White Paper
FUJITSU Server PRIMERGY
Performance Report PRIMERGY GX2460 M1

This document provides an overview of benchmarks executed on the FUJITSU Server PRIMERGY GX2460 M1.

Explains PRIMERGY GX2460 M1 performance data in comparison to other PRIMERGY models. In addition to the benchmark results, the explanation for each benchmark and benchmark environment are also included.

Version
1.1
2021/07/28

Contents

Document history

Version 1.0 (2021/05/28)

New:

- Technical data
- SPECcpu2017, STREAM, LINPACK
  Measured with AMD EPYC 7002 Series Processors

Version 1.1 (2021/07/28)

Updated:

- Contact information and URLs
  Updated to the latest one
- Minor correction

Technical data

PRIMERGY GX2460 M1

In this document, the power of 10 (example: 1 GB = 10^9 bytes) is used to indicate the capacity of the internal storage, and the power of 2 (example: 1 GB = 2^{30} bytes) is used to indicate the capacity of the cache or memory module. Any other exceptional notation will be specified separately.

モデル	PRIMERGY GX2460 M1
Model versions	PY GX2460 M1
Form factor	Rack server
Number of sockets	2
Number of processors orderable	2
Processor type	AMD EPYC 7002 Series Processors
Number of memory slots	16
Maximum memory configuration	1024GB
Onboard HDD controller	SATA Controller embedded on CPU
Max. number of internal hard disks	8 x BC-SATA HDD/ PCIe SSD
PCI slots	6 x PCI-Express 4.0(x16)

Processor							
Processor model	Number of cores	Number of threads	Cache [MB]	Rated frequency [GHz]	Maximu turbo frequency [GHz]	Maximum memory frequency [MHz]	TDP [W]
EPYC 7502	32	64	128	2.50	3.35	3200	180
EPYC 7452	32	64	128	2.35	3.35	3200	155
EPYC 7402	24	48	128	2.80	3.35	3200	180
EPYC 7352	24	48	128	2.30	3.20	3200	155
EPYC 7302	16	32	128	3.00	3.30	3200	155
EPYC 7282	16	32	64	2.80	3.20	3200	120
EPYC 7262	8	16	128	3.20	3.40	3200	155
EPYC 7252	8	16	64	3.10	3.20	3200	120

All processors that can be ordered with the PRIMERGY GX2460 M1 support AMD Turbo Core Technology. This technology allows you to operate the processor with higher frequencies than the rated frequency. The maximum frequency that can actually be achieved depends on the type of applications and the processing load.
The turbo functionality can be set in the BIOS option. Generally, Fujitsu generally recommends leaving the [Core Performance Boost] option set at the standard setting of [Auto], as performance is substantially increased by the higher frequencies. However, the higher frequencies depend on the operating conditions mentioned above and is not always guaranteed. If you need stable performance or want to reduce power consumption, it may be beneficial to set [Core Performance Boost] to disable to disable Turbo function.

Memory modules							
Type	Capacity [GB]	Number of ranks	Bit width of the memory chips	Frequency [MHz]	Load reduced	Registered	ECC
16 GB (1x16 GB) 1Rx4 DDR4-3200 RDIMM	16	1	4	3200		✓	✓
16 GB (1x16 GB) 2Rx8 DDR4-3200 RDIMM	16	1	4	3200		✓	✓
32 GB (1x32 GB) 2Rx4 DDR4-3200 RDIMM	32	2	4	3200		✓	✓
64 GB (1x64 GB) 2Rx4 DDR4-3200 RDIMM	32	2	4	3200		✓	✓
64 GB (1x64 GB) 4Rx4 DDR4-3200 LRDIMM	32	4	4	3200	✓	✓	✓

Power supplies	Maximum number
Standard PSU 2200W	2

Includes components that will be supported after the system release. Also, some components may not be available in all countries or sales regions.

Detailed technical information is available in the data sheet PRIMERGY GX2460 M1.

SPECcpu2017

Benchmark description

SPECcpu2017 is a benchmark which measures the system efficiency with integer and floating-point operations. It consists of an integer test suite (SPECrate 2017 Integer, SPECspeed 2017 Integer) containing 10 applications and a floating-point test suite (SPECrate 2017 Floating Point, SPECspeed 2017 Floating Point) containing 14 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.

SPECcpu2017 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.

SPECcpu2017 contains two different performance measurement methods. The first method (SPECspeed 2017 Integer or SPECspeed 2017 Floating Point) determines the time which is required to process a single task. The second method (SPECrate 2017 Integer or SPECrate 2017 Floating Point) 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." They differ in the use of compiler optimization. When publishing the results, the base values are always used and the peak values are optional.

Benchmark	Number of single benchmarks	Arithmetic	Compiler optimization	Measurement result	
SPECspeed2017_int_peak	10	Integer	Aggressive (peak)	Speed	Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_peak	10		Conservative (base)		Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_peak	10		Aggressive (peak)	Throughput	Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_peak	10		Conservative (base)		Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_peak	10	Floating point	Aggressive (peak)	Speed	Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_peak	10		Conservative (base)		Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_peak	13		Aggressive (peak)	Throughput	Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_peak	13		Conservative (base)		Performance
SPECspeed2017_int_energy_peak					Power efficiency

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. For example, value "1" was defined for the SPECspeed2017_int_base, SPECrate2017_int_base, SPECspeed2017_fp_base, and SPECrate2017_fp_base results of the reference system. A SPECspeed2017_int_base value of 2 means that the measuring system has handled this benchmark twice as fast as the reference system. A SPECrate2017_fp_base value of 4 means that the measuring system has handled this benchmark about 4/[# base copies] times faster than the reference system. "# base copies" specifies how many parallel instances of the benchmark have been executed.

Not every SPECcpu2017 measurement is submitted by Fujitsu for publication at SPEC. This is why the SPEC web pages do not have every result. As Fujitsu archives the log files for all measurements, it is possible to prove the correct implementation of the measurements at any time.

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMERGY GX2460 M1
Processor	2 x AMD EPYC 7002 Series Processors
Memory	16 x 32GB (1x32GB) 2Rx4 PC4-3200AA-L
Software	
BIOS settings	Determinism Slider = Power NUMA nodes per socket = NPS4 SVM Mode = Disabled “cTDP” and “Package Power Limit” were set: EPYC 7502: 200 EPYC 7452: 180 EPYC 7402: 200 EPYC 7352: 180 EPYC 7302: 180 EPYC 7282: 150 EPYC 7262: 180 EPYC 7252: 150
Operating system	SUSE Linux Enterprise Server 15 SP2 5.3.18-22-default
Compiler	C/C++/Fortran: Version 2.0.0 of AOCC

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

Benchmark results

In terms of processors, the benchmark results depend primarily on the size of the processor cache, 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.

The results with "est." are estimated values.

Processors	Number of cores	Cache [MB]	Rated freqyency [GHz]	Maximum memory frequency [MHz]	Number of memory channels	Number of chips	SPECrate2017 int_base	SPECrate2017 fp_base
EPYC 7502	32	128	2.50	3200	8	2	**413**	**362**
EPYC 7452	32	128	2.35	3200	8	2	**412**	**361**
EPYC 7402	24	128	2.80	3200	8	2	**343**	**336**
EPYC 7352	24	128	2.30	3200	8	2	**331**	**329**
EPYC 7302	16	128	3.00	3200	8	2	**242**	**284**
EPYC 7282	16	64	2.80	3200	8	2	**214**	**199**
EPYC 7262	8	128	3.20	3200	8	*2*	**134**	**181**
EPYC 7252	8	64	3.10	3200	8	2	**119**	**148**
[Reference] *1								
Xeon Platinum 8280	28	38.5	2.70	2933	6	2	**342**	**283**
Xeon Platinum 8268	24	35.8	2.90	2933	6	2	**304**	**265**
Xeon Gold 6242	16	22.0	2.80	2933	6	2	**214**	**208**
Xeon Gold 6244	8	24.8	3.60	2933	6	2	**131**	**150**

*1: The scores in RX2540 M5

The PRIMERGY GX2460 M1 with EPYC 7502 scored 21% higher on SPECrate2017_int_base and 28% higher on SPECrate2017_fp_base than the PRIMERGY RX2540 M5 with Xeon Platinum 8280. Compared to the Xeon Platinum 8280, EPCY 7502 has 14% more cores, more than three times the cache size, 33% more memory channels, and supports 3200 MHz memory frequency. These differences have a significant impact on CPU throughput performance.

SPECrate2017: Comparison of RPIMERGY GX2450 M1 and PRIMERGY RX2540 M5

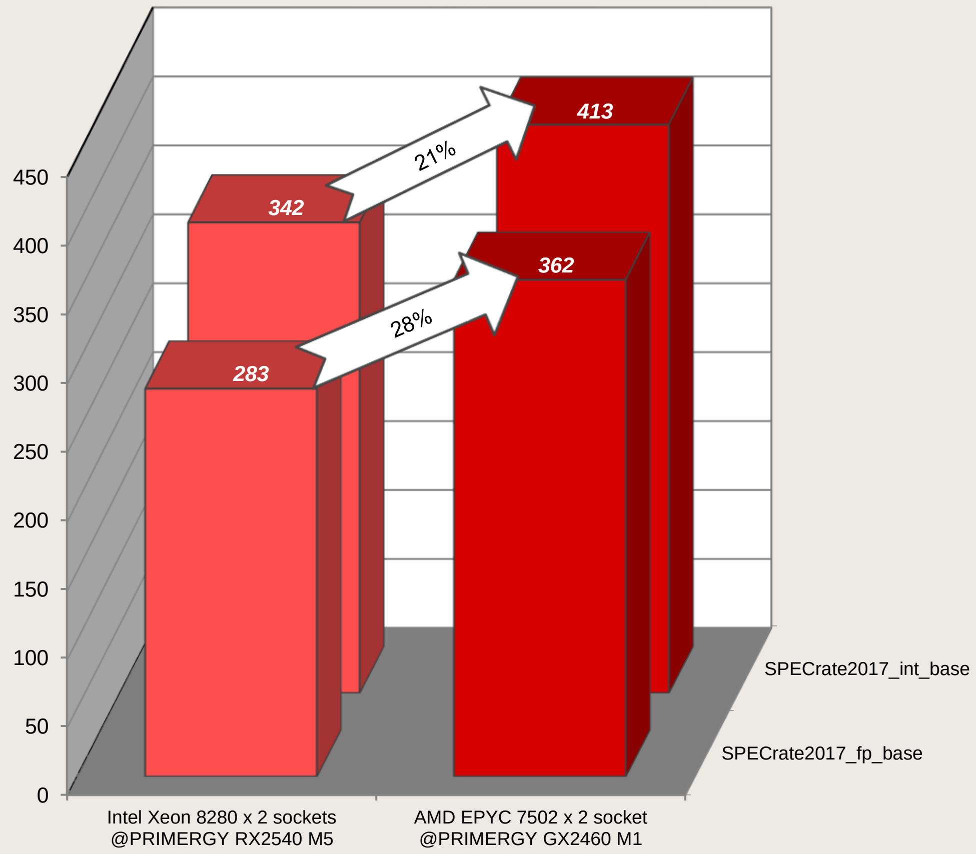

STREAM

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. This provides optimal load distribution for the available processor cores.

In the STREAM benchmark, a data area consisting of 8-byte elements is continuously copied to four operation types. Arithmetic operations are also performed on operation types other than COPY.

Arithmeticstype	Arithmetics	Bytes per step	Floating-point calculations 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.

In this chapter, throughputs are indicated as a power of 10. (1 GB/s = 10^9 Byte/s)

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMERGY GX2460 M1
Processor	AMD EPYC 7002 Series Processors
Memory	16 × 32 GB (1 x 32 GB) 2Rx4 PC4-3200AA-L
Software	
BIOS settings	Default
Operating system	SUSE Linux Enterprise Server 15 SP2 5.3.18-22-default

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

Benchmark results

Processor	Memory frequency [MHz]	Number of memory channels	Maximum memory bandwidth [GB/s]	Number of cores	Rated processor frequency [GHz]	Number of processors	TRIAD [GB/s]	Efficiency [%]
EPYC 7502	3200	8	204.8	32	2.50	2	286	69.8
EPYC 7452	3200	8	204.8	32	2.35	2	287	70.1
EPYC 7402	3200	8	204.8	24	2.80	2	255	62.3
EPYC 7352	3200	8	204.8	24	2.30	2	250	61.0
EPYC 7302	3200	8	204.8	16	3.00	2	251	61.3
EPYC 7282	3200	8	85.3 *1	16	2.80	2	168	98.4
EPYC 7262	3200	8	204.8	8	3.20	2	290	70.8
EPYC 7252	3200	8	85.3 *1	8	3.10	2	167	97.8
[Reference] *2								
Xeon Platinum 8280	2933	6	140.8	28	2.7	2	242	85.9
Xeon Platinum 8268	2933	6	140.8	24	2.9	2	243	86.3
Xeon Gold 6242	2933	6	140.8	16	2.8	2	223	79.2
Xeon Gold 6244	2933	6	140.8	8	3.3	2	161	57.2

*1: Since EPYC 7282 and EPYC 7252 have four memory channels and are optimized for a memory frequency of 2667 MHz, the maximum memory bandwidth is different from other processors.
*2: The scores in RX2540 M5.

The following diagram illustrates the memory throughput of PRIMERGY GX2460 M1 with AMD EPYC 7002 Series Processors in comparison to PRIMERGY RX2540 M5 with 2nd Generation Intel Xeon Scalable Processors. With 33% more memory channels and 3200 MHz memory frequency support, the EPYC 7502 delivers 18% better performance than the Xeon Platinum 8280.

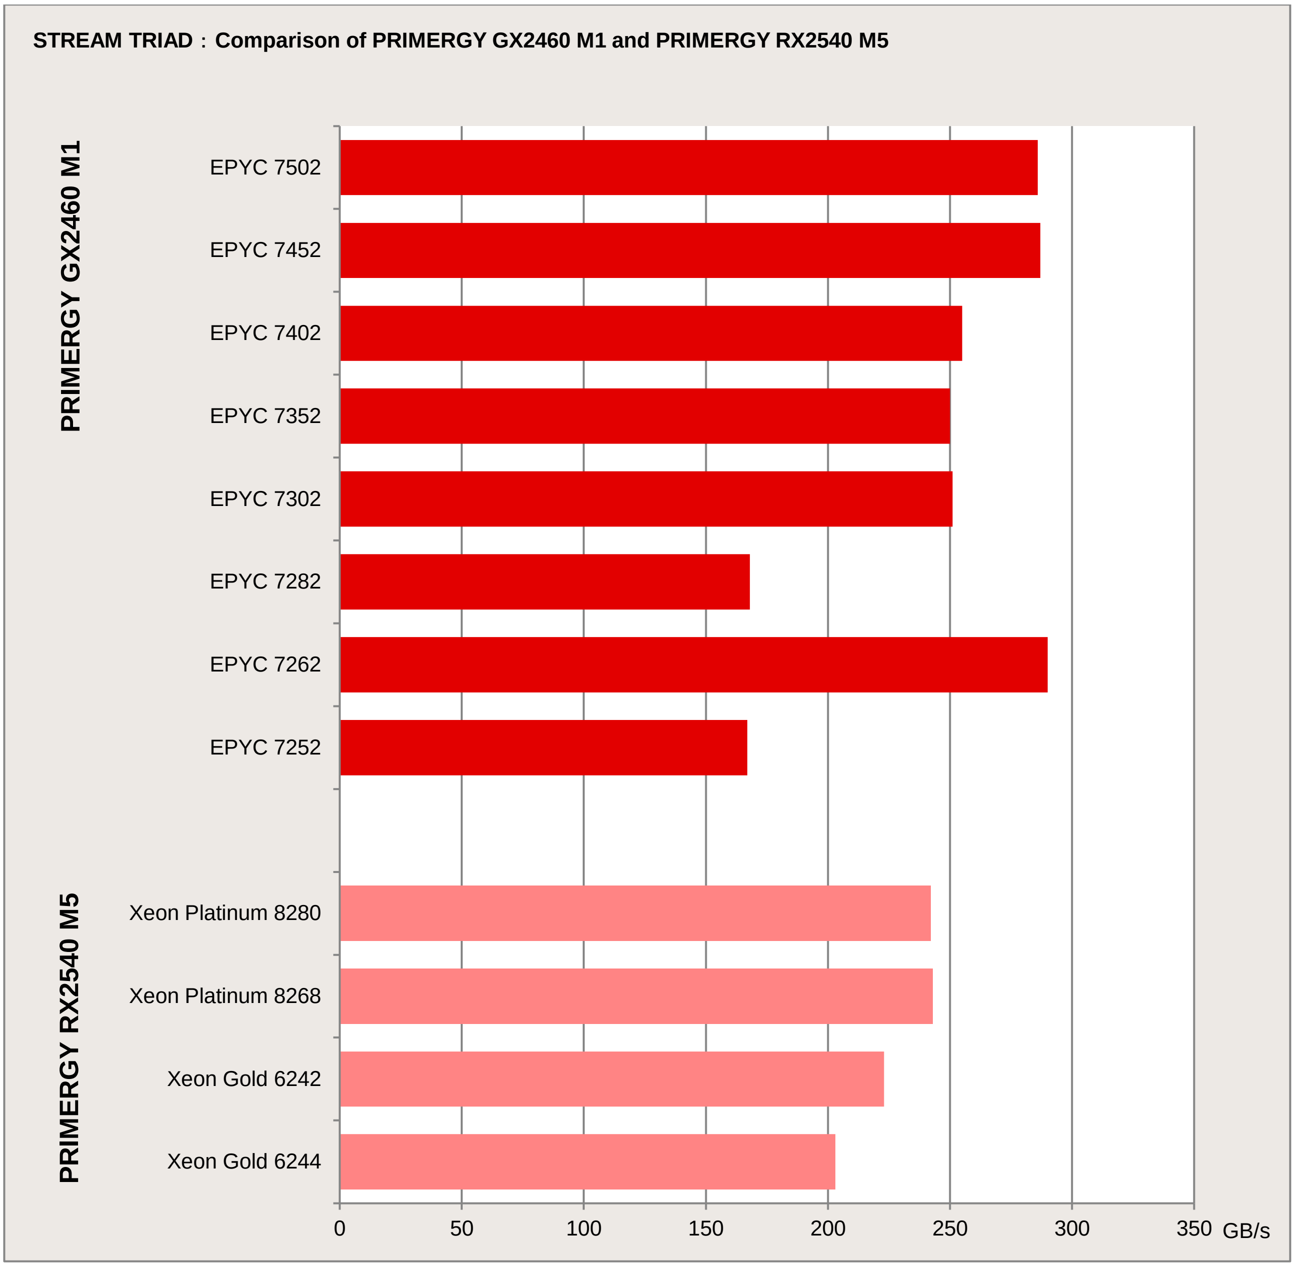
STREAM TRIAD : Comparison of PRIMERGY GX2460 M1 and PRIMERGY RX2540 M5
PRIMERGY GX2460 M1
EPYC 7502
EPYC 7452
EPYC 7402
EPYC 7352
EPYC 7302
EPYC 7282
EPYC 7262
EPYC 7252
PRIMERGY RX2540 M5
Xeon Platinum 8280
Xeon Platinum 8268
Xeon Gold 6242
Xeon Gold 6244
0
50
100
150
200
250
300
350
GB/s

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. The description can be found in the following document.
https://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 x 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 x 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. In other words, if n is doubled, the measurement time will be approximately eight times longer. The size of n also has an influence on the measurement result itself. As n increases, the measured value asymptotically approaches its 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. he 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: 1 billion floating point operations/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
x Number of processor cores of the computer
x 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 https://top500.org/. This requires using an HPL-based LINPACK version (see https://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 https://software.intel.com/content/www/us/en/develop/articles/intel-mkl-benchmarks-suite.html.

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 GX2460 M1
Processor	2x AMD EPYC 7002 Series Processors
Memory	16x 32GB (1x32GB) 2Rx4 PC4-3200AA-L
Software	
BIOS settings	DRAM scrub time = disabled Determinism Slider = Performance “cTDP” and “Package Power Limit” were set: EPYC 7502: 200 EPYC 7402: 200 EPYC 7452: 180 EPYC 7352: 180 EPYC 7302: 180 EPYC 7252: 150 EPYC 7282: 150 EPYC 7262: 180
Operating system	Ubuntu 20.04 5.0.0-13-generic

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

Benchmark results

Processor	Number of cores	Rated processor frequency [GHz]	Number of Processors	Rpeak [GFlops]	Rmax [GFlops]	Efficiency [%]
EPYC 7502	32	2.50	2	2560	2296	89.7
EPYC 7452	32	2.35	2	2406	2120	88.1
EPYC 7402	24	2.80	2	2150	2000	93.0
EPYC 7352	24	2.30	2	1766	1782	100.9
EPYC 7302	16	3.00	2	1536	1411	91.9
EPYC 7282	16	2.80	2	1434	1209	84.3
EPYC 7262	8	3.20	2	819	790	96.4
EPYC 7252	8	3.10	2	794	734	92.5
[Reference] *1						
Xeon Platinum 8280	28	2.70	2	4838	3522	72.8
Xeon Platinum 8268	24	2.90	2	4454	3096	70.0
Xeon Gold 6242	16	2.80	2	2867	2253	79.0
Xeon Gold 6244	8	3.30	2	1690	1325	78.0

*1: Scores in RX2540 M5

The following diagram illustrates the scores of LINPACK of PRIMERGY GX2460 M1 with AMD EPYC 7002 Series Processors in comparison to PRIMERGY RX2540 M5 with 2nd Generation Intel Xeon Scalable Processors. Xeon Platinum 8280 supports AVX -512 instructions, while the EPYC 7502 supports only AVX - 256 instructions, but has 14% more cores than Xeon Platinum 8280, resulting in 53% worse performance than Xeon Platinum 8280.

Rpeak values in the table above were calculated by the rated frequency of each processor. Since we enabled Turbo mode in the measurements, the average Turbo frequency exceeded the rated frequency for some processors. That's why some are more than 100% efficient.

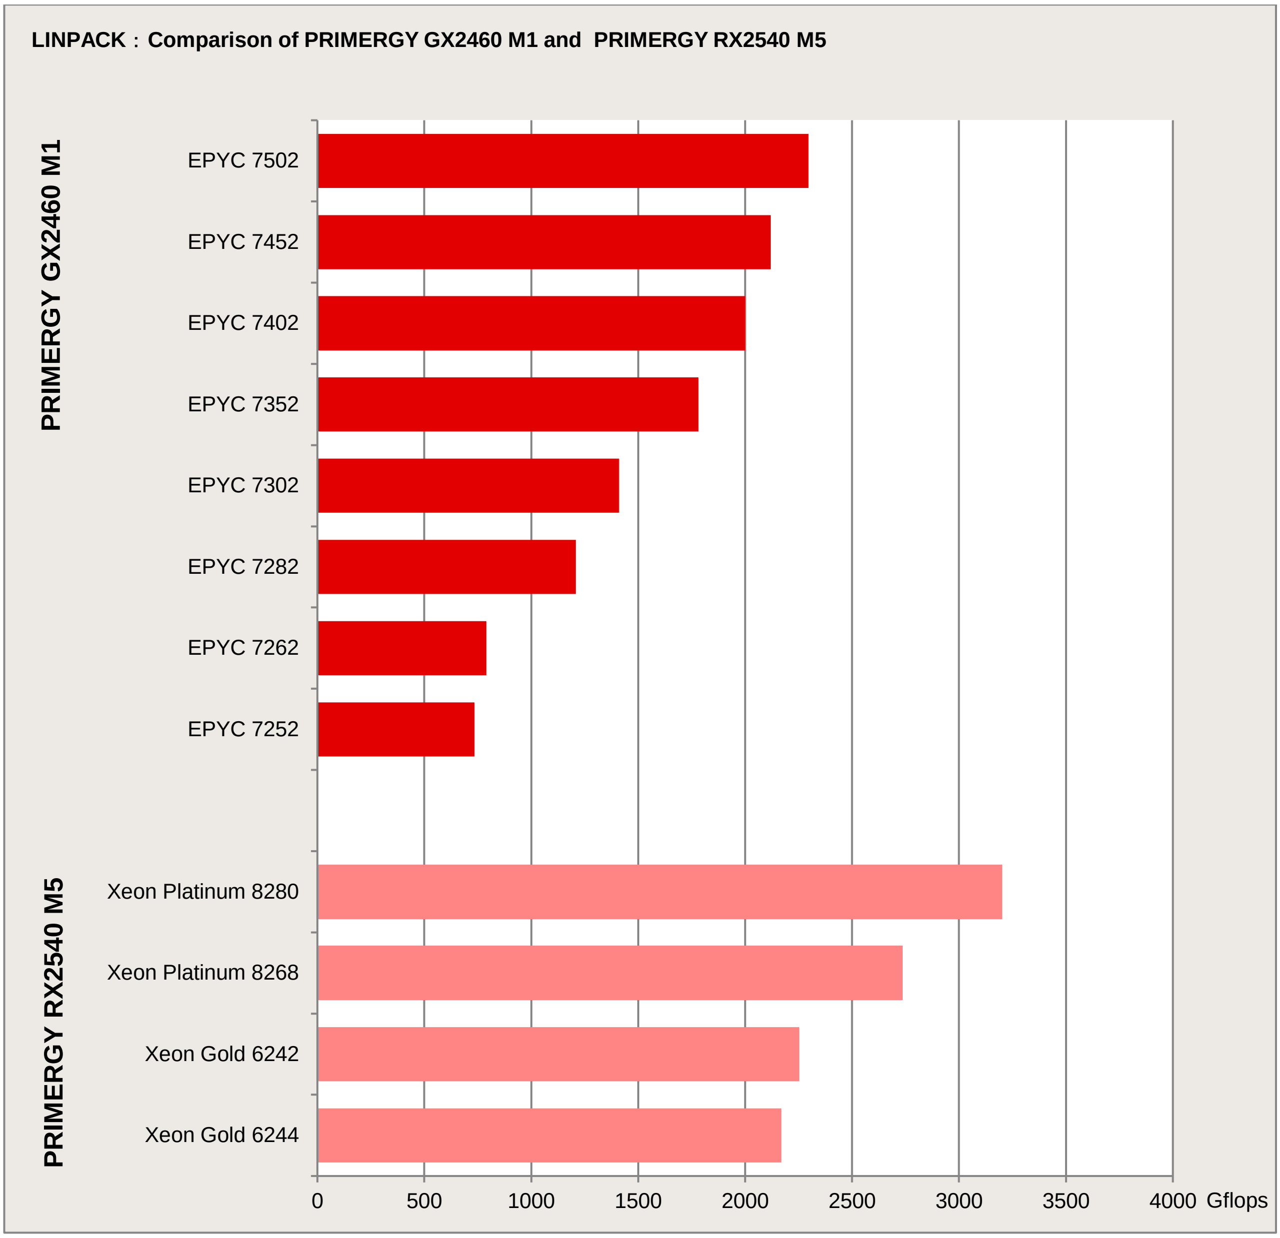
LINPACK : Comparison of PRIMERGY GX2460 M1 and PRIMERGY RX2540 M5
PRIMERGY GX2460 M1
EPYC 7502
EPYC 7452
EPYC 7402
EPYC 7352
EPYC 7302
EPYC 7282
EPYC 7262
EPYC 7252
PRIMERGY RX2540 M5
Xeon Platinum 8280
Xeon Platinum 8268
Xeon Gold 6242
Xeon Gold 6244
0
500
1000
1500
2000
2500
3000
3500
4000
Gflops

Literature

PRIMERGY Servers

https://www.fujitsu.com/global/products/computing/servers/primergy/

PRIMERGY GX2460 M1

This whitepaper:

https://docs.ts.fujitsu.com/dl.aspx?id=ea851fb9-1951-43a3-868c-a48bf3bda13a

https://docs.ts.fujitsu.com/dl.aspx?id=157df055-b6bc-4631-9a75-1f3044d2248b

Data sheet:
https://sp.ts.fujitsu.com/dmsp/Publications/public/ds-py-gx2460-m1.pdf

PRIMERGY Performance

https://www.fujitsu.com/global/products/computing/servers/primergy/benchmarks/

SPECcpu2017

https://www.spec.org/osg/cpu2017/

Benchmark Overview SPECcpu2017
https://docs.ts.fujitsu.com/dl.aspx?id=20f1f4e2-5b3c-454a-947f-c169fca51eb1

STREAM

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

LINPACK

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

TOP500
https://top500.org/

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

Contact

FUJITSU

Website: https://www.fujitsu.com/global/

PRIMERGY Performance and Benchmarks

mailto:fj-benchmark@dl.jp.fujitsu.com