FUJITSU

White Paper
FUJITSU Server PRIMERGY
Performance Report PRIMERGY CX2570 M1

This document contains a summary of the benchmarks executed for the FUJITSU Server PRIMERGY CX2570 M1.

The PRIMERGY CX2570 M1 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
1.2 2015-06-18

Contents

Document history

Version 1.0 (2014-12-18)

New:

- Technical data
- SPECcpu2006
  Measurements with Intel® Xeon® Processor E5-2600 v3 Product Family
- Disk I/O: Performance of storage media
  Results for 2.5" storage media
- Disk I/O: Performance of RAID controllers
  Measurements with “LSI SW RAID on Intel C610 (Onboard SATA)”, “PRAID CP400i”, “PRAID EP400i” and “PRAID EP420i” controllers
- LINPACK
  Measurements with Intel® Xeon® Processor E5-2600 v3 Product Family

Version 1.1 (2015-01-22)

Updated:

- Technical data
  Table of GPGPUs/coprocessors added
- SPECcpu2006
  Additional Measurements with Intel® Xeon® Processor E5-2600 v3 Product Family

Version 1.2 (2015-06-18)

New:

- STREAM
  Measurements with Intel® Xeon® Processor E5-2600 v3 Product Family

Updated:

- LINPACK
  Additional measurements with Intel® Xeon® Processor E5-2600 v3 Product Family

Technical data

PRIMERGY CX2570 M1

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 CX2570 M1
Form factor	Server node
Chipset	Intel® C610
Number of sockets	2
Number of processors orderable	2
Processor type	Intel® Xeon® Processor E5-2600 v3 Product Family
Number of memory slots	16 (8 per processor)
Maximum memory configuration	512 GB
Onboard HDD controller	Controller with RAID 0, RAID 1 and RAID 10 for up to 6 × 2.5" SATA HDDs
PCI slots	4 × PCI-Express 3.0 x16

Processors (since system release)								
Processor	Cores	Threads	Cache [MB]	QPI Speed [GT/s]	Rated Frequency [Ghz]	Max. Turbo Frequency [Ghz]	Max. Memory Frequency [MHz]	TDP [Watt]
Xeon E5-2623 v3	4	8	10	8.00	3.00	3.50	1866	105
Xeon E5-2637 v3	4	8	15	9.60	3.50	3.70	2133	135
Xeon E5-2603 v3	6	6	15	6.40	1.60	n/a	1600	85
Xeon E5-2609 v3	6	6	15	6.40	1.90	n/a	1600	85
Xeon E5-2620 v3	6	12	15	8.00	2.40	3.20	1866	85
Xeon E5-2643 v3	6	12	20	9.60	3.40	3.70	2133	135
Xeon E5-2630L v3	8	16	20	8.00	1.80	2.90	1866	55
Xeon E5-2630 v3	8	16	20	8.00	2.40	3.20	1866	85
Xeon E5-2640 v3	8	16	20	8.00	2.60	3.40	1866	90
Xeon E5-2667 v3	8	16	20	9.60	3.20	3.60	2133	135
Xeon E5-2650 v3	10	20	25	9.60	2.30	3.00	2133	105
Xeon E5-2660 v3	10	20	25	9.60	2.60	3.3	2133	105
Xeon E5-2687W v3	10	20	25	9.60	3.10	3.50	2133	160
Xeon E5-2650L v3	12	24	30	9.60	1.80	2.50	2133	65
Xeon E5-2670 v3	12	24	30	9.60	2.30	3.10	2133	120
Xeon E5-2680 v3	12	24	30	9.60	2.50	3.30	2133	120
Xeon E5-2690 v3	12	24	30	9.60	2.60	3.50	2133	135
Xeon E5-2683 v3	14	28	35	9.60	2.00	3.00	2133	120
Xeon E5-2695 v3	14	28	35	9.60	2.30	3.30	2133	120
Xeon E5-2697 v3	14	28	35	9.60	2.60	3.60	2133	145
Xeon E5-2698 v3	16	32	40	9.60	2.30	3.60	2133	135
Xeon E5-2699 v3	18	36	45	9.60	2.30	3.60	2133	145

All the processors that can be ordered with the PRIMERGY CX2570 M1, apart from Xeon E5-2603 v3 and Xeon E5-2609 v3, 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 frequency maximum 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 will 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 "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**
8GB (1x8GB) 1Rx4 DDR4-2133 R ECC	8	1	4	2133			✓	✓
8GB (1x8GB) 2Rx8 DDR4-2133 R ECC	8	2	8	2133			✓	✓
16GB (1x16GB) 2Rx4 DDR4-2133 R ECC	16	2	4	2133			✓	✓
32GB (1x32GB) 4Rx4 DDR4-2133 LR ECC	32	4	4	2133		✓	✓	✓

GPGPUs/coprocessors (since system release)			
GPGPU/coprocessor	**Cores**	**Peak double precision floating point performance [GFlops]**	**Max. number of GPGPUs**
PY NVIDIA Tesla K20 GPGPU w/o brackets	2496	1170	2
PY NVIDIA Tesla K20X GPGPU w/o brackets	2688	1310	2
PY NVIDIA Tesla K40 GPGPU	2880	1430	2
Intel Xeon Phi Co-Processor 3120P	57	1001	2
Intel Xeon Phi Co-Processor 31S1P	57	1003	2
Intel Xeon Phi Co-Processor 5110P	60	1011	2
Intel Xeon Phi Co-Processor 7120P	61	1208	2

Some components may not be available in all countries or sales regions.
Detailed technical information is available in the data sheet PRIMERGY CX2570 M1.

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 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; the peak values are optional.

Benchmark	Arithmetics	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 favour 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” specify 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

All results have been measured on a PRIMERGY CX2550 M1. The PRIMERGY CX2550 M1 and the PRIMERGY CX2570 M1 are electronically equivalent.

System Under Test (SUT)	
Hardware	
Model	PRIMERGY CX2550 M1
Processor	2 processors of Intel® Xeon® Processor E5-2600 v3 Product Family
Memory	16 × 16GB (1x16GB) 2Rx4 DDR4-2133 R ECC
Software	
Operating system	SPECint_base2006, SPECint2006: Xeon E5-2640 v3: Red Hat Enterprise Linux Server release 6.5 All others: Red Hat Enterprise Linux Server release 7.0 SPECfp_base2006, SPECfp2006: Red Hat Enterprise Linux Server release 7.0 SPECint_rate_base2006, SPECint_rate2006: Xeon E5-2603 v3, E5-2630 v3, E5-2650 v3, E5-2650L v3, E5-2660 v3, E5-2695 v3, E5-2698 v3: Red Hat Enterprise Linux Server release 7.0 All others: Red Hat Enterprise Linux Server release 6.5 SPECfp_rate_base2006, SPECfp_rate2006: Xeon E5-2690 v3: Red Hat Enterprise Linux Server release 6.5 All others: Red Hat Enterprise Linux Server release 7.0
Operating system settings	echo always > /sys/kernel/mm/redhat_transparent_hugepage/enabled
Compiler	SPECint_base2006, SPECint2006: Xeon E5-2640 v3: C/C++: Version 14.0.0.080 of Intel C++ Studio XE for Linux All others: C/C++: Version 15.0.0.090 of Intel C++ Studio XE for Linux SPECint_rate_base2006, SPECint_rate2006: Xeon E5-2603 v3, E5-2630 v3, E5-2650 v3, E5-2650L v3, E5-2660 v3, E5-2695 v3, E5-2698 v3: C/C++: Version 15.0.0.090 of Intel C++ Studio XE for Linux All others: C/C++: Version 14.0.0.080 of Intel C++ Studio XE for Linux SPECfp_base2006, SPECfp2006, SPECfp_rate_base2006, SPECfp_rate2006: C/C++: Version 15.0.0.090 of Intel C++ Studio XE for Linux Fortran: Version 15.0.0.090 of Intel Fortran Studio XE 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 on 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.

Processor	Number of processors	SPECint_base2006	SPECint2006	SPECfp_base2006	SPECfp2006	SPECint_rate_base2006	SPECint_rate2006	SPECfp_rate_base2006	SPECfp_rate2006
Xeon E5-2623 v3	2	56.6	58.7	96.2	99.7	411	425	374	385
Xeon E5-2637 v3	2	61.8	64.3	106	109	451	468	427	441
Xeon E5-2603 v3	2	29.3	30.5	55.5	57.2	267	275	283	290
Xeon E5-2609 v3	2	33.7	35.1	62.5	64.5	307	318	327	335
Xeon E5-2620 v3	2	54.1	56.4	95.0	100	506	523	464	475
Xeon E5-2643 v3	2	64.4	67.2	114	118	669	693	576	592
Xeon E5-2630L v3	2	51.4	53.5	89.8	95.9	558	577	470	482
Xeon E5-2630 v3	2	56.2	59.0	102	108	658	686	565	580
Xeon E5-2640 v3	2	58.8	62.3	104	109	704	729	581	594
Xeon E5-2667 v3	2	62.3	65.1	116	121	815	843	652	669
Xeon E5-2650 v3	2	54.0	56.3	103	107	818	853	679	698
Xeon E5-2660 v3	2	57.6	59.7	110	114	883	916	700	719
Xeon E5-2687W v3	2	61.4	64.1	115	121	954	982	743	764
Xeon E5-2650L v3	2	46.2	47.9	88.1	92.0	777	810	636	655
Xeon E5-2670 v3	2	56.6	58.9	105	110	970	1000	752	773
Xeon E5-2680 v3	2	60.0	62.3	111	115	1030	1060	768	789
Xeon E5-2690 v3	2	62.6	65.2	114	119	1080	1110	790	814
Xeon E5-2683 v3	2	53.5	55.8	99.9	104	1070	1100	796	820
Xeon E5-2695 v3	2	58.5	61.0	105	110	1110	1150	807	830
Xeon E5-2697 v3	2	62.8	65.6	111	117	1200	1240	844	868
Xeon E5-2698 v3	2	62.9	65.7	109	116	1230	1280	865	892
Xeon E5-2699 v3	2	63.9	66.1	109	116	1350	1390	905	933

Disk I/O: Performance of storage media

Benchmark description

Performance measurements of disk subsystems for PRIMERGY servers are used to assess their performance and enable a comparison of the different storage connections for PRIMERGY servers. As standard, these performance measurements are carried out with a defined measurement method, which models the accesses of real application scenarios on the basis of specifications.

The essential specifications are:

- Share of random accesses / sequential accesses
- Share of read / write access types
- Block size (kB)
- Number of parallel accesses (# of outstanding I/Os)

A given value combination of these specifications is known as “load profile”. The following five standard load profiles can be allocated to typical application scenarios:

Standard load profile	Access	Type of access		Block size [kB]	Application
		read	write		
File copy	random	50%	50%	64	Copying of files
File server	random	67%	33%	64	File server
Database	random	67%	33%	8	Database (data transfer) Mail server
Streaming	sequential	100%	0%	64	Database (log file), Data backup; Video streaming (partial)
Restore	sequential	0%	100%	64	Restoring of files

In order to model applications that access in parallel with a different load intensity, the “# of Outstanding I/Os” is increased, starting with 1, 3, 8 and going up to 512 (from 8 onwards in increments to the power of two).

The measurements of this document are based on these standard load profiles.

The main results of a measurement are:

- Throughput [MB/s] Throughput in megabytes per second
- Transactions [IO/s] Transaction rate in I/O operations per second
- Latency [ms] Average response time in ms

The data throughput has established itself as the normal measurement variable for sequential load profiles, whereas the measurement variable “transaction rate” is mostly used for random load profiles with their small block sizes. Data throughput and transaction rate are directly proportional to each other and can be transferred to each other according to the formula

Data throughput [MB/s]	*= Transaction rate [IO/s] × Block size [MB]*
Transaction rate [IO/s]	*= Data throughput [MB/s] / Block size [MB]*

This section specifies capacities of storage media on a basis of 10 (1 TB = 10^{12} bytes) while all other capacities, file sizes, block sizes and throughputs are specified on a basis of 2 (1 MB/s = 2^{20} bytes/s).

All the details of the measurement method and the basics of disk I/O performance are described in the white paper “Basics of Disk I/O Performance”.

Benchmark environment

All the measurement results discussed in this section apply for the hardware and software components listed below:

System Under Test (SUT)		
Hardware		
Model	PRIMERGY CX2570 M1	
Controller	1 × RAID Ctrl SAS 6G 0/1 (D2607)	
Storage media	SSD	HDD
	Intel SSDSC2BA100G3C Intel SSDSC2BA200G3C Intel SSDSC2BA400G3C Intel SSDSC2BA800G3C	HGST HUC101212CSS600 Seagate ST300MM0006 Seagate ST450MM0006 Seagate ST600MM0006 Seagate ST900MM0006 Seagate ST91000640NS Seagate ST9250610NS Seagate ST9500620NS Western Digital WD3001BKHG Western Digital WD4501BKHG Western Digital WD6001BKHG Western Digital WD9001BKHG
	Innodisk DESMH-32D07RC1DCF-A88 Innodisk DESMH-64D07RC1DCF-A88 Innodisk DESMH-A28D07RC1DCF-A88	
Software		
Operating system	Microsoft Windows Server 2008 Enterprise x64 Edition SP2	
Administration software	ServerView RAID Manager 5.7.2	
File system	NTFS	
Measuring tool	Iometer 2006.07.27	
Measurement data	32 GB measurement file	

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

Benchmark results

The results shown here are intended to help you select the appropriate storage media for the PRIMERGY CX2570 M1 under the aspect of disk-I/O performance. For this purpose, a single storage medium was measured in the configuration specified in the subsection Benchmark environment. The measurements were made using controllers which have their main features listed in the table below:

Storage medium	Controller name	Cache	Supported interfaces		RAID levels
SSD/HDD	RAID Ctrl SAS 6G 0/1 (D2607)	-	SATA 3G/6G SAS 3G/6G	PCIe 2.0 x8	0, 1, 1E, 10
DOM	AHCI SATA on Intel C610	-	SATA 6G	-	-

Storage media

When selecting the type and number of storage media you can move the weighting in the direction of storage capacity, performance, security or price. The following types of storage media can be used in the PRIMERGY CX2570 M1:

Storage medium type	Interface	Form factor
HDD	SATA 6G	2.5"
HDD	SAS 6G	2.5"
SSD	SATA 6G	2.5"
DOM	SATA 6G	Disk on module

HDDs and SSDs are operated via host bus adapters, usually RAID controllers, with a SATA or SAS interface. The interface of the RAID controller to the chipset of the systemboard is typically PCIe or, in the case of the integrated onboard controllers, an internal bus interface of the systemboard. DOM stands for "Disk on module". It is an extremely space and energy-saving flash memory that is particularly used as a boot drive in servers. The memory technology is equivalent to that of SSDs. For a range of PRIMERGY servers Fujitsu offers a DOM with a SATA 6G interface, which can be inserted directly into the SATA port of the system board.

Of all the storage medium types SSDs offer by far the highest transaction rates for random load profiles as well as the shortest access times. In return, however, the price per gigabyte of storage capacity is substantially higher.

Cache settings

In most cases, the cache of HDDs has a great influence on disk-I/O performance. It is frequently regarded as a security problem in case of power failure and is thus switched off. On the other hand, it was integrated by hard disk manufacturers for the good reason of increasing the write performance. For performance reasons it is therefore advisable to enable the hard disk cache. To prevent data loss in case of power failure you are recommended to equip the system with a UPS.

For the purpose of easy and reliable handling of the settings for RAID controllers and hard disks it is advisable to use the RAID-Manager software “ServerView RAID” that is supplied for PRIMERGY servers. All the cache settings for controllers and hard disks can usually be made en bloc – specifically for the application – by using the pre-defined modi “Performance” or “Data Protection”. The “Performance” mode ensures the best possible performance settings for the majority of the application scenarios.

Performance values

The performance values of the PRIMERGY CX2570 M1 are summarized in the following tables, in each case specifically for a single storage medium and with various access types and block sizes. The established measurement variables, as already mentioned in the subsection Benchmark description, are used here. Thus, transaction rate is specified for random accesses and data throughput for sequential accesses. To avoid any confusion among the measurement units the tables have been separated for the two access types.

The table cells contain the maximum achievable values. This means that each value is the maximum achievable value of the whole range of load intensities (# of Outstanding I/Os). In order to also visualize the numerical values each table cell is highlighted with a horizontal bar, the length of which is proportional to the numerical value in the table cell. All bars shown in the same scale of length have the same color. In other words, a visual comparison only makes sense for table cells with the same colored bars. Since the horizontal bars in the table cells depict the maximum achievable performance values, they are shown by the color getting lighter as you move from left to right. The light shade of color at the right end of the bar tells you that the value is a maximum value and can only be achieved under optimal prerequisites. The darker the shade becomes as you move to the left, the more frequently it will be possible to achieve the corresponding value in practice.

SSDs in comparison with the most powerful HDD

Random accesses (maximum performance values in IO/s):

PRIMERGY CX2570 M1					
Capacity [GB]	Storage device	Interface	Transactions [IO/s]		
			Database	Fileserver	Filecopy
1200	HGST HUC101212CSS600 (HDD)	SAS 6G	638	539	539
800	Intel SSDSC2BA800G3C	SATA 6G	35120	5554	5313
400	Intel SSDSC2BA400G3C	SATA 6G	36667	5453	5338
200	Intel SSDSC2BA200G3C	SATA 6G	35023	4903	4466
128	Innodisk DESMH-A28D07RC1DCF-A88	SATA 6G	28	24	21
100	Intel SSDSC2BA100G3C	SATA 6G	28535	3371	3128
64	Innodisk DESMH-64D07RC1DCF-A88	SATA 6G	31	25	24
32	Innodisk DESMH-32D07RC1DCF-A88	SATA 6G	61	21	18

Sequential accesses (maximum performance values in MB/s):

PRIMERGY CX2570 M1				
Capacity [GB]	Storage device	Interface	Throughput [MB/s]	
			Streaming	Restore
1200	HGST HUC101212CSS600 (HDD)	SAS 6G	191	191
800	Intel SSDSC2BA800G3C	SATA 6G	382	342
400	Intel SSDSC2BA400G3C	SATA 6G	434	341
200	Intel SSDSC2BA200G3C	SATA 6G	410	330
128	Innodisk DESMH-A28D07RC1DCF-A88	SATA 6G	255	126
100	Intel SSDSC2BA100G3C	SATA 6G	434	196
64	Innodisk DESMH-64D07RC1DCF-A88	SATA 6G	261	114
32	Innodisk DESMH-32D07RC1DCF-A88	SATA 6G	211	73

HDDs

Random accesses (maximum performance values in IO/s):

PRIMERGY CX2570 M1					
Capacity [GB]	Storage device	Interface	Transactions [IO/s]		
			Database	Fileserver	Filecopy
1200	HGST HUC101212CSS600	SAS 6G	638	539	539
1000	Seagate ST91000640NS	SATA 6G	372	317	314
900	Western Digital WD9001BKHG	SAS 6G	568	485	462
900	Seagate ST900MM0006	SAS 6G	502	436	422
600	Western Digital WD6001BKHG	SAS 6G	572	484	460
600	Seagate ST600MM0006	SAS 6G	551	471	456
500	Seagate ST9500620NS	SATA 6G	240	215	221
450	Seagate ST450MM0006	SAS 6G	533	453	437
450	Western Digital WD4501BKHG	SAS 6G	515	462	450
300	Seagate ST300MM0006	SAS 6G	535	460	445
300	Western Digital WD3001BKHG	SAS 6G	521	447	420
250	Seagate ST9250610NS	SATA 6G	318	275	272

Sequential accesses (maximum performance values in MB/s):

PRIMERGY CX2570 M1				
Capacity [GB]	Storage device	Interface	Throughput [MB/s]	
			Streaming	Restore
1200	HGST HUC101212CSS600	SAS 6G	191	191
1000	Seagate ST91000640NS	SATA 6G	108	107
900	Western Digital WD9001BKHG	SAS 6G	192	193
900	Seagate ST900MM0006	SAS 6G	191	191
600	Western Digital WD6001BKHG	SAS 6G	193	193
600	Seagate ST600MM0006	SAS 6G	194	193
500	Seagate ST9500620NS	SATA 6G	110	109
450	Seagate ST450MM0006	SAS 6G	188	188
450	Western Digital WD4501BKHG	SAS 6G	192	192
300	Seagate ST300MM0006	SAS 6G	190	190
300	Western Digital WD3001BKHG	SAS 6G	192	193
250	Seagate ST9250610NS	SATA 6G	114	113

Disk I/O: Performance of RAID controllers

Benchmark description

Performance measurements of disk subsystems for PRIMERGY servers are used to assess their performance and enable a comparison of the different storage connections for PRIMERGY servers. As standard, these performance measurements are carried out with a defined measurement method, which models the accesses of real application scenarios on the basis of specifications.

The essential specifications are:

- Share of random accesses / sequential accesses
- Share of read / write access types
- Block size (kB)
- Number of parallel accesses (# of outstanding I/Os)

A given value combination of these specifications is known as "load profile". The following five standard load profiles can be allocated to typical application scenarios:

Standard load profile	Access	Type of access		Block size [kB]	Application
		read	write		
File copy	random	50%	50%	64	Copying of files
File server	random	67%	33%	64	File server
Database	random	67%	33%	8	Database (data transfer) Mail server
Streaming	sequential	100%	0%	64	Database (log file), Data backup; Video streaming (partial)
Restore	sequential	0%	100%	64	Restoring of files

In order to model applications that access in parallel with a different load intensity, the "# of Outstanding I/Os" is increased, starting with 1, 3, 8 and going up to 512 (from 8 onwards in increments to the power of two).

The measurements of this document are based on these standard load profiles.

The main results of a measurement are:

- Throughput [MB/s] Throughput in megabytes per second
- Transactions [IO/s] Transaction rate in I/O operations per second
- Latency [ms] Average response time in ms

The data throughput has established itself as the normal measurement variable for sequential load profiles, whereas the measurement variable "transaction rate" is mostly used for random load profiles with their small block sizes. Data throughput and transaction rate are directly proportional to each other and can be transferred to each other according to the formula

Data throughput [MB/s]	*= Transaction rate [IO/s] × Block size [MB]*
Transaction rate [IO/s]	*= Data throughput [MB/s] / Block size [MB]*

This section specifies capacities of storage media on a basis of 10 (1 TB = 10^{12} bytes) while all other capacities, file sizes, block sizes and throughputs are specified on a basis of 2 (1 MB/s = 2^{20} bytes/s).

All the details of the measurement method and the basics of disk I/O performance are described in the white paper "Basics of Disk I/O Performance".

Benchmark environment

All the measurement results discussed in this chapter were determined using the hardware and software components listed below:

System Under Test (SUT)	
Hardware	
Controller	1 × “LSI SW RAID on Intel C610 (Onboard SATA)” 1 × “PRAID CP400i” 1 × “PRAID EP400i” 1 × “PRAID EP420i”
Drive	6 × 2.5" SATA SSD Intel SSDSC2BA400G3C 6 × 2.5" SATA HDD Seagate ST91000640NS 6 × 2.5" SAS SSD Toshiba PX02SMF040 6 × 2.5" SAS HDD Toshiba AL13SXB450N
Software	
BIOS settings	Intel Virtualization Technology = Disabled VT-d = Disabled Energy Performance = Performance Utilization Profile = Unbalanced CPU C6 Report = Disabled
Operating system	Microsoft Windows Server 2012 Standard
Operating system settings	Choose or customize a power plan: High performance For the processes that create disk I/Os: set the AFFINITY to the CPU node to which the PCIe slot of the RAID controller is connected
Administration software	ServerView RAID Manager 5.7.2
Initialization of RAID arrays	RAID arrays are initialized before the measurement with an elementary block size of 64 kB (“stripe size”)
File system	NTFS
Measuring tool	Iometer 2006.07.27
Measurement data	Measurement files of 32 GB with 1 – 8 hard disks; 64 GB with 9 – 16 hard disks; 128 GB with 17 or more hard disks

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

Benchmark results

The results presented here are designed to help you choose the right solution from the various configuration options of the PRIMERGY CX2570 M1 in the light of disk-I/O performance. Various combinations of RAID controllers and storage media will be analyzed below. Information on the selection of storage media themselves is to be found in the section “Disk I/O: Performance of storage media”.

Hard disks

The hard disks are the first essential component. If there is a reference below to “hard disks”, this is meant as the generic term for HDDs (“hard disk drives”, in other words conventional hard disks) and SSDs (“solid state drives”, i.e. non-volatile electronic storage media).

Mixed drive configurations of SAS and SATA hard disks in one system are permitted, unless they are excluded in the configurator for special hard disk types.

More detailed performance statements about hard disk types are available in the section “Disk I/O: Performance of storage media” in this performance report.

Model versions

The maximum number of hard disks in the system depends on the system configuration. The following table lists the essential cases.

Form factor	Interface	Connection type	Number of PCIe controllers	Maximum number of hard disks
2.5"	SATA 6G	direct	0	6
2.5"	SATA 6G, SAS 12G	direct	1	6

RAID controller

In addition to the hard disks the RAID controller is the second performance-determining key component. In the case of these controllers the "modular RAID" concept of the PRIMERGY servers offers a plethora of options to meet the various requirements of a wide range of different application scenarios.

The following table summarizes the most important features of the available RAID controllers of the PRIMERGY CX2570 M1. A short alias is specified here for each controller, which is used in the subsequent list of the performance values.

Controller name	Alias	Cache	Supported interfaces		In the system		BBU/ FBU
					Max. # disks per controller	RAID levels	
LSI SW RAID on Intel C610 (Onboard SATA)	Onboard C610	-	SATA 6G	-	4 × 2.5"	0, 1, 10	-/-
PRAID CP400i	PRAID CP400i	-	SATA 6G SAS 12G	PCIe 3.0 x8	6 × 2.5"	0, 1, 1E, 5, 10, 50	-/-
PRAID EP400i	PRAID EP400i	1 GB	SATA 6G SAS 12G	PCIe 3.0 x8	6 × 2.5"	0, 1, 1E, 5, 6, 10, 50, 60	-/✓
PRAID EP420i	PRAID EP420i	2 GB	SATA 6G SAS 12G	PCIe 3.0 x8	6 × 2.5"	0, 1, 1E, 5, 6, 10, 50, 60	-/✓

The onboard RAID controller is implemented in the chipset Intel C610 on the system board of the server and uses the CPU of the server for the RAID functionality. This controller is a simple solution that does not require a PCIe slot.

System-specific interfaces

The interfaces of a controller in CPU direction (PCIe or in the event of onboard controllers "Direct Media Interface", DMI in short) and in the direction of hard disks (SAS or SATA) have in each case specific limits for data throughput. These limits are listed in the following table. The minimum of these two values is a definite limit, which cannot be exceeded. This value is highlighted in bold in the following table.

Controller alias	Effective in the configuration				Connection via expander
	# Disk-side data channels	Limit for throughput of disk interface	# CPU-side data channels	Limit for throughput of CPU-side interface	
Onboard C610	6 × SATA 6G	3090 MB/s	4 × DMI 2.0	**1716 MB/s**	-
PRAID CP400i	6 × SAS 12G	**6180 MB/s**	8 × PCIe 3.0	6761 MB/s	-
PRAID EP400i	6 × SAS 12G	**6180 MB/s**	8 × PCIe 3.0	6761 MB/s	-
PRAID EP420i	6 × SAS 12G	**6180 MB/s**	8 × PCIe 3.0	6761 MB/s	-

More details about the RAID controllers of the PRIMERGY systems are available in the white paper "RAID Controller Performance".

Settings

In most cases, the cache of HDDs has a great influence on disk-I/O performance. It is frequently regarded as a security problem in case of power failure and is thus switched off. On the other hand, it was integrated by hard disk manufacturers for the good reason of increasing the write performance. For performance reasons it is therefore advisable to enable the hard disk cache. To prevent data loss in case of power failure you are recommended to equip the system with a UPS.

In the case of controllers with a cache there are several parameters that can be set. The optimal settings can depend on the RAID level, the application scenario and the type of data medium. In the case of RAID levels 5 and 6 in particular (and the more complex RAID level combinations 50 and 60) it is obligatory to enable the controller cache for application scenarios with write share. If the controller cache is enabled, the data temporarily stored in the cache should be safeguarded against loss in case of power failure. Suitable accessories are available for this purpose (e.g. a BBU or FBU).

For the purpose of easy and reliable handling of the settings for RAID controllers and hard disks it is advisable to use the RAID-Manager software "ServerView RAID" that is supplied for PRIMERGY servers. All the cache settings for controllers and hard disks can usually be made en bloc – specifically for the application – by using the pre-defined modi "Performance" or "Data Protection". The "Performance" mode ensures the best possible performance settings for the majority of the application scenarios.

More information about the setting options of the controller cache is available in the white paper "RAID Controller Performance".

Performance values

In general, disk-I/O performance of a RAID array depends on the type and number of hard disks, on the RAID level and on the RAID controller. If the limits of the system-specific interfaces are not exceeded, the statements on disk-I/O performance are therefore valid for all PRIMERGY systems. This is why all the performance statements of the document "RAID Controller Performance" also apply for the PRIMERGY CX2570 M1 if the configurations measured there are also supported by this system.

The performance values of the PRIMERGY CX2570 M1 are listed in table form below, specifically for different RAID levels, access types and block sizes. Substantially different configuration versions are dealt with separately. The established measurement variables, as already mentioned in the subsection Benchmark description, are used here. Thus, transaction rate is specified for random accesses and data throughput for sequential accesses. To avoid any confusion among the measurement units the tables have been separated for the two access types.

The table cells contain the maximum achievable values. This has three implications: On the one hand hard disks with optimal performance were used (the components used are described in more detail in the subsection Benchmark environment). Furthermore, cache settings of controllers and hard disks, which are optimal for the respective access scenario and the RAID level, are used as a basis. And ultimately each value is the maximum value for the entire load intensity range (# of outstanding I/Os).

In order to also visualize the numerical values each table cell is highlighted with a horizontal bar, the length of which is proportional to the numerical value in the table cell. All bars shown in the same scale of length have the same color. In other words, a visual comparison only makes sense for table cells with the same colored bars.

Since the horizontal bars in the table cells depict the maximum achievable performance values, they are shown by the color getting lighter as you move from left to right. The light shade of color at the right end of the bar tells you that the value is a maximum value and can only be achieved under optimal prerequisites. The darker the shade becomes as you move to the left, the more frequently it will be possible to achieve the corresponding value in practice.

2.5" - Random accesses (maximum performance values in IO/s):

PRIMERGY CX2570 M1 Model version PY CX2570 M1								
Configuration version			RAID level	HDDs random 8 kB blocks 67% read [IO/s]	HDDs random 64 kB blocks 67% read [IO/s]	SSDs random 8 kB blocks 67% read [IO/s]	SSDs random 64 kB blocks 67% read [IO/s]
RAID Controller	Hard disk type	#Disks					
Onboard C610	ST91000640NS SATA HDD SSDSC2BA400G3C SATA SSD	2	1	464	389	47337	7870
		6	0	1351	726	87164	18447
		6	10	874	469	68074	15090
PRAID CP400i	AL13SXB450N SAS HDD PX02SMF040 SAS SSD	2	1	1078	874	75925	12445
		6	10	3374	1660	101480	43040
		6	0	4831	2373	136523	59689
		6	5	2357	1136	30175	16445
PRAID EP400i	AL13SXB450N SAS HDD PX02SMF040 SAS SSD	2	1	1340	866	78733	12318
		6	10	3822	1876	112461	41368
		6	0	5080	2349	131019	61411
		6	5	2800	1362	52636	19099
PRAID EP420i	AL13SXB450N SAS HDD PX02SMF040 SAS SSD	2	1	1777	1058	80178	12460
		6	10	4071	2118	113567	44360
		6	0	5472	2538	131040	61042
		6	5	3680	1452	52773	20629

2.5" - Sequential accesses (maximum performance values in MB/s):

PRIMERGY CX2570 M1 Model version PY CX2570 M1							
Configuration version							
RAID Controller	Hard disk type	#Disks	RAID level	HDDs sequential 64 kB blocks 100% read [MB/s]	HDDs sequential 64 kB blocks 100% write [MB/s]	SSDs sequential 64 kB blocks 100% read [MB/s]	SSDs sequential 64 kB blocks 100% write [MB/s]
Onboard C610	ST91000640NS SATA HDD SSDSC2BA400G3C SATA SSD	2	1	112	108	726	443
		6	0	622	567	1494	1232
		6	10	319	288	1270	618
PRAID CP400i	AL13SXB450N SAS HDD PX02SMF040 SAS SSD	2	1	379	232	1603	421
		6	10	735	650	4851	1243
		6	0	1317	1291	4764	2482
		6	5	1098	1062	4756	1766
PRAID EP400i	AL13SXB450N SAS HDD PX02SMF040 SAS SSD	2	1	396	231	1596	420
		6	10	748	664	4862	1242
		6	0	1339	1290	4758	2490
		6	5	1098	1061	4740	2082
PRAID EP420i	AL13SXB450N SAS HDD PX02SMF040 SAS SSD	2	1	384	253	1595	421
		6	10	769	665	4875	1248
		6	0	1346	1322	4707	2469
		6	5	1138	1079	4695	2068

Conclusion

At full configuration with powerful hard disks the PRIMERGY CX2570 M1 achieves a throughput of up to 4875 MB/s for sequential load profiles and a transaction rate of up to 136523 IO/s for typical, random application scenarios.

For best possible performance we recommend one of the plug-in PCIe controllers. To operate SSDs within the maximum performance range the PRAID CP400i is already suited for the simpler RAID levels 0, 1 and 10, and a controller with cache is to be preferred for RAID 5.

In the event of HDDs the controller cache for random load profiles with a significant write share has performance advantages for all RAID levels.

STREAM

Benchmark description

STREAM is a synthetic benchmark that has been used for many years to determine memory throughput and which 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, 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	
Enclosure	PRIMERGY CX400 M1
Model	PRIMERGY CX2570 M1
Processor	2 processors of Intel® Xeon® Processor E5-2600 v3 Product Family
Memory	16 × 16GB (1x16GB) 2Rx4 DDR4-2133 R ECC
Software	
BIOS settings	EnergyPerformance = Performance Xeon E5-2603 v3, E5-2609 v3, E5-2630 v3, E5-2637 v3, E5-2667 v3, E5-2699 v3: COD Enable = disabled, Early Snoop = disabled Xeon E5-2623 v3, E5-2630L v3, E5-2640 v3: COD Enable = disabled, Early Snoop = enabled All others: COD Enable = enabled, Early Snoop = disabled
Operating system	Red Hat Enterprise Linux Server release 7.0
Operating system settings	Transparent Huge Pages inactivated
Compiler	Intel C++ Composer XE 2013 SP1 for Linux Update 1
Benchmark	Stream.c Version 5.9

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

Benchmark results

Processor	Memory Frequency [MHz]	Max. Memory Bandwidth [GB/s]	Cores	Processor Frequency [GHz]	Number of Processors	TRIAD [GB/s]
Xeon E5-2603 v3	1600	51	6	1.60	2	**47.8**
Xeon E5-2609 v3	1600	51	6	1.90	2	**58.4**
Xeon E5-2623 v3	1866	59	4	3.00	2	**73.5**
Xeon E5-2620 v3	1866	59	6	2.40	2	88.2 (est.)
Xeon E5-2630L v3	1866	59	8	1.80	2	**86.4**
Xeon E5-2630 v3	1866	59	8	2.40	2	**89.9**
Xeon E5-2640 v3	1866	59	8	2.60	2	**90.4**
Xeon E5-2637 v3	2133	68	4	3.50	2	**90.1**
Xeon E5-2643 v3	2133	68	6	3.40	2	90.1 (est.)
Xeon E5-2667 v3	2133	68	8	3.20	2	**91.2**
Xeon E5-2650 v3	2133	68	10	2.30	2	115 (est.)
Xeon E5-2660 v3	2133	68	10	2.60	2	**114**
Xeon E5-2687W v3	2133	68	10	3.10	2	**114**
Xeon E5-2650L v3	2133	68	12	1.80	2	**115**
Xeon E5-2670 v3	2133	68	12	2.30	2	**117**
Xeon E5-2680 v3	2133	68	12	2.50	2	**117**
Xeon E5-2690 v3	2133	68	12	2.60	2	118 (est.)
Xeon E5-2683 v3	2133	68	14	2.00	2	117 (est.)
Xeon E5-2695 v3	2133	68	14	2.30	2	117 (est.)
Xeon E5-2697 v3	2133	68	14	2.60	2	**117**
Xeon E5-2698 v3	2133	68	16	2.30	2	**117**
Xeon E5-2699 v3	2133	68	18	2.30	2	**116**

The results marked (est.) are estimates.

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 ${}^{2}/{}_{3}n^3 + 2n^2$. 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 CX2570 M1
Processor	2 processors of Intel® Xeon® Processor E5-2600 v3 Product Family
Memory	8 × 16GB (1x16GB) 2Rx4 DDR4-2133 R ECC
Software	
BIOS settings	EnergyPerformance = Performance COD Enable = disabled Early Snoop = disabled All processors apart from Xeon E5-2603 v3 and E5-2609 v3: Turbo Mode = Enabled (default) = Disabled Hyper Threading = Disabled
Operating system	Red Hat Enterprise Linux Server release 6.6
Benchmark	Shared memory version: Intel Optimized LINPACK Benchmark 11.2 for Linux OS

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

Benchmark results

Processor	Cores	Rated Frequency [Ghz]	Number of processors	Rpeak [GFlops]	Rmax (without Turbo Mode) [GFlops]	Efficiency	Rmax (with Turbo Mode) [GFlops]	Efficiency
Xeon E5-2623 v3	4	3.00	2	384	**359**	**93%**	394	103%
Xeon E5-2637 v3	4	3.50	2	448	**419**	**94%**	419	94%
Xeon E5-2603 v3	6	1.60	2	307	**289**	**94%**		
Xeon E5-2609 v3	6	1.90	2	365	**342**	**94%**		
Xeon E5-2620 v3	6	2.40	2	461	**434**	**94%**	451	98%
Xeon E5-2643 v3	6	3.40	2	653	**611**	**94%**	613	94%
Xeon E5-2630L v3	8	1.80	2	461	**403**	**87%**	403	87%
Xeon E5-2630 v3	8	2.40	2	614	**540**	**88%**	540	88%
Xeon E5-2640 v3	8	2.60	2	666	**601**	**90%**	600	90%
Xeon E5-2667 v3	8	3.20	2	819	**738**	**90%**	737	90%
Xeon E5-2650 v3	10	2.30	2	736	**681**	**93%**	683	93%
Xeon E5-2660 v3	10	2.60	2	832	**717**	**86%**	715	86%
Xeon E5-2687W v3	10	3.10	2	992	**837**	**84%**	837	84%
Xeon E5-2650L v3	12	1.80	2	691	**527**	**76%**	522	76%
Xeon E5-2670 v3	12	2.30	2	883	**800**	**91%**	801	91%
Xeon E5-2680 v3	12	2.50	2	960	**817**	**85%**	819	85%
Xeon E5-2690 v3	12	2.60	2	998	**864**	**87%**	864	87%
Xeon E5-2683 v3	14	2.00	2	896	**836**	**93%**	842	94%
Xeon E5-2695 v3	14	2.30	2	1030	**895**	**87%**	895	87%
Xeon E5-2697 v3	14	2.60	2	1165	**995**	**85%**	995	85%
Xeon E5-2698 v3	16	2.30	2	1178	**1049**	**89%**	1048	89%
Xeon E5-2699 v3	18	2.30	2	1325	**1152**	**87%**	1154	87%

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 does not as a matter of principle 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.

System comparison

The following diagram illustrates the throughput of the PRIMERGY CX2570 M1 in comparison to its predecessor, the PRIMERGY CX270 S2, in their most performant configuration.

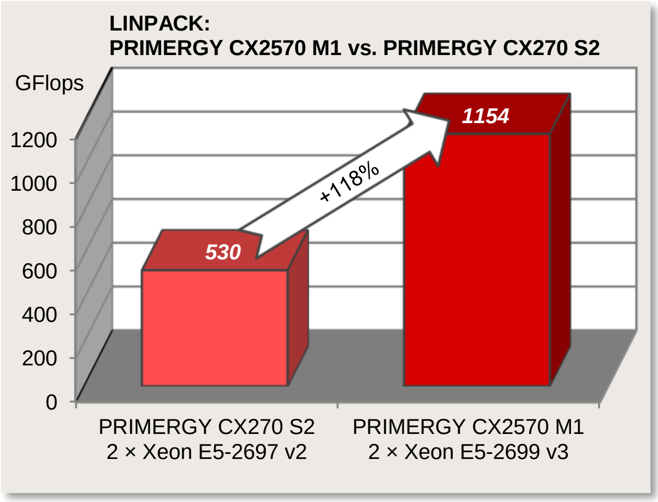

Literature

PRIMERGY Servers

http://primergy.com/

PRIMERGY CX2570 M1

This White Paper:

- http://docs.ts.fujitsu.com/dl.aspx?id=dbee567a-708b-48b9-972c-730de94a338f
- http://docs.ts.fujitsu.com/dl.aspx?id=d6ebf220-8238-4099-8679-51a100eeddec
- http://docs.ts.fujitsu.com/dl.aspx?id=23fb06a0-cf81-4423-8aa5-0f40453628f0

Data sheet
http://docs.ts.fujitsu.com/dl.aspx?id=ca7443e3-5735-4e6a-8c1a-b85c5f610aef

PRIMERGY Performance

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

Performance of Server Components

http://www.fujitsu.com/fts/products/computing/servers/mission-critical/benchmarks/x86-components.html

BIOS optimizations for Xeon E5-2600 v3 based systems
http://docs.ts.fujitsu.com/dl.aspx?id=f154aca6-d799-487c-8411-e5b4e558c88b

RAID Controller Performance
http://docs.ts.fujitsu.com/dl.aspx?id=e2489893-cab7-44f6-bff2-7aeea97c5aef

Disk I/O: Performance of storage media and RAID controllers

Basics of Disk I/O Performance
http://docs.ts.fujitsu.com/dl.aspx?id=65781a00-556f-4a98-90a7-7022feacc602

Information about Iometer
http://www.iometer.org

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/

SPECcpu2006

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

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

STREAM

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

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

2015-06-18 WW EN