White Paper
FUJITSU Server PRIMERGY
Performance Report PRIMERGY RX2510 M2

This document contains a summary of the benchmarks executed for the FUJITSU Server PRIMERGY RX2510 M2.

The PRIMERGY RX2510 M2 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
1.0
2016-07-06

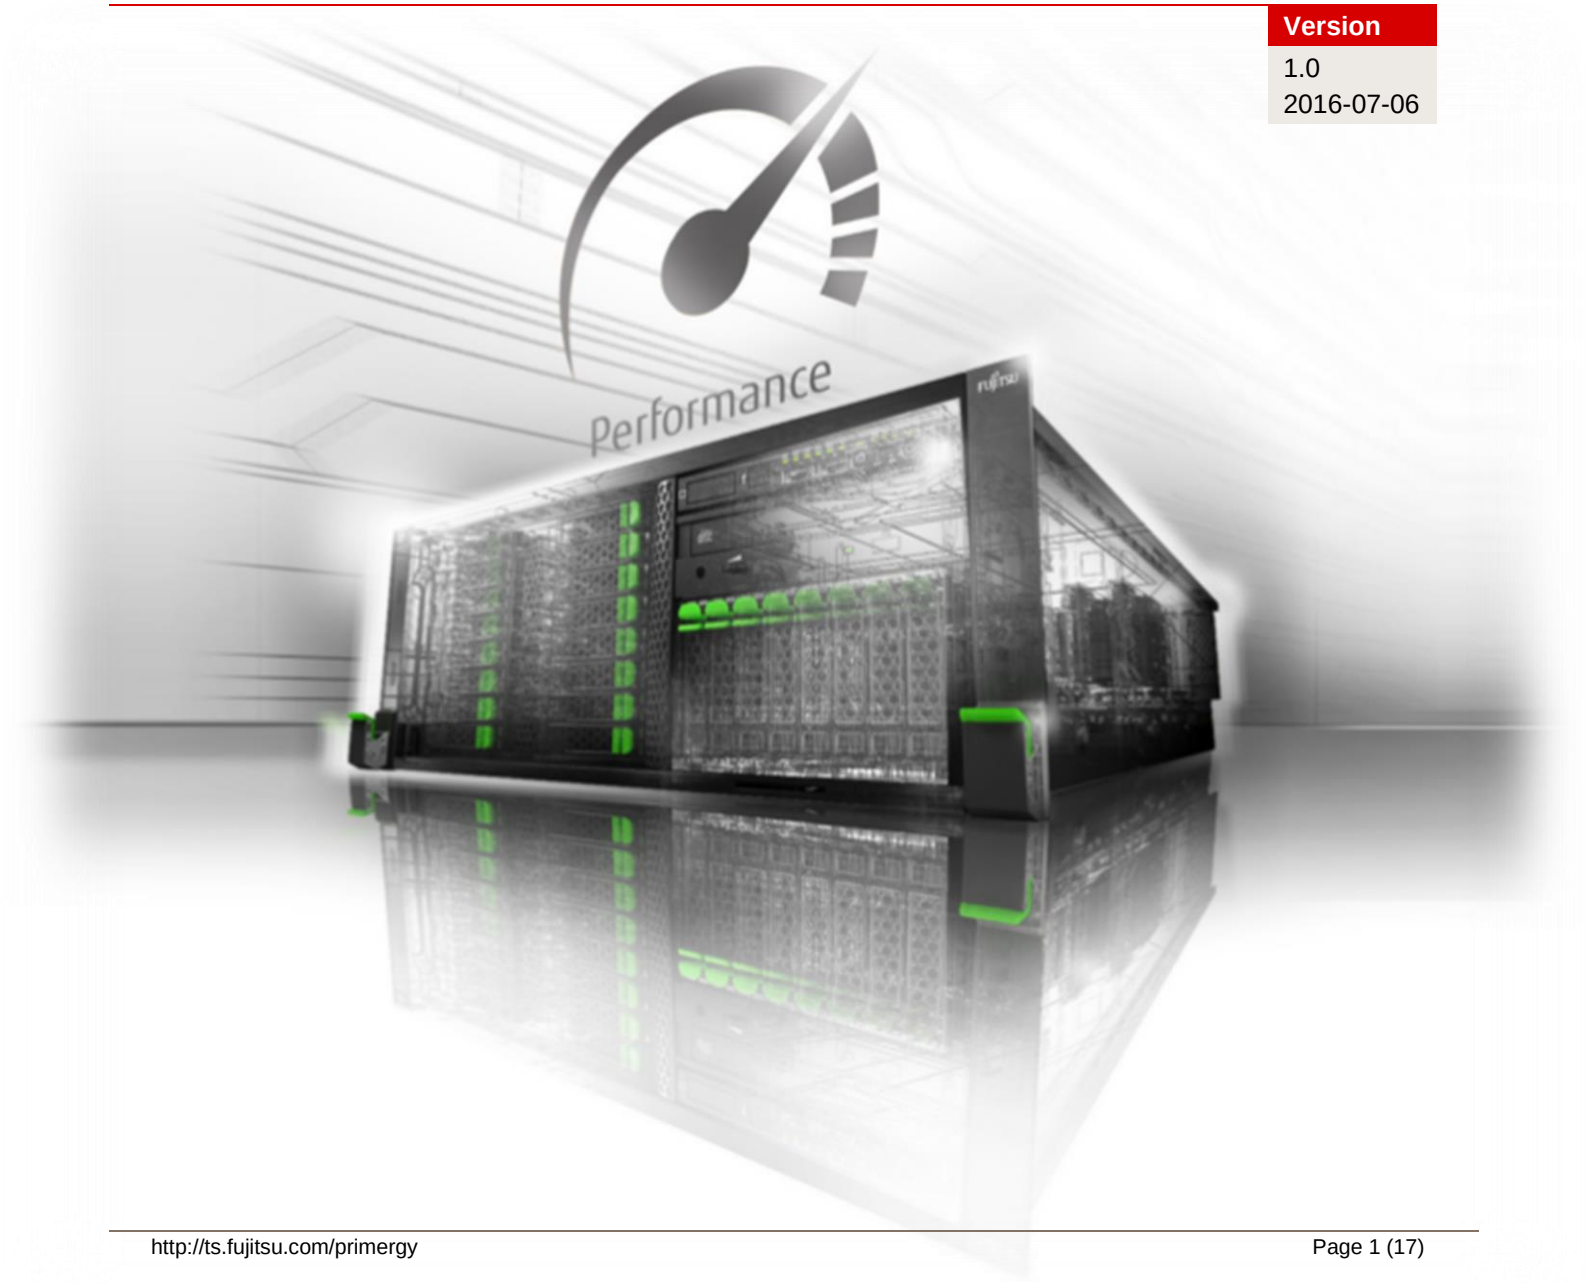

Contents

Document history 2
Technical data 3
SPECcpu2006 5
Disk I/O: Performance of RAID controllers 8
STREAM 15
Literature 17
Contact 17

Document history

Version 1.0 (2016-07-06)

New:

- Technical data
- SPECcpu2006
  Measurements with Intel® Xeon® Processor E5-2600 v4 Product Family
- 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
- STREAM
  Measurements with Intel® Xeon® Processor E5-2600 v4 Product Family

Technical data

PRIMERGY RX2510 M2

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 RX2510 M2
Form factor	Rack server
Chipset	Intel® C612
Number of sockets	2
Number of processors orderable	1 or 2
Processor type	Intel® Xeon® Processor E5-2600 v4 Product 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 4 SATA HDDs
PCI slots	2 × PCI-Express 3.0 x8 2 × PCI-Express 3.0 x16
Max. number of internal hard disks	4

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 v4	4	8	10	8.00	2.60	3.20	2133	85
Xeon E5-2603 v4	6	6	15	6.40	1.70	n/a	1866	85
Xeon E5-2609 v4	8	8	20	6.40	1.70	n/a	1866	85
Xeon E5-2620 v4	8	16	20	8.00	2.10	3.00	2133	85
Xeon E5-2630L v4	10	20	25	8.00	1.80	2.90	2133	55
Xeon E5-2630 v4	10	20	25	8.00	2.20	3.10	2133	85
Xeon E5-2640 v4	10	20	25	8.00	2.40	3.40	2133	90
Xeon E5-2650 v4	12	24	30	9.60	2.20	2.90	2400	105
Xeon E5-2650L v4	14	28	35	9.60	1.70	2.50	2400	65
Xeon E5-2660 v4	14	28	35	9.60	2.00	3.20	2400	105

All the processors that can be ordered with the PRIMERGY RX2510 M2, apart from Xeon E5-2603 v4 and Xeon E5-2609 v4, 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]	Load reduced	Registered	ECC
4GB (1x4GB) 1Rx8 DDR4-2400 R ECC	4	1	8	2400		✓	✓
8GB (1x8GB) 1Rx4 DDR4-2400 R ECC	8	1	4	2400		✓	✓
16GB (1x16GB) 2Rx4 DDR4-2400 R ECC	16	2	4	2400		✓	✓
32GB (1x32GB) 2Rx4 DDR4-2400 R ECC	32	2	4	2400		✓	✓

Power supplies (since system release)	Max. number
Modular PSU 450W platinum hp	2
Modular PSU 800W 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 RX2510 M2.

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

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX2510 M2
Processor	2 processors of Intel® Xeon® Processor E5-2600 v4 Product Family
Memory	8 × 16GB (1x16GB) 2Rx4 DDR4-2400 R ECC
Software	
BIOS settings	Energy Performance = Performance Utilization Profile = Unbalanced CPU C1E Support = Disabled Home Snoop Dir OSB = Disabled COD Enable = Disabled Early Snoop = Disabled Home Snoop Dir OSB = Enabled
Operating system	SUSE Linux Enterprise Server 12 SP1 (x86_64)
Operating system settings	echo always > /sys/kernel/mm/transparent_hugepage/enabled
Compiler	C/C++: Version 16.0.0.101 of Intel C++ Studio XE for Linux Fortran: Version 16.0.0.101 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_rate_base2006	SPECint_rate2006	SPECfp_rate_base2006	SPECfp_rate2006
Xeon E5-2623 v4	2	**371**	**390**	**319**	**326**
Xeon E5-2603 v4	2	**306**	**320**	**291**	**297**
Xeon E5-2609 v4	2	**404**	**422**	**342**	**350**
Xeon E5-2620 v4	2	**599**	**625**	**434**	**445**
Xeon E5-2630L v4	2	**648**	**674**	**453**	**467**
Xeon E5-2630 v4	2	**730**	**758**	**483**	**498**
Xeon E5-2640 v4	2	**766**	**796**	**499**	**513**
Xeon E5-2650 v4	2	**880**	**914**	**563**	**580**
Xeon E5-2650L v4	2	**838**	**869**	**525**	**543**
Xeon E5-2660 v4	2	**956**	**985**	**588**	**606**

Disk I/O: Performance of RAID controllers

Benchmark description

Performance measurements of disk subsystems for PRIMERGY and PRIMEQUEST servers are used to assess their performance and enable a comparison of the different storage connections for these 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 from 1 to 512 (in steps 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		
Processor	2 × Xeon E5-2623 v4 @ 2.60GHz	
Controller	1 × “LSI SW RAID on Intel C610 (Onboard SATA)” Intel C610 PCH, Code name Wellsburg Driver name: megasr1.sys, Driver version: 16.02.2014.0811 BIOS version: A.14.02121826R 1 × “PRAID CP400i”, “PRAID EP400i”, “PRAID EP420i”: Driver name: megasas2.sys, Driver version: 6.706.06 Firmware package: 24.7.0-0061	
Storage media	SSDs	HDDs
	Samsung MZ7KM240HAGR	HGST HUC156045CSS204 Seagate ST1000NM0033
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 R2 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 6.2.1	
Benchmark version	3.0	
Stripe size	Controller default	
Measuring tool	Iometer 1.1.0	
Measurement area	The first 10% of the usable LBA area is used for sequential accesses; the next 25% for random accesses.	
File system	raw	
Total number of Iometer workers	1	
Alignment of Iometer accesses	Aligned to whole multiples of 4096 bytes	

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 RX2510 M2 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. Only the highest supported version is named for all the interfaces we have dealt with in this section.

Form factor	Interface	Connection type	Number of PCIe controllers	Maximum number of hard disks
3.5"	SATA 6G	direct	0	4
3.5"	SATA 6G, SAS 12G	direct	1	4

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 RX2510 M2. 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		FBU
					Max. # disks per controller	RAID levels	
LSI SW RAID on Intel C610 (Onboard SATA)	Onboard C610	-	SATA 6G	-	4 × 3.5"	0, 1, 10	-
PRAID CP400i	PRAID CP400i	-	SATA 6G SAS 12G	PCIe 3.0 x8	4 × 3.5"	0, 1, 1E, 5, 10, 50	-
PRAID EP400i	PRAID EP400i	1 GB	SATA 6G SAS 12G	PCIe 3.0 x8	4 × 3.5"	0, 1, 1E, 5, 6, 10, 50, 60	✓
PRAID EP420i	PRAID EP420i	2 GB	SATA 6G SAS 12G	PCIe 3.0 x8	4 × 3.5"	0, 1, 1E, 5, 6, 10, 50, 60	✓

Onboard RAID controllers are implemented in the chipset on the server system board and use the CPU of the server for the RAID functionality. This simple solution does not require a PCIe slot. The chipset does not communicate with the CPU via PCIe, but via the “Direct Media Interface”, in short DMI. The PRIMERGY RX2510 M2 has the Intel C610 chipset.

System-specific interfaces

The interfaces of a controller in CPU direction (DMI or PCIe) 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	
1 × Onboard C610	1 × 4 × SATA 6G	2060 MB/s	4 × DMI 2.0	**1716 MB/s**	-
PRAID CP400i	4 × SAS 12G	**4120 MB/s**	8 × PCIe 3.0	6761 MB/s	-
PRAID EP400i	4 × SAS 12G	**4120 MB/s**	8 × PCIe 3.0	6761 MB/s	-
PRAID EP420i	4 × SAS 12G	**4120 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. an FBU).

For the purpose of easy and reliable handling of the settings for RAID controllers and hard disks it is advisable to use the software “ServerView RAID Manager” that is supplied for the server. 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”, “Data Protection” or “Fast Path optimum”. The “Performance” mode ensures the best possible performance settings for the majority of the application scenarios with HDDs. In connection with the “FastPath” RAID controller option, the “Fast Path optimum” mode should be selected if maximum transaction rates are to be achieved with SSDs for random accesses with small blocks (≤ 8 kB, e. g. OLTP operation of databases).

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 logical drive 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 RX2510 M2 if the configurations measured there are also supported by this system.

The performance values of the PRIMERGY RX2510 M2 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.

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

PRIMERGY RX2510 M2							
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	ST1000NM0033 SATA HDD MZ7KM240HAGR SATA SSD	2	1	349	319	43156	8392
		4	0	646	368	74488	17622
		4	10	545	315	60497	14045
PRAID CP400i	HUC156045CSS204 SAS HDD MZ7KM240HAGR SATA SSD	2	1	1153	971	66251	8407
		4	10	2080	1073	110800	16837
		4	0	2605	1361	141045	25091
		4	5	1353	764	27337	12009
PRAID EP400i	HUC156045CSS204 SAS HDD MZ7KM240HAGR SATA SSD	2	1	1590	888	66320	8412
		4	10	2479	1532	130194	16602
		4	0	2778	1914	173987	26273
		4	5	2387	1034	62412	10074
PRAID EP420i	HUC156045CSS204 SAS HDD MZ7KM240HAGR SATA SSD	2	1	1544	994	65552	8356
		4	10	2453	1614	130655	16622
		4	0	2799	1918	174254	26293
		4	5	2438	1112	61456	10039

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

PRIMERGY RX2510 M2							
Configuration version			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]
RAID Controller	Hard disk type	#Disks					
Onboard C610	ST1000NM0033 SATA HDD MZ7KM240HAGR SATA SSD	2	1	185	180	950	492
		4	0	716	712	1541	1262
		4	10	368	356	1375	641
PRAID CP400i	HUC156045CSS204 SAS HDD MZ7KM240HAGR SATA SSD	2	1	327	230	1019	476
		4	10	584	457	2028	947
		4	0	919	915	2044	1893
		4	5	684	685	2031	1405
PRAID EP400i	HUC156045CSS204 SAS HDD MZ7KM240HAGR SATA SSD	2	1	350	232	1018	476
		4	10	627	437	2028	948
		4	0	957	945	2040	1894
		4	5	730	709	2041	1491
PRAID EP420i	HUC156045CSS204 SAS HDD MZ7KM240HAGR SATA SSD	2	1	375	232	1014	476
		4	10	606	471	2025	949
		4	0	977	948	2038	1893
		4	5	751	709	2055	1490

Conclusion

At full configuration with powerful hard disks the PRIMERGY RX2510 M2 achieves a throughput of up to 2055 MB/s for sequential load profiles and a transaction rate of up to 174254 IO/s for typical, random application scenarios.

For best possible performance we recommend one of the plug-in PRAID 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 PRAID 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	
Model	PRIMERGY RX2510 M2
Processor	2 processors of Intel® Xeon® Processor E5-2600 v4 Product Family
Memory	8 × 16GB (1x16GB) 2Rx4 DDR4-2400 R ECC
Software	
BIOS settings	Energy Performance = Performance Utilization Profile = Unbalanced COD Enable = Disabled Early Snoop = Disabled Home Snoop Dir OSB = Enabled All processors except Xeon E5-2603 v4 and E5-2609 v4: Hyper-Threading = Disabled
Operating system	SUSE Linux Enterprise Server 12 (x86_64)
Operating system settings	Transparent Huge Pages inactivated
Compiler	Intel C++ Composer XE 2016 for Linux
Benchmark	STREAM Version 5.10

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 v4	1866	59.7	6	1.70	2	53.9
Xeon E5-2609 v4	1866	59.7	8	1.70	2	54.3
Xeon E5-2623 v4	2133	68.3	4	2.60	2	55.7
Xeon E5-2620 v4	2133	68.3	8	2.10	2	61.6
Xeon E5-2630L v4	2133	68.3	10	1.80	2	61.7
Xeon E5-2630 v4	2133	68.3	10	2.20	2	61.4
Xeon E5-2640 v4	2133	68.3	10	2.40	2	61.8
Xeon E5-2650 v4	2400	76.8	12	2.20	2	68.7
Xeon E5-2650L v4	2400	76.8	14	1.70	2	68.2
Xeon E5-2660 v4	2400	76.8	14	2.00	2	68.5

Further information about memory performance can be found in the White Paper Memory performance of Xeon E5-2600 v4 (Broadwell-EP) based systems.

Literature

PRIMERGY Servers

http://primergy.com/

PRIMERGY RX2510 M2

This White Paper:
http://docs.ts.fujitsu.com/dl.aspx?id=6f040412-794d-4933-8280-2db730846f6f
http://docs.ts.fujitsu.com/dl.aspx?id=ed6c949e-20fe-4f88-b418-c22351abcb98
http://docs.ts.fujitsu.com/dl.aspx?id=f0281f18-d45a-4c98-bc65-f7b08cb61477

Data sheet
http://docs.ts.fujitsu.com/dl.aspx?id=9ef8807f-4dd8-4aa2-b29e-048e551fd67c

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 v4 based systems
http://docs.ts.fujitsu.com/dl.aspx?id=eb90c352-8d98-4f5a-9eed-b5aade5ccae1

Memory performance of Xeon E5-2600 v4 (Broadwell-EP) based systems
http://docs.ts.fujitsu.com/dl.aspx?id=8f372445-ee63-4369-8683-da9557673357

RAID Controller Performance
http://docs.ts.fujitsu.com/dl.aspx?id=9845be50-7d4f-4ef7-ac61-bbde399c1014

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

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

2016-07-06 WW EN