# White Paper
# FUJITSU Server PRIMERGY
# Performance Report PRIMERGY RX2520 M4

This document contains a summary of the benchmarks executed for the FUJITSU Server PRIMERGY RX2520 M4

The PRIMERGY RX2520 M4 performance data are compared with the data of other PRIMERGY models and discussed. In addition to the benchmark results, an explanation has been included for each benchmark and for the benchmark environment.

**Version**
1.1
2018/04/10

# Contents



# Document history

## *Version 1.0 (2017/12/08)*

New:

- Technical data
- SPECcpu2006
  Measurements with Intel® Xeon® Processor Scalable Family
- vServCon
  Result for Intel® Xeon® Processor Scalable Family
- STREAM
  Measurements with Intel® Xeon® Processor Scalable Family
- LINPACK
  Measurements with Intel® Xeon® Processor Scalable Family

## *Version 1.1 (2018/04/10)*

Updated:

- vServCon
  Additional measurements with Intel® Xeon® Processor Scalable Family

# Technical data

**PRIMERGY RX2520 M4**

**PRIMERGY RX2520 M4**
PY RX2520 M4 12 x 3.5'

**PRIMERGY RX2520 M4**
PY RX2520 M4 8x 2.5' expandable

Decimal prefixes according to the SI standard are used for measurement units in this white paper (e.g. 1 GB = $10^9$ bytes). In contrast, these prefixes should be interpreted as binary prefixes (e.g. 1 GB = $2^{30}$ bytes) for the capacities of caches and memory modules. Separate reference will be made to any further exceptions where applicable.

| Model | PRIMERGY RX2520 M4 |
|---|---|
| Model versions | PY RX2520 M4 4x 3.5' expandable<br>PY RX2520 M4 12x 3.5'<br>PY RX2520 M4 8x 2.5' expandable<br>PY RX2520 M4 16 x 2.5'<br>PY RX2520 M4 24x 2.5' |
| Form factor | Rack server |
| Chipset | Intel® C624 |
| Number of sockets | 2 |
| Number of processors orderable | 1 or 2 |
| Processor type | Intel® Xeon® Processor Scalable Family |
| Number of memory slots | 12 (6 per processor) |
| Maximum memory configuration | 384 GB |
| Onboard HDD controller | Controller with RAID 0, RAID 1 or RAID 10 for up to 8 SATA HDDs |
| PCI slots | PCI-Express 3.0 x8 × 3<br>PCI-Express 3.0 x16 × 3 |
| Max. number of internal hard disks | PY RX2520 M4 4 x 3.5' expandable : 8<br>PY RX2520 M4 12 x 3.5' : 12<br>PY RX2520 M4 8 x 2.5' expandable : 24<br>PY RX2520 M4 16 x 2.5' : 16<br>PY RX2520 M4 24 x 2.5' : 24 |

| Processors (since system release) | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| Processor | Cores | Threads | Cache [MB] | UPI Speed [GT/s] | Rated Frequency [Ghz] | Max. Turbo Frequency [Ghz] | Max. Memory Frequency [MHz] | TDP [Watt] |
| Xeon Bronze 3104 | 6 | 6 | 8.3 | 9.6 | 1.7 | n/a | 2133 | 85 |
| Xeon Bronze 3106 | 8 | 8 | 11.0 | 9.6 | 1.7 | n/a | 2133 | 85 |
| Xeon Silver 4108 | 8 | 16 | 11.0 | 9.6 | 1.8 | 3.0 | 2400 | 85 |
| Xeon Silver 4110 | 8 | 16 | 11.0 | 9.6 | 2.1 | 3.0 | 2400 | 85 |
| Xeon Silver 4114 | 10 | 20 | 13.8 | 9.6 | 2.2 | 3.0 | 2400 | 85 |
| Xeon Silver 4116 | 12 | 24 | 16.5 | 9.6 | 2.1 | 3.0 | 2400 | 85 |
| Xeon Gold 5115 | 10 | 20 | 13.8 | 9.6 | 2.4 | 3.2 | 2400 | 85 |
| Xeon Gold 5118 | 12 | 24 | 16.5 | 9.6 | 2.3 | 3.2 | 2400 | 105 |
| Xeon Gold 5120 | 14 | 28 | 19.3 | 9.6 | 2.2 | 3.2 | 2400 | 105 |
| Xeon Silver 4112 | 4 | 8 | 8.3 | 9.6 | 2.6 | 3.0 | 2400 | 85 |
| Xeon Gold 5122 | 4 | 8 | 16.5 | 10.4 | 3.6 | 3.7 | 2666 | 105 |

All the processors that can be ordered with the PRIMERGY RX2520 M4, apart from Xeon Bronze 3104 and Xeon Bronze 3106, support Intel® Turbo Boost Technology 2.0. This technology allows you to operate the processor with higher frequencies than the nominal frequency. Listed in the processor table is "Max. Turbo Frequency" for the theoretical maximum frequency with only one active core per processor. The maximum frequency that can actually be achieved depends on the number of active cores, the current consumption, electrical power consumption, and the temperature of the processor.
As a matter of principle, Intel does not guarantee that the maximum turbo frequency can be reached. This is related to manufacturing tolerances, which result in a variance regarding the performance of various examples of a processor model. The range of the variance covers the entire scope between the nominal frequency and the maximum turbo frequency.
The turbo functionality can be set via BIOS option. Fujitsu generally recommends leaving the "Turbo Mode" option set at the standard setting of "Enabled", as performance is substantially increased by the higher frequencies. However, since the higher frequencies depend on general conditions and are not always guaranteed, it can be advantageous to disable the "Turbo Mode" option for application scenarios with intensive use of AVX instructions and a high number of instructions per clock unit, as well as for those that require constant performance or lower electrical power consumption.

| Memory modules (since system release) | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| Memory module | Capacity [GB] | Ranks | Bit width of the memory chips | Frequency [MHz] | Low voltage | Load reduced | Registered | ECC |
| 8 GB (1x8 GB) 1Rx4 DDR4-2666 R ECC | 8 | 1 | 4 | 2666 | | | ✓ | ✓ |
| 16 GB (1x16 GB) 1Rx4 DDR4-2666 R ECC | 16 | 1 | 4 | 2666 | | | ✓ | ✓ |
| 8 GB (1x8 GB) 2Rx8 DDR4-2666 R ECC | 8 | 2 | 8 | 2666 | | | ✓ | ✓ |
| 16 GB (1x16 GB) 2Rx8 DDR4-2666 R ECC | 16 | 2 | 8 | 2666 | | | ✓ | ✓ |
| 16 GB (1x16 GB) 2Rx4 DDR4-2666 R ECC | 16 | 2 | 4 | 2666 | | | ✓ | ✓ |
| 32 GB (1x32 GB) 2Rx4 DDR4-2666 R ECC | 32 | 2 | 4 | 2666 | | | ✓ | ✓ |

| Power supplies (since system release) | Max. number |
|---|---|
| Modular PSU 450 W platinum hp | 2 |
| Modular PSU 800 W platinum hp | 2 |
| Modular PSU 800 W titanium hp | 2 |
| Modular PSU 1200 W platinum hp | 2 |

Some components may not be available in all countries or sales regions.
Detailed technical information is available in the data sheet PRIMERGY RX2520 M4.

# SPECcpu2006

## Benchmark description

SPECcpu2006 is a benchmark which measures the system efficiency with integer and floating-point operations. It consists of an integer test suite (SPECint2006) containing 12 applications and a floating-point test suite (SPECfp2006) containing 17 applications. Both test suites are extremely computing-intensive and concentrate on the CPU and the memory. Other components, such as Disk I/O and network, are not measured by this benchmark.

SPECcpu2006 is not tied to a special operating system. The benchmark is available as source code and is compiled before the actual measurement. The used compiler version and their optimization settings also affect the measurement result.

SPECcpu2006 contains two different performance measurement methods: The first method (SPECint2006 or SPECfp2006) determines the time which is required to process a single task. The second method (SPECint_rate2006 or SPECfp_rate2006) determines the throughput, i.e. the number of tasks that can be handled in parallel. Both methods are also divided into two measurement runs, “base” and “peak”, which differ in the use of compiler optimization. When publishing the results, the base values are always used and the peak values are optional.

| Benchmark | Arithmetic | Type | Compiler optimization | Measurement result | Application |
|---|---|---|---|---|---|
| SPECint2006 | integer | peak | aggressive | Speed | single-threaded |
| SPECint_base2006 | integer | base | conservative | | |
| SPECint_rate2006 | integer | peak | aggressive | Throughput | multi-threaded |
| SPECint_rate_base2006 | integer | base | conservative | | |
| SPECfp2006 | floating point | peak | aggressive | Speed | single-threaded |
| SPECfp_base2006 | floating point | base | conservative | | |
| SPECfp_rate2006 | floating point | peak | aggressive | Throughput | multi-threaded |
| SPECfp_rate_base2006 | floating point | base | conservative | | |

The measurement results are the geometric average from normalized ratio values which have been determined for individual benchmarks. The geometric average - in contrast to the arithmetic average - means that there is a weighting in favor of the lower individual results. Normalized means that the measurement is how fast is the test system compared to a reference system. Value “1” was defined for the SPECint_base2006, SPECint_rate_base2006, SPECfp_base2006, and SPECfp_rate_base2006 results of the reference system. For example, a SPECint_base2006 value of 2 means that the measuring system has handled this benchmark twice as fast as the reference system. A SPECfp_rate_base2006 value of 4 means that the measuring system has handled this benchmark some 4/[# base copies] times faster than the reference system. “# base copies” specifies how many parallel instances of the benchmark have been executed.

Not every SPECcpu2006 measurement is submitted by us for publication at SPEC. This is why the SPEC web pages do not have every result. As we archive the log files for all measurements, we can prove the correct implementation of the measurements at any time.

## Benchmark environment

| System Under Test (SUT) | |
|---|---|
| **Hardware** | |
| Model | PRIMERGY RX2520 M4 |
| Processor | Intel® Xeon® Processor Scalable Family × 2 |
| Memory | 32 GB (1x32 GB) 2Rx4 PC4-2666V R ECC x 24 |
| **Software** | |
| BIOS settings | Xeon Gold 5120<br>HWPM Support = Disabled<br>Intel Virtualization Technology = Disabled<br>Link Frequency Select = 10.4 GT/s<br>Sub NUMA Clustering = Enabled<br>IMC Interleaving = 1-way<br>LLC Dead Line Alloc = Disabled<br>Stale AtoS = Enabled |
| Operating system | SUSE Linux Enterprise Server 12 SP2 (x86_64) |
| Operating system settings | Stack size set to unlimited using "ulimit -s unlimited"<br>cpupower -c all frequency-set -g performance<br>mkdir /home/memory<br>mount -t tmpfs -o size=376g,rw tmpfs /home/memory<br>echo 10000000 > /proc/sys/kernel/sched_min_granularity_ns<br>echo 15000000 > /proc/sys/kernel/sched_wakeup_granularity_ns<br>cpu idle state set with:<br>cpupower idle-set -d 1<br>cpupower idle-set -d 2<br><br>SPECfp_rate_base2006:<br>echo 0 > /proc/sys/kernel/numa_balancing |
| Compiler | SPECint_rate_base2006:<br>C/C++: Version 18.0.0.128 of Intel C/C++ Compiler for Linux<br><br>SPECfp_rate_base2006:<br>C/C++: Version 17.0.3.191 of Intel C/C++ Compiler for Linux<br>Fortran: Version 17.0.3.191 of Intel Fortran Compiler for Linux |

Some components may not be available in all countries or sales regions.

## Benchmark results

In terms of processors, the benchmark result depends primarily on the size of the processor cache, the support for Hyper-Threading, the number of processor cores, and the processor frequency. In the case of processors with Turbo mode, the number of cores, which are loaded by the benchmark, determines the maximum processor frequency that can be achieved. In the case of single-threaded benchmarks, which largely load one core only, the maximum processor frequency that can be achieved is higher than with multi-threaded benchmarks.

This result of cursive is an estimated value.

| Processor | Number of processors | SPECint_base2006 | SPECint2006 | SPECint_rate_base2006 | SPECint_rate2006 |
|---|---|---|---|---|---|
| Xeon Bronze 3104 | 2 | | | *327* | |
| Xeon Bronze 3106 | 2 | | | *437* | |
| Xeon Silver 4108 | 2 | | | *636* | |
| Xeon Silver 4110 | 2 | | | *706* | |
| Xeon Silver 4114 | 2 | | | *904* | |
| Xeon Silver 4116 | 2 | | | *1052* | |
| Xeon Gold 5115 | 2 | | | *973* | |
| Xeon Gold 5118 | 2 | | | *1151* | |
| Xeon Gold 5120 | 2 | | | 1300 | |
| Xeon Silver 4112 | 2 | | | *423* | |
| Xeon Gold 5122 | 2 | | | *543* | |

| Processor | Number of processors | SPECfp_base2006 | SPECfp2006 | SPECfp_rate_base2006 | SPECfp_rate2006 |
|---|---|---|---|---|---|
| Xeon Bronze 3104 | 2 | | | *354* | |
| Xeon Bronze 3106 | 2 | | | *468* | |
| | | | | | |
| Xeon Silver 4108 | 2 | | | *622* | |
| Xeon Silver 4110 | 2 | | | *671* | |
| Xeon Silver 4114 | 2 | | | *815* | |
| Xeon Silver 4116 | 2 | | | *912* | |
| | | | | | |
| Xeon Gold 5115 | 2 | | | *850* | |
| Xeon Gold 5118 | 2 | | | *965* | |
| Xeon Gold 5120 | 2 | | | 1050 | |
| | | | | | |
| Xeon Silver 4112 | 2 | | | *418* | |
| | | | | | |
| Xeon Gold 5122 | 2 | | | *519* | |

The following two diagrams illustrate the throughput of the PRIMERGY RX2520 M4 in comparison to its predecessor PRIMERGYRX2520 M1, in their respective most performant configuration.

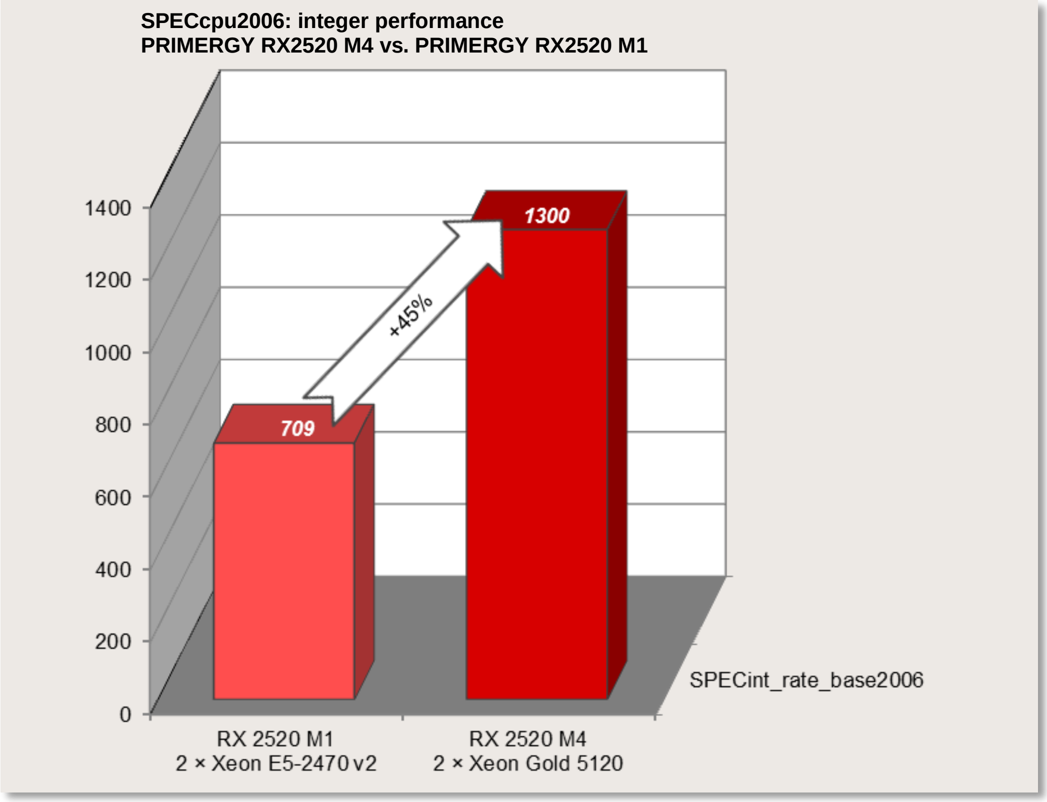


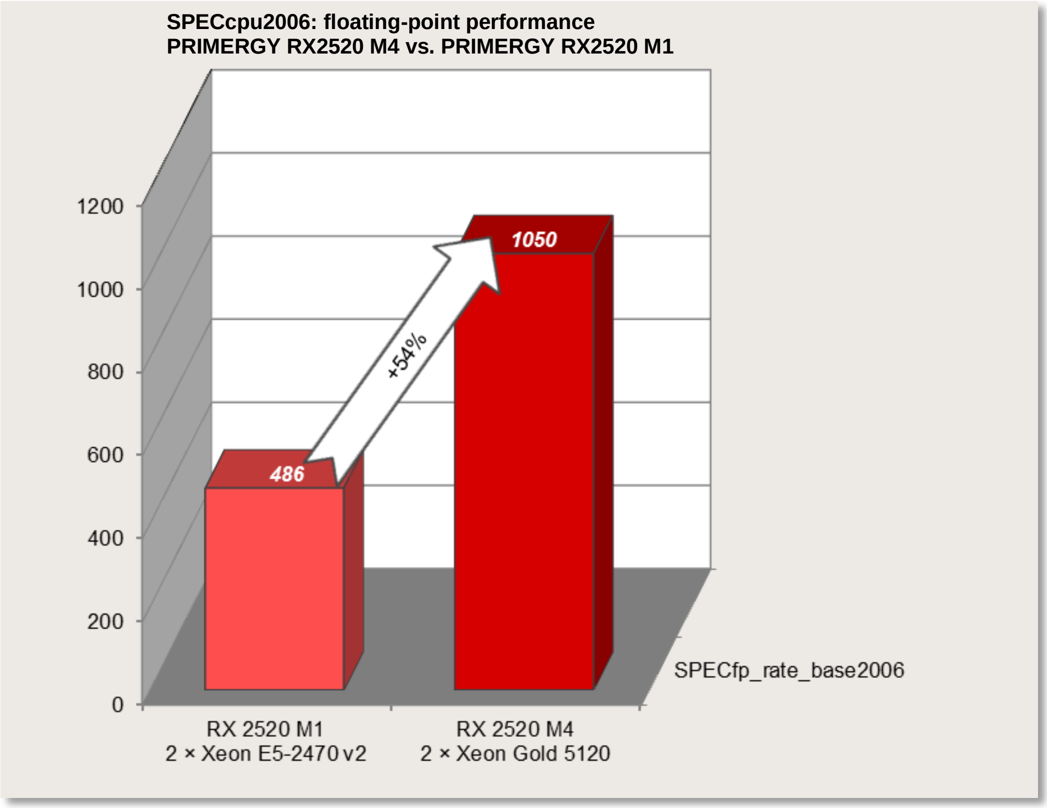

# vServCon

## Benchmark description

vServCon is a benchmark used by Fujitsu to compare server configurations with hypervisor with regard to their suitability for server consolidation. This allows both the comparison of systems, processors and I/O technologies as well as the comparison of hypervisors, virtualization forms, and additional drivers for virtual machines.

vServCon is not a new benchmark in the true sense of the word. It is more a framework that combines already established benchmarks (or in modified form) as workloads in order to reproduce the load of a consolidated and virtualized server environment. Three proven benchmarks are used which cover the application scenarios database, application server, and web server.

| Application scenario | Benchmark | No. of logical CPU cores | Memory |
|---|---|---|---|
| Database | Sysbench (adapted) | 2 | 1.5 GB |
| Java application server | SPECjbb (adapted, with 50% - 60% load) | 2 | 2 GB |
| Web server | WebBench | 1 | 1.5 GB |

Each of the three application scenarios is allocated to a dedicated virtual machine (VM). A fourth machine, the so-called idle VM, is added to these. These four VMs make up a “tile”. Depending on the performance capability of the underlying server hardware, you may as part of a measurement also have to start several identical tiles in parallel in order to achieve a maximum performance score.

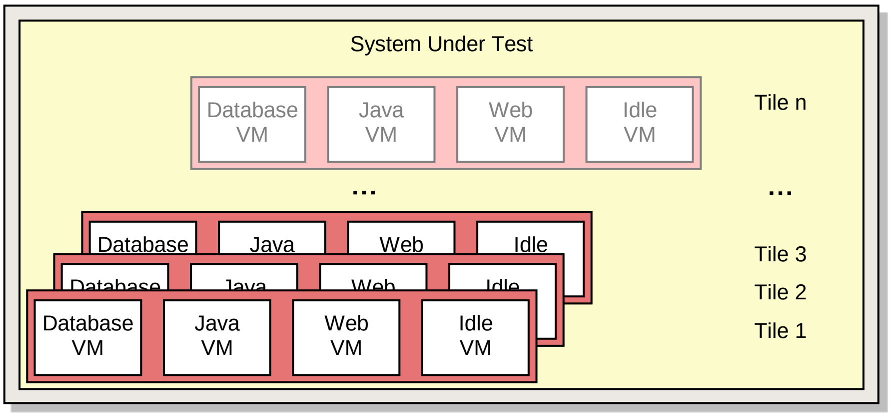


Each of the three vServCon application scenarios provides a specific benchmark result in the form of application-specific transaction rates for the respective VM. In order to derive a normalized score, the individual benchmark result for one tile is put in relation to the respective result of a reference system. The resulting relative performance value is then suitably weighted and finally added up for all VMs and tiles. The outcome is a score for this tile number.

As a general rule, start with one tile, and this procedure is performed for an increasing number of tiles until no further significant increase in this vServCon score occurs. The final vServCon score is then the maximum of the vServCon scores for all tile numbers. This score thus reflects the maximum total throughput that can be achieved by running the mix defined in vServCon that consists of numerous VMs up to the possible full utilization of CPU resources. This is why the measurement environment for vServCon measurements is designed in such a way that only the CPU is the limiting factor and that no limitations occur as a result of other resources.

The progression of the vServCon scores for the tile numbers provides useful information about the scaling behavior of the “System under Test”.

A detailed description of vServCon is in the document: Benchmark Overview vServCon.

## Benchmark environment

The typical measurement set-up is illustrated below:

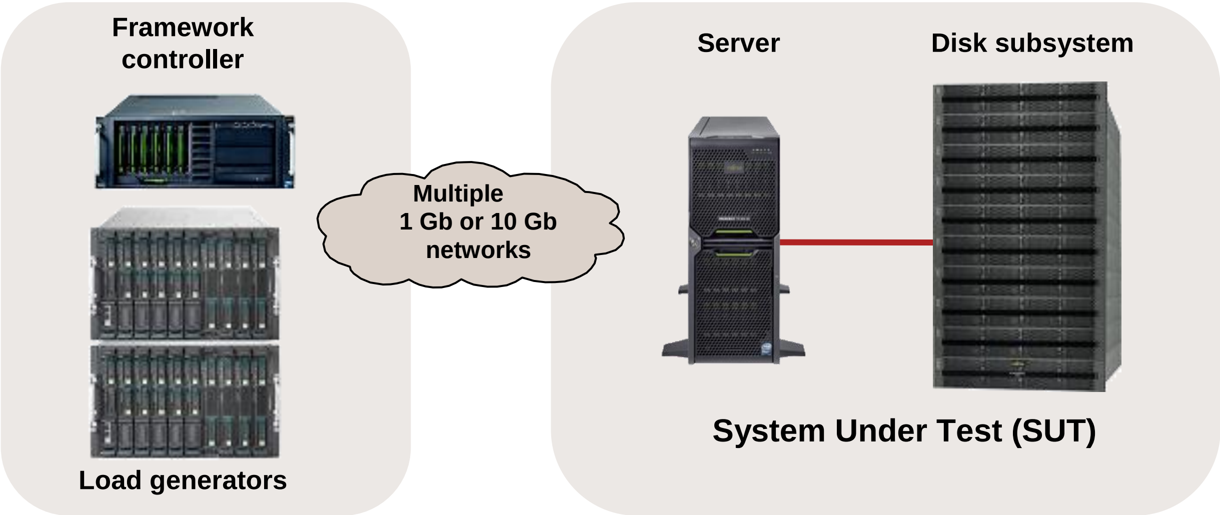


All results were determined by way of example on a PRIMERGY RX2530 M4.

| System Under Test (SUT) | |
|---|---|
| **Hardware** | |
| Processor | 2 × Intel® Xeon® Processor Scalable Family |
| Memory | 12 × 32 GB (1x32 GB) 2Rx4 DDR4-2666 R ECC |
| Network interface | 1 × Emulex OneConnect OCe14000 Dual Port Adapter with 10 Gb SFP+ DynamicLoM interface module |
| Disk subsystem | 1 ×dual-channel FC controller Emulex LPe160021<br>LINUX/LIO based flash storage system |
| **Software** | |
| Operating system | VMware ESXi 6.5.0b Build 5146846 |

| Load generator (incl. Framework controller) | |
|---|---|
| **Hardware (Shared)** | |
| Enclosure | 5 x PRIMERGY RX2530 M2 |
| **Hardware** | |
| Processor | 2 × XeonE5-2683 v4 |
| Memory | 128 GB |
| Network interface | 3 × 1 Gbit LAN |
| **Software** | |
| Operating system | VMware ESXi 6.0.0 U1b Build 3380124 |

| Load generator VM (on various servers) | |
|---|---|
| **Hardware** | |
| Processor | 1 × logical CPU |
| Memory | 4048 MB |
| Network interface | 2 × 1 Gbit/s LAN |
| **Software** | |
| Operating system | Microsoft Windows Server 2008 Standard Edition 32bit |

Some components may not be available in all countries or sales regions.

## Benchmark results

The PRIMERGY dual-socket rack and tower systems dealt with here are based on processors of the Intel® Xeon® Processor Scalable Family. The features of the processors are summarized in the section "Technical data".

The available processors of these systems with their results can be seen in the following table.

The result of cursive is an estimated value.

| Processor | | | Score | #Tiles |
|---|---|---|---|---|
| Intel® Xeon® Processor Scalable Family | 4 Cores Hyper-Threading, Turbo-Modus | Silver 4112 | 9.00 | 8 |
| | | Gold 5122 | 11.5 | 10 |
| | 6 Cores | Bronze 3104 | 6.97 | 7 |
| | 8 Cores | Bronze 3106 | 8.78 | 7 |
| | 8 Cores Hyper-Threading, Turbo-Modus | Silver 4108 | 13.5 | 11 |
| | | Silver 4110 | 16.3 | 12 |
| | 10 Cores Hyper-Threading, Turbo-Modus | Silver 4114 | 19.2 | 16 |
| | | Gold 5115 | 20.8 | 16 |
| | 12 Cores Hyper-Threading, Turbo-Modus | Silver 4116 | 22.4 | 18 |
| | | Gold 5118 | 24.5 | 19 |
| | 14 Cores Hyper-Threading, Turbo-Modus | Gold 5120 | 27.9 | 19 |

The following diagram compares the virtualization performance values that can be achieved with the processors reviewed here.

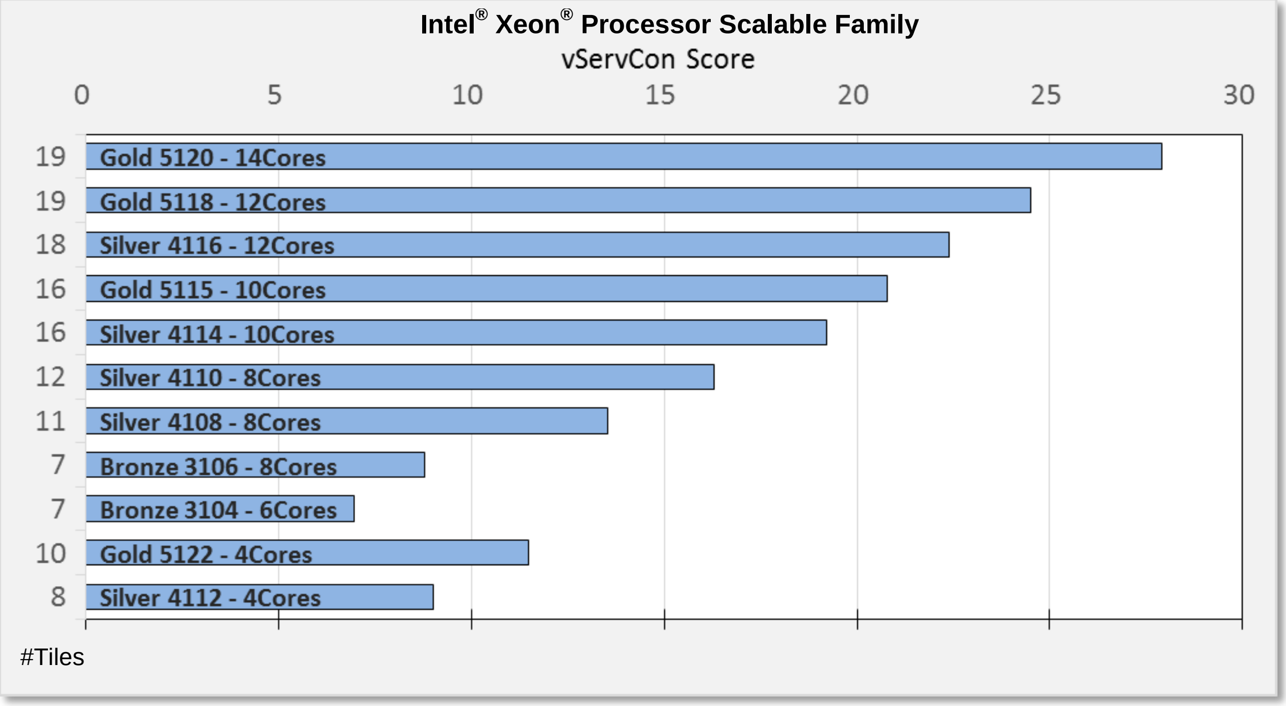


The relatively large performance differences between the processors can be explained by their features. The values scale on the basis of the number of cores, the size of the L3 cache and the CPU clock frequency and as a result of the features of Hyper-Threading and turbo mode, which are available in most processor types. Furthermore, the data transfer rate between processors (“UPI Speed”) also determines performance.

A low performance can be seen in the Xeon Bronze 3104 and Bronze 3106 processors, as they have to manage without Hyper-Threading (HT) and turbo mode (TM). In principle, these weakest processors are only to a limited extent suitable for the virtualization environment.

Within a group of processors with the same number of cores scaling can be seen via the CPU clock frequency.

As a matter of principle, the memory access speed also influences performance. A guideline in the virtualization environment for selecting main memory is that sufficient quantity is more important than the speed of the memory accesses. The vServCon scaling measurements presented here were all performed with a memory access speed – depending on the processor type – of at most 2666 MHz.

The next diagram illustrates the virtualization performance for increasing numbers of VMs based on the Xeon Gold 5120 (14 core) processors.

In addition to the increased number of physical cores, Hyper-Threading, which is supported by almost all processors of the Intel® Xeon® Processor Scalable Product Family, is an additional reason for the high number of VMs that can be operated. As is known, a physical processor core is consequently divided into two logical cores so that the number of cores available for the hypervisor is doubled. This standard feature thus generally increases the virtualization performance of a system.

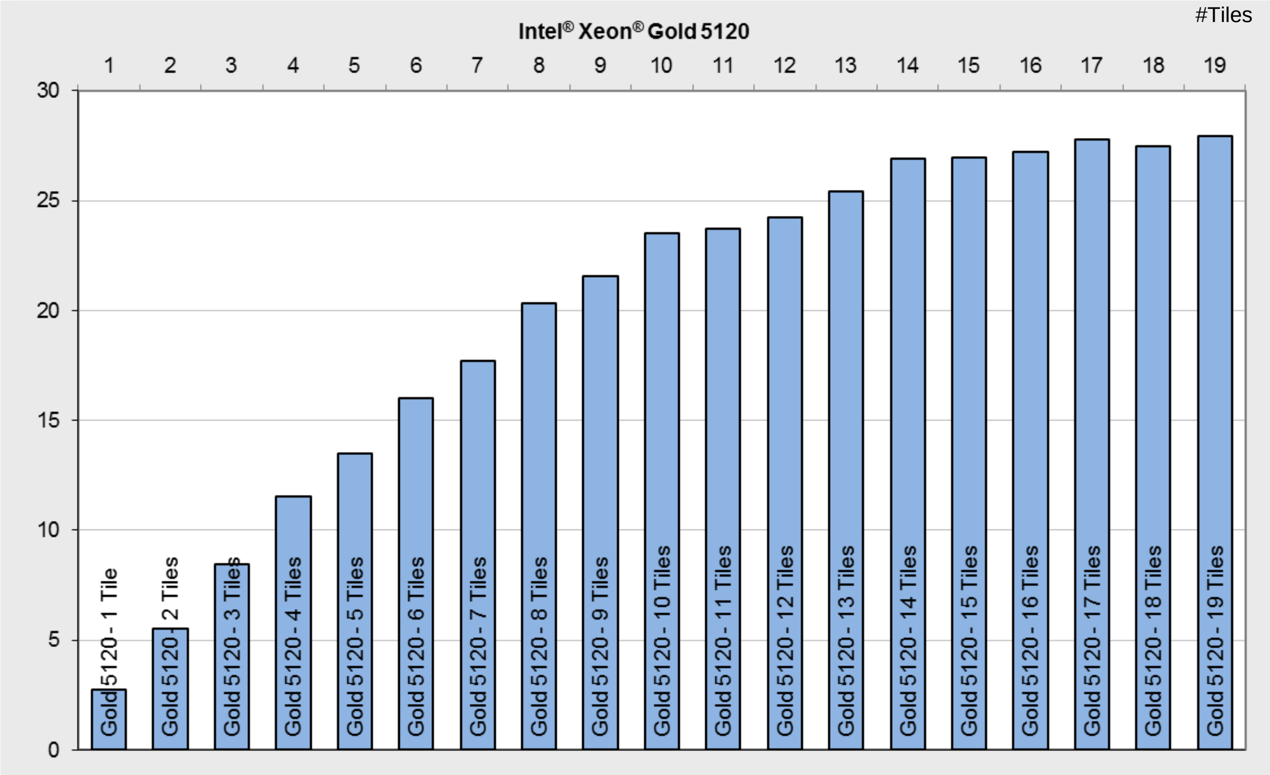


The previous diagram examined the total performance of all application VMs of a host. However, studying the performance from an individual application VM viewpoint is also interesting. This information is in the previous diagram. For example, the total optimum is reached in the above Xeon Gold 5120 situation with 57 application VMs (19 tiles, not including the idle VMs). The low load case is represented by three application VMs (one tile, not including the idle VM). Remember, the vServCon score for one tile is an average value across the three application scenarios in vServCon. This average performance of one tile drops when changing from the low load case to the total optimum of the vServCon score - from 2.77 to 27.9/19=1.47, i.e. to 53%. The individual types of application VMs can react very differently in the high load situation. It is thus clear that in a specific situation the performance requirements of an individual application must be balanced against the overall requirements regarding the numbers of VMs on a virtualization host.

# STREAM

## Benchmark description

STREAM is a synthetic benchmark that has been used for many years to determine memory throughput and was developed by John McCalpin during his professorship at the University of Delaware. Today STREAM is supported at the University of Virginia, where the source code can be downloaded in either Fortran or C. STREAM continues to play an important role in the HPC environment in particular. It is for example an integral part of the HPC Challenge benchmark suite.

The benchmark is designed in such a way that it can be used both on PCs and on server systems. The unit of measurement of the benchmark is GB/s, i.e. the number of gigabytes that can be read and written per second.

STREAM measures the memory throughput for sequential accesses. These can generally be performed more efficiently than accesses that are randomly distributed on the memory, because the processor caches are used for sequential access.

Before execution the source code is adapted to the environment to be measured. Therefore, the size of the data area must be at least 12 times larger than the total of all last-level processor caches so that these have as little influence as possible on the result. The OpenMP program library is used to enable selected parts of the program to be executed in parallel during the runtime of the benchmark, consequently achieving optimal load distribution to the available processor cores.

During implementation the defined data area, consisting of 8 byte elements, it is successively copied to four types, and arithmetic calculations are also performed to some extent.

| Type | Execution | Bytes per step | Floating-point calculation per step |
|---|---|---|---|
| COPY | a(i) = b(i) | 16 | 0 |
| SCALE | a(i) = q × b(i) | 16 | 1 |
| SUM | a(i) = b(i) + c(i) | 24 | 1 |
| TRIAD | a(i) = b(i) + q × c(i) | 24 | 2 |

The throughput is output in GB/s for each type of calculation. The differences between the various values are usually only minor on modern systems. In general, only the determined TRIAD value is used as a comparison.

The measured results primarily depend on the clock frequency of the memory modules; the processors influence the arithmetic calculations.

This chapter specifies throughputs on a basis of 10 (1 GB/s = $10^9$ Byte/s).

## Benchmark environment

| System Under Test (SUT) | |
|---|---|
| **Hardware** | |
| Model | PRIMERGY RX2520 M4 |
| Processor | 2 × Intel® Xeon® Processor Scalable Family |
| Memory | 12 × 32 GB (1x32 GB) 2Rx4 PC4-2666V R ECC |
| **Software** | |
| BIOS settings | HWPM Support = Disabled<br>Intel Virtualization Technology = Disabled<br>Sub NUMA Clustering = Disabled<br>IMC Interleaving = 2-way<br>LLC Dead Line Alloc = Disabled<br>Stale AtoS = Enabled |
| Operating system | SUSE Linux Enterprise Server 12 SP2 (x86_64) |
| Operating system settings | Transparent Huge Pages inactivated<br>sched_cfs_bandwidth_slice_us = 50000<br>sched_latency_ns = 240000000<br>sched_migration_cost_ns = 5000000<br>sched_min_granularity_ns = 100000000<br>sched_wakeup_granularity_ns = 150000000<br>cpupower -c all frequency-set -g performance<br>cpupower idle-set -d 1<br>cpupower idle-set -d 2<br>cpupower idle-set -d 3<br>echo 0 > /proc/sys/kernel/numa_balancing<br>echo 1 >  /proc/sys/vm/drop_caches<br>ulimit -s unlimited<br>nohz_full = 1-xx<br>run with avx2 |
| Compiler | Version 17.0.0.098 of Intel C++ Compiler for Linux |
| Benchmark | Stream.c Version 5.10 |

Some components may not be available in all countries or sales regions.

## Benchmark results

This results in italic are estimated values.

| Processor | Memory Frequency [MHz] | Max. Memory Bandwidth [GB/s] | Cores | Processor Frequency [GHz] | Number of Processors | TRIAD [GB/s] |
|---|---|---|---|---|---|---|
| Xeon Bronce 3104 | 2133 | 102.4 | 6 | 1.7 | 2 | *78.5* |
| Xeon Bronce 3106 | 2133 | 102.4 | 8 | 1.7 | 2 | *106* |
| | | | | | | |
| Xeon Silver 4108 | 2400 | 115.2 | 8 | 1.8 | 2 | *118* |
| Xeon Silver 4110 | 2400 | 115.2 | 8 | 2.1 | 2 | *118* |
| Xeon Silver 4114 | 2400 | 115.2 | 10 | 2.2 | 2 | *139* |
| Xeon Silver 4116 | 2400 | 115.2 | 12 | 2.1 | 2 | *150* |
| | | | | | | |
| Xeon Gold 5115 | 2400 | 115.2 | 10 | 2.4 | 2 | *141* |
| Xeon Gold 5118 | 2400 | 115.2 | 12 | 2.3 | 2 | *155* |
| Xeon Gold 5120 | 2400 | 115.2 | 14 | 2.2 | 2 | 167 |
| | | | | | | |
| Xeon Silver 4112 | 2400 | 115.2 | 4 | 2.6 | 3 | *83.5* |
| | | | | | | |
| Xeon Gold 5122 | 2666 | 128.0 | 4 | 3.6 | 4 | *100* |

The following diagram illustrates the throughput of the PRIMERGY RX2520 M4 in comparison to its predecessor, the PRIMERGY RX2520 M1

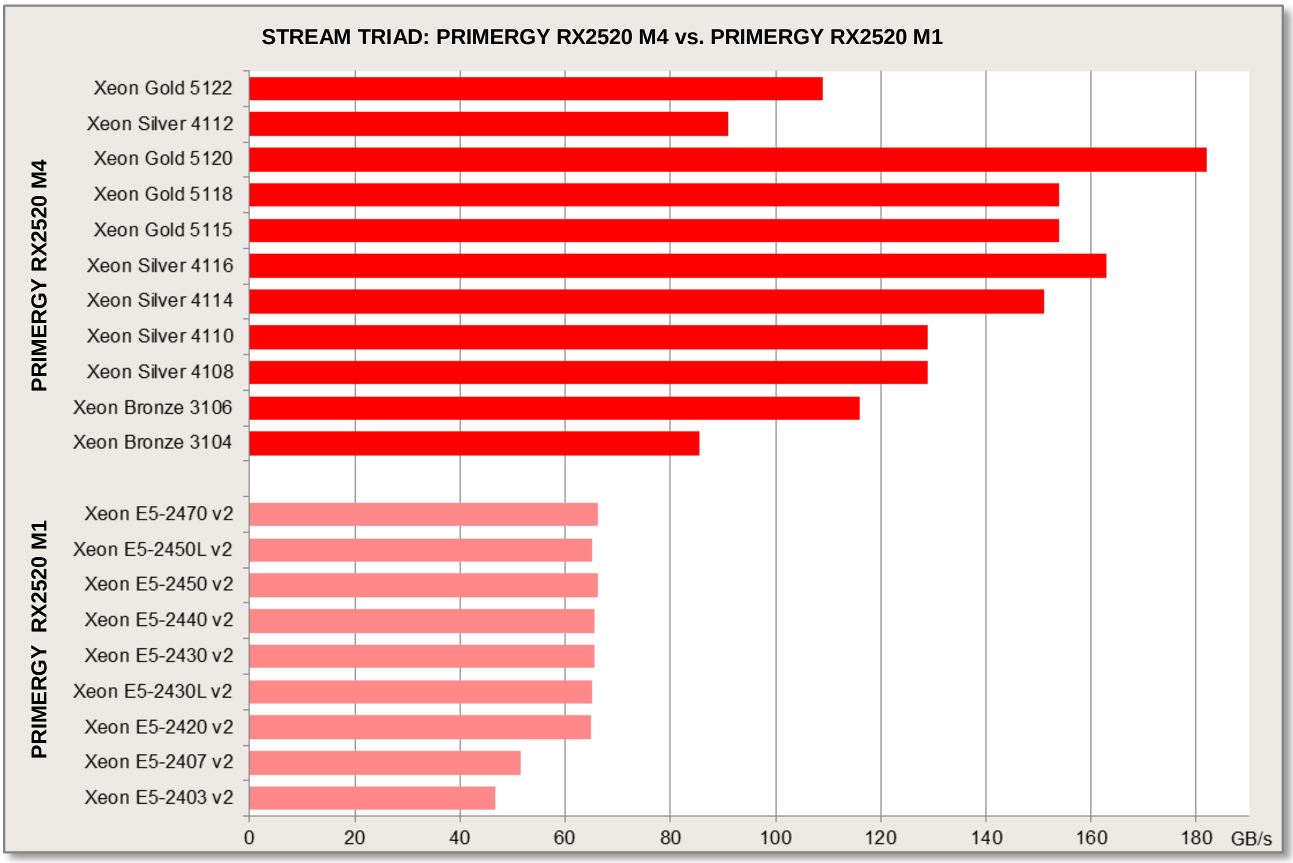

# LINPACK

## Benchmark description

LINPACK was developed in the 1970s by Jack Dongarra and some other people to show the performance of supercomputers. The benchmark consists of a collection of library functions for the analysis and solution of linear system of equations. A description can be found in the document http://www.netlib.org/utk/people/JackDongarra/PAPERS/hplpaper.pdf.

LINPACK can be used to measure the speed of computers when solving a linear equation system. For this purpose, an n × n matrix is set up and filled with random numbers between -2 and +2. The calculation is then performed via LU decomposition with partial pivoting.

A memory of $8n^2$ bytes is required for the matrix. In case of an n × n matrix the number of arithmetic operations required for the solution is $^2/_3n^3 + 2n^2$. Thus, the choice of n determines the duration of the measurement: a doubling of n results in an approximately eight-fold increase in the duration of the measurement. The size of n also has an influence on the measurement result itself. As n increases, the measured value asymptotically approaches a limit. The size of the matrix is therefore usually adapted to the amount of memory available. Furthermore, the memory bandwidth of the system only plays a minor role for the measurement result, but a role that cannot be fully ignored. The processor performance is the decisive factor for the measurement result. Since the algorithm used permits parallel processing, in particular the number of processors used and their processor cores are - in addition to the clock rate - of outstanding significance.

LINPACK is used to measure how many floating point operations were carried out per second. The result is referred to as **Rmax** and specified in GFlops (Giga Floating Point Operations per Second).

An upper limit, referred to as **Rpeak**, for the speed of a computer can be calculated from the maximum number of floating point operations that its processor cores could theoretically carry out in one clock cycle.

*Rpeak = Maximum number of floating point operations per clock cycle*
*× Number of processor cores of the computer*
*× Rated processor frequency [GHz]*

LINPACK is classed as one of the leading benchmarks in the field of high performance computing (HPC). LINPACK is one of the seven benchmarks currently included in the HPC Challenge benchmark suite, which takes other performance aspects in the HPC environment into account.

Manufacturer-independent publication of LINPACK results is possible at http://www.top500.org/. The use of a LINPACK version based on HPL is prerequisite for this (see http://www.netlib.org/benchmark/hpl/).

Intel offers a highly optimized LINPACK version (shared memory version) for individual systems with Intel processors. Parallel processes communicate here via "shared memory", i.e. jointly used memory. Another version provided by Intel is based on HPL (High Performance Linpack). Intercommunication of the LINPACK processes here takes place via OpenMP and MPI (Message Passing Interface). This enables communication between the parallel processes - also from one computer to another. Both versions can be downloaded from http://software.intel.com/en-us/articles/intel-math-kernel-library-linpack-download/.

Manufacturer-specific LINPACK versions also come into play when graphics cards for General Purpose Computation on Graphics Processing Unit (GPGPU) are used. These are based on HPL and include extensions which are needed for communication with the graphics cards.

## Benchmark environment

| System Under Test (SUT) | |
|---|---|
| **Hardware** | |
| Model | PRIMERGY RX2520 M4 |
| Processor | Intel® Xeon® Processor Scalable Family × 2 |
| Memory | 32 GB (1x32 GB) 2Rx4 PC4-2666V R ECC × 12 |
| **Software** | |
| BIOS settings | HyperThreading = Disabled<br>Link Frequency Select = 10.4 GT/s<br>HWPM Support = Disabled<br>Intel Virtualization Technology = Disabled<br>Sub NUMA Clustering = Disabled<br>IMC Interleaving = 1-way<br>LLC Dead Line Alloc = Disabled<br>Stale AtoS = Enabled |
| Operating system | SUSE Linux Enterprise Server 12 SP2 (x86_64) |
| Operating system settings | Transparent Huge Pages inactivated<br>sched_cfs_bandwidth_slice_us = 50000<br>sched_latency_ns = 240000000<br>sched_migration_cost_ns = 5000000<br>sched_min_granularity_ns = 100000000<br>sched_wakeup_granularity_ns = 150000000<br>cpupower -c all frequency-set -g performance<br>aio-max-nr = 1048576<br>ulimit -s unlimited<br>nohz_full = 1-xx<br>run with avx2 |
| Benchmark | MPI version: Intel® Math Kernel Library Benchmarks for Linux OS (l_mklb_p_2017.3.017) |

Some components may not be available in all countries or sales regions.

## Benchmark results

This result of cursive is  an estimated value from the result of RX2530 M4.

| **Processor** | **Cores** | **Processor Frequency [GHz]** | **Number of Processors** | **Rpeak [GFlops]** | **Rmax [GFlops]** | **Efficiency** |
|---|---|---|---|---|---|---|
| Xeon Bronze 3104 | 6 | 1.7 | 2 | 326 | ***239*** | 73% |
| Xeon Bronze 3106 | 8 | 1.7 | 2 | 435 | ***319*** | 73% |
| Xeon Silver 4108 | 8 | 1.8 | 2 | 461 | ***299*** | 65% |
| Xeon Silver 4110 | 8 | 2.1 | 2 | 538 | ***514*** | 96% |
| Xeon Silver 4114 | 10 | 2.2 | 2 | 704 | ***673*** | 96% |
| Xeon Silver 4116 | 12 | 2.1 | 2 | 806 | ***768*** | 95% |
| Xeon Gold 5115 | 10 | 2.4 | 2 | 768 | ***686*** | 89% |
| Xeon Gold 5118 | 12 | 2.3 | 2 | 883 | ***842*** | 95% |
| Xeon Gold 5120 | 14 | 2.2 | 2 | 986 | **705** | 72% |
| Xeon Silver 4112 | 4 | 2.6 | 2 | 333 | ***316*** | 95% |
| Xeon Gold 5122 | 4 | 3.6 | 2 | 922 | ***739*** | 80% |

*Rmax = Measurement result*

*Rpeak = Maximum number of floating point operations per clock cycle*
*× Number of processor cores of the computer*
*× Rated frequency [GHz]*

As explained in the section "Technical Data", Intel generally does not guarantee that the maximum turbo frequency can be reached in the processor models due to manufacturing tolerances. A further restriction applies for workloads, such as those generated by LINPACK, with intensive use of AVX instructions and a high number of instructions per clock unit. Here the frequency of a core can also be limited if the upper limits of the processor for power consumption and temperature are reached before the upper limit for the current consumption. This can result in the achievement of a lower performance with turbo mode than without turbo mode. In such cases, you should disable the turbo functionality via BIOS option.

# Literature

## PRIMERGY Servers

http://primergy.com/

## PRIMERGY RX2520 M4

This White Paper:

- http://docs.ts.fujitsu.com/dl.aspx?id=4b1a7fc1-dc02-41fc-a2d3-d10e6d545a88
- http://docs.ts.fujitsu.com/dl.aspx?id=590ef6ea-2bad-45ab-ae80-9ff2c66608a7

Data sheet
http://docs.ts.fujitsu.com/dl.aspx?id=57157f8d-4b3d-412f-b262-0687398c8aa6

## PRIMERGY Performance

http://www.fujitsu.com/fts/x86-server-benchmarks

## SPECcpu2006

http://www.spec.org/osg/cpu2006

Benchmark overview SPECcpu2006
http://docs.ts.fujitsu.com/dl.aspx?id=1a427c16-12bf-41b0-9ca3-4cc360ef14ce

## vServCon

Benchmark Overview vServCon
http://docs.ts.fujitsu.com/dl.aspx?id=b953d1f3-6f98-4b93-95f5-8c8ba3db4e59

## STREAM

http://www.cs.virginia.edu/stream/

## LINPACK

The LINPACK Benchmark: Past, Present, and Future
http://www.netlib.org/utk/people/JackDongarra/PAPERS/hplpaper.pdf

TOP500
http://www.top500.org/

HPL - A Portable Implementation of the High-Performance Linpack Benchmark for Distributed-Memory Computers
http://www.netlib.org/benchmark/hpl/

Intel Math Kernel Library – LINPACK Download
http://software.intel.com/en-us/articles/intel-math-kernel-library-linpack-download/

# Contact

| **FUJITSU** |
| --- |
| Website: http://www.fujitsu.com/ |
| **PRIMERGY Product Marketing** |
| mailto:Primergy-PM@ts.fujitsu.com |
| **PRIMERGY Performance and Benchmarks** |
| mailto:primergy.benchmark@ts.fujitsu.com |



2018/04/10 WW EN