SPEC CPU®2017 Integer Rate Result

Copyright 2017-2020 Standard Performance Evaluation Corporation

Huawei
(Test Sponsor: Peng Cheng Laboratory)

Huawei TaiShan 200 Server (Model 2480)
(2.6 GHz,Huawei Kunpeng 920 7260)

$SPECrate®2017\_int\_base = \quad 628$

$SPECrate®2017\_int\_peak = Not Run$

CPU2017 License: 5036

Test Sponsor: Peng Cheng Laboratory

Tested by: Peng Cheng Laboratory

Test Date: May-2020

Hardware Availability: Jan-2020

Software Availability: Jul-2020

Benchmark	Copies	Result
500.perlbench_r	256	753
502.gcc_r	256	511
505.mcf_r	256	320
520.omnetpp_r	256	329
523.xalancbmk_r	256	615
525.x264_r	256	1320
531.deepsjeng_r	256	784
541.leela_r	256	736
548.exchange2_r	256	1230
557.xz_r	256	412

SPECrate®2017_int_base (628)

Hardware

CPU Name: Huawei Kunpeng 920 7260
 Max MHz: 2600
 Nominal: 2600
Enabled: 256 cores, 4 chips
Orderable: 1,2,3,4 chips
Cache L1: 64 KB I + 64 KB D on chip per core
 L2: 512 KB I+D on chip per core
 L3: 64 MB I+D on chip per chip
 Other: None
Memory: 1 TB (32 x 32 GB 2Rx4 PC4-2933Y-R)
Storage: 1 x 960 GB SAS SSD
Other: None

Software

OS: kylin release 10 (Azalea)
 4.19.90-5.ky10.aarch64
Compiler: C/C++/Fortran: Version 9.1.0 of GCC, the GNU Compiler Collection
Parallel: No
Firmware: Huawei Corp. Version 1.20 released Apr-2020
File System: xfs
System State: Run level 5 (multi-user graphical)
Base Pointers: 64-bit
Peak Pointers: Not Applicable
Other: jemalloc memory allocator V5.2.1
Power Management: BIOS set to prefer performance at the cost of additional power usage

Standard Performance Evaluation Corporation (info@spec.org)   https://www.spec.org/

SPEC CPU®2017 Integer Rate Result

Copyright 2017-2020 Standard Performance Evaluation Corporation

Huawei
(Test Sponsor: Peng Cheng Laboratory)

Huawei TaiShan 200 Server (Model 2480)
(2.6 GHz,Huawei Kunpeng 920 7260)

SPECrate®2017_int_base =   628

SPECrate®2017_int_peak = Not Run

CPU2017 License: 5036	**Test Date:** May-2020
Test Sponsor: Peng Cheng Laboratory	**Hardware Availability:** Jan-2020
Tested by: Peng Cheng Laboratory	**Software Availability:** Jul-2020

Results Table

	Base							Peak						
Benchmark	Copies	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio	Copies	Seconds	Ratio	Seconds	Ratio	Seconds	Ratio
500.perlbench_r	256	542	752	**542**	**753**	541	753							
502.gcc_r	256	708	512	711	510	**709**	**511**							
505.mcf_r	256	1292	320	**1292**	**320**	1290	321							
520.omnetpp_r	256	**1020**	**329**	1020	329	1019	330							
523.xalancbmk_r	256	441	613	**439**	**615**	439	615							
525.x264_r	256	339	1320	**340**	**1320**	340	1320							
531.deepsjeng_r	256	**374**	**784**	374	785	374	783							
541.leela_r	256	**576**	**736**	576	736	576	736							
548.exchange2_r	256	507	1320	548	1220	**547**	**1230**							
557.xz_r	256	671	412	670	412	**670**	**412**							
SPECrate®2017_int_base =		**628**												
SPECrate®2017_int_peak =		**Not Run**												

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

Submit Notes

The taskset mechanism was used to bind copies to processors. The config file option 'submit'
was used to generate taskset 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"

Environment Variables Notes

Environment variables set by runcpu before the start of the run:
LD_LIBRARY_PATH =
     "/usr/local/gcc-9.1.0/lib64/:/usr/local/gcc-9.1.0/lib/:/lib64::/home/jem
     alloc-5.2.1-setup/lib"

General Notes

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
jemalloc: configured and built at default for 64bit targets;
jemalloc: built with the kylin V10, and the system compiler gcc 7.3.0;
jemalloc: sources available via jemalloc.net or https://github.com/jemalloc/jemalloc/releases

(Continued on next page)

SPEC CPU®2017 Integer Rate Result

Copyright 2017-2020 Standard Performance Evaluation Corporation

Huawei

(Test Sponsor: Peng Cheng Laboratory)

Huawei TaiShan 200 Server (Model 2480)
(2.6 GHz,Huawei Kunpeng 920 7260)

SPECrate®2017_int_base =   628

SPECrate®2017_int_peak = Not Run

CPU2017 License: 5036
Test Sponsor: Peng Cheng Laboratory
Tested by: Peng Cheng Laboratory

Test Date: May-2020
Hardware Availability: Jan-2020
Software Availability: Jul-2020

General Notes (Continued)

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.
NA: 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.

Platform Notes

```
BIOS configuration:
Power Policy Set to Performance
Custom Refresh Rate Set to 64ms
CPU Prefetcher Set to Enabled
L3 Cache Model Set to in :private out:private

Sysinfo program /home/cpu2017/bin/sysinfo
Rev: r6365 of 2019-08-21 295195f888a3d7edb1e6e46a485a0011
running on localhost.localdomain Fri May 29 09:13:26 2020

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
*
* 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.
*
     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.)

From lscpu:
     Architecture:               aarch64
     CPU op-mode(s):             64-bit
     Byte Order:                 Little Endian
     CPU(s):                     256
     On-line CPU(s) list:        0-255
     Thread(s) per core:         1
     Core(s) per socket:         64
```

(Continued on next page)

Huawei

(Test Sponsor: Peng Cheng Laboratory)

Huawei TaiShan 200 Server (Model 2480)
(2.6 GHz,Huawei Kunpeng 920 7260)

SPECrate®2017_int_base =   628

SPECrate®2017_int_peak = Not Run

CPU2017 License: 5036   **Test Date:**   May-2020
Test Sponsor: Peng Cheng Laboratory   **Hardware Availability:** Jan-2020
Tested by: Peng Cheng Laboratory   **Software Availability:** Jul-2020

Platform Notes (Continued)

```
      Socket(s):                     4
      NUMA node(s):                  8
      Vendor ID:                     HiSilicon
      Model:                         0
      Model name:                    Kunpeng-920
      Stepping:                      0x1
      BogoMIPS:                      200.00
      L1d cache:                     16 MiB
      L1i cache:                     16 MiB
      L2 cache:                      128 MiB
      L3 cache:                      512 MiB
      NUMA node0 CPU(s):             0-31
      NUMA node1 CPU(s):             32-63
      NUMA node2 CPU(s):             64-95
      NUMA node3 CPU(s):             96-127
      NUMA node4 CPU(s):             128-159
      NUMA node5 CPU(s):             160-191
      NUMA node6 CPU(s):             192-223
      NUMA node7 CPU(s):             224-255
      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; __user pointer sanitization
      Vulnerability Spectre v2:      Not affected
      Vulnerability Tsx async abort:  Not affected
      Flags:                          fp asimd evtstrm aes pmull sha1 sha2 crc32 atomics
      fphp asimdhp cpuid asimdrdm jscvt fcma dcpop asimddp asimdfhm ssbs

 From numactl --hardware  WARNING: a numactl 'node' might or might not correspond to a
 physical chip.

 From /proc/meminfo
    MemTotal:       1071619840 kB
    HugePages_Total:   100000
    Hugepagesize:       2048 kB

 From /etc/*release* /etc/*version*
    kylin-release: kylin release 10 (Azalea)
    os-release:
      NAME="kylin"
      VERSION="10 (Azalea)"
      ID="kylin"
      VERSION_ID="10"
      PRETTY_NAME="kylin 10 (Azalea)"
```

(Continued on next page)

Huawei
(Test Sponsor: Peng Cheng Laboratory)

Huawei TaiShan 200 Server (Model 2480)
(2.6 GHz,Huawei Kunpeng 920 7260)

SPECrate®2017_int_base = 628

SPECrate®2017_int_peak = Not Run

CPU2017 License: 5036
Test Sponsor: Peng Cheng Laboratory
Tested by: Peng Cheng Laboratory

Test Date: May-2020
Hardware Availability: Jan-2020
Software Availability: Jul-2020

Platform Notes (Continued)

```
      ANSI_COLOR="0;31"


    system-release: kylin release 10 (Azalea)
    system-release-cpe: cpe:/o:kylin:kylin:10:ga:server


  uname -a:
    Linux localhost.localdomain 4.19.90-5.ky10.aarch64 #1 SMP Wed Apr 8 09:34:13 CST 2020
    aarch64 aarch64 aarch64 GNU/Linux


  Kernel self-reported vulnerability status:

  itlb_multihit:                        Not affected
  CVE-2018-3620 (L1 Terminal Fault):    Not affected
  Microarchitectural Data Sampling:     Not affected
  CVE-2017-5754 (Meltdown):             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
  tsx_async_abort:                      Not affected

  run-level 5 May 29 09:08

  SPEC is set to: /home/cpu2017
    Filesystem              Type  Size  Used Avail Use% Mounted on
    /dev/mapper/klas00-home xfs   838G   16G  822G   2% /home

  From /sys/devices/virtual/dmi/id
    BIOS:    Huawei Corp. 1.20 04/14/2020
    Vendor:  Huawei
    Product: TaiShan 200 (Model 2480)
    Serial:  2102312UXX10KC000007

  Additional information from dmidecode 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:
      32x Hynix HMA84GR7CJR4N-WM 32 GB 2 rank 2933


  (End of data from sysinfo program)
The sysinfo is missing the cpu name, the processor under test is Huawei Kunpeng 920 7260.
The L3 capacity is 64MB per processor for Huawei Kunpeng 920 7260 processor for a SUT
total of 256 MiB.
```

Standard Performance Evaluation Corporation (info@spec.org)   https://www.spec.org/

SPEC CPU®2017 Integer Rate Result

Copyright 2017-2020 Standard Performance Evaluation Corporation

Huawei

(Test Sponsor: Peng Cheng Laboratory)
Huawei TaiShan 200 Server (Model 2480)
(2.6 GHz,Huawei Kunpeng 920 7260)

SPECrate®2017_int_base =   628

SPECrate®2017_int_peak = Not Run

CPU2017 License: 5036
Test Sponsor: Peng Cheng Laboratory
Tested by: Peng Cheng Laboratory

Test Date: May-2020
Hardware Availability: Jan-2020
Software Availability: Jul-2020

Compiler Version Notes

```
==============================================================================
C        | 500.perlbench_r(base) 502.gcc_r(base) 505.mcf_r(base)
         | 525.x264_r(base) 557.xz_r(base)
------------------------------------------------------------------------------
Using built-in specs.
COLLECT_GCC=/usr/local/gcc-9.1.0/bin/gcc
COLLECT_LTO_WRAPPER=/usr/local/gcc-9.1.0/libexec/gcc/aarch64-unknown-linux-gnu/9.1.0/lto-wrapper
Target: aarch64-unknown-linux-gnu
Configured with: ../configure --enable-checking=release
  --enable-languages=c,c++,fortran --disable-multilib
  --prefix=/usr/local/gcc-9.1.0
Thread model: posix
gcc version 9.1.0 (GCC)
------------------------------------------------------------------------------


==============================================================================
C++      | 520.omnetpp_r(base) 523.xalancbmk_r(base) 531.deepsjeng_r(base)
         | 541.leela_r(base)
------------------------------------------------------------------------------
Using built-in specs.
COLLECT_GCC=/usr/local/gcc-9.1.0/bin/g++
COLLECT_LTO_WRAPPER=/usr/local/gcc-9.1.0/libexec/gcc/aarch64-unknown-linux-gnu/9.1.0/lto-wrapper
Target: aarch64-unknown-linux-gnu
Configured with: ../configure --enable-checking=release
  --enable-languages=c,c++,fortran --disable-multilib
  --prefix=/usr/local/gcc-9.1.0
Thread model: posix
gcc version 9.1.0 (GCC)
------------------------------------------------------------------------------


==============================================================================
Fortran  | 548.exchange2_r(base)
------------------------------------------------------------------------------
Using built-in specs.
COLLECT_GCC=/usr/local/gcc-9.1.0/bin/gfortran
COLLECT_LTO_WRAPPER=/usr/local/gcc-9.1.0/libexec/gcc/aarch64-unknown-linux-gnu/9.1.0/lto-wrapper
Target: aarch64-unknown-linux-gnu
Configured with: ../configure --enable-checking=release
  --enable-languages=c,c++,fortran --disable-multilib
  --prefix=/usr/local/gcc-9.1.0
Thread model: posix
gcc version 9.1.0 (GCC)
------------------------------------------------------------------------------
```

SPEC CPU®2017 Integer Rate Result

Huawei
(Test Sponsor: Peng Cheng Laboratory)
Huawei TaiShan 200 Server (Model 2480)
(2.6 GHz,Huawei Kunpeng 920 7260)

SPECrate®2017_int_base =   628

SPECrate®2017_int_peak = Not Run

CPU2017 License: 5036
Test Sponsor: Peng Cheng Laboratory
Tested by: Peng Cheng Laboratory

Test Date: May-2020
Hardware Availability: Jan-2020
Software Availability: Jul-2020

Base Compiler Invocation

C benchmarks:
`gcc`

C++ benchmarks:
`g++`

Fortran benchmarks:
`gfortran`

Base Portability Flags

500.perlbench_r: `-DSPEC_LINUX_AARCH64 -DSPEC_LP64`
502.gcc_r: `-DSPEC_LP64`
505.mcf_r: `-DSPEC_LP64`
520.omnetpp_r: `-DSPEC_LP64`
523.xalancbmk_r: `-DSPEC_LINUX -DSPEC_LP64`
525.x264_r: `-DSPEC_LP64`
531.deepsjeng_r: `-DSPEC_LP64`
541.leela_r: `-DSPEC_LP64`
548.exchange2_r: `-DSPEC_LP64`
557.xz_r: `-DSPEC_LP64`

Base Optimization Flags

C benchmarks:
```
-mabi=lp64 -std=c99 -z muldefs -O3 -march=armv8.2-a+lse -fno-PIE
-no-pie -fomit-frame-pointer -funroll-loops -fno-strict-aliasing
-fgnu89-inline -L/home/jemalloc-5.2.1-setup/lib -ljemalloc
```

C++ benchmarks:
```
-mabi=lp64 -std=c++03 -O3 -march=armv8.2-a+lse -fno-PIE -no-pie
-fomit-frame-pointer -funroll-loops -L/home/jemalloc-5.2.1-setup/lib
-ljemalloc
```

Fortran benchmarks:
```
-mabi=lp64 -O3 -march=armv8.2-a+lse -fno-PIE -no-pie
-fomit-frame-pointer -funroll-loops -L/home/jemalloc-5.2.1-setup/lib
-ljemalloc
```

SPEC CPU®2017 Integer Rate Result

Copyright 2017-2020 Standard Performance Evaluation Corporation

Huawei

(Test Sponsor: Peng Cheng Laboratory)

Huawei TaiShan 200 Server (Model 2480)
(2.6 GHz,Huawei Kunpeng 920 7260)

SPECrate®2017_int_base =   628

SPECrate®2017_int_peak = Not Run

CPU2017 License: 5036

Test Sponsor: Peng Cheng Laboratory

Tested by: Peng Cheng Laboratory

Test Date: May-2020

Hardware Availability: Jan-2020

Software Availability: Jul-2020

The flags files that were used to format this result can be browsed at

http://www.spec.org/cpu2017/flags/gcc.2020-06-30.html
http://www.spec.org/cpu2017/flags/PCL-Platform-Settings-Kunpeng-V1.0-revF.html

You can also download the XML flags sources by saving the following links:

http://www.spec.org/cpu2017/flags/gcc.2020-06-30.xml
http://www.spec.org/cpu2017/flags/PCL-Platform-Settings-Kunpeng-V1.0-revF.xml

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

Tested with SPEC CPU®2017 v1.1.0 on 2020-05-28 21:13:25-0400.
Report generated on 2020-06-30 14:41:29 by CPU2017 PDF formatter v6255.
Originally published on 2020-06-30.