SPEC CPU®2017 Integer Speed Result

Copyright 2017-2024 Standard Performance Evaluation Corporation

Hewlett Packard Enterprise

(Test Sponsor: HPE)

ProLiant DL380 Gen11
(2.40 GHz, Intel Xeon Gold 6448H)

SPECspeed®2017_int_base = 15.9

SPECspeed®2017_int_peak = 16.1

CPU2017 License: 3
Test Sponsor: HPE
Tested by: HPE

Test Date: Jun-2023
Hardware Availability: May-2023
Software Availability: Dec-2022

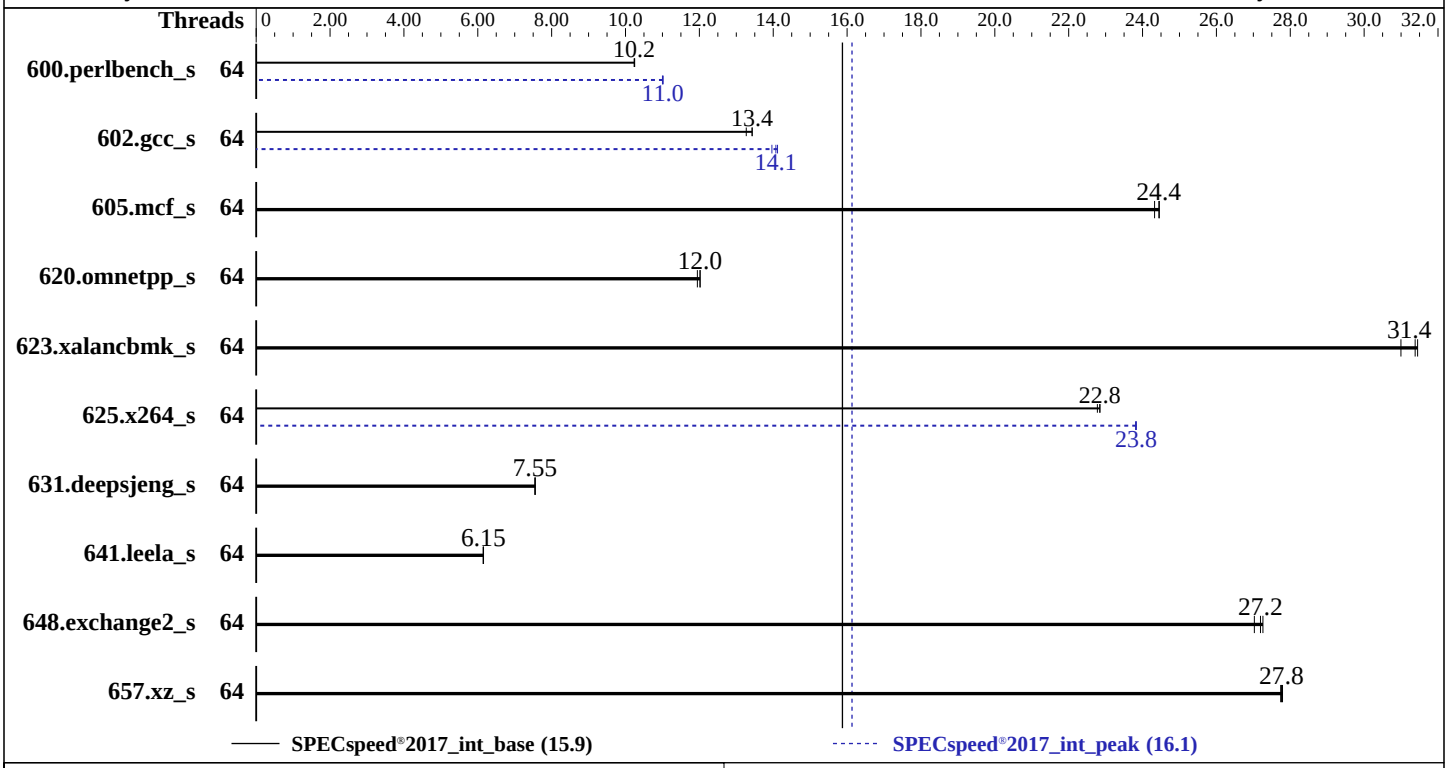

Hardware

CPU Name:	Intel Xeon Gold 6448H
Max MHz:	4100
Nominal:	2400
Enabled:	64 cores, 2 chips
Orderable:	1, 2 chip(s)
Cache L1:	32 KB I + 48 KB D on chip per core
L2:	2 MB I+D on chip per core
L3:	60 MB I+D on chip per chip
Other:	None
Memory:	512 GB (16 x 32 GB 2Rx8 PC5-4800B-R)
Storage:	1 x 480 GB SATA SSD
Other:	None

Software

OS:	Red Hat Enterprise Linux release 9.0 (Plow) Kernel 5.14.0-70.13.1.el9_0.x86_64
Compiler:	C/C++: Version 2023.0 of Intel oneAPI DPC++/C++ Compiler for Linux; Fortran: Version 2023.0 of Intel Fortran Compiler for Linux
Parallel:	Yes
Firmware:	HPE BIOS Version v1.40 05/18/2023 released May-2023
File System:	xfs
System State:	Run level 3 (multi-user)
Base Pointers:	64-bit
Peak Pointers:	64-bit
Other:	jemalloc memory allocator V5.0.1
Power Management:	BIOS and OS set to prefer performance at the cost of additional power usage

SPEC CPU®2017 Integer Speed Result

Copyright 2017-2024 Standard Performance Evaluation Corporation

Hewlett Packard Enterprise
(Test Sponsor: HPE)
ProLiant DL380 Gen11
(2.40 GHz, Intel Xeon Gold 6448H)

SPECspeed®2017_int_base = 15.9

SPECspeed®2017_int_peak = 16.1

CPU2017 License: 3
Test Sponsor: HPE
Tested by: HPE

Test Date: Jun-2023
Hardware Availability: May-2023
Software Availability: Dec-2022

Results Table

Benchmark	Base							Peak						
	Threads	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Threads	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
600.perlbench_s	64	173	10.2	**173**	**10.2**	173	10.2	64	161	11.0	**161**	**11.0**	161	11.0
602.gcc_s	64	297	13.4	300	13.3	**297**	**13.4**	64	282	14.1	285	14.0	**283**	**14.1**
605.mcf_s	64	193	24.5	194	24.3	**193**	**24.4**	64	193	24.5	194	24.3	**193**	**24.4**
620.omnetpp_s	64	137	11.9	136	12.0	**136**	**12.0**	64	137	11.9	136	12.0	**136**	**12.0**
623.xalancbmk_s	64	45.1	31.4	45.7	31.0	**45.2**	**31.4**	64	45.1	31.4	45.7	31.0	**45.2**	**31.4**
625.x264_s	64	77.4	22.8	**77.2**	**22.8**	77.2	22.8	64	**74.0**	**23.8**	74.1	23.8	74.0	23.8
631.deepsjeng_s	64	190	7.56	190	7.55	**190**	**7.55**	64	190	7.56	190	7.55	**190**	**7.55**
641.leela_s	64	277	6.16	**277**	**6.15**	278	6.14	64	277	6.16	**277**	**6.15**	278	6.14
648.exchange2_s	64	**108**	**27.2**	109	27.0	108	27.3	64	**108**	**27.2**	109	27.0	108	27.3
657.xz_s	64	222	27.8	223	27.7	**223**	**27.8**	64	222	27.8	223	27.7	**223**	**27.8**

SPECspeed®2017_int_base = 15.9

SPECspeed®2017_int_peak = 16.1

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

Compiler Notes

```
SPEC has ruled that the compiler used for this result was performing a compilation
that specifically improves the performance of the 523.xalancbmk_r / 623.xalancbmk_s
benchmarks using a priori knowledge of the SPEC code and dataset to perform a
transformation that has narrow applicability.

In order to encourage optimizations that have wide applicability (see rule 1.4
https://www.spec.org/cpu2017/Docs/runrules.html#rule_1.4), SPEC will no longer
publish results using this optimization.

This result is left in the SPEC results database for historical reference.
```

Submit Notes

```
The numactl mechanism was used to bind copies to processors. The config file option 'submit'
was used to generate numactl commands to bind each copy to a specific processor.
For details, please see the config file.
```

Operating System Notes

```
Stack size set to unlimited using "ulimit -s unlimited"
Transparent Huge Pages enabled by default
Prior to runcpu invocation
Filesystem page cache synced and cleared with:
sync; echo 3>       /proc/sys/vm/drop_caches
runcpu command invoked through numactl i.e.:
numactl --interleave=all runcpu <etc>
tuned-adm profile was set to Throughput-Performance using "tuned-adm profile throughput-performance"
```

Environment Variables Notes

```
Environment variables set by runcpu before the start of the run:
KMP_AFFINITY = "granularity=fine,scatter"
LD_LIBRARY_PATH = "/home/cpu2017/lib/intel64:/home/cpu2017/je5.0.1-64"
MALLOC_CONF = "retain:true"
OMP_STACKSIZE = "192M"
```

General Notes

```
 Binaries compiled on a system with 2x Intel Xeon Platinum 8280M CPU + 384GB RAM
 memory using Red Hat Enterprise Linux 8.0
NA: 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.
 jemalloc, a general purpose malloc implementation
 built with the RedHat Enterprise 7.5, and the system compiler gcc 4.8.5
 sources available from jemalloc.net or https://github.com/jemalloc/jemalloc/releases
```

Platform Notes

```
The system ROM used for this result contains Intel microcode version 0x2b0004a1 for
the Intel Xeon Gold 6448H processor.
BIOS Configuration:
 Workload Profile set to General Peak Frequency Compute
 Thermal Configuration set to Maximum Cooling
 Intel Hyper-Threading set to Disabled
 Memory Patrol Scrubbing set to Disabled
 Last Level Cache (LLC) Prefetch set to Enabled
 Last Level Cache (LLC) Dead Line Allocation set to Disabled
 Enhanced Processor Performance Profile set to Aggressive
 Dead Block Predictor set to Enabled
 Sub-NUMA Clustering set to Enabled SNC2(2-clusters)
 Workload Profile set to Custom
  Adjacent Sector Prefetch set to Disabled
  Minimum Processor Idle Power Package C-State set to No Package State

 Sysinfo program /home/cpu2017/bin/sysinfo
 Rev: r6732 of 2022-11-07 fe91c89b7ed5c36ae2c92cc097bec197
 running on localhost.localdomain Wed Jun 14 04:05:12 2023

 SUT (System Under Test) info as seen by some common utilities.

 ----------------------------------------------------------
 Table of contents
 ----------------------------------------------------------
  1. uname -a
  2. w
  3. Username
  4. ulimit -a
  5. sysinfo process ancestry
  6. /proc/cpuinfo
  7. lscpu
  8. numactl --hardware
  9. /proc/meminfo
 10. who -r
```

(Continued on next page)

Platform Notes (Continued)

```
 11. Systemd service manager version: systemd 250 (250-6.el9_0)
 12. Failed units, from systemctl list-units --state=failed
 13. Services, from systemctl list-unit-files
 14. Linux kernel boot-time arguments, from /proc/cmdline
 15. cpupower frequency-info
 16. tuned-adm active
 17. sysctl
 18. /sys/kernel/mm/transparent_hugepage
 19. /sys/kernel/mm/transparent_hugepage/khugepaged
 20. OS release
 21. Disk information
 22. /sys/devices/virtual/dmi/id
 23. dmidecode
 24. BIOS
------------------------------------------------------------

------------------------------------------------------------
1. uname -a
  Linux localhost.localdomain 5.14.0-70.13.1.el9_0.x86_64 #1 SMP PREEMPT Thu Apr 14 12:42:38 EDT 2022 x86_64
  x86_64 x86_64 GNU/Linux

------------------------------------------------------------
2. w
   04:05:12 up 0 min,  0 users,  load average: 1.62, 0.61, 0.22
  USER     TTY        LOGIN@   IDLE   JCPU   PCPU WHAT

------------------------------------------------------------
3. Username
  From environment variable $USER:  root

------------------------------------------------------------
4. ulimit -a
  real-time non-blocking time  (microseconds, -R) unlimited
  core file size              (blocks, -c) 0
  data seg size               (kbytes, -d) unlimited
  scheduling priority                 (-e) 0
  file size                   (blocks, -f) unlimited
  pending signals                     (-i) 2062754
  max locked memory           (kbytes, -l) 64
  max memory size             (kbytes, -m) unlimited
  open files                          (-n) 1024
  pipe size                (512 bytes, -p) 8
  POSIX message queues         (bytes, -q) 819200
  real-time priority                  (-r) 0
  stack size                  (kbytes, -s) unlimited
  cpu time                   (seconds, -t) unlimited
  max user processes                  (-u) 2062754
  virtual memory              (kbytes, -v) unlimited
  file locks                          (-x) unlimited

------------------------------------------------------------
5. sysinfo process ancestry
 /usr/lib/systemd/systemd --switched-root --system --deserialize 27
 sshd: /usr/sbin/sshd -D [listener] 0 of 10-100 startups
 sshd: root [priv]
 sshd: root@notty
 bash -c cd $SPEC/ && $SPEC/intspeed_spr.sh
 runcpu --nobuild --action validate --define default-platform-flags -c
   ic2023.0-lin-sapphirerapids-speed-20221201.cfg --define cores=64 --tune base,peak -o all --define
   intspeedaffinity --define drop_caches intspeed
```

(Continued on next page)

Platform Notes (Continued)

```
 runcpu --nobuild --action validate --define default-platform-flags --configfile
   ic2023.0-lin-sapphirerapids-speed-20221201.cfg --define cores=64 --tune base,peak --output_format all
   --define intspeedaffinity --define drop_caches --nopower --runmode speed --tune base:peak --size refspeed
   intspeed --nopreenv --note-preenv --logfile $SPEC/tmp/CPU2017.001/templogs/preenv.intspeed.001.0.log
   --lognum 001.0 --from_runcpu 2
 specperl $SPEC/bin/sysinfo
$SPEC = /home/cpu2017

------------------------------------------------------------
6. /proc/cpuinfo
    model name      : Intel(R) Xeon(R) Gold 6448H
    vendor_id       : GenuineIntel
    cpu family      : 6
    model           : 143
    stepping        : 7
    microcode       : 0x2b0004a1
    bugs            : spectre_v1 spectre_v2 spec_store_bypass swapgs
    cpu cores       : 32
    siblings        : 32
    2 physical ids (chips)
    64 processors (hardware threads)
    physical id 0: core ids 0-31
    physical id 1: core ids 0-31
    physical id 0: apicids
    0,2,4,6,8,10,12,14,16,18,20,22,24,26,28,30,32,34,36,38,40,42,44,46,48,50,52,54,56,58,60,62
    physical id 1: apicids
    128,130,132,134,136,138,140,142,144,146,148,150,152,154,156,158,160,162,164,166,168,170,172,174,176,178,1
    80,182,184,186,188,190
  Caution: /proc/cpuinfo data regarding chips, cores, and threads is not necessarily reliable, especially for
  virtualized systems.  Use the above data carefully.

------------------------------------------------------------
7. lscpu

From lscpu from util-linux 2.37.4:
  Architecture:                    x86_64
  CPU op-mode(s):                  32-bit, 64-bit
  Address sizes:                   46 bits physical, 57 bits virtual
  Byte Order:                      Little Endian
  CPU(s):                          64
  On-line CPU(s) list:             0-63
  Vendor ID:                       GenuineIntel
  BIOS Vendor ID:                  Intel(R) Corporation
  Model name:                      Intel(R) Xeon(R) Gold 6448H
  BIOS Model name:                 Intel(R) Xeon(R) Gold 6448H
  CPU family:                      6
  Model:                           143
  Thread(s) per core:              1
  Core(s) per socket:              32
  Socket(s):                       2
  Stepping:                        7
  BogoMIPS:                        4800.00
  Flags:                           fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36
                                   clflush dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx pdpe1gb rdtscp
                                   lm constant_tsc art arch_perfmon pebs bts rep_good nopl xtopology
                                   nonstop_tsc cpuid aperfmperf tsc_known_freq pni pclmulqdq dtes64 monitor
                                   ds_cpl vmx smx est tm2 ssse3 sdbg fma cx16 xtpr pdcm pcid dca sse4_1
                                   sse4_2 x2apic movbe popcnt tsc_deadline_timer aes xsave avx f16c rdrand
                                   lahf_lm abm 3dnowprefetch cpuid_fault epb cat_l3 cat_l2 cdp_l3
                                   invpcid_single cdp_l2 ssbd mba ibrs ibpb stibp ibrs_enhanced tpr_shadow
```

(Continued on next page)

Platform Notes (Continued)

```
                                  vnmi flexpriority ept vpid ept_ad fsgsbase tsc_adjust bmi1 avx2 smep bmi2
                                  erms invpcid cqm rdt_a avx512f avx512dq rdseed adx smap avx512ifma
                                  clflushopt clwb intel_pt avx512cd sha_ni avx512bw avx512vl xsaveopt xsavec
                                  xgetbv1 xsaves cqm_llc cqm_occup_llc cqm_mbm_total cqm_mbm_local
                                  split_lock_detect avx_vnni avx512_bf16 wbnoinvd dtherm ida arat pln pts
                                  avx512vbmi umip pku ospke waitpkg avx512_vbmi2 gfni vaes vpclmulqdq
                                  avx512_vnni avx512_bitalg tme avx512_vpopcntdq la57 rdpid bus_lock_detect
                                  cldemote movdiri movdir64b enqcmd fsrm md_clear serialize tsxldtrk pconfig
                                  arch_lbr avx512_fp16 amx_tile flush_l1d arch_capabilities
  Virtualization:                 VT-x
  L1d cache:                      3 MiB (64 instances)
  L1i cache:                      2 MiB (64 instances)
  L2 cache:                       128 MiB (64 instances)
  L3 cache:                       120 MiB (2 instances)
  NUMA node(s):                   4
  NUMA node0 CPU(s):              0-7,32-39
  NUMA node1 CPU(s):              8-15,40-47
  NUMA node2 CPU(s):              16-23,48-55
  NUMA node3 CPU(s):              24-31,56-63
  Vulnerability Itlb multihit:    Not affected
  Vulnerability L1tf:             Not affected
  Vulnerability Mds:              Not affected
  Vulnerability Meltdown:         Not affected
  Vulnerability Spec store bypass: Mitigation; Speculative Store Bypass disabled via prctl
  Vulnerability Spectre v1:       Mitigation; usercopy/swapgs barriers and __user pointer sanitization
  Vulnerability Spectre v2:       Mitigation; Enhanced IBRS, IBPB conditional, RSB filling
  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       48K       3M   12 Data            1    64        1             64
     L1i       32K       2M    8 Instruction     1    64        1             64
     L2         2M     128M   16 Unified         2  2048        1             64
     L3        60M     120M   15 Unified         3 65536        1             64

------------------------------------------------------------
8. numactl --hardware
NOTE: a numactl 'node' might or might not correspond to a physical chip.
  available: 4 nodes (0-3)
  node 0 cpus: 0-7,32-39
  node 0 size: 128714 MB
  node 0 free: 128073 MB
  node 1 cpus: 8-15,40-47
  node 1 size: 129020 MB
  node 1 free: 128673 MB
  node 2 cpus: 16-23,48-55
  node 2 size: 128984 MB
  node 2 free: 128290 MB
  node 3 cpus: 24-31,56-63
  node 3 size: 129009 MB
  node 3 free: 128655 MB
  node distances:
  node   0   1   2   3
    0:  10  20  30  30
    1:  20  10  30  30
    2:  30  30  10  20
    3:  30  30  20  10

------------------------------------------------------------
```

(Continued on next page)

Platform Notes (Continued)

```
 9. /proc/meminfo
    MemTotal:       528105800 kB

------------------------------------------------------------
10. who -r
  run-level 3 Jun 14 04:04

------------------------------------------------------------
11. Systemd service manager version: systemd 250 (250-6.el9_0)
  Default Target  Status
  multi-user      degraded

------------------------------------------------------------
12. Failed units, from systemctl list-units --state=failed
    UNIT         LOAD   ACTIVE SUB    DESCRIPTION
  * sep5.service loaded failed failed systemd script to load sep5 driver at boot time

------------------------------------------------------------
13. Services, from systemctl list-unit-files
  STATE            UNIT FILES
  enabled          NetworkManager NetworkManager-dispatcher NetworkManager-wait-online auditd crond
                   dbus-broker firewalld getty@ irqbalance kdump lvm2-monitor mdmonitor microcode
                   nis-domainname rhsmcertd rsyslog selinux-autorelabel-mark sep5 sshd sssd
                   systemd-network-generator tuned udisks2 upower
  enabled-runtime  systemd-remount-fs
  disabled         blk-availability canberra-system-bootup canberra-system-shutdown
                   canberra-system-shutdown-reboot chrony-wait chronyd console-getty cpupower debug-shell
                   hwloc-dump-hwdata ipsec kvm_stat man-db-restart-cache-update nftables powertop rdisc rhsm
                   rhsm-facts rpmdb-rebuild serial-getty@ sshd-keygen@ systemd-boot-check-no-failures
                   systemd-pstore systemd-sysext
  indirect         sssd-autofs sssd-kcm sssd-nss sssd-pac sssd-pam sssd-ssh sssd-sudo

------------------------------------------------------------
14. Linux kernel boot-time arguments, from /proc/cmdline
  BOOT_IMAGE=(hd0,gpt2)/vmlinuz-5.14.0-70.13.1.el9_0.x86_64
  root=/dev/mapper/rhel-root
  ro
  resume=/dev/mapper/rhel-swap
  rd.lvm.lv=rhel/root
  rd.lvm.lv=rhel/swap

------------------------------------------------------------
15. cpupower frequency-info
  analyzing CPU 0:
    Unable to determine current policy
    boost state support:
      Supported: yes
      Active: yes

------------------------------------------------------------
16. tuned-adm active
  Current active profile: throughput-performance

------------------------------------------------------------
17. sysctl
  kernel.numa_balancing               1
  kernel.randomize_va_space           2
  vm.compaction_proactiveness        20
  vm.dirty_background_bytes           0
  vm.dirty_background_ratio          10
```

(Continued on next page)

Platform Notes (Continued)

```
  vm.dirty_bytes                    0
  vm.dirty_expire_centisecs      3000
  vm.dirty_ratio                   40
  vm.dirty_writeback_centisecs    500
  vm.dirtytime_expire_seconds   43200
  vm.extfrag_threshold            500
  vm.min_unmapped_ratio             1
  vm.nr_hugepages                   0
  vm.nr_hugepages_mempolicy         0
  vm.nr_overcommit_hugepages        0
  vm.swappiness                    10
  vm.watermark_boost_factor     15000
  vm.watermark_scale_factor        10
  vm.zone_reclaim_mode              0

------------------------------------------------------------
18. /sys/kernel/mm/transparent_hugepage
  defrag          always defer defer+madvise [madvise] never
  enabled         [always] madvise never
  hpage_pmd_size  2097152
  shmem_enabled   always within_size advise [never] deny force

------------------------------------------------------------
19. /sys/kernel/mm/transparent_hugepage/khugepaged
  alloc_sleep_millisecs   60000
  defrag                      1
  max_ptes_none             511
  max_ptes_shared           256
  max_ptes_swap              64
  pages_to_scan            4096
  scan_sleep_millisecs    10000

------------------------------------------------------------
20. OS release
  From /etc/*-release /etc/*-version
  os-release     Red Hat Enterprise Linux 9.0 (Plow)
  redhat-release Red Hat Enterprise Linux release 9.0 (Plow)
  system-release Red Hat Enterprise Linux release 9.0 (Plow)

------------------------------------------------------------
21. Disk information
SPEC is set to: /home/cpu2017
  Filesystem            Type  Size  Used Avail Use% Mounted on
  /dev/mapper/rhel-home xfs   372G  111G  262G  30% /home

------------------------------------------------------------
22. /sys/devices/virtual/dmi/id
    Vendor:         HPE
    Product:        ProLiant DL380 Gen11
    Product Family: ProLiant
    Serial:         CNX21000G8

------------------------------------------------------------
23. dmidecode
  Additional information from dmidecode 3.3 follows.  WARNING: Use caution when you interpret this section.
  The 'dmidecode' program reads system data which is "intended to allow hardware to be accurately
  determined", but the intent may not be met, as there are frequent changes to hardware, firmware, and the
  "DMTF SMBIOS" standard.
  Memory:
    1x Hynix HMCG88AEBRA168N 32 GB 2 rank 4800
```

(Continued on next page)

Platform Notes (Continued)

```
     15x Hynix HMCG88MEBRA113N 32 GB 2 rank 4800


 ------------------------------------------------------------
 24. BIOS
 (This section combines info from /sys/devices and dmidecode.)
    BIOS Vendor:       HPE
    BIOS Version:      1.40
    BIOS Date:         05/18/2023
    BIOS Revision:     1.40
    Firmware Revision: 1.10
```

Compiler Version Notes

```
==============================================================================================================
C       | 600.perlbench_s(base, peak) 602.gcc_s(base, peak) 605.mcf_s(base, peak) 625.x264_s(base, peak)
        | 657.xz_s(base, peak)
--------------------------------------------------------------------------------------------------------------
Intel(R) oneAPI DPC++/C++ Compiler for applications running on Intel(R) 64, Version 2023.0.0 Build 20221201
Copyright (C) 1985-2022 Intel Corporation. All rights reserved.
--------------------------------------------------------------------------------------------------------------

==============================================================================================================
C++     | 620.omnetpp_s(base, peak) 623.xalancbmk_s(base, peak) 631.deepsjeng_s(base, peak)
        | 641.leela_s(base, peak)
--------------------------------------------------------------------------------------------------------------
Intel(R) oneAPI DPC++/C++ Compiler for applications running on Intel(R) 64, Version 2023.0.0 Build 20221201
Copyright (C) 1985-2022 Intel Corporation. All rights reserved.
--------------------------------------------------------------------------------------------------------------

==============================================================================================================
Fortran | 648.exchange2_s(base, peak)
--------------------------------------------------------------------------------------------------------------
Intel(R) Fortran Compiler for applications running on Intel(R) 64, Version 2023.0.0 Build 20221201
Copyright (C) 1985-2022 Intel Corporation. All rights reserved.
--------------------------------------------------------------------------------------------------------------
```

Base Compiler Invocation

C benchmarks:
icx

C++ benchmarks:
icpx

Fortran benchmarks:
ifx

Base Portability Flags

600.perlbench_s: `-DSPEC_LP64 -DSPEC_LINUX_X64`
602.gcc_s: `-DSPEC_LP64`
605.mcf_s: `-DSPEC_LP64`
620.omnetpp_s: `-DSPEC_LP64`
623.xalancbmk_s: `-DSPEC_LP64 -DSPEC_LINUX`
625.x264_s: `-DSPEC_LP64`
631.deepsjeng_s: `-DSPEC_LP64`
641.leela_s: `-DSPEC_LP64`
648.exchange2_s: `-DSPEC_LP64`
657.xz_s: `-DSPEC_LP64`

Base Optimization Flags

C benchmarks:
`-m64 -std=c11 -Wl,-z,muldefs -xsapphirerapids -O3 -ffast-math -flto -mfpmath=sse -funroll-loops -qopt-mem-layout-trans=4 -fiopenmp -DSPEC_OPENMP -L/usr/local/jemalloc64-5.0.1/lib -ljemalloc`

C++ benchmarks:
`-m64 -std=c++14 -Wl,-z,muldefs -xsapphirerapids -O3 -ffast-math -flto -mfpmath=sse -funroll-loops -qopt-mem-layout-trans=4 -L/usr/local/jemalloc64-5.0.1/lib -ljemalloc`

Fortran benchmarks:
`-m64 -Wl,-z,muldefs -xsapphirerapids -O3 -ffast-math -flto -mfpmath=sse -funroll-loops -qopt-mem-layout-trans=4 -nostandard-realloc-lhs -align array32byte -L/usr/local/jemalloc64-5.0.1/lib -ljemalloc`

Peak Compiler Invocation

C benchmarks:
`icx`

C++ benchmarks:
`icpx`

Fortran benchmarks:
`ifx`

Peak Portability Flags

Same as Base Portability Flags

Peak Optimization Flags

C benchmarks:

600.perlbench_s: -m64 -std=c11 -Wl,-z,muldefs -fprofile-generate(pass 1)
-fprofile-use=default.profdata(pass 2) -xCORE-AVX2(pass 1)
-flto -Ofast(pass 1) -xCORE-AVX512 -O3 -ffast-math
-mfpmath=sse -funroll-loops -qopt-mem-layout-trans=4
-fiopenmp -DSPEC_OPENMP -fno-strict-overflow
-L/usr/local/jemalloc64-5.0.1/lib -ljemalloc

602.gcc_s: -m64 -std=c11 -Wl,-z,muldefs -fprofile-generate(pass 1)
-fprofile-use=default.profdata(pass 2) -xCORE-AVX2(pass 1)
-flto -Ofast(pass 1) -xCORE-AVX512 -O3 -ffast-math
-mfpmath=sse -funroll-loops -qopt-mem-layout-trans=4
-fiopenmp -DSPEC_OPENMP -L/usr/local/jemalloc64-5.0.1/lib
-ljemalloc

605.mcf_s: basepeak = yes

625.x264_s: -m64 -std=c11 -Wl,-z,muldefs -xsapphirerapids -O3
-ffast-math -flto -mfpmath=sse -funroll-loops
-qopt-mem-layout-trans=4 -fiopenmp -DSPEC_OPENMP
-fno-alias -L/usr/local/jemalloc64-5.0.1/lib -ljemalloc

657.xz_s: basepeak = yes

C++ benchmarks:

620.omnetpp_s: basepeak = yes

623.xalancbmk_s: basepeak = yes

631.deepsjeng_s: basepeak = yes

641.leela_s: basepeak = yes

Fortran benchmarks:

648.exchange2_s: basepeak = yes

The flags files that were used to format this result can be browsed at
http://www.spec.org/cpu2017/flags/Intel-ic2023-official-linux64.html
http://www.spec.org/cpu2017/flags/HPE-Platform-Flags-Intel-SPR-rev2.1.html

You can also download the XML flags sources by saving the following links:
http://www.spec.org/cpu2017/flags/Intel-ic2023-official-linux64.xml
http://www.spec.org/cpu2017/flags/HPE-Platform-Flags-Intel-SPR-rev2.1.xml

SPEC CPU and SPECspeed are registered trademarks 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 SPEC CPU®2017 v1.1.9 on 2023-06-13 18:35:12-0400.
Report generated on 2024-01-29 17:53:40 by CPU2017 PS formatter v6716.
Originally published on 2023-07-04.