SPECaccel®2023 Result

NVIDIA Corporation
Tesla A100-SXM-80GB
DGX-A100

SPECaccel 2023_base = 2.63

SPECaccel 2023_peak = 2.63

accel2023 License: 9045
Test Sponsor: NVIDIA Corporation
Tested by: NVIDIA Corporation

Test Date: Oct-2023
Hardware Availability: Jul-2020
Software Availability: Nov-2023

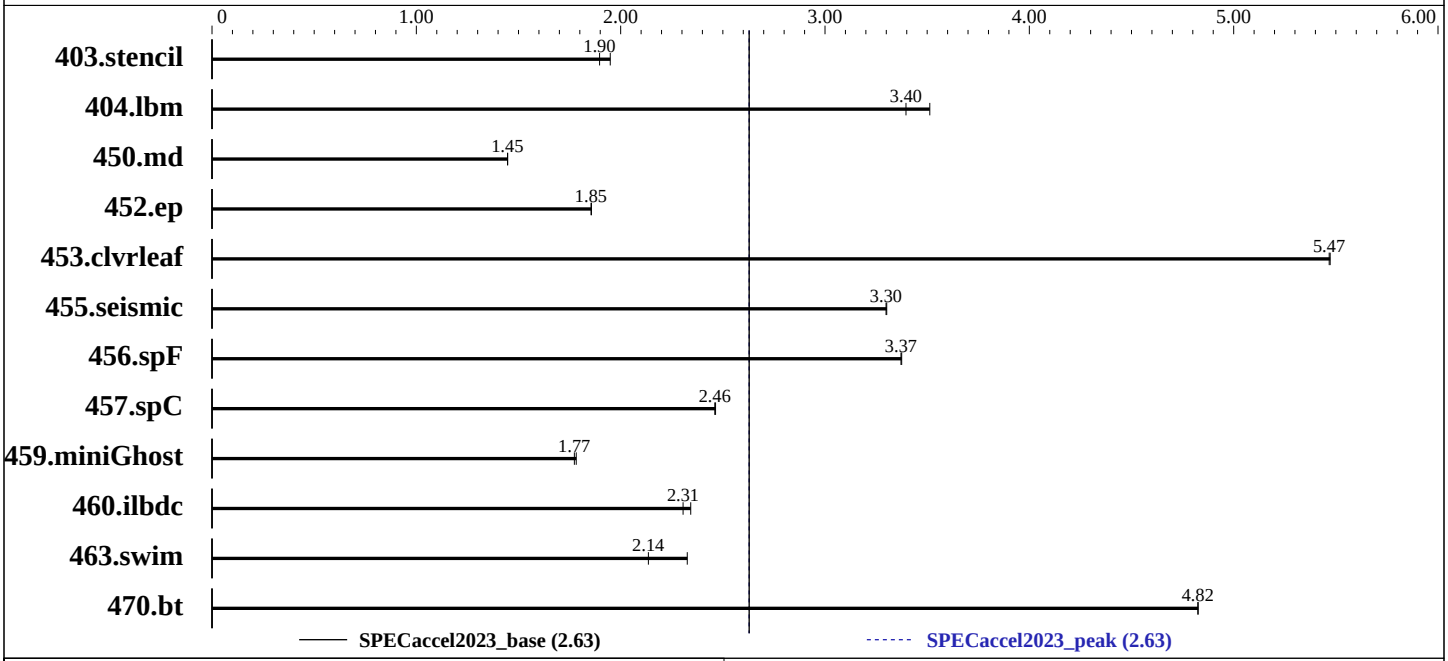

Hardware

CPU Name:	AMD EPYC 7742
Max MHz.:	3400
Nominal:	2250
Enabled:	128 cores, 2 chips, 2 threads/core
Orderable:	2 chips
Cache L1:	32 KB I + 32 KB D on chip per core
L2:	512 KB I+D on chip per core
L3:	256 MB I+D on chip per chip 16 MB shared / 4 cores
Other:	None
Memory:	2 TB (32 x 64 GB 2Rx8 PC4-3200AA-R)
Storage:	OS: 2TB U.2 NVMe SSD drive Internal Storage: 30TB (8x 3.84TB U.2 NVMe SSD drives)
Other:	None
Base Threads Run:	1
Min. Peak Threads:	1
Max. Peak Threads:	1

Accelerator

Accel Model Name:	A100-SXM-80GB
Accel Vendor:	NVIDIA Corporation
Accel Name:	Tesla A100-SXM-80GB
Type of Accel:	GPU
Accel Connection:	NVLINK 3.0, NVSWITCH 2.0 600GB/s
Does Accel Use ECC:	Yes
Accel Description:	See Notes
Accel Driver:	NVIDIA UNIX x86_64 Kernel Module 535.54.03

Software

OS:	Ubuntu 22.04.3 LTS 5.15.0-1031-nvidia
Compiler:	C/Fortran: Version 23.11 of the NVHPC SDK
Firmware:	American Megatrends 1.21
File System:	ext4
System State:	Run level 5 (multi-user)
Other:	None
Base Parallel Model:	LOP

(Continued on next page)

Software (Continued)

Base Threads Run: 1
Peak Parallel Models: LOP
Max. Peak Threads: 1
Min. Peak Threads: 1

Results Table

Benchmark	Base							Peak						
	Model	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Model	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
403.stencil	LOP	226	1.95	**232**	**1.90**			LOP	226	1.95	**232**	**1.90**		
404.lbm	LOP	**134**	**3.40**	130	3.51			LOP	**134**	**3.40**	130	3.51		
450.md	LOP	**415**	**1.45**	415	1.45			LOP	**415**	**1.45**	415	1.45		
452.ep	LOP	223	1.86	**224**	**1.85**			LOP	223	1.86	**224**	**1.85**		
453.clvrleaf	LOP	183	5.47	**183**	**5.47**			LOP	183	5.47	**183**	**5.47**		
455.seismic	LOP	236	3.30	**236**	**3.30**			LOP	236	3.30	**236**	**3.30**		
456.spF	LOP	141	3.38	**141**	**3.37**			LOP	141	3.38	**141**	**3.37**		
457.spC	LOP	**219**	**2.46**	219	2.46			LOP	**219**	**2.46**	219	2.46		
459.miniGhost	LOP	**333**	**1.77**	331	1.78			LOP	**333**	**1.77**	331	1.78		
460.ilbdc	LOP	**241**	**2.31**	237	2.34			LOP	**241**	**2.31**	237	2.34		
463.swim	LOP	**206**	**2.14**	189	2.33			LOP	**206**	**2.14**	189	2.33		
470.bt	LOP	219	4.83	**219**	**4.82**			LOP	219	4.83	**219**	**4.82**		

SPEC accel2023_base = 2.63

SPEC accel2023_peak = 2.63

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

Operating System Notes

```
Shell stacksize set to unlimited via "limit stacksize unlimited"
```

Platform Notes

```
Information from nvaccelinfo
CUDA Driver Version:           12020
NVRM version:                  NVIDIA UNIX x86_64 Kernel Module  535.54.03  Tue Jun  6 22:20:39 UTC 2023
Device Number:                 0
Device Name:                   NVIDIA A100-SXM4-80GB
Device Revision Number:        8.0
Global Memory Size:            84987740160
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
```

(Continued on next page)

Platform Notes (Continued)

```
 Maximum Block Dimensions:      1024, 1024, 64
 Maximum Grid Dimensions:       2147483647 x 65535 x 65535
 Maximum Memory Pitch:          2147483647B
 Texture Alignment:             512B
 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:             1593 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
 Default Target:                cc80

 Sysinfo program /local/home/mcolgrove/ACCELV2/bin/sysinfo
 Rev: r6622 of 2021-04-07 b1a7d5f8f71be5aff70a755cad7211a0
 running on luna Wed Oct 25 14:50:10 2023

 SUT (System Under Test) info as seen by some common utilities.
 For more information on this section, see
    https://www.spec.org/cpu2017/Docs/config.html#sysinfo

 From /proc/cpuinfo
    model name : AMD EPYC 7742 64-Core Processor
       2  "physical id"s (chips)
       256 "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 : 64
       siblings  : 128
       physical 0: cores 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
       25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52
       53 54 55 56 57 58 59 60 61 62 63
       physical 1: cores 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
       25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52
       53 54 55 56 57 58 59 60 61 62 63

 From lscpu from util-linux 2.37.2:
```

(Continued on next page)

Platform Notes (Continued)

```
Architecture:                    x86_64
CPU op-mode(s):                  32-bit, 64-bit
Address sizes:                   43 bits physical, 48 bits virtual
Byte Order:                      Little Endian
CPU(s):                          256
On-line CPU(s) list:             0-255
Vendor ID:                       AuthenticAMD
Model name:                      AMD EPYC 7742 64-Core Processor
CPU family:                      23
Model:                           49
Thread(s) per core:              2
Core(s) per socket:              64
Socket(s):                       2
Stepping:                        0
Frequency boost:                 enabled
CPU max MHz:                     2250.0000
CPU min MHz:                     1500.0000
BogoMIPS:                        4491.45
Flags:                           fpu vme de pse tsc msr pae mce cx8 apic sep mtrr
pge mca cmov pat pse36 clflush mmx fxsr sse sse2 ht syscall nx mmxext fxsr_opt
pdpe1gb rdtscp lm constant_tsc rep_good nopl nonstop_tsc cpuid extd_apicid
aperfmperf rapl pni pclmulqdq monitor ssse3 fma cx16 sse4_1 sse4_2 x2apic movbe
popcnt aes xsave avx f16c rdrand lahf_lm cmp_legacy svm extapic cr8_legacy abm sse4a
misalignsse 3dnowprefetch osvw ibs skinit wdt tce topoext perfctr_core perfctr_nb
bpext perfctr_llc mwaitx cpb cat_l3 cdp_l3 hw_pstate ssbd mba ibrs ibpb stibp
vmmcall fsgsbase bmi1 avx2 smep bmi2 cqm rdt_a rdseed adx smap clflushopt clwb
sha_ni xsaveopt xsavec xgetbv1 cqm_llc cqm_occup_llc cqm_mbm_total cqm_mbm_local
clzero irperf xsaveerptr rdpru wbnoinvd amd_ppin arat npt lbrv svm_lock nrip_save
tsc_scale vmcb_clean flushbyasid decodeassists pausefilter pfthreshold avic
v_vmsave_vmload vgif v_spec_ctrl umip rdpid overflow_recov succor smca sme sev
sev_es
Virtualization:                  AMD-V
L1d cache:                       4 MiB (128 instances)
L1i cache:                       4 MiB (128 instances)
L2 cache:                        64 MiB (128 instances)
L3 cache:                        512 MiB (32 instances)
NUMA node(s):                    8
NUMA node0 CPU(s):               0-15,128-143
NUMA node1 CPU(s):               16-31,144-159
NUMA node2 CPU(s):               32-47,160-175
NUMA node3 CPU(s):               48-63,176-191
NUMA node4 CPU(s):               64-79,192-207
NUMA node5 CPU(s):               80-95,208-223
NUMA node6 CPU(s):               96-111,224-239
NUMA node7 CPU(s):               112-127,240-255
Vulnerability Gather data sampling: Not affected
Vulnerability Itlb multihit:     Not affected
```

(Continued on next page)

Platform Notes (Continued)

```
      Vulnerability L1tf:                Not affected
      Vulnerability Mds:                 Not affected
      Vulnerability Meltdown:            Not affected
      Vulnerability Mmio stale data:     Not affected
      Vulnerability Retbleed:            Vulnerable
      Vulnerability Spec store bypass:   Vulnerable
      Vulnerability Spectre v1:          Vulnerable: __user pointer sanitization and
      usercopy barriers only; no swapgs barriers
      Vulnerability Spectre v2:          Vulnerable, IBPB: disabled, STIBP: disabled,
      PBRSB-eIBRS: Not affected
      Vulnerability Srbds:               Not affected
      Vulnerability Tsx async abort:     Not affected

  From lscpu --cache:
      NAME ONE-SIZE ALL-SIZE WAYS TYPE        LEVEL  SETS PHY-LINE COHERENCY-SIZE
      L1d       32K       4M    8 Data            1    64        1             64
      L1i       32K       4M    8 Instruction     1    64        1             64
      L2       512K      64M    8 Unified         2  1024        1             64
      L3        16M     512M   16 Unified         3 16384        1             64

  /proc/cpuinfo cache data
     cache size : 512 KB

  From numactl --hardware
  WARNING: a numactl 'node' might or might not correspond to a physical chip.
    available: 8 nodes (0-7)
    node 0 cpus: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 128 129 130 131 132 133 134 135 136
    137 138 139 140 141 142 143
    node 0 size: 257831 MB
    node 0 free: 111273 MB
    node 1 cpus: 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 144 145 146 147 148 149
    150 151 152 153 154 155 156 157 158 159
    node 1 size: 257991 MB
    node 1 free: 214768 MB
    node 2 cpus: 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 160 161 162 163 164 165
    166 167 168 169 170 171 172 173 174 175
    node 2 size: 258039 MB
    node 2 free: 214538 MB
    node 3 cpus: 48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63 176 177 178 179 180 181
    182 183 184 185 186 187 188 189 190 191
    node 3 size: 258027 MB
    node 3 free: 218280 MB
    node 4 cpus: 64 65 66 67 68 69 70 71 72 73 74 75 76 77 78 79 192 193 194 195 196 197
    198 199 200 201 202 203 204 205 206 207
    node 4 size: 258039 MB
    node 4 free: 124948 MB
    node 5 cpus: 80 81 82 83 84 85 86 87 88 89 90 91 92 93 94 95 208 209 210 211 212 213
```

(Continued on next page)

Platform Notes (Continued)

```
  214 215 216 217 218 219 220 221 222 223
  node 5 size: 258039 MB
  node 5 free: 213907 MB
  node 6 cpus: 96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111 224 225 226
  227 228 229 230 231 232 233 234 235 236 237 238 239
  node 6 size: 258039 MB
  node 6 free: 193828 MB
  node 7 cpus: 112 113 114 115 116 117 118 119 120 121 122 123 124 125 126 127 240 241
  242 243 244 245 246 247 248 249 250 251 252 253 254 255
  node 7 size: 258027 MB
  node 7 free: 154417 MB
  node distances:
  node   0   1   2   3   4   5   6   7
    0:  10  12  12  12  32  32  32  32
    1:  12  10  12  12  32  32  32  32
    2:  12  12  10  12  32  32  32  32
    3:  12  12  12  10  32  32  32  32
    4:  32  32  32  32  10  12  12  12
    5:  32  32  32  32  12  10  12  12
    6:  32  32  32  32  12  12  10  12
    7:  32  32  32  32  12  12  12  10

From /proc/meminfo
   MemTotal:       2113571332 kB
   HugePages_Total:       0
   Hugepagesize:       2048 kB

/sys/devices/system/cpu/cpu*/cpufreq/scaling_governor has
   performance

/usr/bin/lsb_release -d
   Ubuntu 22.04.3 LTS

From /etc/*release* /etc/*version*
   debian_version: bookworm/sid
   dgx-release:
      DGX_NAME="DGX Server"
      DGX_PRETTY_NAME="NVIDIA DGX Server"
      DGX_SWBUILD_DATE="2020-10-26-11-53-11"
      DGX_SWBUILD_VERSION="5.0.0"
      DGX_COMMIT_ID="7501dff"
      DGX_PLATFORM="DGX Server for DGX A100"
      DGX_SERIAL_NUMBER="1663521001239"

   ec2_version: Ubuntu 20.04.1 LTS (Focal Fossa)
   os-release:
      PRETTY_NAME="Ubuntu 22.04.3 LTS"
```

(Continued on next page)

Platform Notes (Continued)

```
      NAME="Ubuntu"
      VERSION_ID="22.04"
      VERSION="22.04.3 LTS (Jammy Jellyfish)"
      VERSION_CODENAME=jammy
      ID=ubuntu
      ID_LIKE=debian
      HOME_URL="https://www.ubuntu.com/"

uname -a:
   Linux luna 5.15.0-1031-nvidia #31-Ubuntu SMP Tue Aug 15 23:56:08 UTC 2023 x86_64
   x86_64 x86_64 GNU/Linux

Kernel self-reported vulnerability status:

gather_data_sampling:                                  Not affected
CVE-2018-12207 (iTLB Multihit):                        Not affected
CVE-2018-3620 (L1 Terminal Fault):                     Not affected
Microarchitectural Data Sampling:                      Not affected
CVE-2017-5754 (Meltdown):                              Not affected
mmio_stale_data:                                       Not affected
retbleed:                                              Vulnerable
CVE-2018-3639 (Speculative Store Bypass):              Vulnerable
CVE-2017-5753 (Spectre variant 1):                     Vulnerable: __user pointer
                                                     sanitization and usercopy barriers
                                                       only; no swapgs barriers
CVE-2017-5715 (Spectre variant 2):                   Vulnerable, IBPB: disabled, STIBP:
                                                       disabled, PBRSB-eIBRS: Not
                                                       affected
CVE-2020-0543 (Special Register Buffer Data Sampling): Not affected
CVE-2019-11135 (TSX Asynchronous Abort):               Not affected

run-level 5 Sep 8 09:26

SPEC is set to: /local/home/mcolgrove/ACCELV2
   Filesystem     Type  Size  Used Avail Use% Mounted on
   /dev/md0       ext4  1.8T  1.1T  597G  65% /

From /sys/devices/virtual/dmi/id
    Vendor:         NVIDIA
    Product:        DGXA100 920-23687-2530-000
    Product Family: DGX

Cannot run dmidecode; consider saying (as root)
   chmod +s /usr/sbin/dmidecode

BIOS:
   BIOS Vendor:       American Megatrends Inc.
```

(Continued on next page)

Platform Notes (Continued)

```
    BIOS Version:      1.21
    BIOS Date:         03/09/2023

 (End of data from sysinfo program)
```

Compiler Version Notes

```
==============================================================================
C          | 457.spC(base)
------------------------------------------------------------------------------
/usr/bin/ld: /usr/lib/x86_64-linux-gnu/crt1.o: in function `_start':
(.text+0x1b): undefined reference to `main'
pgacclnk: child process exit status 1: /usr/bin/ld
nvc 23.11-0 64-bit target on x86-64 Linux -tp znver2
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
C          | 403.stencil(base) 404.lbm(base) 452.ep(base) 470.bt(base)
------------------------------------------------------------------------------
nvc 23.11-0 64-bit target on x86-64 Linux -tp znver2
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
C          | 457.spC(base)
------------------------------------------------------------------------------
/usr/bin/ld: /usr/lib/x86_64-linux-gnu/crt1.o: in function `_start':
(.text+0x1b): undefined reference to `main'
pgacclnk: child process exit status 1: /usr/bin/ld
nvc 23.11-0 64-bit target on x86-64 Linux -tp znver2
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
C          | 403.stencil(base) 404.lbm(base) 452.ep(base) 470.bt(base)
------------------------------------------------------------------------------
nvc 23.11-0 64-bit target on x86-64 Linux -tp znver2
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
```

(Continued on next page)

Compiler Version Notes (Continued)

```
Fortran    | 450.md(base) 455.seismic(base) 456.spF(base) 460.ilbdc(base)
           | 463.swim(base)
------------------------------------------------------------------------------
nvfortran 23.11-0 64-bit target on x86-64 Linux -tp znver2
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
Fortran, C | 453.clvrleaf(base) 459.miniGhost(base)
------------------------------------------------------------------------------
nvfortran 23.11-0 64-bit target on x86-64 Linux -tp znver2
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
nvc 23.11-0 64-bit target on x86-64 Linux -tp znver2
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------
```

Base Compiler Invocation

C benchmarks:
nvc

Fortran benchmarks:
nvfortran

Benchmarks using both Fortran and C:
nvfortran nvc

Base Portability Flags

403.stencil: -DSPEC_NO_NOTHING
457.spC: -mcmodel=medium -Wl,--no-relax

Base Optimization Flags

C benchmarks:
-Ofast -mp=gpu -Mfprelaxed -Mstack_arrays -static-nvidia

(Continued on next page)

Base Optimization Flags (Continued)

Fortran benchmarks:
`-Ofast -mp=gpu -Mfprelaxed -Mstack_arrays -static-nvidia`

Benchmarks using both Fortran and C:

453.clvrleaf: `-Ofast -mp=gpu -Mfprelaxed -Mstack_arrays -static-nvidia`

459.miniGhost: `-Mnomain -Ofast -mp=gpu -Mfprelaxed -Mstack_arrays -static-nvidia`

Peak Optimization Flags

C benchmarks:

403.stencil: `basepeak = yes`

404.lbm: `basepeak = yes`

452.ep: `basepeak = yes`

457.spC: `basepeak = yes`

470.bt: `basepeak = yes`

Fortran benchmarks:

450.md: `basepeak = yes`

455.seismic: `basepeak = yes`

456.spF: `basepeak = yes`

460.ilbdc: `basepeak = yes`

463.swim: `basepeak = yes`

Benchmarks using both Fortran and C:

453.clvrleaf: `basepeak = yes`

459.miniGhost: `basepeak = yes`

The flags file that was used to format this result can be browsed at
`http://www.spec.org/accel2023/flags/nv2023_flags_v2.html`

You can also download the XML flags source by saving the following link:
`http://www.spec.org/accel2023/flags/nv2023_flags_v2.xml`

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

Tested with SPECaccel2023 v2.0.17 on 2023-10-25 17:50:10-0400.
Report generated on 2023-12-06 13:07:07 by accel2023 PDF formatter v112.
Originally published on 2023-11-08.