# SPECaccel®2023 Result



## NVIDIA Corporation
## Tesla H100 96GB
## MGX-GH200

**SPECaccel 2023_base = 4.49**

**SPECaccel 2023_peak = 4.49**

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

**Test Date:** Oct-2023
**Hardware Availability:** Sep-2023
**Software Availability:** Nov-2023

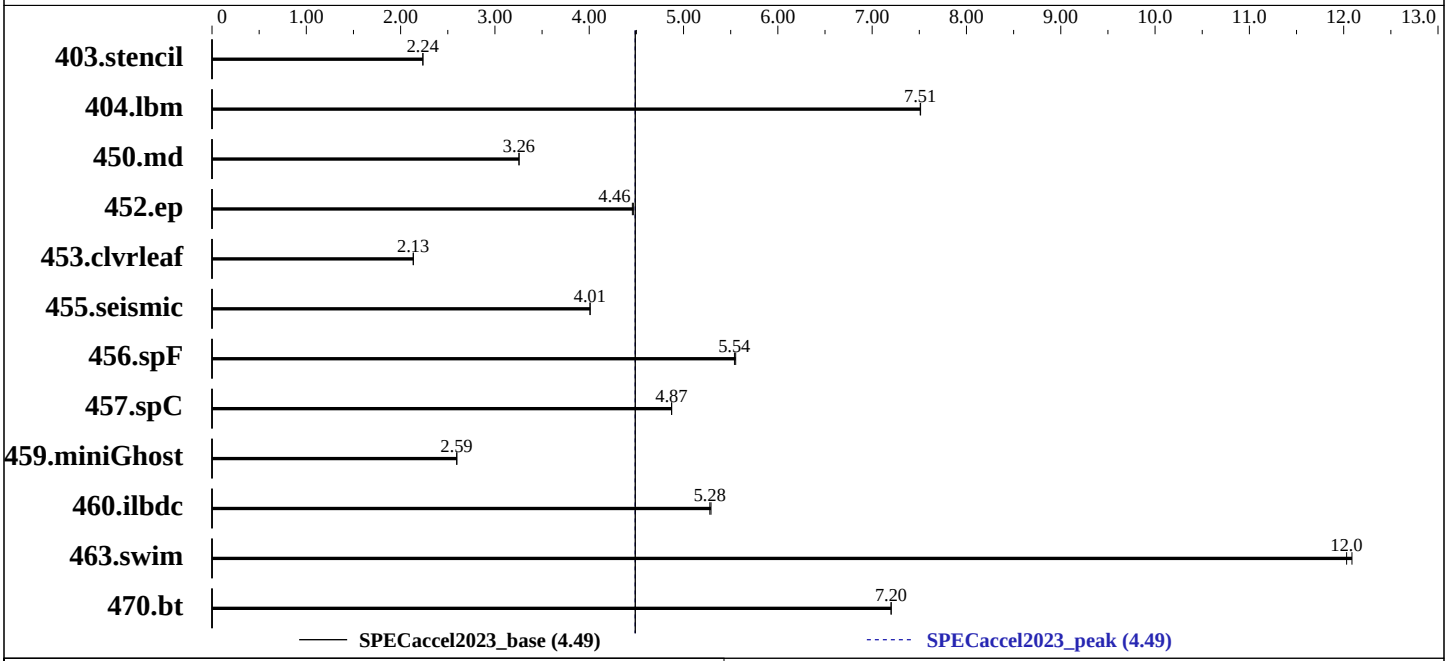


### Hardware

| | |
|---|---|
| CPU Name: | Grace |
| Max MHz.: | 3100 |
| Nominal: | 2900 |
| Enabled: | 72 cores, 1 chip |
| Orderable: | 1 chip |
| Cache L1: | 64 KB I + 64 KB D on chip per core |
| L2: | 1 MB I+D on chip per core |
| L3: | 117 MB I+D on chip per chip |
| Other: | None |
| Memory: | 480 GB (1 x 480 GB LPDDR5) |
| Storage: | 1.8 TB NVMe |
| Other: | None |
| Base Threads Run: | 1 |
| Min. Peak Threads: | 1 |
| Max. Peak Threads: | 1 |

### Accelerator

| | |
|---|---|
| Accel Model Name: | H100 96GB |
| Accel Vendor: | NVIDIA Corporation |
| Accel Name: | Tesla H100 96GB |
| Type of Accel: | GPU |
| Accel Connection: | NVLink-C2C |
| Does Accel Use ECC: | Yes |
| Accel Description: | Grace Hopper Superchip w/ 96GB device memory |
| Accel Driver: | NVIDIA UNIX Open Kernel Module for aarch64 525.105.17 |

### Software

| | |
|---|---|
| OS: | Ubuntu 22.04.3 LTS<br>6.2.0-1010-nvidia-64k |
| Compiler: | C/Fortran: Version 23.11 NVHPC SDK |
| Firmware: | NVIDIA 00010001 09/21/2023 |
| File System: | ext4 |
| System State: | Run level 3 (multi-user) |
| Other: | None |
| Base Parallel Model: | TGT |
| Base Threads Run: | 1 |
| Peak Parallel Models: | TGT |

**(Continued on next page)**

### Software (Continued)

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 | TGT | **197** | **2.24** | 197 | 2.24 | | | TGT | **197** | **2.24** | 197 | 2.24 | | |
| 404.lbm | TGT | **60.6** | **7.51** | 60.6 | 7.51 | | | TGT | **60.6** | **7.51** | 60.6 | 7.51 | | |
| 450.md | TGT | 184 | 3.26 | **184** | **3.26** | | | TGT | 184 | 3.26 | **184** | **3.26** | | |
| 452.ep | TGT | **93.1** | **4.46** | 92.9 | 4.47 | | | TGT | **93.1** | **4.46** | 92.9 | 4.47 | | |
| 453.clvrleaf | TGT | **469** | **2.13** | 468 | 2.13 | | | TGT | **469** | **2.13** | 468 | 2.13 | | |
| 455.seismic | TGT | 195 | 4.01 | **195** | **4.01** | | | TGT | 195 | 4.01 | **195** | **4.01** | | |
| 456.spF | TGT | 85.5 | 5.55 | **85.7** | **5.54** | | | TGT | 85.5 | 5.55 | **85.7** | **5.54** | | |
| 457.spC | TGT | 111 | 4.87 | **111** | **4.87** | | | TGT | 111 | 4.87 | **111** | **4.87** | | |
| 459.miniGhost | TGT | 227 | 2.60 | **227** | **2.59** | | | TGT | 227 | 2.60 | **227** | **2.59** | | |
| 460.ilbdc | TGT | 105 | 5.29 | **105** | **5.28** | | | TGT | 105 | 5.29 | **105** | **5.28** | | |
| 463.swim | TGT | 36.4 | 12.1 | **36.6** | **12.0** | | | TGT | 36.4 | 12.1 | **36.6** | **12.0** | | |
| 470.bt | TGT | **147** | **7.20** | 147 | 7.20 | | | TGT | **147** | **7.20** | 147 | 7.20 | | |

**SPEC accel2023_base = 4.49**

**SPEC accel2023_peak = 4.49**

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"
```

## General Notes

```
Environment variables set by runaccel before the start of the run:
LD_LIBRARY_PATH = "/var/data0/sandbox/nvuser/mcolgrove/nvhpc_2311_lib"
Set to the location of the NVHPC compiler runtime libraries.
```

## Platform Notes

```
 Sysinfo program /var/data0/sandbox/nvuser/mcolgrove/ACCELv2/bin/sysinfo
 Rev: r6622 of 2021-04-07 b1a7d5f8f71be5aff70a755cad7211a0
 running on LegoCG1-96GB-QS-102 Tue Oct 17 02:41:09 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
```

**(Continued on next page)**

## Platform Notes (Continued)

```
 *
 * Did not identify cpu model.  If you would
 * like to write your own sysinfo program, see
 * www.spec.org/cpu2017/config.html#sysinfo
 *
 *
 * 0 "physical id" tags found.  Perhaps this is an older system,
 * or a virtualized system.  Not attempting to guess how to
 * count chips/cores for this system.
 *
       72 "processors"
    cores, siblings (Caution: counting these is hw and system dependent. The following
    excerpts from /proc/cpuinfo might not be reliable.  Use with caution.)

 From lscpu from util-linux 2.37.2:
      Architecture:                         aarch64
      CPU op-mode(s):                       64-bit
      Byte Order:                           Little Endian
      CPU(s):                               72
      On-line CPU(s) list:                  0-71
      Vendor ID:                            ARM
      Model:                                0
      Thread(s) per core:                   1
      Core(s) per socket:                   72
      Socket(s):                            1
      Stepping:                             r0p0
      Frequency boost:                      disabled
      CPU max MHz:                          3429.0000
      CPU min MHz:                          81.0000
      BogoMIPS:                             2000.00
      Flags:                                fp asimd evtstrm aes pmull sha1 sha2 crc32
      atomics fphp asimdhp cpuid asimdrdm jscvt fcma lrcpc dcpop sha3 sm3 sm4 asimddp
      sha512 sve asimdfhm dit uscat ilrcpc flagm ssbs sb paca pacg dcpodp sve2 sveaes
      svepmull svebitperm svesha3 svesm4 flagm2 frint svei8mm svebf16 i8mm bf16 dgh bti
      L1d cache:                            4.5 MiB (72 instances)
      L1i cache:                            4.5 MiB (72 instances)
      L2 cache:                             72 MiB (72 instances)
      L3 cache:                             114 MiB (1 instance)
      NUMA node(s):                         9
      NUMA node0 CPU(s):                    0-71
      NUMA node1 CPU(s):
      NUMA node2 CPU(s):
      NUMA node3 CPU(s):
      NUMA node4 CPU(s):
      NUMA node5 CPU(s):
      NUMA node6 CPU(s):
      NUMA node7 CPU(s):
```

**(Continued on next page)**

## Platform Notes (Continued)

```
      NUMA node8 CPU(s):
      Vulnerability Gather data sampling: Not affected
      Vulnerability Itlb multihit:        Not affected
      Vulnerability L1tf:                 Not affected
      Vulnerability Mds:                  Not affected
      Vulnerability Meltdown:             Not affected
      Vulnerability Mmio stale data:      Not affected
      Vulnerability Retbleed:             Not affected
      Vulnerability Spec store bypass:    Mitigation; Speculative Store Bypass disabled
      via prctl
      Vulnerability Spectre v1:           Mitigation; __user pointer sanitization
      Vulnerability Spectre v2:           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       64K     4.5M    4 Data            1    256                    64
      L1i       64K     4.5M    4 Instruction     1    256                    64
      L2         1M      72M    8 Unified         2   2048                    64
      L3       114M     114M   12 Unified         3 155648                    64

  From numactl --hardware
  WARNING: a numactl 'node' might or might not correspond to a physical chip.
    available: 9 nodes (0-8)
    node 0 cpus: 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 64 65 66 67 68 69 70 71
    node 0 size: 490308 MB
    node 0 free: 76199 MB
    node 1 cpus:
    node 1 size: 97280 MB
    node 1 free: 96959 MB
    node 2 cpus:
    node 2 size: 0 MB
    node 2 free: 0 MB
    node 3 cpus:
    node 3 size: 0 MB
    node 3 free: 0 MB
    node 4 cpus:
    node 4 size: 0 MB
    node 4 free: 0 MB
    node 5 cpus:
    node 5 size: 0 MB
    node 5 free: 0 MB
    node 6 cpus:
    node 6 size: 0 MB
```

**(Continued on next page)**

## Platform Notes (Continued)

```
   node 6 free: 0 MB
   node 7 cpus:
   node 7 size: 0 MB
   node 7 free: 0 MB
   node 8 cpus:
   node 8 size: 0 MB
   node 8 free: 0 MB
   node distances:
   node   0   1   2   3   4   5   6   7   8
     0:  10  80  80  80  80  80  80  80  80
     1:  80  10  255  255  255  255  255  255  255
     2:  80  255  10  255  255  255  255  255  255
     3:  80  255  255  10  255  255  255  255  255
     4:  80  255  255  255  10  255  255  255  255
     5:  80  255  255  255  255  10  255  255  255
     6:  80  255  255  255  255  255  10  255  255
     7:  80  255  255  255  255  255  255  10  255
     8:  80  255  255  255  255  255  255  255  10

 From /proc/meminfo
    MemTotal:       601690240 kB
    HugePages_Total:       0
    Hugepagesize:     524288 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
    l4t-release:
       L4T_NAME="L4T Server"
       L4T_PRETTY_NAME="NVIDIA L4T Server"
       L4T_SWBUILD_DATE="2023-09-12-19-18-30"
       L4T_SWBUILD_VERSION="6.0.0-nvidia"
       L4T_COMMIT_ID="3244539"
       L4T_PLATFORM=""
       L4T_SERIAL_NUMBER=""
    os-release:
       PRETTY_NAME="Ubuntu 22.04.3 LTS"
       NAME="Ubuntu"
       VERSION_ID="22.04"
       VERSION="22.04.3 LTS (Jammy Jellyfish)"
       VERSION_CODENAME=jammy
       ID=ubuntu
```

**(Continued on next page)**

## Platform Notes (Continued)

```
      ID_LIKE=debian
      HOME_URL="https://www.ubuntu.com/"

uname -a:
   Linux LegoCG1-96GB-QS-102 6.2.0-1010-nvidia-64k #10-Ubuntu SMP PREEMPT_DYNAMIC Wed Aug
   30 06:23:50 UTC 2023 aarch64 aarch64 aarch64 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:                                              Not affected
CVE-2018-3639 (Speculative Store Bypass):              Mitigation: Speculative Store
                                                       Bypass disabled via prctl
CVE-2017-5753 (Spectre variant 1):                     Mitigation: __user pointer
                                                       sanitization
CVE-2017-5715 (Spectre variant 2):                     Not affected
CVE-2020-0543 (Special Register Buffer Data Sampling): Not affected
CVE-2019-11135 (TSX Asynchronous Abort):               Not affected

run-level 3 Oct 14 00:07

SPEC is set to: /var/data0/sandbox/nvuser/mcolgrove/ACCELv2
   Filesystem     Type  Size  Used Avail Use% Mounted on
   /dev/nvme0n1p2 ext4  1.8T  1.5T  199G  89% /

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

BIOS:
   BIOS Vendor:       NVIDIA
   BIOS Version:              00010001
   BIOS Date:         20230921

(End of data from sysinfo program)
```

## Compiler Version Notes

```
==============================================================================
C          | 403.stencil(base) 404.lbm(base) 452.ep(base) 470.bt(base)
------------------------------------------------------------------------------
nvc Rel Dev-r238862 linuxarm64 target on aarch64 Linux -tp neoverse-v2
```

**(Continued on next page)**

## Compiler Version Notes (Continued)

```
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
C           | 457.spC(base)
------------------------------------------------------------------------------
nvc-Warning-Only small, tiny and large code models are allowed on AArch64.
  The code model large will be used.
/usr/bin/ld: /usr/lib/aarch64-linux-gnu/crt1.o: in function `__wrap_main':
(.text+0x38): undefined reference to `main'
pgacclnk: child process exit status 1: /usr/bin/ld
nvc Rel Dev-r238862 linuxarm64 target on aarch64 Linux -tp neoverse-v2
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 Rel Dev-r238862 linuxarm64 target on aarch64 Linux -tp neoverse-v2
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
C           | 457.spC(base)
------------------------------------------------------------------------------
nvc-Warning-Only small, tiny and large code models are allowed on AArch64.
  The code model large will be used.
/usr/bin/ld: /usr/lib/aarch64-linux-gnu/crt1.o: in function `__wrap_main':
(.text+0x38): undefined reference to `main'
pgacclnk: child process exit status 1: /usr/bin/ld
nvc Rel Dev-r238862 linuxarm64 target on aarch64 Linux -tp neoverse-v2
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------

==============================================================================
Fortran     | 450.md(base) 455.seismic(base) 456.spF(base) 460.ilbdc(base)
            | 463.swim(base)
------------------------------------------------------------------------------
nvfortran Rel Dev-r238862 linuxarm64 target on aarch64 Linux -tp neoverse-v2
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
------------------------------------------------------------------------------
```

**(Continued on next page)**

## Compiler Version Notes (Continued)

```
==============================================================================
Fortran, C | 453.clvrleaf(base) 459.miniGhost(base)
------------------------------------------------------------------------------
nvfortran Rel Dev-r238862 linuxarm64 target on aarch64 Linux -tp neoverse-v2
NVIDIA Compilers and Tools
Copyright (c) 2023, NVIDIA CORPORATION & AFFILIATES.  All rights reserved.
nvc Rel Dev-r238862 linuxarm64 target on aarch64 Linux -tp neoverse-v2
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
-march=neoverse-v2

Fortran benchmarks:
-Ofast -mp=gpu -Mfprelaxed -Mstack_arrays -static-nvidia
-march=neoverse-v2

Benchmarks using both Fortran and C:

453.clvrleaf: -Ofast -mp=gpu -Mfprelaxed -Mstack_arrays -static-nvidia
-march=neoverse-v2

**(Continued on next page)**

## Base Optimization Flags (Continued)

459.miniGhost: `-Mnomain -Ofast -mp=gpu -Mfprelaxed -Mstack_arrays -static-nvidia -march=neoverse-v2`

## 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_arm.html`

You can also download the XML flags source by saving the following link:
`http://www.spec.org/accel2023/flags/nv2023_flags_v2_arm.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 2023-10-16 22:41:08-0400.
Report generated on 2023-12-06 13:07:20 by accel2023 PDF formatter v112.
Originally published on 2023-11-08.