

# White Paper

## FUJITSU Server PRIMERGY & PRIMEQUEST

### Benchmark Overview SAP Server Power Standard Application Benchmark

This document is an explanation of the SAP Server Power Standard Application Benchmark which was released in February 2011. The benchmark supplements the family of SAP Standard Application Benchmarks by an instrument for the assessment of energy efficiency in SAP environments.

#### Version

1.1  
2015-02-05



## Contents

Document history .....	2
SAP Benchmarks – An overview .....	3
a) 2-tier configuration architecture .....	4
b) 3-tier configuration architecture .....	5
SAP Benchmarks – Terminology .....	6
Benchmark Users and Average Dialog Response Time .....	6
Throughput Measurement in SAPS .....	6
Benchmark Toolkit .....	6
The SAP Server Power Standard Application Benchmark .....	7
Features of the SAP Server Power Benchmark .....	7
Metrics of the SAP Server Power Benchmark .....	9
Literature .....	11
Contact .....	11
SAP Copyrights .....	11

## Document history

### **Version 1.0 (September 2012)**

- Original version

### **Version 1.1 (February 2015)**

- New layout

## SAP Benchmarks – An overview

Since 1993 the SAP Standard Application Benchmarks have been developed by SAP to provide basic information for configuring, sizing and for platform comparison. The first benchmarks available (for SAP R/3 Release 1.1H) were targeted for FI (Financial Accounting), SD (Sales and Distribution), and MM (Materials Management) followed by ATO (Assemble-To-Order), PP (Production Planning), WM (Warehouse Management), BW (Business Information Warehouse) and now many more.

The SAP Benchmark Council (established in 1995 and consisting of representatives of SAP as well as hardware, logo and technology partners involved in benchmarking) define and control the content of the benchmarks and establish rules that encompass the testing procedures. The procedures involve the hardware companies running most of the benchmarks and sending the results to SAP. On request SAP certifies the results.

An SAP Standard Application Benchmark consists of script files that simulate the most typical transactions and workflow of a user scenario. It has a predefined SAP client database that contains sample company data against which the benchmark is run. The benchmark transactions of each component usually reflect the data throughput of an installation (for example orders, number of goods movements, etc.). However, benchmark transactions do not reflect reporting, because the resource consumption of a customer-defined report depends on the volume of data sought and is therefore not comparable. An exception is the BW benchmark which is mainly (but well defined) reporting activity.

Application components are customized for a benchmark run in such way that the system resource requirements are minimized while still representing an economic reality. Comparable customizing settings (buffer sizes, number of work processes, etc.) can be found in real-world customer installations that need high data throughput.

In general each benchmark user has his own master data such as material, vendor or customer master data to avoid data locking situations. For most benchmarks a maximum of 1000 parallel benchmark users can be simulated simultaneously per client. The multi-tier client/server architecture consists of a presentation, application and database layer. The presentation layer – in real life the PCs of the logged-in users – is handled by one (or occasionally more) machines dubbed 'benchmark driver(s)'. There is no way for the application layer to tell whether it is driven by real users or a simulation environment.




Figure 1: Multi-tier client/server architecture

Possible configurations for a benchmark simulation are:

### a) 2-tier configuration architecture

Database and application layer reside on a single system – the simulation is driven by the presentation server (aka benchmark driver).

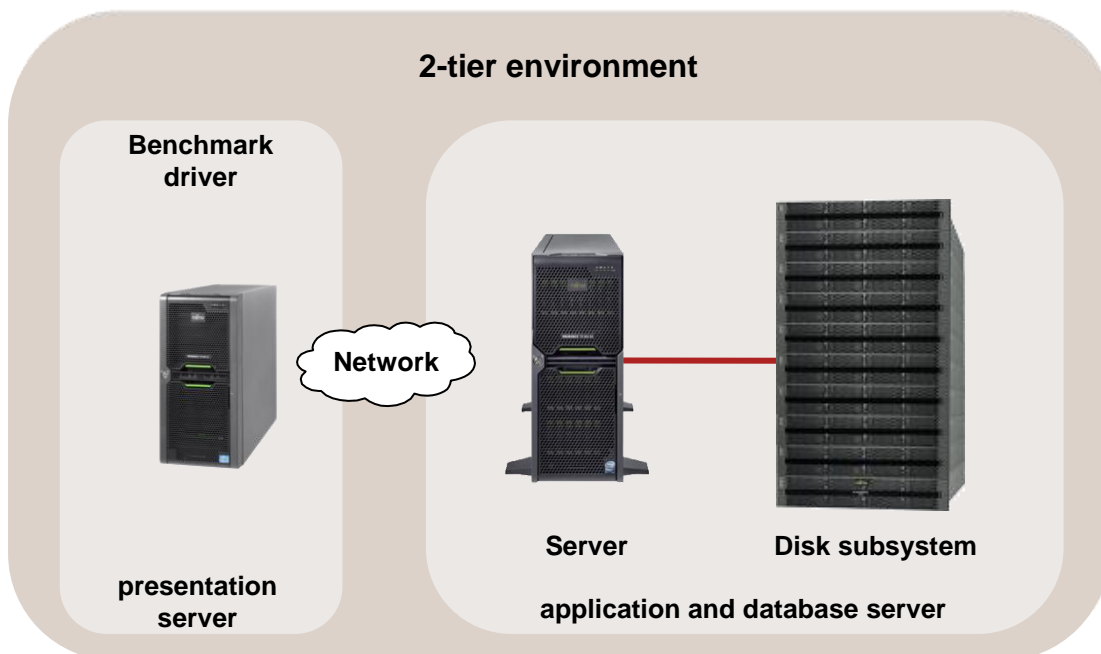


Figure 2: 2-tier configuration

## b) 3-tier configuration architecture

Database and application layer reside on different systems – the simulation is driven by a presentation server (aka benchmark driver).

Based on the architecture possible configurations are:

- 1 database server (or more - using parallel database techniques)
- n application servers with dedicated enqueue, update, message and dialog functions
- n presentation servers (benchmark driver)

An impressive degree of scalability can be achieved with such a configuration.

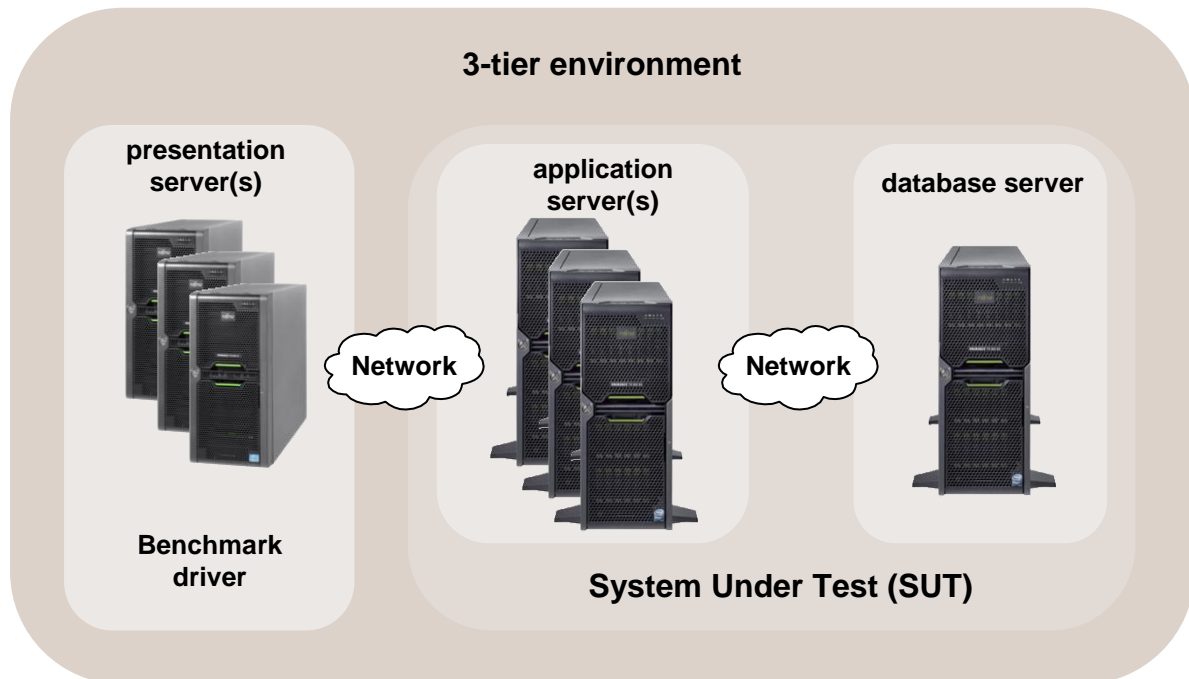


Figure 3: 3-tier configuration

The SAP Standard Application Benchmarks measure all performance-relevant metrics such as database request times, wait times, CPU utilization, average dialog response times by a given number of benchmark users (with a fixed think time of 10 seconds between each dialog step) and the achieved throughput. The most significant parameters have to be part of every publication about SAP benchmarks initiated by the platform partners.

The following information must be part of a benchmark press release:

- The SAP Business Suite component (e.g. "SAP ERP 6.0 with Enhancement Pack 5")
- RDBMS and operating system release
- Tested standard SAP Business Suite components (FI, PP, SD or a combination of these)
- Number of tested benchmark users (if applicable)
- Average dialog response time in 'n.nn sec' (if applicable)
- Achieved throughput in dialog steps / hour or business indicators such as "accounts balanced"
- Type of client/server configuration
- A detailed description of the hardware configuration (type, size of main memory, average CPU utilization and functions of the individual servers involved in the benchmark)
- Confirmation that the benchmark is certified by SAP (e.g. "This benchmark fully complies with SAP's issued benchmark regulations and has been audited and certified by SAP.")
- Reference where readers can get more information (e.g. "Details regarding this benchmark are available upon request from the hardware partner or SAP AG").

## SAP Benchmarks – Terminology

### Benchmark Users and Average Dialog Response Time

A benchmark can only be certified if the average dialog response time is less than a fixed amount of time (think about it as system reaction time). More and more benchmark users are added to the system until the average response time is outside the granted time frame.

Only SAP audited and certified benchmarks may be published by partners to ensure results that can be fairly compared with each other. A typical result would read like '2550 SD benchmark users with an average dialog response time of 0.96 seconds'.

### Throughput Measurement in SAPS

SAP has defined a unit for measuring throughput in a SAP Business Suite environment: **SAPS** (SAP Application Benchmark Performance Standard).

100 SAPS are defined as 2,000 fully processed order items per hour in the standard SD application benchmark. This throughput is achieved by processing 6,000 dialog steps (screen changes) and 2,000 postings per hour or processing 2,400 SAP transactions in the SD benchmark.

In the SD standard benchmark 'fully processed' means the full workflow of an order item (creating the order, creating a delivery note for this order, displaying the order, changing the delivery, posting a goods issue, listing orders and creating an invoice) has completed.

### Benchmark Toolkit

In order to have a benchmark environment which enables fairly easy usage and reproducible results, a continuously maintained and updated toolkit is publicly available. It makes heavy use of Perl scripts plus some C source code which is either precompiled by SAP or by the user for the specific target platform.

## The SAP Server Power Standard Application Benchmark

The family of SAP Standard Application Benchmarks was extended to include two new members in February 2011. The extension was motivated by the increased significance of energy efficiency for server selection. SAP applications need large IT infrastructures. It is obvious that energy efficiency has also become a key selection criterion in this environment.

The two newly introduced benchmarks are called:

- SAP Server Power Standard Application Benchmark
- SAP System Power Standard Application Benchmark

Both benchmarks focus on measuring energy consumption with the aid of power metering equipment during benchmark runs and also on incorporating the power, expressed in watts, in the measurement result.

The new SAP Benchmarks differ according to the extent of components involved in the power measurement. In the case of the System Power Benchmark external disk subsystems and network components, such as switches, are involved in the power measurement; the measurement with the Server Power Benchmark is restricted to the application and database server or servers. Due to the greater simplicity, and because the energy efficiency of very different components is not mixed, the Server Power Benchmark has been preferred in the previous activities of SAP partners.

Fujitsu was closely involved in the development of the SAP Power Benchmarks within the framework of the SAP Benchmark Council.

### Features of the SAP Server Power Benchmark

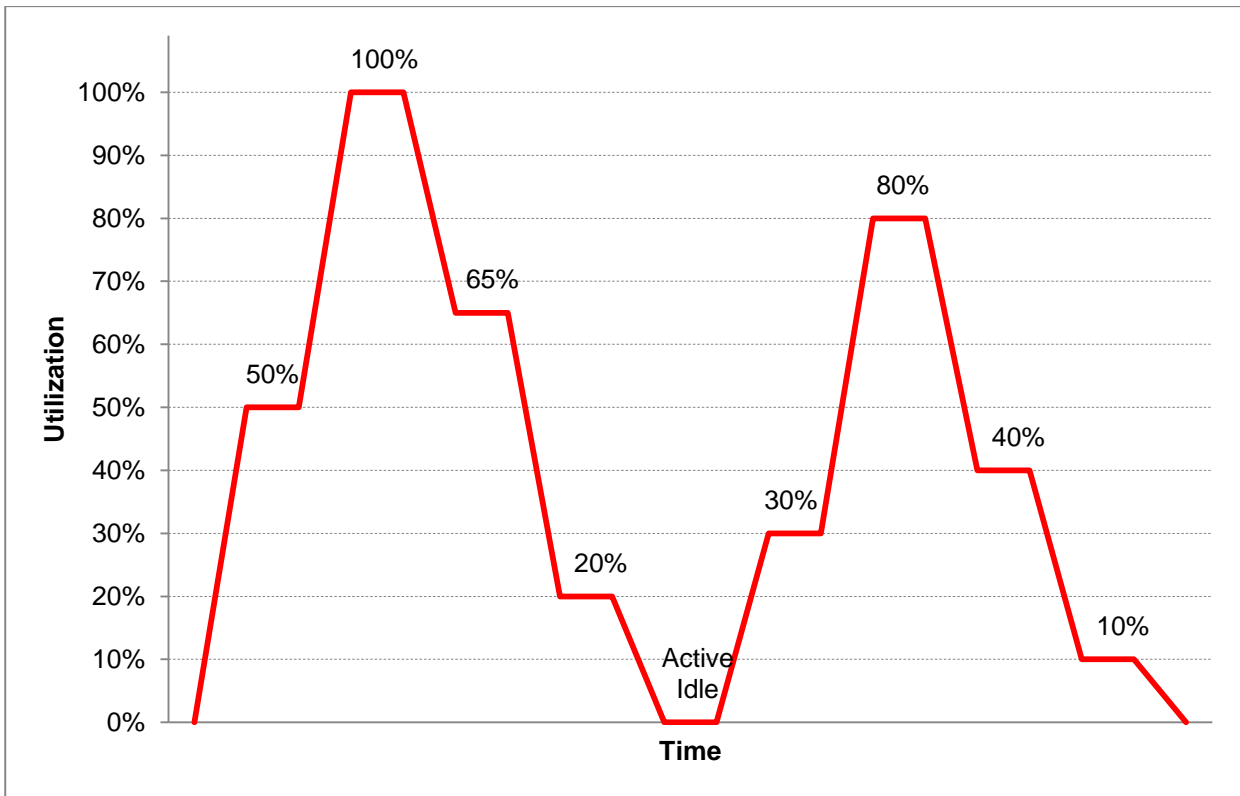
The benchmark is based on the SAP Sales & Distribution (SD) load profile [L4]. The SD benchmark is one of the most well-known and frequently used SAP Standard Application Benchmarks. It is a dialog benchmark, i.e. a larger number (with 4 to 5 digits) of users are simulated, which perform a specific sequence of SAP transactions in loops. The transactions map a sale-from-stock scenario. From a technical viewpoint, this load profile is CPU and memory hungry, with a moderate proportion of disk and network I/O.

The SD benchmark is available in the 2-tier form described above with a single application and database server, and in the 3-tier form with several application servers and one database server. Precisely these two configuration options are also available with the SAP Server Power Benchmark. In each case, it is necessary to measure the power consumption of the application and database server or servers. Otherwise, the load profile of the classic SD benchmark runs during the measurement.


The particular feature of the SAP Server Power Benchmark is that the load profile is performed in several load levels. In the case of classic SD measurement there is only one measurement interval with maximum CPU utilization. The Server Power Benchmark has nine measurement intervals with load levels between Active Idle and 100% (maximum CPU utilization). The first of the two diagrams below shows these load levels and in particular their sequence. The horizontal sections are the actual measurement intervals. During these phases the QoS (Quality of Service) requirements of the SD benchmark have to be fulfilled - especially the requirement for the average response time per dialog step to be less than one second. As with the SD benchmark, the same requirement of at least 15 minutes also applies for the length of the measurement interval. Between the horizontal phases the number of simulated users is adapted to the required load level in each case.

The number of measurement intervals, their length and the necessity of the interposed adaptation to the simulated number of users mean a comparatively high time requirement per measurement of between four and five hours. The configuration and goal of every measurement should be carefully considered beforehand.

Power measurement for various load levels is a principle of the benchmarks for energy efficiency and are also to be found for example in the SPECpower\_ssj2008 benchmark [L5]. In the category of these benchmarks SPECpower\_ssj2008 is to a certain degree a forerunner. The non-linear descending sequence of the load levels and the significance of Active Idle are the main differences to SPECpower\_ssj2008. In the case of SAP Server Power Benchmark a simulated user runs for each configured SAP instance with Active Idle. This ensures that the SAP system is also on full standby during the idle phase. SAP instance is a term from the SAP software architecture and denotes a pool of work processes that perform SAP transactions. On the other hand, no transactions at all are performed in the case of Active Idle for SPECpower\_ssj2008.



In the meantime servers have sophisticated mechanisms for the regulation of power consumption subject to utilization. A particularly effective example is the reduction in CPU frequency at low utilization. The load level methodology in the benchmarks for energy efficiency aims to test the quality of these mechanisms. The second diagram below is taken from the certificate of an SAP Server Power Benchmark and shows the large range in power consumption between Active Idle and full load. The power consumption in the SAP Server Power Benchmark certificates is also published in this much detail.



For the sake of completeness it should be mentioned that - apart from power consumption - room temperature is also continuously measured in the vicinity of the measuring configuration during the measurement and may at no point in time be below 20° C. The fan-driven server cooling, which is incorporated in the energy budget, should take place under realistic data center conditions.



## Metrics of the SAP Server Power Benchmark

The main metric, which is denoted as the *Power Efficiency Indicator*, is watts/kSAPS. This metric puts power consumption and performance into relation to each other. Performance is expressed by the SAPS throughput measurement, which is explained above. k stands for kilos (1000). The main metric says how much energy (watts) is needed for the set work quantum of 1000 SAPS. In order to calculate this, the arithmetic mean of the nine watt values of the load levels is divided by the arithmetic mean of the corresponding kSAPS values. This metric is all the more remarkable in the light of the fact that the motto “less is better” applies. Otherwise, and this is for example also valid for SPECpower\_ssj2008, benchmark metrics are usually formed according to the “larger is better” principle. However, in the case of the SAP Server Power Benchmark emphasis was placed on the energy actually consumed, and “less is better” applies for this. One alternative, which corresponds to the SPECpower\_ssj2008 metric, would have been the reciprocal value kSAPS/watts.

The efficiency metric of watts/kSAPS reveals little about which absolute performance level of the servers measured is reached. This is why the arithmetic mean of the nine SAPS values obtained is specified as the second metric. Although two servers, A and B, could for example have the same efficiency, one of the two servers could support considerably more SAP users. See the secondary SAPS metric for the latter circumstances.

The minimum room temperature obtained during the measurement is specified as the third metric.

The SAPS metric as the average value of the throughputs obtained in the nine measurement intervals is the only performance measure that is specified. The number of respectively configured benchmark users is not mentioned. In this way, the SAP Server Power Benchmark is clearly separated from the SD benchmark, for which the number of users is the primary metric. It does not make sense to compare SAP Server Power measurements and classic publications with the SD benchmark, because the respective goals of maximum energy efficiency and maximum performance are different. Likewise, the secondary SAPS metric of the SD benchmark cannot be compared with the SAPS metric of the power benchmark, because the latter is a mean value taken from nine load levels and the former is a simple value under maximum load.

The following page shows as an example the full certificate of the SAP Server Power Benchmark of the PRIMERGY RX300 S7.



# CERTIFICATION

## SAP® Standard Application Benchmarks

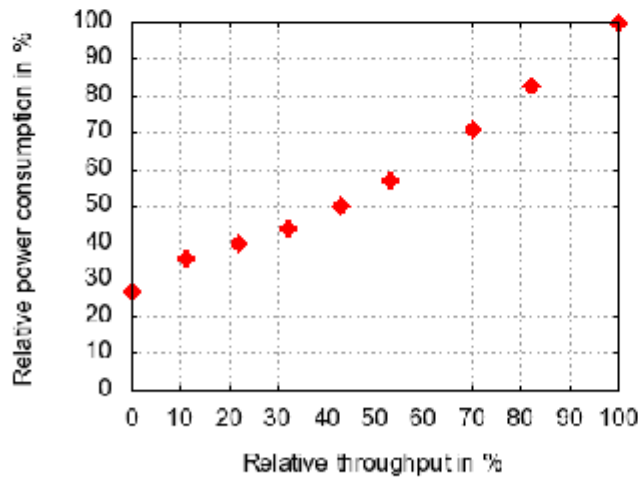
The SAP Server Power Standard Application Benchmark performed on February 7, 2012 by Fujitsu in Paderborn, Germany, was certified on March 6, 2012, with the following data:

<b>Power Efficiency Indicator – Server (watts/kSAPS):</b>	10.5
<b>Power Efficiency Indicator – Storage (watts/kSAPS):</b>	not measured
<b>Power Efficiency Indicator – Server &amp; Storage (watts/kSAPS) :</b>	not measured
<b>Average throughput over all load levels (SAPS):</b>	15,480
<b>Minimum ambient temperature (degrees Celsius):</b>	22.0
<b>Operating system, central server:</b>	Windows 2008 R2 Enterprise Edition
<b>RDBMS:</b>	SQL Server 2008
<b>SAP Business Suite software:</b>	SAP enhancement package 4 for SAP ERP 6.0


**Configuration:**

Component	Usage	Hardware
1	DB and Application server	Fujitsu PRIMERGY RX300 S7, 2 processors / 16 cores / 32 threads, Intel Xeon Processor E5-2660, 2.20 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 20 MB L3 cache per processor, 128 GB main memory, 450W PSU Module Platinum, 1 x 250 GB 2.5" SATA HDD

**Power Characteristics:**



**Scope of power measurement:**



Certification number: 2012007



## Literature

### [L1] PRIMERGY & PRIMEQUEST Servers

<http://www.fujitsu.com/fts/products/computing/servers/>


### [L2] PRIMERGY & PRIMEQUEST Performance


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


### [L3] Benchmark descriptions

<http://www.fujitsu.com/fts/products/computing/servers/primergy/benchmarks/benchmark-descriptions.html>

This White Paper:

 <http://docs.ts.fujitsu.com/dl.aspx?id=9a500709-589c-4a36-9a5d-bcf28debad7>

 <http://docs.ts.fujitsu.com/dl.aspx?id=cef058a1-7032-4187-b628-285a971f8fea>

 <http://docs.ts.fujitsu.com/dl.aspx?id=2dcfc52b-896a-49ba-9131-946188fc4e68>

### [L4] Benchmark Overview SAP SD Standard Application Benchmark

<http://docs.ts.fujitsu.com/dl.aspx?id=0a1e69a6-e366-4fd1-a1a6-0dd93148ea10>

### [L5] Benchmark Overview SPECpower\_ss2008

<http://docs.ts.fujitsu.com/dl.aspx?id=166f8497-4bf0-4190-91a1-884b90850ee0>

### [L6] SAP Standard Application Benchmarks

<http://www.sap.com/benchmark>

## Contact

### FUJITSU

Website: <http://www.fujitsu.com/>

### PRIMERGY & PRIMEQUEST Product Marketing

<mailto:Primergy-PM@ts.fujitsu.com>

### PRIMERGY Performance and Benchmarks

<mailto:primergy.benchmark@ts.fujitsu.com>

## SAP Copyrights

■SAP, SAP Logo, mySAP.com, mySAP.com Marketplace, mySAP.com Workplace, mySAP.com Business Scenarios, mySAP.com Application Hosting, WebFlow, R/2, R/3, RIVA, ABAP, SAP Business Workflow, SAP EarlyWatch, SAP ArchiveLink, BAPI, SAPHIRE, Management Cockpit, SEM, are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world. All other products mentioned are trademarks or registered trademarks of their respective companies.