SPEChpc™ 2021 Small Result

Copyright 2021-2023 Standard Performance Evaluation Corporation

xFusion

FusionServer G5500 V6 (Intel Xeon Platinum 8380, Nvidia A100-PCIE-80G)

SPEChpc 2021_sml_base = 6.58

SPEChpc 2021_sml_peak = 6.79

hpc2021 License: 6488
Test Sponsor: xFusion
Tested by: xFusion

Test Date: Jul-2022
Hardware Availability: Apr-2021
Software Availability: May-2022

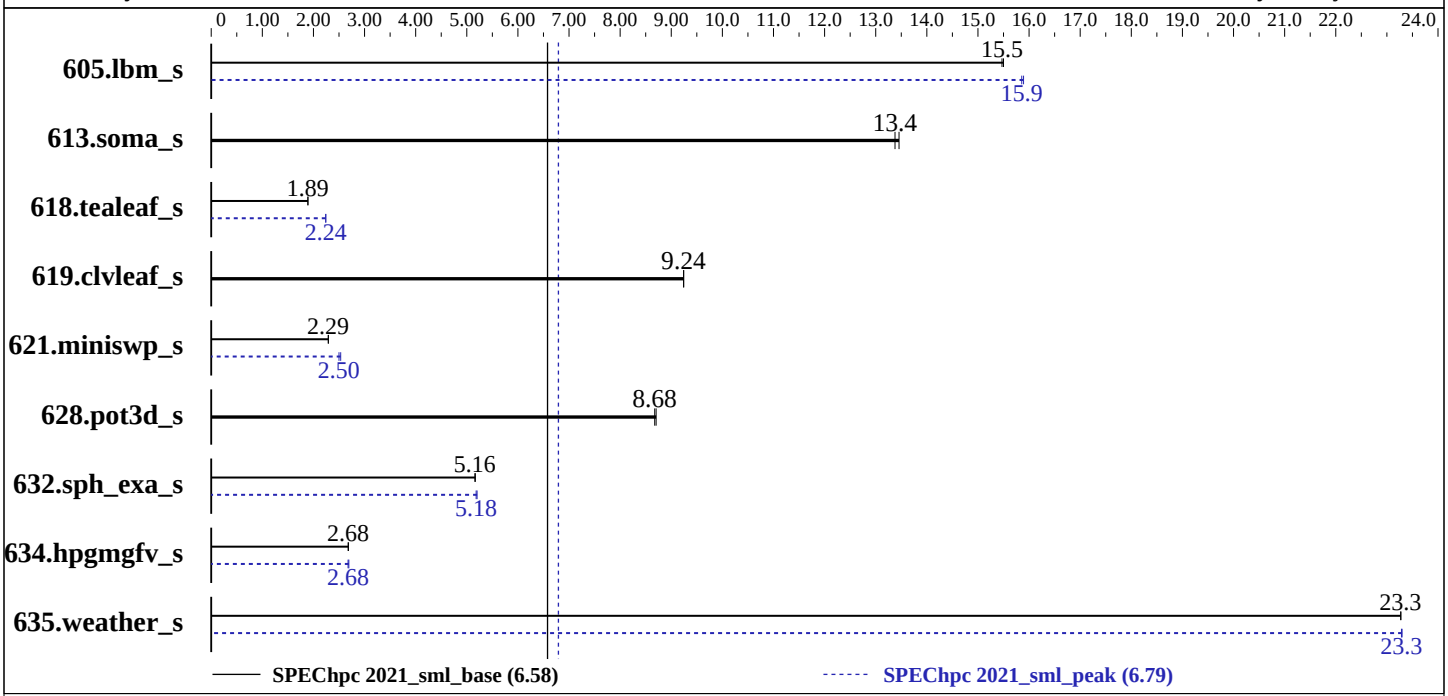

Results Table

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	7	1	100	15.5	**100**	**15.5**			ACC	7	1	97.5	15.9	**97.7**	**15.9**		
613.soma_s	ACC	7	1	**120**	**13.4**	119	13.5			ACC	7	1	**120**	**13.4**	119	13.5		
618.tealeaf_s	ACC	7	1	**1084**	**1.89**	1083	1.89			ACC	7	1	915	2.24	**915**	**2.24**		
619.clvleaf_s	ACC	7	1	**179**	**9.24**	179	9.24			ACC	7	1	**179**	**9.24**	179	9.24		
621.miniswp_s	ACC	7	1	**480**	**2.29**	480	2.29			ACC	7	1	434	2.53	**440**	**2.50**		
628.pot3d_s	ACC	7	1	192	8.70	**193**	**8.68**			ACC	7	1	192	8.70	**193**	**8.68**		
632.sph_exa_s	ACC	7	1	**446**	**5.16**	445	5.17			ACC	7	1	442	5.20	**444**	**5.18**		
634.hpgmgfv_s	ACC	7	1	**364**	**2.68**	363	2.68			ACC	7	1	**364**	**2.68**	362	2.69		
635.weather_s	ACC	7	1	112	23.3	**112**	**23.3**			ACC	7	1	112	23.3	**112**	**23.3**		

SPEChpc 2021_sml_base = 6.58

SPEChpc 2021_sml_peak = 6.79

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

xFusion

FusionServer G5500 V6 (Intel Xeon Platinum 8380, Nvidia A100-PCIE-80G)

SPEChpc 2021_sml_base = 6.58

SPEChpc 2021_sml_peak = 6.79

hpc2021 License: 6488
Test Sponsor: xFusion
Tested by: xFusion

Test Date: Jul-2022
Hardware Availability: Apr-2021
Software Availability: May-2022

Hardware Summary

Type of System:	SMP
Compute Node:	FusionServer G5500 V6
Interconnect:	None
Compute Nodes Used:	1
Total Chips:	2
Total Cores:	80
Total Threads:	80
Total Memory:	1 TB
Max. Peak Threads:	1

Software Summary

Compiler:	Nvidia HPC SDK 22.5
MPI Library:	OpenMPI Version 4.0.5, included with NVHPC SDK
Other MPI Info:	--
Other Software:	--
Base Parallel Model:	ACC
Base Ranks Run:	7
Base Threads Run:	1
Peak Parallel Models:	ACC
Minimum Peak Ranks:	7
Maximum Peak Ranks:	7
Max. Peak Threads:	1
Min. Peak Threads:	1

Node Description: FusionServer G5500 V6

Hardware

Number of nodes:	1
Uses of the node:	compute
Vendor:	xFusion
Model:	FusionServer G5500 V6
CPU Name:	Intel Xeon Platinum 8380
CPU(s) orderable:	1, 2 chips
Chips enabled:	2
Cores enabled:	80
Cores per chip:	40
Threads per core:	1
CPU Characteristics:	Intel Turbo Boost Technology up to 3.4 GHz
CPU MHz:	2300
Primary Cache:	32 KB I + 48 KB D on chip per core
Secondary Cache:	1.25 MB I+D on chip per core
L3 Cache:	60 MB I+D on chip per chip
Other Cache:	None
Memory:	1 TB (16 x 64 GB 2Rx4 PC4-3200A-R)
Disk Subsystem:	1 x 3.2 TB NVMe SSD
Other Hardware:	None
Accel Count:	8
Accel Model:	Tesla A100 PCIe 80GB
Accel Vendor:	Nvidia Corporation
Accel Type:	GPU
Accel Connection:	PCIe Gen4 x16
Accel ECC enabled:	Yes
Accel Description:	Nvidia Tesla A100 PCIe 80GB
Adapter:	None
Number of Adapters:	0
Slot Type:	None
Data Rate:	None
Ports Used:	0
Interconnect Type:	None

Software

Accelerator Driver:	NVIDIA UNIX x86_64 Kernel Module 515.43.04
Adapter:	None
Adapter Driver:	None
Adapter Firmware:	None
Operating System:	CentOS Linux release 8.2.2004 4.18.0-193.el8.x86_644
Local File System:	xfs
Shared File System:	None
System State:	Multi-user, run level 3
Other Software:	None

Interconnect Description: None

Hardware		Software
Vendor:	None	: --
Model:	None	
Switch Model:	None	
Number of Switches:	0	
Number of Ports:	0	
Data Rate:	None	
Firmware:	None	
Topology:	None	
Primary Use:	None	

Submit Notes

```
The config file option 'submit' was used.
MPIRUN_OPTS = --allow-run-as-root --bind-to none
submit = mpirun --allow-run-as-root -x UCX_MEMTYPE_CACHE=n -np $ranks perl $[top]/bind.pl $command
```

Compiler Version Notes

```
==============================================================================
 CC  605.lbm_s(base, peak) 613.soma_s(base, peak) 618.tealeaf_s(base, peak)
     621.miniswp_s(base, peak) 634.hpgmgfv_s(base, peak)
------------------------------------------------------------------------------
nvc 22.5-0 64-bit target on x86-64 Linux -tp skylake-avx512
NVIDIA Compilers and Tools
Copyright (c) 2022, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
 CXXC 632.sph_exa_s(base, peak)
------------------------------------------------------------------------------
nvc++ 22.5-0 64-bit target on x86-64 Linux -tp skylake-avx512
NVIDIA Compilers and Tools
Copyright (c) 2022, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
 FC  619.clvleaf_s(base, peak) 628.pot3d_s(base, peak) 635.weather_s(base,
     peak)
------------------------------------------------------------------------------
nvfortran 22.5-0 64-bit target on x86-64 Linux -tp skylake-avx512
NVIDIA Compilers and Tools
Copyright (c) 2022, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------
```

Base Compiler Invocation

C benchmarks:
`mpicc`

C++ benchmarks:
`mpicxx`

Fortran benchmarks:
`mpif90`

Base Portability Flags

632.sph_exa_s: `--c++17`

Base Optimization Flags

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

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

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

Base Other Flags

C benchmarks:
`-w`

C++ benchmarks:
`-w`

Fortran benchmarks:
`-w`

Peak Compiler Invocation

C benchmarks:
`mpicc`

C++ benchmarks:
`mpicxx`

Fortran benchmarks:
`mpif90`

Peak Optimization Flags

C benchmarks:

605.lbm_s: `-fast -acc=gpu -O3 -Mfprelaxed -Mnouniform -DSPEC_ACCEL_AWARE_MPI`

613.soma_s: `basepeak = yes`

618.tealeaf_s: `-fast -acc=gpu -Msafeptr -DSPEC_ACCEL_AWARE_MPI`

621.miniswp_s: `-fast -acc=gpu -gpu=pinned`

634.hpgmgfv_s: `-fast -acc=gpu -static-nvidia -DSPEC_ACCEL_AWARE_MPI`

C++ benchmarks:
`-fast -acc=gpu -O3 -Mfprelaxed -Mnouniform -Mstack_arrays -static-nvidia -DSPEC_ACCEL_AWARE_MPI`

Fortran benchmarks:

619.clvleaf_s: `basepeak = yes`

628.pot3d_s: `basepeak = yes`

635.weather_s: `-DSPEC_ACCEL_AWARE_MPI -fast -acc=gpu -O3 -Mfprelaxed -Mnouniform -Mstack_arrays -static-nvidia`

Peak Other Flags

C benchmarks:
`-w`

(Continued on next page)

Peak Other Flags (Continued)

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_v1.0.3.2022-08-24.html`

You can also download the XML flags source by saving the following link:
`http://www.spec.org/hpc2021/flags/nv2021_flags_v1.0.3.2022-08-24.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 v1.0.3 on 2022-07-29 07:45:44-0400.
Report generated on 2023-03-27 12:20:34 by hpc2021 PS formatter v1.0.3.
Originally published on 2022-08-24.