SPEC CPU®2017 Integer Speed Result

Copyright 2017-2024 Standard Performance Evaluation Corporation

IEIT Systems Co., Ltd.

NF5466M7 (Intel Xeon Gold 6548Y+)

SPECspeed®2017_int_base = 15.2

SPECspeed®2017_int_peak = 15.4

CPU2017 License: 3358
Test Sponsor: IEIT Systems Co., Ltd.
Tested by: IEIT Systems Co., Ltd.

Test Date: May-2024
Hardware Availability: Dec-2023
Software Availability: Dec-2023

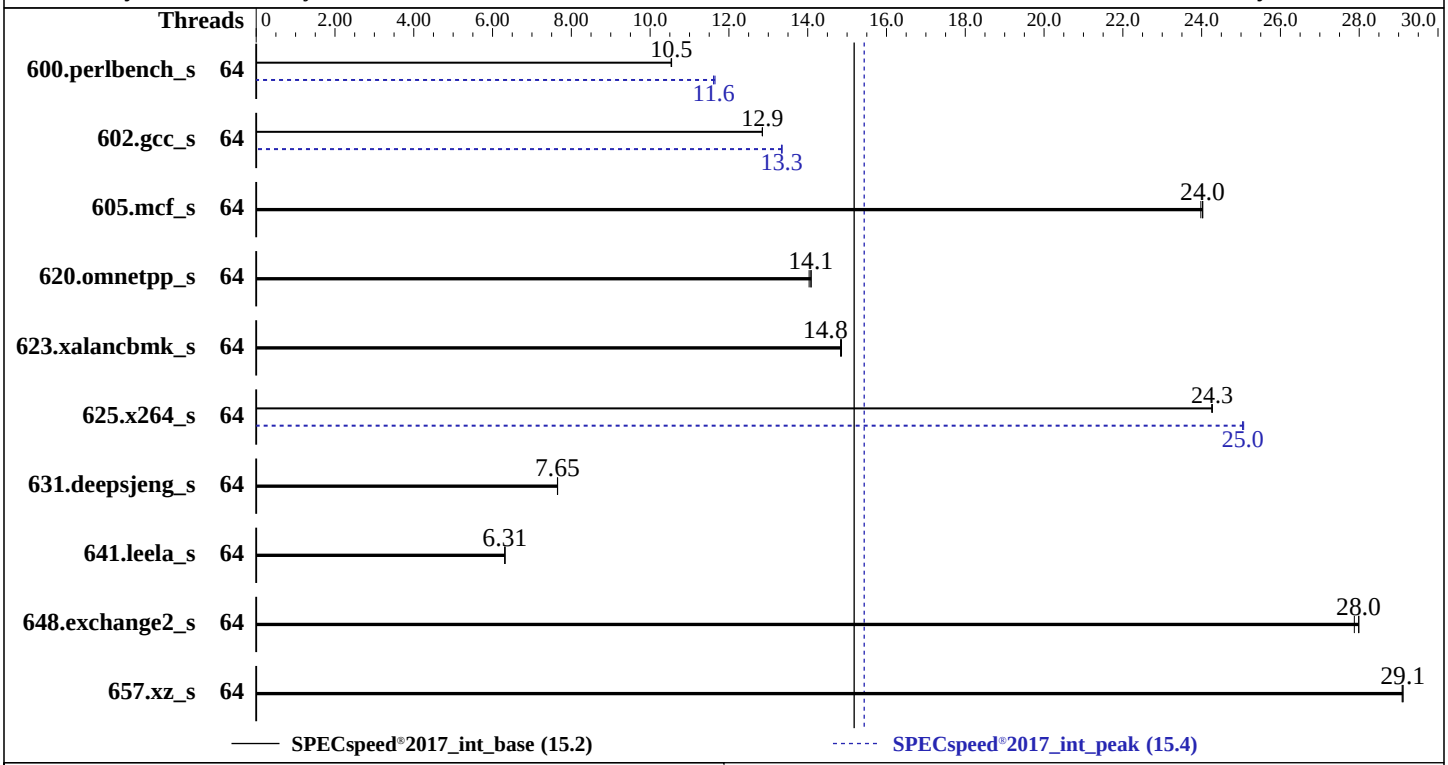

Hardware

CPU Name:	Intel Xeon Gold 6548Y+
Max MHz:	4100
Nominal:	2500
Enabled:	64 cores, 2 chips
Orderable:	1,2 chips
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:	1 TB (16 x 64 GB 2Rx4 PC5-5600B-R, running at 4800)
Storage:	1 x 1.92 TB NVME SSD
Other:	CPU Cooling: Air

Software

OS:	SUSE Linux Enterprise Server 15 SP5 5.14.21-150500.53-default
Compiler:	C/C++: Version 2024.0.2 of Intel oneAPI DPC++/C++ Compiler for Linux; Fortran: Version 2024.0.2 of Intel Fortran Compiler for Linux;
Parallel:	Yes
Firmware:	Version 05.13.00 released Nov-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 set to prefer performance at the cost of additional power usage.

Results Table

Benchmark	Base							Peak						
	Threads	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Threads	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
600.perlbench_s	64	**168**	**10.5**	168	10.5	169	10.5	64	152	11.6	**153**	**11.6**	153	11.6
602.gcc_s	64	**310**	**12.9**	310	12.9	310	12.8	64	**298**	**13.3**	299	13.3	298	13.3
605.mcf_s	64	**197**	**24.0**	196	24.0	197	24.0	64	**197**	**24.0**	196	24.0	197	24.0
620.omnetpp_s	64	116	14.1	**116**	**14.1**	116	14.0	64	116	14.1	**116**	**14.1**	116	14.0
623.xalancbmk_s	64	95.4	14.9	95.5	14.8	**95.5**	**14.8**	64	95.4	14.9	95.5	14.8	**95.5**	**14.8**
625.x264_s	64	72.7	24.3	**72.7**	**24.3**	72.7	24.3	64	70.4	25.0	70.4	25.1	**70.4**	**25.0**
631.deepsjeng_s	64	187	7.65	**187**	**7.65**	187	7.65	64	187	7.65	**187**	**7.65**	187	7.65
641.leela_s	64	270	6.31	270	6.31	**270**	**6.31**	64	270	6.31	270	6.31	**270**	**6.31**
648.exchange2_s	64	**105**	**28.0**	105	27.9	105	28.0	64	**105**	**28.0**	105	27.9	105	28.0
657.xz_s	64	**212**	**29.1**	212	29.1	213	29.1	64	**212**	**29.1**	212	29.1	213	29.1

SPECspeed®2017_int_base = **15.2**

SPECspeed®2017_int_peak = **15.4**

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

Operating System Notes

```
Stack size set to unlimited using "ulimit -s unlimited"
```

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 Redhat Enterprise Linux 8.0
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

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

```
BIOS configuration:
ENERGY_PERF_BIAS_CFG mode set to Performance
Hardware Prefetch set to Disable
VT Support set to Disable
Sub NUMA Cluster (SNC) set to Disable
Hyper Threading set to Disable

Sysinfo program /home/CPU2017/bin/sysinfo
Rev: r6732 of 2022-11-07 fe91c89b7ed5c36ae2c92cc097bec197
running on localhost Fri May 17 13:24:05 2024

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
11. Systemd service manager version: systemd 249 (249.16+suse.171.gdad0071f15)
12. Services, from systemctl list-unit-files
13. Linux kernel boot-time arguments, from /proc/cmdline
14. cpupower frequency-info
15. tuned-adm active
16. sysctl
17. /sys/kernel/mm/transparent_hugepage
18. /sys/kernel/mm/transparent_hugepage/khugepaged
19. OS release
20. Disk information
21. /sys/devices/virtual/dmi/id
22. dmidecode
23. BIOS
------------------------------------------------------------

------------------------------------------------------------
1. uname -a
  Linux localhost 5.14.21-150500.53-default #1 SMP PREEMPT_DYNAMIC Wed May 10 07:56:26 UTC 2023 (b630043)
  x86_64 x86_64 x86_64 GNU/Linux

------------------------------------------------------------
2. w
   13:24:05 up 1 min,  1 user,  load average: 0.69, 0.28, 0.10
  USER     TTY      FROM             LOGIN@   IDLE   JCPU   PCPU WHAT
  root     tty1     -                13:23   13.00s  1.92s  0.00s sh
  reportable-ic2024.0.2-lin-sapphirerapids-speed-smt-off-20231213.sh

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

------------------------------------------------------------
4. ulimit -a
  core file size          (blocks, -c) unlimited
```

(Continued on next page)

Platform Notes (Continued)

```
  data seg size           (kbytes, -d) unlimited
  scheduling priority             (-e) 0
  file size               (blocks, -f) unlimited
  pending signals                 (-i) 4124593
  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) 4124593
  virtual memory          (kbytes, -v) unlimited
  file locks                      (-x) unlimited

------------------------------------------------------------
5. sysinfo process ancestry
 /usr/lib/systemd/systemd --switched-root --system --deserialize 30
 login -- root
 -bash
 sh reportable-ic2024.0.2-lin-sapphirerapids-speed-smt-off-20231213.sh
 runcpu --nobuild --action validate --define default-platform-flags -c
   ic2024.0.2-lin-sapphirerapids-speed-20231213.cfg --define cores=64 --tune base,peak -o all --define
   intspeedaffinity --define drop_caches intspeed
 runcpu --nobuild --action validate --define default-platform-flags --configfile
   ic2024.0.2-lin-sapphirerapids-speed-20231213.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/CPU2017.009/templogs/preenv.intspeed.009.0.log
   --lognum 009.0 --from_runcpu 2
 specperl $SPEC/bin/sysinfo
$SPEC = /home/CPU2017

------------------------------------------------------------
6. /proc/cpuinfo
    model name      : INTEL(R) XEON(R) GOLD 6548Y+
    vendor_id       : GenuineIntel
    cpu family      : 6
    model           : 207
    stepping        : 2
    microcode       : 0x21000161
    bugs            : spectre_v1 spectre_v2 spec_store_bypass swapgs eibrs_pbrsb
    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:
```

(Continued on next page)

Platform Notes (Continued)

```
  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
  Model name:                      INTEL(R) XEON(R) GOLD 6548Y+
  CPU family:                      6
  Model:                           207
  Thread(s) per core:              1
  Core(s) per socket:              32
  Socket(s):                       2
  Stepping:                        2
  CPU max MHz:                     4100.0000
  CPU min MHz:                     800.0000
  BogoMIPS:                        3800.00
  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 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 fsgsbase tsc_adjust bmi1 hle
                                   avx2 smep bmi2 erms invpcid rtm 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 avx_vnni avx512_bf16 wbnoinvd dtherm ida arat pln pts hwp
                                   hwp_act_window hwp_epp hwp_pkg_req hfi 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
  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):                    2
  NUMA node0 CPU(s):               0-31
  NUMA node1 CPU(s):               32-63
  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 and seccomp
  Vulnerability Spectre v1:        Mitigation; usercopy/swapgs barriers and __user pointer sanitization
  Vulnerability Spectre v2:        Mitigation; Enhanced IBRS, IBPB conditional, RSB filling, PBRSB-eIBRS SW
                                   sequence
  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
```

(Continued on next page)

Platform Notes (Continued)

```
------------------------------------------------------------
8. numactl --hardware
NOTE: a numactl 'node' might or might not correspond to a physical chip.
  available: 2 nodes (0-1)
  node 0 cpus: 0-31
  node 0 size: 515558 MB
  node 0 free: 514295 MB
  node 1 cpus: 32-63
  node 1 size: 515619 MB
  node 1 free: 514799 MB
  node distances:
  node   0   1
    0:  10  21
    1:  21  10

------------------------------------------------------------
9. /proc/meminfo
   MemTotal:       1055926352 kB

------------------------------------------------------------
10. who -r
  run-level 3 May 17 13:23

------------------------------------------------------------
11. Systemd service manager version: systemd 249 (249.16+suse.171.gdad0071f15)
  Default Target  Status
  multi-user      running

------------------------------------------------------------
12. Services, from systemctl list-unit-files
  STATE            UNIT FILES
  enabled          YaST2-Firstboot YaST2-Second-Stage apparmor auditd cron display-manager getty@ irqbalance
                   issue-generator kbdsettings klog lvm2-monitor nscd nvmefc-boot-connections postfix
                   purge-kernels rollback rsyslog smartd sshd systemd-pstore
  enabled-runtime  systemd-remount-fs
  disabled         autofs autoyast-initscripts blk-availability boot-sysctl ca-certificates chrony-wait
                   chronyd console-getty cups cups-browsed debug-shell ebtables exchange-bmc-os-info
                   firewalld gpm grub2-once haveged haveged-switch-root ipmi ipmievd issue-add-ssh-keys kdump
                   kdump-early kexec-load lunmask man-db-create multipathd nfs nfs-blkmap nvmf-autoconnect
                   rpcbind rpmconfigcheck rsyncd serial-getty@ smartd_generate_opts snmpd snmptrapd
                   systemd-boot-check-no-failures systemd-network-generator systemd-sysext
                   systemd-time-wait-sync systemd-timesyncd tuned udisks2 vncserver@ wicked wickedd-auto4
                   wickedd-dhcp4 wickedd-dhcp6 wickedd-nanny
  indirect         wickedd

------------------------------------------------------------
13. Linux kernel boot-time arguments, from /proc/cmdline
  BOOT_IMAGE=/boot/vmlinuz-5.14.21-150500.53-default
  root=UUID=0a76fbd6-6541-4f13-928e-2a3caa636f4c
  splash=silent
  resume=/dev/disk/by-uuid/e6a5bcc3-5591-4cd4-a737-ee6cc3cdc083
  mitigations=auto
  quiet
  security=apparmor
  crashkernel=447M,high
  crashkernel=72M,low

------------------------------------------------------------
14. cpupower frequency-info
```

(Continued on next page)

Platform Notes (Continued)

```
  analyzing CPU 0:
    current policy: frequency should be within 800 MHz and 4.10 GHz.
                    The governor "powersave" may decide which speed to use
                    within this range.
    boost state support:
      Supported: yes
      Active: yes

------------------------------------------------------------
15. tuned-adm active
  It seems that tuned daemon is not running, preset profile is not activated.
  Preset profile: throughput-performance

------------------------------------------------------------
16. sysctl
  kernel.numa_balancing               1
  kernel.randomize_va_space           2
  vm.compaction_proactiveness        20
  vm.dirty_background_bytes           0
  vm.dirty_background_ratio          10
  vm.dirty_bytes                      0
  vm.dirty_expire_centisecs        3000
  vm.dirty_ratio                     20
  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                      60
  vm.watermark_boost_factor       15000
  vm.watermark_scale_factor          10
  vm.zone_reclaim_mode                0

------------------------------------------------------------
17. /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

------------------------------------------------------------
18. /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

------------------------------------------------------------
19. OS release
  From /etc/*-release /etc/*-version
  os-release SUSE Linux Enterprise Server 15 SP5

------------------------------------------------------------
20. Disk information
SPEC is set to: /home/CPU2017
```

(Continued on next page)

Platform Notes (Continued)

```
  Filesystem     Type  Size  Used Avail Use% Mounted on
  /dev/nvme0n1p4 xfs   1.3T   30G  1.2T   3% /home

------------------------------------------------------------
21. /sys/devices/virtual/dmi/id
    Vendor:         IEIT SYSTEMS
    Product:        NF5466-M7-A0-R0-00
    Product Family: Not specified
    Serial:         000000000

------------------------------------------------------------
22. dmidecode
  Additional information from dmidecode 3.4 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:
    8x Samsung M321R8GA0PB0-CWMKH 64 GB 2 rank 5600, configured at 4800
    8x Samsung M321R8GA0PB0-CWMXH 64 GB 2 rank 5600, configured at 4800


------------------------------------------------------------
23. BIOS
(This section combines info from /sys/devices and dmidecode.)
   BIOS Vendor:       American Megatrends International, LLC.
   BIOS Version:      05.13.00
   BIOS Date:         11/30/2023
```

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 2024.0.2 Build 20231213
Copyright (C) 1985-2023 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 2024.0.2 Build 20231213
Copyright (C) 1985-2023 Intel Corporation. All rights reserved.
--------------------------------------------------------------------------------------------------------------

==============================================================================================================
Fortran | 648.exchange2_s(base, peak)
--------------------------------------------------------------------------------------------------------------
Intel(R) Fortran Compiler for applications running on Intel(R) 64, Version 2024.0.2 Build 20231213
Copyright (C) 1985-2023 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:
-w -std=c11 -m64 -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:
-w -std=c++14 -m64 -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:
-w -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: -w -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: -w -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: -w -std=c11 -m64 -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

(Continued on next page)

Peak Optimization Flags (Continued)

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-ic2024-official-linux64.html`
`http://www.spec.org/cpu2017/flags/IEIT-Platform-Settings-intel-V1.0.html`

You can also download the XML flags sources by saving the following links:
`http://www.spec.org/cpu2017/flags/Intel-ic2024-official-linux64.xml`
`http://www.spec.org/cpu2017/flags/IEIT-Platform-Settings-intel-V1.0.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 2024-05-17 13:24:04-0400.
Report generated on 2024-07-17 11:44:51 by CPU2017 PS formatter v6716.
Originally published on 2024-07-16.