

White Paper

FUJITSU Server PRIMERGY & PRIMEQUEST Benchmark Overview vServCon

Conventional benchmarks are less suited for the assessment of virtualized operating systems and applications, which is why special virtualization benchmarks exist. The “vServCon” benchmark developed by Fujitsu for internal purposes allows virtualization solutions and PRIMERGY and PRIMEQUEST servers to be measured and assessed.

This document describes the problems concerning benchmarks for virtualized environments as well as the fundamentals of the “vServCon” benchmark and its use at Fujitsu.

Version

1.6a

2015-02-25



Contents

Introduction	3
vServCon	4
vServCon Benchmark	4
vServCon Score	5
Load profile and hardware environment	7
Literature.....	9
Contact	9

Introduction

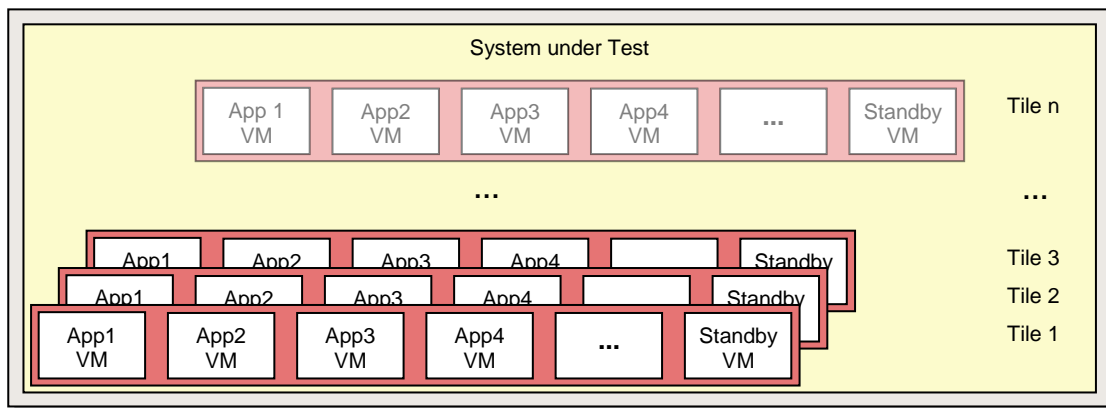
Server virtualization is an increasingly deployed scenario in the implementation of IT infrastructures. On the one hand, virtualization enables more efficient hardware to be used in an optimal way, and on the other hand the dependency on hardware specifics is being reduced. Contrary to the classic server with only one operating system and its applications, with virtualized environments, several operating systems and applications are run in parallel creating heterogeneous environments on one server. The handling of virtual machines is implemented by a virtualization layer, also called hypervisor.

Classic benchmarks are scarcely suited to measure and assess the performance capability of virtualized environments. For this purpose, it is necessary to utilize the hardware resources of a server with simultaneously working virtual machines with different workloads. The one aim of virtualization benchmarks is server consolidation. In this case, the throughput of a set collection of virtual machines is maximized on a single virtualization host by means of suitable replication. vConsolidate (Intel), VMmark V1 (VMware) and SPECvirt_sc2010 (SPEC) fall under this category. Fujitsu carries out scaling measurements of virtual environments using its internal benchmark “vServCon” (based on ideas from “vConsolidate”). The other aim of virtualization benchmarks is data center operations. A server consolidation scenario for several virtualization hosts is assumed in this case. In addition to the throughputs of the virtual machines, the benchmark metric then contains ratios that reflect the efficiency of typical data center operations, such as the relocation of virtual machines. These benchmarks include VMmark V2 (VMware).

For a virtualization benchmark to fulfill its objective, it must map the real world of a data center regarding server consolidation; in other words it must consider existing servers with those application scenarios that are normally virtualized. These servers have weak utilization levels and the aim is thus to consolidate as many of them as possible as virtual machines (VMs). Therefore, such a benchmark must assess for a virtualization host both the suitably determined overall throughput across the various application VMs as well as the number of efficiently operable VMs.

The following solution concept has been established for these two objectives: a representative group of application scenarios is selected in the benchmark. They are started simultaneously as a group of VMs on a virtualization host when making a measurement. Each of these VMs is operated with a suitable load tool at a defined lower load level. All known virtualization benchmarks are thus based on a mixed approach of operating system and applications - plus usually an “idle” or “standby” VM which represents the inactive phases of a virtualization environment and simultaneously increases the number of VMs to be managed by the hypervisor. The term “tile” is the name for such a unit of virtual machines.

It must be possible to increase this well-defined load created by this group of virtual machines on a step-by-step basis until the considered system has reached its performance limit. The following illustration shows the growth of VM load on a system under test by operating several tiles.



An application is executed in each virtual machine where the applications are put under stress via established benchmarks. If necessary, there may also be further infrastructure components. All the individual results are then suitably summarized in one overall result. This score is an indication for the performance capability of a virtualized environment.

vServCon

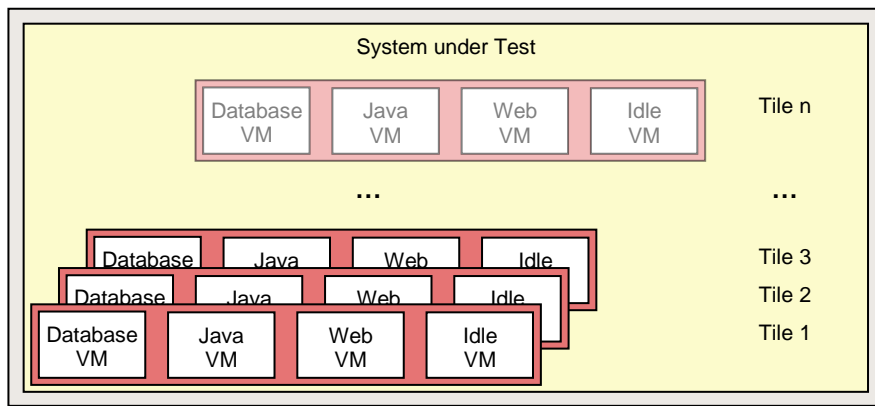
Scalability measurements of virtualized environments at Fujitsu are currently accomplished by means of the internal benchmark “vServCon” (based on ideas from Intel’s “vConsolidate”), which is described below. The abbreviation “vServCon” stands for: “Virtualization enables SERVer CONSolidation”.

vServCon Benchmark

vServCon is not a new benchmark in the actual sense. It is in fact a framework that consolidates already established benchmarks, as workloads, if necessary in modified form in order to simulate the load of a virtualized consolidated server environment. Three proven benchmarks are used, which cover the application scenarios database, application server and web server.

Application scenario	Load tool	# VMs
Database	Sysbench	1
Java-Applikationsserver	SPECjbb2005	1
Web-Server	WebBench	1
(Idle)	-	1

Each of the three application scenarios is assigned to one dedicated virtual machine (VM). Then add to these a fourth VM called the ‘idle VM’. These four VMs form a “tile”. In the terminology of “vConsolidate” this would be a “consolidation stack unit” (CSU). Because of the performance capability of the underlying server hardware, it is usually necessary to have started several identical tiles in parallel as part of a measurement in order to achieve a maximum overall performance.



vServCon requires external load generators: exactly one client system per tile for the web server VM. Load generators and “system under test” are connected via a suitable number of networks.

The execution of the individual load tools is controlled by a separate system, the so-called framework controller. This system monitors the measurements and collects the individual performance data of the VMs.

Two of the standard benchmarks, Sysbench and SPECjbb, had to be adapted for use in a virtualized environment where they could then be measured within the vServCon framework.

vServCon Score

The result of vServCon is a number, known as a “score”, which provides information about the performance of the measured virtualization host. According to the objective such a score has to reflect the maximum total throughput that can be achieved by running a defined mix that consists of numerous application VMs.

The score is determined from the individual results of the VMs. Each of the three vServCon application scenarios provides a specific benchmark result in the form of application-specific transaction rates for the respective VM. In order to derive a normalized score the individual benchmark results for one tile are put in relation to the respective results of an always identical reference system (PRIMERGY RX300 S3 with two Xeon 5130 processors and 16 GB RAM). The resulting relative performance values are then suitably weighted and finally added up for all VMs and tiles. The outcome is the vServCon score for this tile number (see [below](#) for details of the score calculation).

Starting as a rule with one tile, this procedure is performed for an increasing number of tiles until no further significant increase in this vServCon score occurs. The final vServCon score is then the maximum of the vServCon scores for all tile numbers.

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

Optionally, the total CPU load of the host (VMs and all other CPU activities) and, if possible, electrical power consumption can be recorded during a vServCon measurement.

In-depth information: Details on how to calculate a score

The score calculation, as described in detail below, is not absolutely necessary to understand the rest of the document. You can omit it if you do not have a special interest in this topic.

The score for a current tile number n is denoted below with $Score_n$. It is made up from the total of the individual scores of all tiles that were determined for the individual VMs.

$$Score_n = \sum_i \sum_{VM} Score_{VM_i}$$

VM denotes one of the 3 benchmark VMs of a tile (without the idle VM)
i numbers the tiles from 1 to n

Note that it is not possible to simply use the measured value of the benchmark as the score of a VM. This is because the individual benchmarks provide measured values at entirely different dimensions and in different units of measurement. For this reason, the individual benchmark results are qualified by putting them in relation to the result of one single running tile on the reference system. Thus the latter is given the score 1. The scores of all benchmarks now have the same dimension, are free from units of measurement, and can therefore be added up.

The relative score

$$\frac{R_{VM_i}}{R_{VM_{ref}}}$$

R_{VM_i} Benchmark result of the measured VM
 $R_{VM_{ref}}$ Benchmark result of the VM when measured with precisely one tile on the reference system

of a VM is multiplied with a weighting factor α_{VM} that is specific to the respective VM type. The score of the VM is formed in this way:

$$Score_{VM_i} = \frac{R_{VM_i}}{R_{VM_{ref}}} \alpha_{VM}$$

α_{VM} Weighting factor of the VM type

The weighting factor α_{VM} of a VM type takes into account and assesses the CPU and memory configuration of the VM in relation to the CPU and memory configuration of the total tile:

$$\alpha_{VM} = \frac{CPU_{VM}}{CPU_{CSU}} \times 0.8 + \frac{MEM_{VM}}{MEM_{CSU}} \times 0.2$$

$$CPU_{CSU} = \sum_{VM} CPU_{VM} \quad MEM_{CSU} = \sum_{VM} MEM_{VM}$$

CPU_{VM} : Number of configured virtual CPUs of the virtual machine VM

CPU_{Tile} : Number of virtual CPUs in a tile

MEM_{VM} : Size of the configured memory of the virtual machine VM

MEM_{Tile} : Size of the overall memory in a tile

Regarding CPU configuration, please note that a virtual CPU (vCPU) is not absolutely identical with a physical CPU of the virtualization host. A virtual CPU always consists of only one single logical execution unit of a physical CPU. For example, this is the case when using a single core of a multicore CPU, or when using one thread of a CPU with Hyperthreading. The configuration of the virtual systems does not directly refer to the inner structure of a physical CPU. It is the job of the respective virtualization hypervisor to ensure that during runtime the dynamic CPU scheduling takes place in most efficient way possible.

Load profile and hardware environment

In principle, it would be possible to optimize each virtual machine with regard to CPU and memory resources as well as the operating system and application software in such a way that the highest possible score is achieved. However, this would complicate comparability, which is the most important basis for scaling statements, to a large extent. Therefore, Fujitsu defines a standardized benchmark environment and uses a so-called “profile” within the “vServCon” framework, in which the resources and software versions of the operating system and the application software are specified precisely. The compilation of guest operating systems and applications and their specified resource requirements are to be understood as a representative selection for a complex virtualization environment. Even if there are newer and perhaps even more high-performance versions, continuity of the load profile must be maintained for as long a time as possible in order to ensure comparability. The following profile is used for the PRIMERGY and PRIMEQUEST “vServCon” measurements.

	Resource	Web WebBench	Database Sysbench	Java SPECjbb	Idle
Fujitsu Profile v1.0a	# vCPU	1	2	2	1
	Memory	1.5 GB	1.5 GB	2.0 GB	0.4 GB
	OS	Linux 32-bit SUSE SLES 10	Windows Server 2003 R2 EE SP2 64-bit	Linux 64-bit RedHat RHEL 4.5	Windows Server 2003 R2 EE SP2 32-bit
	Application	Apache	Microsoft SQL Server 2005 SP2	BEA JRockit(6)	
	Benchmark	WebBench 5.0	Sysbench v0.4.0 with Intel Binaries	SPECjbb 2005	-
	Disk Subsystem	10 GB	10 GB	5 GB	5 GB

As to the operating system mix the popular Linux distributions in Europe have been considered.

We have also defined the disk subsystem with not only the logical design regarding size and RAID level, but also the physical implementation. Local disk subsystems are usually not a sensible solution for virtualization and consolidation; that is why SAN-based disk subsystems are used. Since the performance capability of the virtualization solution and of the server hardware is to be determined, the disk-I/O subsystems are designed in such a way to avoid being a bottleneck in the longer term.


Neither should the main memory of the “System under Test” be a bottleneck for the vServCon measurement. Therefore, an adequate quantity is equipped so that even with a maximum tile number no swap activities take place at host level.

The benchmark “WebBench” requires external load generators. Load generators and the system under test are connected by a dedicated LAN with sufficient bandwidth.

A second dedicated network is used to control the “vServCon” benchmark.

An additional PRIMERGY system is used as the framework controller.

The measurement set-up is symbolically illustrated below:



The configuration of the individual VMs is always done identically according to a defined specification. Specifically, that is to say neither the operating system nor the respective benchmark software is optimized. Although optimization would by all means be possible to gain a higher score, it is not sensible within the context of comparability. Modifications required due to the hypervisor used can always be possible and are documented accordingly.

Outside of the VMs, at BIOS level or in the virtualization software, settings can be adapted to achieve a better score for the performance of the virtualization platform, if this is urgently required with regard to the performance optimization of the virtualization platform. Any settings recommended or stipulated by the Hypervisor manufacturer are also taken into consideration for vServCon. The standard settings are retained if sensible. This usually represents a very good compromise between performance and energy-efficiency.

Literature

PRIMERGY & PRIMEQUEST Servers

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


PRIMERGY & PRIMEQUEST Performance


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


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=b953d1f3-6f98-4b93-95f5-8c8ba3db4e59>

 <http://docs.ts.fujitsu.com/dl.aspx?id=214ee9dc-9239-4985-86e4-f0f9ac78ea25>

 <http://docs.ts.fujitsu.com/dl.aspx?id=c3d5ce5d-5610-43c6-86b4-051549940a71>

Benchmarks

SPECjbb2005

<http://www.spec.org/jbb2005/>

SPECvirtualization

<http://www.spec.org/benchmarks.html#virtual>

VMmark

<http://www.vmmark.com/>

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>