SPEChpc™ 2021 Small Result

Copyright 2021-2023 Standard Performance Evaluation Corporation

Transtec

(Test Sponsor: Helmholtz-Zentrum Dresden - Rossendorf)

Hemera: Supermicro SuperServer 1029GQ-TXRT (Intel Xeon Gold 6136, Tesla P100-SXM2-16GB)

SPEChpc 2021_sml_base = 9.75

SPEChpc 2021_sml_peak = Not Run

hpc2021 License: 065A
Test Sponsor: Helmholtz-Zentrum Dresden - Rossendorf
Tested by: Helmholtz-Zentrum Dresden - Rossendorf

Test Date: Sep-2021
Hardware Availability: Jul-2017
Software Availability: Jul-2021

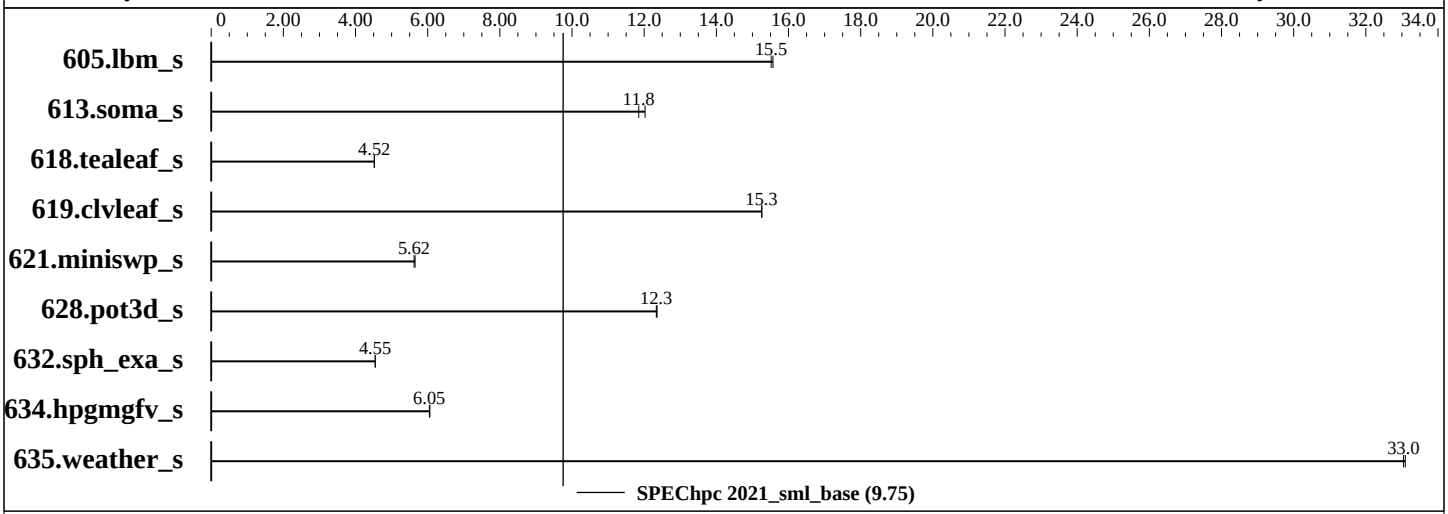

Results Table

Benchmark	Base									Peak								
	Model	Ranks	Thrds/Rnk	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Model	Ranks	Thrds/Rnk	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
605.lbm_s	ACC	32	1	**99.9**	**15.5**	99.6	15.6											
613.soma_s	ACC	32	1	133	12.0	**135**	**11.8**											
618.tealeaf_s	ACC	32	1	453	4.52	**453**	**4.52**											
619.clvleaf_s	ACC	32	1	108	15.3	**108**	**15.3**											
621.miniswp_s	ACC	32	1	195	5.65	**196**	**5.62**											
628.pot3d_s	ACC	32	1	**136**	**12.3**	136	12.4											
632.sph_exa_s	ACC	32	1	**506**	**4.55**	506	4.55											
634.hpgmgfv_s	ACC	32	1	**161**	**6.05**	161	6.06											
635.weather_s	ACC	32	1	78.6	33.1	**78.7**	**33.0**											

SPEChpc 2021_sml_base = 9.75

SPEChpc 2021_sml_peak = Not Run

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

SPEChpc™ 2021 Small Result

Copyright 2021-2023 Standard Performance Evaluation Corporation

Transtec

(Test Sponsor: Helmholtz-Zentrum Dresden - Rossendorf)

Hemera: Supermicro SuperServer 1029GQ-TXRT (Intel Xeon Gold 6136, Tesla P100-SXM2-16GB)

SPEChpc 2021_sml_base = 9.75

SPEChpc 2021_sml_peak = Not Run

hpc2021 License: 065A
Test Sponsor: Helmholtz-Zentrum Dresden - Rossendorf
Tested by: Helmholtz-Zentrum Dresden - Rossendorf

Test Date: Sep-2021
Hardware Availability: Jul-2017
Software Availability: Jul-2021

Hardware Summary

Type of System:	Homogenous Cluster
Compute Node:	Compute Node
Interconnect:	Infiniband (EDR)
Compute Nodes Used:	8
Total Chips:	8
Total Cores:	96
Total Threads:	96
Total Memory:	3 TB
Max. Peak Threads:	--

Software Summary

Compiler:	C/C++/Fortran: Version 21.7 of NVIDIA HPC SDK for Linux
MPI Library:	OpenMPI Version 4.0.5
Other MPI Info:	None
Other Software:	None
Base Parallel Model:	ACC
Base Ranks Run:	32
Base Threads Run:	1
Peak Parallel Models:	Not Run
Minimum Peak Ranks:	--
Maximum Peak Ranks:	--
Max. Peak Threads:	--
Min. Peak Threads:	--

Node Description: Compute Node

Hardware

Number of nodes:	8
Uses of the node:	compute
Vendor:	Intel
Model:	SuperServer 1029GQ-TXRT
CPU Name:	Intel Xeon Gold 6136
CPU(s) orderable:	1 chips
Chips enabled:	1
Cores enabled:	12
Cores per chip:	12
Threads per core:	1
CPU Characteristics:	Intel Turbo Boost Technology up to 3.7 GHz
CPU MHz:	3000
Primary Cache:	32 KB I + 32 KB D on chip per core
Secondary Cache:	1 MB I+D on chip per core
L3 Cache:	25344 KB I+D on chip per chip
Other Cache:	None
Memory:	384 GB (12 x 32GB 2Rx4 PC4-2666V-RB2-12)
Disk Subsystem:	1 x 500 GB
Other Hardware:	None
Accel Count:	4
Accel Model:	Tesla P100-SXM2-16GB
Accel Vendor:	NVIDIA Corporation
Accel Type:	GPU
Accel Connection:	PCIe 3.0 16x
Accel ECC enabled:	Yes
Accel Description:	--
Adapter:	Mellanox MT4115
Number of Adapters:	2
Slot Type:	PCI-Express 3.0 x16
Data Rate:	100 Gb/s
Ports Used:	2

(Continued on next page)

Software

Accelerator Driver:	--
Adapter:	Mellanox MT4115
Adapter Driver:	--
Adapter Firmware:	12.28.2006
Operating System:	CentOS Linux release 7.9.2009 (Core) 3.10.0-1160.6.1.el7.x86_64
Local File System:	xfs
Shared File System:	GPFS Version 5.0.5.0 6 NSD (vendor: NEC) 5 building blocks (vendor: NetApp): 2x (240 x 8 TB HDD) 1x (180 x 12 TB HDD) 1x (240 x 16 TB HDD) 1x (120 x 16 TB HDD)
System State:	Multi-user, run level 3
Other Software:	None

Node Description: Compute Node

Hardware (Continued)

Interconnect Type: EDR Infiniband

Interconnect Description: Infiniband (EDR)

Hardware

Vendor:	Mellanox Technologies
Model:	Mellanox SB7790
Switch Model:	36 x EDR 100 Gb/s
Number of Switches:	2
Number of Ports:	36
Data Rate:	100 Gb/s
Firmware:	--
Topology:	Mesh (blocking factor: 8:1)
Primary Use:	MPI Traffic, GPFS

Software

: --

Submit Notes

```
The config file option 'submit' was used.
  MPI startup command:
    mpirun --bind-to socket -np $ranks $[top]/mpirunCUDA.sh $command
  contents of $[top]/mpirunCUDA.sh
    #!/bin/bash
    export CUDA_VISIBLE_DEVICES=$OMPI_COMM_WORLD_LOCAL_RANK
    $@
```

Compiler Version Notes

```
==============================================================================
 CC  605.lbm_s(base) 613.soma_s(base) 618.tealeaf_s(base) 621.miniswp_s(base)
     634.hpgmgfv_s(base)
------------------------------------------------------------------------------
nvc 21.7-0 64-bit target on x86-64 Linux -tp skylake
NVIDIA Compilers and Tools
Copyright (c) 2021, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
 CXXC 632.sph_exa_s(base)
------------------------------------------------------------------------------
nvc++ 21.7-0 64-bit target on x86-64 Linux -tp skylake
NVIDIA Compilers and Tools
Copyright (c) 2021, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
```

(Continued on next page)

Compiler Version Notes (Continued)

```
------------------------------------------------------------------------------

==============================================================================
 FC  619.clvleaf_s(base) 628.pot3d_s(base) 635.weather_s(base)
------------------------------------------------------------------------------
nvfortran 21.7-0 64-bit target on x86-64 Linux -tp skylake
NVIDIA Compilers and Tools
Copyright (c) 2021, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------
```

Base Compiler Invocation

C benchmarks:
`mpicc`

C++ benchmarks:
`mpicxx`

Fortran benchmarks:
`mpif90`

Base Portability Flags

621.miniswp_s: `-DUSE_KBA -DUSE_ACCELDIR`
632.sph_exa_s: `-DSPEC_USE_LT_IN_KERNELS --c++17`

Base Optimization Flags

C benchmarks:
`-Mfprelaxed -Mnouniform -Mstack_arrays -fast -acc=gpu -Minfo=accel -DSPEC_ACCEL_AWARE_MPI`

C++ benchmarks:
`-Mfprelaxed -Mnouniform -Mstack_arrays -fast -acc=gpu -Minfo=accel -DSPEC_ACCEL_AWARE_MPI`

Fortran benchmarks:
`-DSPEC_ACCEL_AWARE_MPI -Mfprelaxed -Mnouniform -Mstack_arrays -fast -acc=gpu -Minfo=accel`

Base 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.html`

You can also download the XML flags source by saving the following link:
`http://www.spec.org/hpc2021/flags/nv2021_flags.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 2021-09-11 14:12:13-0400.
Report generated on 2023-08-25 19:03:01 by hpc2021 PS formatter v1.0.3.
Originally published on 2021-10-27.