SPEChpc™ 2021 Tiny Result

Copyright 2021-2023 Standard Performance Evaluation Corporation

Lenovo Global Technology

ThinkSystem SR670 V2 (Intel Xeon Platinum 8380, Nvidia A100-PCIE-40G)

SPEChpc 2021_tny_base = 40.5

SPEChpc 2021_tny_peak = Not Run

hpc2021 License: 28
Test Sponsor: Lenovo Global Technology
Tested by: Lenovo Global Technology

Test Date: Aug-2021
Hardware Availability: Aug-2021
Software Availability: Aug-2021

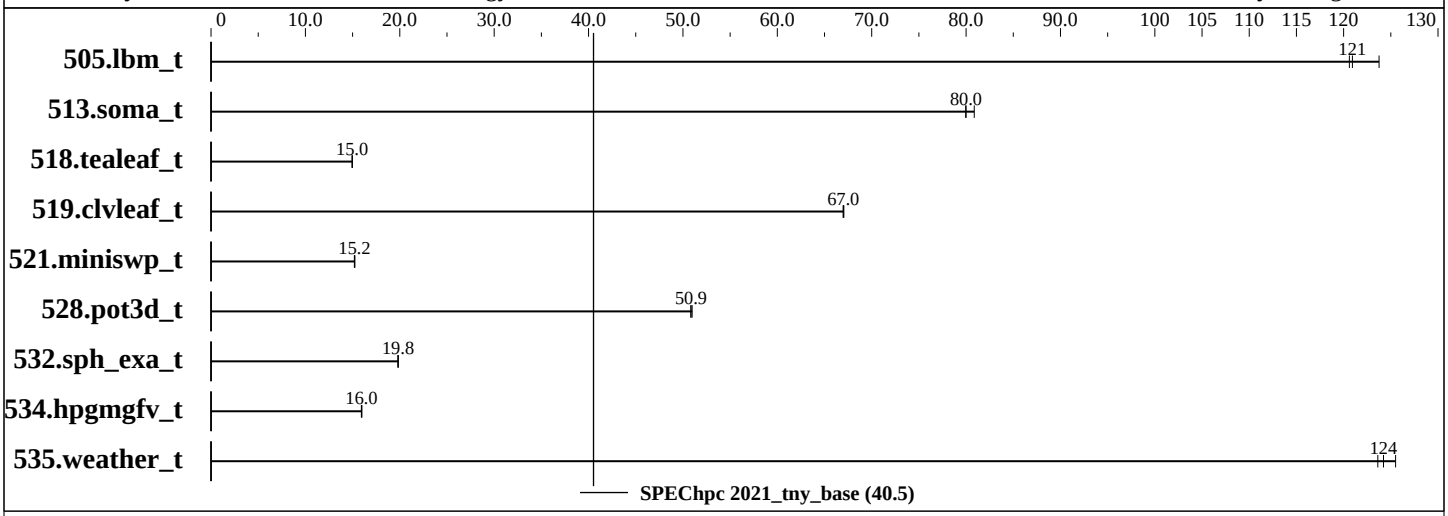

Results Table

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	6	1	18.2	124	**18.6**	**121**	18.7	121									
513.soma_t	ACC	6	1	45.8	80.9	46.3	80.0	**46.2**	**80.0**									
518.tealeaf_t	ACC	6	1	110	15.0	110	15.0	**110**	**15.0**									
519.clvleaf_t	ACC	6	1	24.6	67.0	**24.6**	**67.0**	24.6	67.0									
521.miniswp_t	ACC	6	1	105	15.2	105	15.2	**105**	**15.2**									
528.pot3d_t	ACC	6	1	41.7	51.0	**41.8**	**50.9**	41.8	50.8									
532.sph_exa_t	ACC	6	1	98.3	19.8	98.4	19.8	**98.4**	**19.8**									
534.hpgmgfv_t	ACC	6	1	73.5	16.0	73.8	15.9	**73.6**	**16.0**									
535.weather_t	ACC	6	1	**26.0**	**124**	26.1	124	25.7	126									

SPEChpc 2021_tny_base = 40.5

SPEChpc 2021_tny_peak = Not Run

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

Lenovo Global Technology

ThinkSystem SR670 V2 (Intel Xeon Platinum 8380, Nvidia A100-PCIE-40G)

SPEChpc 2021_tny_base = 40.5

SPEChpc 2021_tny_peak = Not Run

hpc2021 License: 28
Test Sponsor: Lenovo Global Technology
Tested by: Lenovo Global Technology

Test Date: Aug-2021
Hardware Availability: Aug-2021
Software Availability: Aug-2021

Hardware Summary

Type of System:	Homogenous
Compute Node:	ThinkSystem SR670 V2
Interconnect:	None
File Server Node:	ThinkSystem SR670 V2
Compute Nodes Used:	1
Total Chips:	2
Total Cores:	80
Total Threads:	80
Total Memory:	512 GB
Max. Peak Threads:	--

Software Summary

Compiler:	Nvidia HPC SDK 21.5
MPI Library:	Open MPI 4.0.5
Other MPI Info:	None
Other Software:	--
Base Parallel Model:	ACC
Base Ranks Run:	6
Base Threads Run:	1
Peak Parallel Models:	Not Run
Minimum Peak Ranks:	--
Maximum Peak Ranks:	--
Max. Peak Threads:	--
Min. Peak Threads:	--

Node Description: ThinkSystem SR670 V2

Hardware

Number of nodes:	1
Uses of the node:	compute
Vendor:	Lenovo Global Technology
Model:	ThinkSystem SR670 V2
CPU Name:	Intel Xeon Platinum 8380
CPU(s) orderable:	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:	1280 KB I+D on chip per core
L3 Cache:	60 MB I+D on chip per chip
Other Cache:	None
Memory:	512 GB (16 x 32 GB 2Rx8 PC4-3200A-R)
Disk Subsystem:	1 x 4 TB NVMe SSD
Other Hardware:	None
Accel Count:	8
Accel Model:	Tesla A100 PCIe 40GB
Accel Vendor:	Nvidia Corporation
Accel Type:	GPU
Accel Connection:	PCIe Gen4 x16
Accel ECC enabled:	Yes
Accel Description:	Nvidia Tesla A100 PCIe 40GB
Adapter:	Mellanox ConnectX-6 HDR
Number of Adapters:	1
Slot Type:	PCI-Express 4.0 x16
Data Rate:	200 Gb/s
Ports Used:	1
Interconnect Type:	Nvidia Mellanox ConnectX-6 HDR

Software

Accelerator Driver:	470.42.01
Adapter:	Mellanox ConnectX-6 HDR
Adapter Driver:	5.2-1.0.4
Adapter Firmware:	20.28.1002
Operating System:	Red Hat Enterprise Linux Server release 8.3, Kernel 4.18.0-193.el8.x86_64
Local File System:	xfs
Shared File System:	XFS
System State:	Multi-user, run level 3
Other Software:	None

Node Description: ThinkSystem SR670 V2

Hardware

Number of nodes:	1
Uses of the node:	Fileserver
Vendor:	Lenovo Global Technology
Model:	ThinkSystem SR670 V2
CPU Name:	Intel Xeon Platinum 8380
CPU(s) orderable:	2 chips
Chips enabled:	2
Cores enabled:	80
Cores per chip:	40
Threads per core:	1
CPU Characteristics:	Turbo up to 3.4 GHz
CPU MHz:	2300
Primary Cache:	32 KB I + 48 KB D on chip per core
Secondary Cache:	1280 KB I+D on chip per core
L3 Cache:	60 MB I+D on chip per chip
Other Cache:	None
Memory:	512 GB (16 x 32 GB 2Rx8 PC4-3200A-R)
Disk Subsystem:	1 x 4 TB NVMe SSD
Other Hardware:	None
Accel Count:	8
Accel Model:	Tesla A100 PCIe 40GB
Accel Vendor:	Nvidia
Accel Type:	GPU
Accel Connection:	Nvidia Tesla A100 PCIe 40GB
Accel ECC enabled:	Yes
Accel Description:	Nvidia Tesla A100 PCIe 40GB
Adapter:	Mellanox ConnectX-6 HDR
Number of Adapters:	1
Slot Type:	PCI-Express 4.0 x16
Data Rate:	200 Gb/s
Ports Used:	1
Interconnect Type:	Nvidia Mellanox ConnectX-6 HDR

Software

Accelerator Driver:	None
Adapter:	Mellanox ConnectX-6 HDR
Adapter Driver:	5.2-1.0.4
Adapter Firmware:	20.28.1002
Operating System:	Red Hat Enterprise Linux Server release 8.3
Local File System:	xfs
Shared File System:	None
System State:	Multi-User, run level 3
Other Software:	None

Interconnect Description: None

Hardware

Vendor:	None
Model:	None
Switch Model:	None
Number of Switches:	0
Number of Ports:	0
Data Rate:	None
Firmware:	None
Topology:	None

(Continued on next page)

Software

: --

Interconnect Description: None

Hardware (Continued)

Primary Use: None

Submit Notes

```
Indiviual Ranks were bound to the CPU cores on the same NUMA node as
the GPU using 'numactl' within the following "bind.pl" perl script:
---- Start bind.pl ------
my %bind;
$bind{0} = "1-3";
$bind{1} = "4-7";
$bind{2} = "8-10";
$bind{3} = "11-14";
$bind{4} = "41-43";
$bind{5} = "44-47";
$bind{6} = "61-63";
$bind{7} = "64-67";
my $rank = $ENV{OMPI_COMM_WORLD_LOCAL_RANK};
my $cmd = "taskset -c $bind{$rank} ";
while (my $arg = shift) {
 $cmd .= "$arg ";
}
my $rc = system($cmd);
exit($rc);
---- End bind.pl ------
The config file option 'submit' was used.
submit = mpirun --allow-run-as-root -x UCX_MEMTYPE_CACHE=n
-host localhost:8 -np $ranks perl $[top]/bind.pl $command
```

General Notes

```
Environment variables set by runhpc before the start of the run:
UCX_MEMTYPE_CACHE = "n"
UCX_TLS = "self,shm,cuda_copy"
```

Compiler Version Notes

```
==============================================================================
 CC  505.lbm_t(base) 513.soma_t(base) 518.tealeaf_t(base) 521.miniswp_t(base)
      534.hpgmgfv_t(base)
------------------------------------------------------------------------------
nvc 21.5-0 LLVM 64-bit target on x86-64 Linux -tp skylake
NVIDIA Compilers and Tools
```

(Continued on next page)

Compiler Version Notes (Continued)

```
Copyright (c) 2021, NVIDIA CORPORATION.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
 CXXC 532.sph_exa_t(base)
------------------------------------------------------------------------------
nvc++ 21.5-0 LLVM 64-bit target on x86-64 Linux -tp skylake
NVIDIA Compilers and Tools
Copyright (c) 2021, NVIDIA CORPORATION.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
 FC  519.clvleaf_t(base) 528.pot3d_t(base) 535.weather_t(base)
------------------------------------------------------------------------------
nvfortran 21.5-0 LLVM 64-bit target on x86-64 Linux -tp skylake
NVIDIA Compilers and Tools
Copyright (c) 2021, NVIDIA CORPORATION.  All rights reserved.
------------------------------------------------------------------------------
```

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`
`-DSPEC_ACCEL_AWARE_MPI`

(Continued on next page)

Base Optimization Flags (Continued)

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

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

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 v1.0.1 on 2021-08-23 07:37:04-0400.
Report generated on 2023-08-25 18:57:11 by hpc2021 PS formatter v1.0.3.
Originally published on 2021-10-20.