2023 Standard Performance Evaluation Corporation

GIGA-BYTE TECHNOLOGY CO., LTD

(Test Sponsor: NVIDIA Corporation)

SPEChpc 2021_tny_base = 19.8

SPEChpc 2021_tny_peak = 23.9

GIGA-BYTE G242-P31 (Ampere Altra Q80-33, Tesla A100-PCIE-40GB)

hpc2021 License: 019
Test Sponsor: NVIDIA Corporation
Tested by: NVIDIA Corporation

Test Date: Sep-2021
Hardware Availability: Jun-2021
Software Availability: Sep-2021

Hardware Summary

Type of System: SMP
Compute Node: Ampere Altra
Interconnect: None
Compute Nodes Used: 1
Total Chips: 1
Total Cores: 80
Total Threads: 80
Total Memory: 256 GB
Max. Peak Threads: 1

Software Summary

Compiler: C/C++/Fortran: Version 21.9 of NVIDIA HPC SDK for Linux
MPI Library: OpenMPI Version 4.0.5, included with NVHPC SDK
Other MPI Info: None
Other Software: None
Base Parallel Model: ACC
Base Ranks Run: 2
Base Threads Run: 1
Peak Parallel Models: ACC
Minimum Peak Ranks: 2
Maximum Peak Ranks: 16
Max. Peak Threads: 1
Min. Peak Threads: 1

Node Description: Ampere Altra

Hardware

Number of nodes: 1
Uses of the node: compute
Vendor: GIGA-BYTE TECHNOLOGY CO., LTD
Model: G242-P31
CPU Name: Ampere Altra Q80-33
CPU(s) orderable: 1 chips
Chips enabled: 1
Cores enabled: 80
Cores per chip: 80
Threads per core: 1
CPU Characteristics: Max Frequency 3300Mhz
CPU MHz: 3000
Primary Cache: 64 KB I + 64 KB D on chip per core
Secondary Cache: 1 MB I+D on chip per core
L3 Cache: 32 MB I+D on chip per core
Other Cache: None
Memory: 256 GB (16 x 16 GB 2Rx8 PC4-3200AA-R)
Disk Subsystem: 1 x 960 GB, NVME, M.2, PCIe Gen3
Other Hardware: None
Accel Count: 2
Accel Model: Tesla A100-PCIE-40GB
Accel Vendor: NVIDIA Corporation
Accel Type: GPU
Accel Connection: PCIe 3.0 16x
Accel ECC enabled: Yes
Accel Description: See Notes
Adapter: None
Number of Adapters: 0
Slot Type: None
Data Rate: None
Ports Used: 0

Software

Accelerator Driver: NVIDIA UNIX aarch64 Kernel Module 460.32.03
Adapter: None
Adapter Driver: None
Adapter Firmware: None
Operating System: CentOS 8.3-2011
Local File System: xfs
Shared File System: None
System State: Multi-user, run level 3
Other Software: None

(Continued on next page)



SPEChpc™ 2021 Tiny Result

Copyright 2021-2023 Standard Performance Evaluation Corporation

GIGA-BYTE TECHNOLOGY CO., LTD
(Test Sponsor: NVIDIA Corporation)

SPEChpc 2021_tny_base = 19.8

SPEChpc 2021_tny_peak = 23.9

GIGA-BYTE G242-P31 (Ampere Altra Q80-33, Tesla A100-PCIE-40GB)

hpc2021 License: 019
Test Sponsor: NVIDIA Corporation
Tested by: NVIDIA Corporation

Test Date: Sep-2021
Hardware Availability: Jun-2021
Software Availability: Sep-2021

Node Description: Ampere Altra

Hardware (Continued)

Interconnect Type: None

Interconnect Description: None

Hardware

Vendor: N/A
Model: N/A
Switch Model: N/A
Number of Switches: 0
Number of Ports: 0
Data Rate: 0
Firmware: 0
Topology: N/A
Primary Use: N/A

Software

: --

Submit Notes

The config file option 'submit' was used.
MPI startup command:
mpirun command was used to start MPI jobs.

Platform Notes

Information from nvaccelinfo
CUDA Driver Version: 11020
NVRM version: NVIDIA UNIX aarch64 Kernel Module 460.32.03
Device Number: 0
Device Name: A100-PCIE-40GB
Device Revision Number: 8.0
Global Memory Size: 42505273344
Number of Multiprocessors: 108
Concurrent Copy and Execution: Yes
Total Constant Memory: 65536
Total Shared Memory per Block: 49152
Registers per Block: 65536
Warp Size: 32
Maximum Threads per Block: 1024
Maximum Block Dimensions: 1024, 1024, 64
Maximum Grid Dimensions: 2147483647 x 65535 x 65535
Maximum Memory Pitch: 2147483647B
Texture Alignment: 512B

(Continued on next page)



SPEChpc™ 2021 Tiny Result

Copyright 2021-2023 Standard Performance Evaluation Corporation

GIGA-BYTE TECHNOLOGY CO., LTD
(Test Sponsor: NVIDIA Corporation)

SPEChpc 2021_tny_base = 19.8

SPEChpc 2021_tny_peak = 23.9

GIGA-BYTE G242-P31 (Ampere Altra Q80-33, Tesla A100-PCIE-40GB)

hpc2021 License: 019
Test Sponsor: NVIDIA Corporation
Tested by: NVIDIA Corporation

Test Date: Sep-2021
Hardware Availability: Jun-2021
Software Availability: Sep-2021

Platform Notes (Continued)

Clock Rate:	1410 MHz
Execution Timeout:	No
Integrated Device:	No
Can Map Host Memory:	Yes
Compute Mode:	default
Concurrent Kernels:	Yes
ECC Enabled:	Yes
Memory Clock Rate:	1215 MHz
Memory Bus Width:	5120 bits
L2 Cache Size:	41943040 bytes
Max Threads Per SMP:	2048
Async Engines:	3
Unified Addressing:	Yes
Managed Memory:	Yes
Concurrent Managed Memory:	Yes
Preemption Supported:	Yes
Cooperative Launch:	Yes
Multi-Device:	Yes
Default Target:	cc80

Compiler Version Notes

```
=====
CC 505.lbm_t(base, peak) 513.soma_t(base, peak) 518.tealeaf_t(base, peak)
   521.miniswp_t(base, peak) 534.hpgmgfv_t(base, peak)
-----
```

```
nvc 21.9-0 linuxarm64 target on aarch64 Linux
NVIDIA Compilers and Tools
Copyright (c) 2021, NVIDIA CORPORATION & AFFILIATES. All rights reserved.
-----
```

```
=====
CXXC 532.sph_exa_t(base, peak)
-----
```

```
nvc++ 21.9-0 linuxarm64 target on aarch64 Linux
NVIDIA Compilers and Tools
Copyright (c) 2021, NVIDIA CORPORATION & AFFILIATES. All rights reserved.
-----
```

```
=====
FC 519.clvleaf_t(base, peak) 528.pot3d_t(base, peak) 535.weather_t(base,
   peak)
-----
```

```
nvfortran 21.9-0 linuxarm64 target on aarch64 Linux
NVIDIA Compilers and Tools
Copyright (c) 2021, NVIDIA CORPORATION & AFFILIATES. All rights reserved.
```

(Continued on next page)



SPEChpc™ 2021 Tiny Result

Copyright 2021-2023 Standard Performance Evaluation Corporation

GIGA-BYTE TECHNOLOGY CO., LTD
(Test Sponsor: NVIDIA Corporation)

SPEChpc 2021_tny_base = 19.8

SPEChpc 2021_tny_peak = 23.9

GIGA-BYTE G242-P31 (Ampere Altra Q80-33, Tesla A100-PCIE-40GB)

hpc2021 License: 019
Test Sponsor: NVIDIA Corporation
Tested by: NVIDIA Corporation

Test Date: Sep-2021
Hardware Availability: Jun-2021
Software Availability: Sep-2021

Compiler Version Notes (Continued)

Base Compiler Invocation

C benchmarks:
mpicc

C++ benchmarks:
mpicxx

Fortran benchmarks:
mpif90

Base Portability Flags

521.miniswp_t: -DUSE_KBA -DUSE_ACCELDIR
532.sph_exa_t: -DSPEC_USE_LT_IN_KERNELS --c++17

Base Optimization Flags

C benchmarks:
-Mfprelaxed -Mnouniform -Mstack_arrays -fast -acc=gpu

C++ benchmarks:
-Mfprelaxed -Mnouniform -Mstack_arrays -fast -acc=gpu

Fortran benchmarks:
-Mfprelaxed -Mnouniform -Mstack_arrays -fast -acc=gpu

Base Other Flags

C benchmarks:
-w

C++ benchmarks:
-w

Fortran benchmarks:
-w



SPEChpc™ 2021 Tiny Result

Copyright 2021-2023 Standard Performance Evaluation Corporation

GIGA-BYTE TECHNOLOGY CO., LTD
(Test Sponsor: NVIDIA Corporation)

SPEChpc 2021_tny_base = 19.8

SPEChpc 2021_tny_peak = 23.9

GIGA-BYTE G242-P31 (Ampere Altra Q80-33, Tesla A100-PCIE-40GB)

hpc2021 License: 019
Test Sponsor: NVIDIA Corporation
Tested by: NVIDIA Corporation

Test Date: Sep-2021
Hardware Availability: Jun-2021
Software Availability: Sep-2021

Peak Compiler Invocation

C benchmarks:

mpicc

C++ benchmarks:

mpicxx

Fortran benchmarks:

mpif90

Peak Portability Flags

521.miniswp_t: -DUSE_KBA -DUSE_ACCELDIR
532.sph_exa_t: -DSPEC_USE_LT_IN_KERNELS

Peak Optimization Flags

C benchmarks:

505.lbm_t: basepeak = yes

513.soma_t: -fast -O3 -acc=gpu -gpu=pinned

518.tealeaf_t: -fast -Msafeptr -acc=gpu

521.miniswp_t: -Mfprelaxed -Mnouniform -Mstack_arrays -fast -acc=gpu
-gpu=pinned

534.hpgmgfv_t: -fast -acc=gpu -gpu=pinned -static-nvidia

C++ benchmarks:

-Mfprelaxed -Mnouniform -Mstack_arrays -fast -acc=gpu

Fortran benchmarks:

519.clvleaf_t: -Mfprelaxed -fast -acc=gpu -gpu=pinned

528.pot3d_t: -Mstack_arrays -fast -acc=gpu

535.weather_t: -Mfprelaxed -Mnouniform -Mstack_arrays -fast -acc=gpu



SPEChpc™ 2021 Tiny Result

Copyright 2021-2023 Standard Performance Evaluation Corporation

GIGA-BYTE TECHNOLOGY CO., LTD
(Test Sponsor: NVIDIA Corporation)

SPEChpc 2021_tny_base = 19.8

SPEChpc 2021_tny_peak = 23.9

GIGA-BYTE G242-P31 (Ampere Altra Q80-33, Tesla A100-PCIE-40GB)

hpc2021 License: 019
Test Sponsor: NVIDIA Corporation
Tested by: NVIDIA Corporation

Test Date: Sep-2021
Hardware Availability: Jun-2021
Software Availability: Sep-2021

Peak Other Flags

C benchmarks:

-w

C++ benchmarks:

-w

Fortran benchmarks:

-w

The flags file that was used to format this result can be browsed at
http://www.spec.org/hpc2021/flags/nv2021_flags_v1.0.3.html

You can also download the XML flags source by saving the following link:
http://www.spec.org/hpc2021/flags/nv2021_flags_v1.0.3.xml

SPEChpc is a trademark of the Standard Performance Evaluation Corporation. All other brand and product names appearing in this result are trademarks or registered trademarks of their respective holders.

For questions about this result, please contact the tester. For other inquiries, please contact info@spec.org.

Tested with SPEChpc2021 v1.0.2 on 2021-09-13 19:11:37-0400.
Report generated on 2023-08-25 18:58:32 by hpc2021 PDF formatter v1.0.3.
Originally published on 2021-10-20.