SPEC ACCEL™ ACC Result

Copyright 2015-2019 Standard Performance Evaluation Corporation

Dell
(Test Sponsor: Indiana University)

Tesla V100X-8Q

PowerEdge C4140 Server (KVM virtual machine)

SPECaccel_acc_peak = Not Run

SPECaccel_acc_base = 7.18

ACCEL license: 3440A
Test sponsor: Indiana University
Tested by: Indiana University

Test date: Jul-2019
Hardware Availability: May-2019
Software Availability: Apr-2019

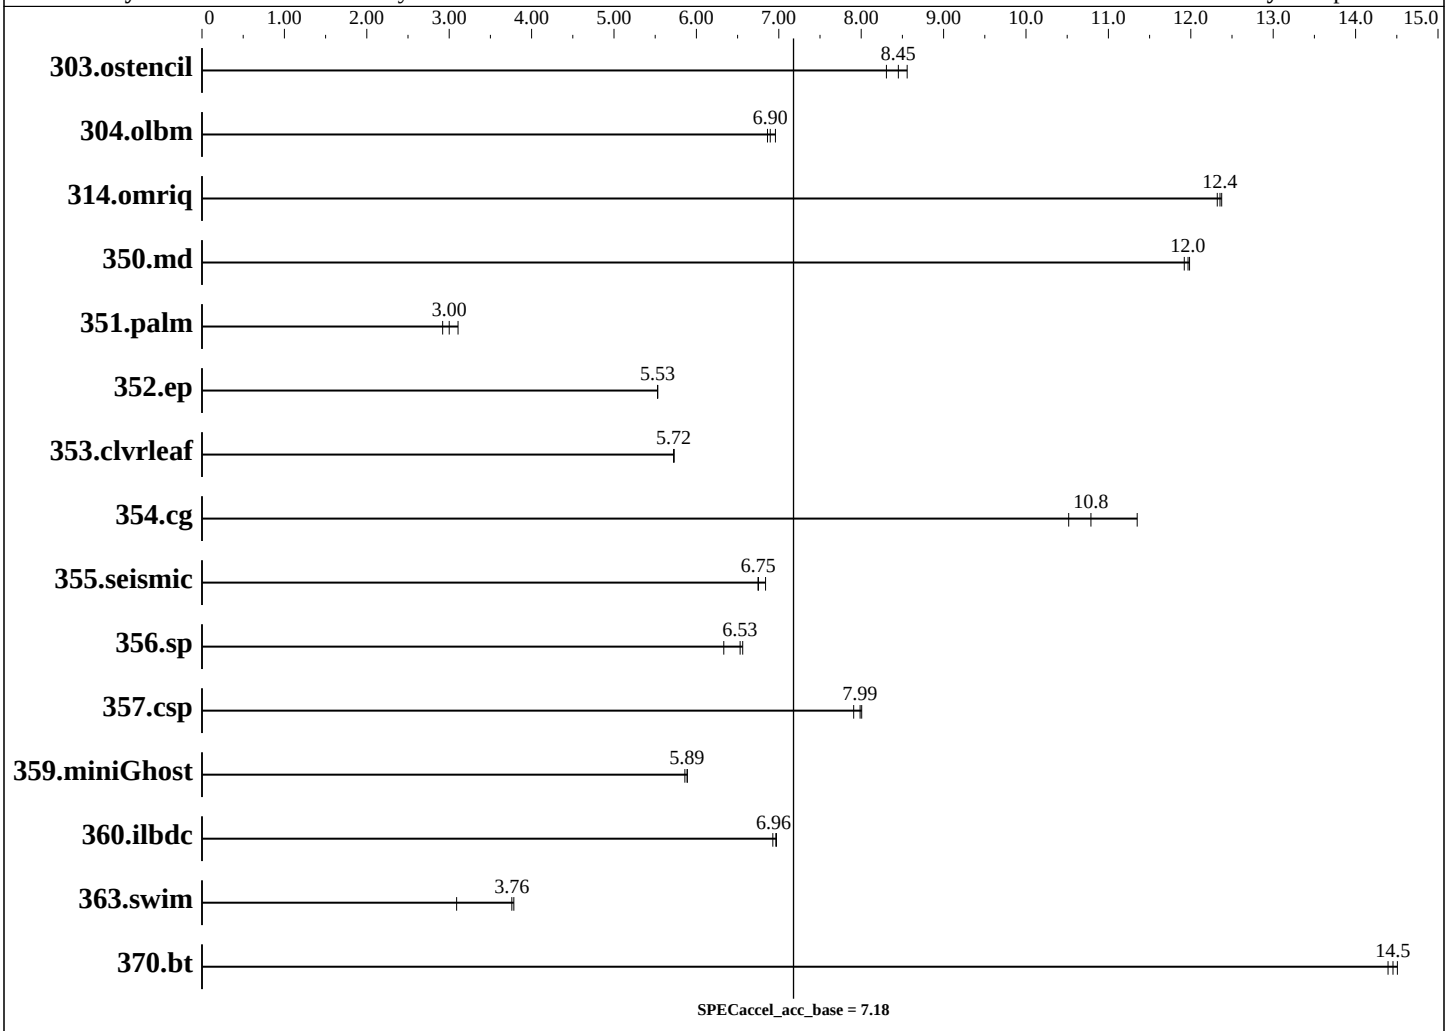

Hardware

CPU Name:	Intel Xeon Gold 6130
CPU Characteristics:	Intel Turbo Boost on, SMT off. 4 of 32 cores allocated to KVM virtual machine.
CPU MHz:	2100
CPU MHz Maximum:	3700
FPU:	Integrated
CPU(s) enabled:	32 cores, 2 chips, 16 cores/chip
CPU(s) orderable:	1,2 chips
Primary Cache:	32 KB I + 32 KB D on chip per core
Secondary Cache:	1 MB I+D on chip per core
L3 Cache:	22 MB I+D on chip per chip

Continued on next page

Accelerator

Accel Model Name:	Tesla V100
Accel Vendor:	NVIDIA Corporation
Accel Name:	Tesla V100X-8Q
Type of Accel:	GPU
Accel Connection:	PCIe
Does Accel Use ECC:	Yes
Accel Description:	half of virtualized V100-SMX2-16GB with NVLink (Persistence Mode enabled) allocated to KVM virtual machine.
Accel Driver:	NVIDIA UNIX x86_64 Kernel Module 418.70

SPEC ACCEL ACC Result

Copyright 2015-2019 Standard Performance Evaluation Corporation

Dell
(Test Sponsor: Indiana University)

Tesla V100X-8Q

PowerEdge C4140 Server (KVM virtual machine)

SPECaccel_acc_peak = Not Run

SPECaccel_acc_base = 7.18

ACCEL license: 3440A
Test sponsor: Indiana University
Tested by: Indiana University

Test date: Jul-2019
Hardware Availability: May-2019
Software Availability: Apr-2019

Hardware (Continued)

Other Cache: None
Memory: 256 GB (16 x 16 GB 2Rx8 PC4-2666V-R)
21 GB (21/256 GB Allocated for KVM)
Disk Subsystem: None
Other Hardware: None

Software

Operating System: VM:
CentOS Linux release 7.6.1810 (Core)
3.10.0-957.12.1.el7.x86_64
Host:
Red Hat Enterprise Linux Server release 7.6
(Maipo) 3.10.0-957.21.3.el7.x86_64
Compiler: PGI Community Edition, Release 19.4
File System: cephfs nfsv4 (ganesha) over 100Gbits/s Ethernet
System State: Run level 3 (multi-user)
Other Software: KVM Version 2.12, CUDA 10.1

Results Table

Benchmark	Base						Peak					
	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
303.ostencil	17.5	8.31	**17.2**	**8.45**	16.9	8.56						
304.olbm	**66.0**	**6.90**	66.3	6.86	65.4	6.96						
314.omriq	**77.4**	**12.4**	77.6	12.3	77.3	12.4						
350.md	21.1	11.9	21.0	12.0	**21.1**	**12.0**						
351.palm	127	2.92	**123**	**3.00**	119	3.11						
352.ep	95.9	5.53	95.8	5.53	**95.9**	**5.53**						
353.clvrleaf	77.7	5.72	77.7	5.73	**77.7**	**5.72**						
354.cg	**37.8**	**10.8**	35.9	11.3	38.8	10.5						
355.seismic	54.1	6.84	**54.8**	**6.75**	54.8	6.75						
356.sp	42.1	6.56	**42.3**	**6.53**	43.6	6.33						
357.csp	33.7	8.01	**33.8**	**7.99**	34.1	7.91						
359.miniGhost	**62.7**	**5.89**	62.7	5.89	63.0	5.86						
360.ilbdc	52.6	6.97	**52.7**	**6.96**	53.0	6.93						
363.swim	60.8	3.78	74.4	3.09	**61.2**	**3.76**						
370.bt	15.4	14.5	**15.4**	**14.5**	15.5	14.4						

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

Copyright 2015-2019 Standard Performance Evaluation Corporation

Dell
(Test Sponsor: Indiana University)

Tesla V100X-8Q

PowerEdge C4140 Server (KVM virtual machine)

SPECaccel_acc_peak = Not Run

SPECaccel_acc_base = 7.18

ACCEL license: 3440A
Test sponsor: Indiana University
Tested by: Indiana University

Test date: Jul-2019
Hardware Availability: May-2019
Software Availability: Apr-2019

Platform Notes

```
Sysinfo program /home/lijunj/junjie_benchmarks/spec/accel-1.2-8q/Docs/sysinfo
$Rev: 6965 $ $Date:: 2015-04-21 #$ c05a7f14b1b1765e3fe1df68447e8a35
running on v100x-8q.novalocal Thu Jul 11 01:01:01 2019

This section contains SUT (System Under Test) info as seen by
some common utilities.  To remove or add to this section, see:
  http://www.spec.org/accel/Docs/config.html#sysinfo

From /proc/cpuinfo
   model name : Intel(R) Xeon(R) Gold 6130 CPU @ 2.10GHz
      4 "physical id"s (chips)
      4 "processors"
   cores, siblings (Caution: counting these is hw and system dependent.  The
   following excerpts from /proc/cpuinfo might not be reliable.  Use with
   caution.)
      cpu cores : 1
      siblings  : 1
      physical 0: cores 0
      physical 1: cores 0
      physical 2: cores 0
      physical 3: cores 0
   cache size : 16384 KB

From /proc/meminfo
   MemTotal:       22458516 kB
   HugePages_Total:       0
   Hugepagesize:       2048 kB

From /etc/*release* /etc/*version*
   centos-release: CentOS Linux release 7.6.1810 (Core)
   centos-release-upstream: Derived from Red Hat Enterprise Linux 7.6 (Source)
   os-release:
      NAME="CentOS Linux"
      VERSION="7 (Core)"
      ID="centos"
      ID_LIKE="rhel fedora"
      VERSION_ID="7"
      PRETTY_NAME="CentOS Linux 7 (Core)"
      ANSI_COLOR="0;31"
      CPE_NAME="cpe:/o:centos:centos:7"
   redhat-release: CentOS Linux release 7.6.1810 (Core)
   system-release: CentOS Linux release 7.6.1810 (Core)
   system-release-cpe: cpe:/o:centos:centos:7

uname -a:
   Linux v100x-8q.novalocal 3.10.0-957.12.1.el7.x86_64 #1 SMP Mon Apr 29
   14:59:59 UTC 2019 x86_64 x86_64 x86_64 GNU/Linux

run-level 3 Jul 10 22:07
```

Continued on next page

Platform Notes (Continued)

```
 SPEC is set to: /home/lijunj/junjie_benchmarks/spec/accel-1.2-8q
    Filesystem                                                Type  Size
    Used Avail Use% Mounted on
    10.255.0.1:/volumes/_nogroup/24ec4401-f96d-40a5-99a5-e96f73257d2f nfs4  128G
    54G   75G  42% /home/lijunj
 Additional information from dmidecode:

    Warning: Use caution when you interpret this section. The 'dmidecode' program

Information from pgaccelinfo:
CUDA Driver Version:           10010
NVRM version:                  NVIDIA UNIX x86_64 Kernel Module  418.70
Device Number:                 0
Device Name:                   GRID V100X-8Q
Device Revision Number:        7.0
Global Memory Size:            8589934592
Number of Multiprocessors:     80
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
Clock Rate:                    1530 MHz
Execution Timeout:             No
Integrated Device:             No
Can Map Host Memory:           Yes
Compute Mode:                  default
Concurrent Kernels:            Yes
ECC Enabled:                   No
Memory Clock Rate:             877 MHz
Memory Bus Width:              4096 bits
L2 Cache Size:                 6291456 bytes
Max Threads Per SMP:           2048
Async Engines:                 2
Unified Addressing:            Yes
Managed Memory:                No
Preemption Supported:          Yes
Cooperative Launch:            Yes
  Multi-Device:                Yes
PGI Default Target:            -ta=tesla:cc70
```

Copyright 2015-2019 Standard Performance Evaluation Corporation

Dell
(Test Sponsor: Indiana University)

Tesla V100X-8Q

PowerEdge C4140 Server (KVM virtual machine)

SPECaccel_acc_peak = Not Run

SPECaccel_acc_base = 7.18

ACCEL license:	3440A	**Test date:**	Jul-2019
Test sponsor:	Indiana University	**Hardware Availability:**	May-2019
Tested by:	Indiana University	**Software Availability:**	Apr-2019

General Notes

```
Four V100-SMX2-16GB GPUs were installed on the host system,
three of them were idle while only one is active for running SPEC Accel.

CPUs and GPUs are connected via PCIe, while the four GPUs are connected
through NVLink.  The NVlink connection is not used in this test.

Stacksize set to 'unlimited':
     ulimit -s unlimited

Spectre & Meltdown:
  Yes: The test sponsor attests, as of date of publication, that CVE-2017-5754 (Meltdown)
   is mitigated in the system as tested and documented.
  Yes: The test sponsor attests, as of date of publication, that CVE-2017-5753 (Spectre variant 1)
   is mitigated in the system as tested and documented.
  Yes: The test sponsor attests, as of date of publication, that CVE-2017-5715 (Spectre variant 2)
   is mitigated in the system as tested and documented.
```

Base Compiler Invocation

C benchmarks:
```
pgcc
```

Fortran benchmarks:
```
pgfortran
```

Benchmarks using both Fortran and C:
```
pgcc pgfortran
```

Base Optimization Flags

C benchmarks:
```
-fast -Mfprelaxed -acc -ta=tesla:cc70 -ta=tesla:cuda10.1
```

Fortran benchmarks:
```
-fast -Mfprelaxed -acc -ta=tesla:cc70 -ta=tesla:cuda10.1
```

Benchmarks using both Fortran and C:

353.clvrleaf: `-fast -Mfprelaxed -acc -ta=tesla:cc70 -ta=tesla:cuda10.1`

359.miniGhost: `-fast -Mfprelaxed -acc -ta=tesla:cc70 -ta=tesla:cuda10.1 -Mnomain`

Copyright 2015-2019 Standard Performance Evaluation Corporation

Dell
(Test Sponsor: Indiana University)

Tesla V100X-8Q

PowerEdge C4140 Server (KVM virtual machine)

SPECaccel_acc_peak = Not Run

SPECaccel_acc_base = 7.18

ACCEL license:	3440A	**Test date:**	Jul-2019
Test sponsor:	Indiana University	**Hardware Availability:**	May-2019
Tested by:	Indiana University	**Software Availability:**	Apr-2019

The flags file that was used to format this result can be browsed at
`https://www.spec.org/accel/flags/pgi2019_flags.html`

You can also download the XML flags source by saving the following link:
`https://www.spec.org/accel/flags/pgi2019_flags.xml`

SPEC ACCEL 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 webmaster@spec.org.

Tested with SPEC ACCEL v1.2.
Report generated on Thu Oct 24 12:55:42 2019 by SPEC ACCEL PS/PDF formatter v1290.
Originally published on 24 October 2019.