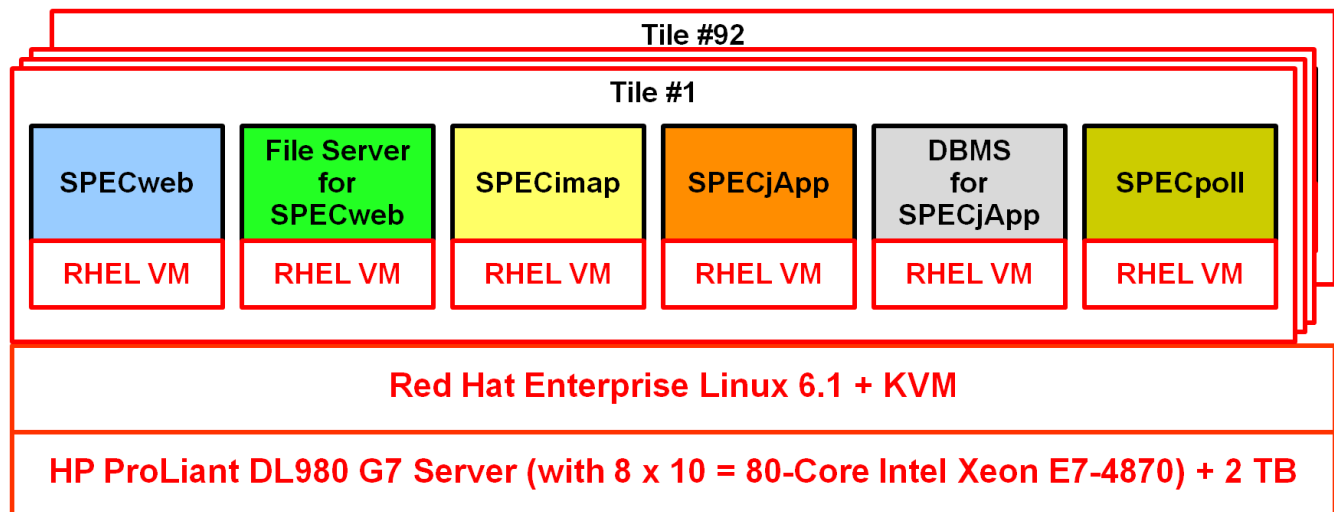




## Red Hat – Industry Standard Benchmarks

# SPECvirt\_sc<sup>®</sup> 2010 Benchmark Results using Red Hat<sup>®</sup> Enterprise Linux<sup>®</sup> and Kernel-based Virtual Machine (KVM)



Version 8.0

April 2012





**SPECvirt\_sc<sup>®</sup> 2010 Benchmark Result using  
Red Hat<sup>®</sup> Enterprise Linux<sup>®</sup> and  
Kernel-based Virtual Machine (KVM)**

1801 Varsity Drive  
Raleigh NC 27606-2072 USA  
Phone: +1 919 754 3700  
Phone: 888 733 4281  
Fax: +1 919 754 3701  
PO Box 13588  
Research Triangle Park NC 27709 USA

Linux is a registered trademark of Linus Torvalds. Red Hat, Red Hat Enterprise Linux and the Red Hat "Shadowman" logo are registered trademarks of Red Hat, Inc. in the United States and other countries.

SPEC<sup>®</sup> and the benchmark name SPECvirt\_sc<sup>®</sup>, SPECimap<sup>®</sup>, SPECweb<sup>®</sup>, SPECjAppServer<sup>®</sup> are registered trademarks of the Standard Performance Evaluation Corporation (SPEC).

All other trademarks referenced herein are the property of their respective owners.

© 2012 by Red Hat, Inc. This material may be distributed only subject to the terms and conditions set forth in the Open Publication License, V1.0 or later (the latest version is presently available at <http://www.opencontent.org/openpub/>).

The information contained herein is subject to change without notice. Red Hat, Inc. shall not be liable for technical or editorial errors or omissions contained herein.

Distribution of modified versions of this document is prohibited without the explicit permission of Red Hat Inc.

Distribution of this work or derivative of this work in any standard (paper) book for commercial purposes is prohibited unless prior permission is obtained from Red Hat Inc.

The GPG fingerprint of the [security@redhat.com](mailto:security@redhat.com) key is:  
CA 20 86 86 2B D6 9D FC 65 F6 EC C4 21 91 80 CD DB 42 A6 0E



# Table of Contents

1. Executive Summary .....	4
1.1 RHEL(KVM) bests VMware ESX wherever head-to-head comparison is possible .....	5
1.2 RHEL was used as the guest OS in all SPECvirt_sc2010 results – regardless of the hypervisor used .....	8
1.3 SPECvirt_sc2010 Results Summary (from <a href="http://www.spec.org/virt_sc2010/results/">http://www.spec.org/virt_sc2010/results/</a> ) .....	9
2.0 SPECvirt_sc2010 Benchmark (from <a href="http://www.spec.org/virt_sc2010/">http://www.spec.org/virt_sc2010/</a> ).....	20
2.1 Workload Design .....	20
2.2 VMs and Tiles.....	22
2.3 Metrics and Submetrics.....	24
2.4 Applications .....	25
2.5 Harness Design .....	25
3.0 Running the Benchmark (One Tile Illustration) .....	27
Appendix A: Fair Use of SPECvirt_sc2010 Results.....	28
Appendix B: References .....	30



# 1. Executive Summary

The non-profit Standard Performance Evaluation Corp. (SPEC) has released **SPECvirt\_sc2010**, the first vendor-neutral benchmark to measure the performance of datacenter servers used for virtualized server consolidation. The new benchmark also includes options for measuring power consumption and power/performance relationships.

SPECvirt\_sc2010 uses a realistic workload and SPEC's proven performance- and power-measurement methodologies to enable vendors, users and researchers to compare system performance across multiple hardware, virtualization platforms, and applications. It was developed by the SPEC virtualization subcommittee, whose members and contributors include AMD, Dell, Fujitsu, HP, IBM, Intel, Oracle, Red Hat, Unisys and VMware.

SPECvirt\_sc2010 is SPEC's first benchmark addressing performance evaluation of datacenter servers used in virtualized server consolidation. SPECvirt\_sc2010 measures the end-to-end performance of all system components including the hardware, virtualization platform, and the virtualized guest operating system and application software. The benchmark supports hardware virtualization, operating system virtualization, and hardware partitioning schemes.

The benchmark utilizes several SPEC workloads representing applications that are common targets of virtualization and server consolidation. SPEC modified each of these standard workloads to match a typical server consolidation scenario of CPU resource requirements, memory, disk I/O, and network utilization for each workload. These workloads are modified versions of SPECweb2005, SPECjAppServer2004, and SPECmail2008. Scaling is achieved by running additional sets of virtual machines (VMs), called "tiles", until overall throughput reaches a peak. All VMs must continue to meet required quality of service (QoS) criteria.



## ***1.1 RHEL(KVM) bests VMware ESX wherever head-to-head comparison is possible***

The industry-leading SPECvirt\_sc2010 scores, shown in Table 1 below, were achieved with Red Hat Enterprise Linux (RHEL) using Kernel-based Virtual Machine (KVM) using features including:

- SR-IOV (Single Root I/O Virtualization) Web, App, and DB VMs for the 10Gbit Ethernet.
- Hugepages
- NUMA (Non-Uniform Memory Access/Architecture)
- Node binding

**RHEL/KVM bests VMware ESX wherever a head-to-head comparison is possible.**

# Cores	Best VMware SPECvirt_sc2010 Score	Best RHEL(KVM) SPECvirt_sc2010 Score
12	1221 @ 78 VMs	1367 @ 84 VMs
16	1570 @ 102 VMs	2388 @ 150 VMs
20	1878 @ 120 VMs	2144 @ 132 VMs
32	2742 @ 168 VMs	N/A
40	3824 @ 234 VMs	4682 @ 288 VMs
64	N/A	5466 @ 336 VMs
80	N/A	8956 @ 552 VMs

*Table 1: VMware versus RHEL(KVM) by # of Cores*



Best SPECvirt\_sc2010 Results by CPU Cores  
(As of April 2012)





Figure 1: Best SPECvirt\_sc2010 Scores by # CPU Cores



# Sockets	Best VMware SPECvirt_sc2010 Score	Best RHEL(KVM) SPECvirt_sc2010 Score
2	1878 @ 120 VMs	2388 @ 150 VMs
4	3824 @ 234 VMs	4682 @ 288 VMs
8	N/A	8956 @ 552 VMs

*Table 2: VMware versus RHEL(KVM) by # of Server Sockets*

**The top RHEL/KVM score beats the top VMware score by a factor of 2.**



*Figure 2: Best SPECvirt\_sc2010 Scores by # Server Sockets*




## ***1.2 RHEL was used as the guest OS in all SPECvirt\_sc2010 results – regardless of the hypervisor used***

As shown in Table 3 below, ALL SPECvirt\_sc2010 results published to date use RHEL as the guest / VM Operating System.

**Figure 3 shows the low virtualization overhead and linear scalability of RHEL(KVM) from the smaller to the largest servers.**

**Clearly, Red Hat Enterprise Linux offers the industry leading platform for virtualization host OS and virtualization guest OS!**



*Figure 3: SPECvirt Tiles per Server Core for RHEL(KVM)*





### 1.3 SPECvirt\_sc2010 Results Summary (from [http://www.spec.org/virt\\_sc2010/results/](http://www.spec.org/virt_sc2010/results/))

Platform	VM Software per Tile	SPECvirt_sc2010 Score		
		# Tiles	#VMs (= 6 * # Tiles)	Score
<b>Hewlett-Packard Company ProLiant DL380 G7</b> Processor: Intel Xeon X5680 (2 chips, <b>12 cores</b> , 6 cores per chip) Memory: 196608 MB (12 x 16 GB 2Rx4 PC3L-10600R at 1333 MHz) <b>VMware Inc. ESX 4.1</b> October 2010	<b>Application Server</b> Red Hat Enterprise Server 5 U5 64-bit (kernel 2.6.18-194.el5) GlassFish Server Open Source Edition 2.1.1 <b>Database Server</b> Red Hat Enterprise Server 5 U5 64-bit (kernel 2.6.18-194.el5) PostgreSQL 8.1.18 <b>Mail Server</b> Red Hat Enterprise Server 5 U5 64-bit (kernel 2.6.18-194.el5) Dovecot 2.0.4 <b>Web Server</b> Red Hat Enterprise Server 5 U5 64-bit (kernel 2.6.18-194.el5) Apache 2.2.16 <b>Infrastructure Server</b> Red Hat Enterprise Server 5 U5 64-bit (kernel 2.6.18-194.el5) Apache 2.2.3 <b>Idle Server</b> Red Hat Enterprise Server 5 U5 64-bit (kernel 2.6.18-194.el5)	13	78	1221 @ 78 VMs



<b>IBM Corporation: IBM BladeCenter HS22V</b> Processor: Intel X5690 (12 cores, 2 chips, 6 cores/chip, 2 threads/core)  Memory: 288 GB VLP DDR-3 (18 x 16GB DIMMs at 800 MHz)  <b>Red Hat Enterprise Linux 6.0 (KVM)</b>  February 2011	<b>Application Server</b> Red Hat Enterprise Linux 6.0 (64-bit) IBM WebSphere Application Server 7.0.0.7  <b>Database Server</b> Red Hat Enterprise Linux 6.0 (64-bit) DB2 v9.7  <b>Mail Server</b> Red Hat Enterprise Linux 6.0 (64-bit) Dovecot 2.0.7  <b>Web Server</b> Red Hat Enterprise Linux 6.0 (64-bit) Apache 2.2.15  <b>Infrastructure Server</b> Red Hat Enterprise Linux 6.0 (64-bit) Apache 2.2.5  <b>Idle Server</b> Red Hat Enterprise Linux 6.0 (64-bit)	14	84	1367 @ 84 VMs
---	--	----	----	---------------



<b>Hewlett-Packard Company ProLiant DL385 G7</b>  Processor: AMD Opteron 6282 SE (2 chips, <b>16 cores</b> , 16 cores per chip)  Memory: 256 GB (16 x 16 GB PC3L-10600R at 1333 MHz)  <b>VMWare Inc. ESXi 5.0</b>  October 2011	<b>Application Server</b> Red Hat Enterprise Server 6 U1 (kernel 2.6.32- 131.0.15.el6.x86_64) GlassFish Server Open Source Edition 3.1  <b>Database Server</b> Red Hat Enterprise Server 6 U1 (kernel 2.6.32- 131.0.15.el6.x86_64) PostgreSQL 9.0.4  <b>Mail Server</b> Red Hat Enterprise Server 6 U1 (kernel 2.6.32- 131.0.15.el6.x86_64) Dovecot 2.0.13  <b>Web Server</b> Red Hat Enterprise Server 6 U1 (kernel 2.6.32- 131.0.15.el6.x86_64) Apache 2.2.19  <b>Infrastructure Server</b> Red Hat Enterprise Server 6 U1 (kernel 2.6.32- 131.0.15.el6.x86_64) Apache 2.2.19  <b>Idle Server</b> Red Hat Enterprise Server 6 U1 (kernel 2.6.32- 131.0.15.el6.x86_64)			
		17	102	1570 @ 120 VMs



<b>Hewlett Packard Company ProLiant DL360p Gen8</b>  Processor: Intel Xeon E5-2690 ( <b>16 cores</b> , 2 chips, 8 cores/chip, 2 threads/core)  Memory: 512 GB (16 x 32 GB PC3L-10600LR at 1333 MHz)  <b>Red Hat Enterprise Linux Server 6.2 (KVM)</b>  April 2012	<b>Application Server</b> Red Hat Enterprise Server 6.2 (kernel 2.6.32-220.el6.x86_64) GlassFish Server Open Source Edition 3.1.1  <b>Database Server</b> Red Hat Enterprise Server 6.2 (kernel 2.6.32-220.el6.x86_64) PostgreSQL 9.0.7  <b>Mail Server</b> Red Hat Enterprise Server 6.2 (kernel 2.6.32-220.el6.x86_64) dovecot 2.0.16  <b>Web Server</b> Red Hat Enterprise Server 6.2 (kernel 2.6.32-220.el6.x86_64) Apache 2.2.15  <b>Infrastructure Server</b> Red Hat Enterprise Server 6.2 (kernel 2.6.32-220.el6.x86_64) Apache 2.2.15  <b>Idle Server</b> Red Hat Enterprise Server 6.2 (kernel 2.6.32-220.el6.x86_64)	25	150	2388 @ 150 VMs
--	---	----	-----	-------------------



<b>Hewlett Packard Company ProLiant BL620c G7</b>  Processor: Intel Xeon E7-2870 ( <b>20 cores</b> , 2 chips, 10 cores/chip, 2 threads/core)  Memory: 512 GB (32 x 16 GB PC3L-10600R at 1066 MHz)  <b>VMware Inc. ESX 4.1 U1</b>  August 2011	<b>Application Server</b> Red Hat Enterprise Server 6 U1 (kernel 2.6.32- 131.0.15.el6.x86_64) GlassFish Server Open Source Edition 3.1  <b>Database Server</b> Red Hat Enterprise Server 6 U1 (kernel 2.6.32- 131.0.15.el6.x86_64) PostgreSQL 9.0.4  <b>Mail Server</b> Red Hat Enterprise Server 6 U1 (kernel 2.6.32- 131.0.15.el6.x86_64) dovecot 2.0.13  <b>Web Server</b> Red Hat Enterprise Server 6 U1 (kernel 2.6.32- 131.0.15.el6.x86_64) Apache 2.2.19  <b>Infrastructure Server</b> Red Hat Enterprise Server 6 U1 (kernel 2.6.32- 131.0.15.el6.x86_64) Apache 2.2.19  <b>Idle Server</b> Red Hat Enterprise Server 6 U1 (kernel 2.6.32- 131.0.15.el6.x86_64)			
		20	120	1878 @ 120 VMs



<b>IBM Corporation: IBM BladeCenter HX5 with MAX5</b> Processor: Intel Xeon E7-4870 ( <b>20 cores</b> , 2 chips, 10 cores/chip, 2 threads/core) Memory: 640 GB VLP DDR-3 (40 x 16GB DIMMs at 1066 Mhz) <b>Red Hat Enterprise Linux 6.0 (KVM)</b> April 2011	<b>Application Server</b> Red Hat Enterprise Linux 6.0 (kernel 2.6.32-71.el6.x86_64) WebShpere Application Server 7.0.0.7 <b>Database Server</b> Red Hat Enterprise Linux 6.0 (kernel 2.6.32-71.el6.x86_64) DB2 v9.7 <b>Mail Server</b> Red Hat Enterprise Linux 6.0 (kernel 2.6.32-71.el6.x86_64) Dovecot 2.0.11 <b>Web Server</b> Red Hat Enterprise Linux 6.0 (kernel 2.6.32-71.el6.x86_64) Apache 2.2.15 <b>Infrastructure Server</b> Red Hat Enterprise Linux 6.0 (kernel 2.6.32-71.el6.x86_64) Apache 2.2.15 <b>Idle Server</b> Red Hat Enterprise Linux 6.0 (kernel 2.6.32-71.el6.x86_64)	22	132	2144 @ 132 VMs
---	--	----	-----	----------------



<b>IBM Corporation:</b> <b>x3850 X5</b> Processor: Intel X7560 ( <b>32 cores</b> , 4 chips, 8 cores/chip, 2 threads/core)  Memory: 1 TB DDR3 (64 x 16 GB Quad Rank x4 PC3-8500 CL7 ECC DDR3 1066Mhz RDIMM)  <b>VMware Inc. ESXi 4.1</b>  February 2911	<b>Application Server</b> Red Hat Enterprise Linux 6.0 (64-bit) IBM WebSphere Application Server 7.0.0.7  <b>Database Server</b> Red Hat Enterprise Linux 6.0 (64-bit) DB2 v9.7  <b>Mail Server</b> Red Hat Enterprise Linux 6.0 (64-bit) Dovecot 2.0.7  <b>Web Server</b> Red Hat Enterprise Linux 6.0 (64-bit) Apache 2.2.15  <b>Infrastructure Server</b> Red Hat Enterprise Linux 6.0 (64-bit) Apache 2.2.5  <b>Idle Server</b> Red Hat Enterprise Linux 6.0 (64-bit)	28	168	2742 @ 168 VMs
---	---	----	-----	-------------------



<b>IBM Corporation: x3850 X5</b>  Processor: Intel E7-8870 ( <b>40 cores</b> , 4 chips, 10 cores/chip, 2 threads/core)  Memory: 1 TB DDR3 (64 x 16 GB Quad Rank x4 PC3-8500 CL7 ECC DDR3 1066Mhz RDIMM)  <b>VMware ESXi 4.1</b>  June 2011	<b>Application Server</b> Red Hat Enterprise Linux 6.0 (64-bit) IBM WebSphere Application Server 7.0.0.7  <b>Database Server</b> Red Hat Enterprise Linux 6.0 (64-bit) DB2 v9.7  <b>Mail Server</b> Red Hat Enterprise Linux 6.0 (64-bit) Dovecot 2.0.11  <b>Web Server</b> Red Hat Enterprise Linux 6.0 (64-bit) Apache 2.2.15  <b>Infrastructure Server</b> Red Hat Enterprise Linux 6.0 (64-bit) Apache 2.2.5  <b>Idle Server</b> Red Hat Enterprise Linux 6.0 (64-bit)	39	234	3824 @ 234 VMs
--	---	----	-----	-------------------





<b>Hewlett Packard Company ProLiant DL580 G7</b>  Processor: Intel Xeon E7-4870 ( <b>40 cores</b> , 4 chips, 10 cores/chip, 2 threads/core)  Memory: 1 TB (64 x 16 GB, Quad Rank x4 PC3- 8500 CL7 ECC DDR3 1066MHz RDIMM)  <b>Red Hat Enterprise Linux Server 6.1 (KVM)</b>  March 2012	<b>Application Server</b> Red Hat Enterprise Linux Server 6.1 (kernel 2.6.32-131.0.15.el6) GlassFish Server Open Source Edition 3.1.1  <b>Database Server</b> Red Hat Enterprise Linux Server 6.1 (kernel 2.6.32-131.0.15.el6) PostgreSQL 9.0.6  <b>Mail Server</b> Red Hat Enterprise Linux Server 6.1 (kernel 2.6.32-131.0.15.el6) dovecot 2.0.9  <b>Web Server</b> Red Hat Enterprise Linux Server 6.1 (kernel 2.6.32-131.0.15.el6) Apache 2.2.15  <b>Infrastructure Server</b> Red Hat Enterprise Linux Server 6.1 (kernel 2.6.32-131.0.15.el6) Apache 2.2.15  <b>Idle Server</b> Red Hat Enterprise Linux Server 6.1 (kernel 2.6.32-131.0.15.el6)	48	288	4682 @ 288 VMs
---	--	----	-----	-------------------



<b>IBM Corporation x3850 X5</b>  Processor: Intel X7560 (8 chips, <b>64 cores</b> , 8 cores per chip)  Memory: 2 TB (128 x 16 GB, Quad Rank x4 PC3-8500 CL7 ECC DDR3 1066MHz LP RDIMM)  <b>Red Hat Enterprise Linux 6.0 (KVM)</b>  November 2010	<b>Application Server</b> <b>Red Hat Enterprise Linux 6.0 (64-bit)</b> WebSphere Application Server 7.0.0.7  <b>Database Server</b> <b>Red Hat Enterprise Linux 6.0 (64-bit)</b> DB2 v9.7  <b>Mail Server</b> <b>Red Hat Enterprise Linux 6.0 (64-bit)</b> Dovecot 2.0.7  <b>Web Server</b> <b>Red Hat Enterprise Linux 6.0 (64-bit)</b> Apache 2.2.15  <b>Infrastructure Server</b> <b>Red Hat Enterprise Linux 6.0 (64-bit)</b> Apache 2.2.15  <b>Idle Server</b> <b>Red Hat Enterprise Linux 6.0 (64-bit)</b>	56	336	5466 @ 336 VMs
--	---	----	-----	----------------



<b>Hewlett Packard Company: ProLiant DL980 G7</b>  Processor: Intel Xeon E7-4870 ( <b>80 cores</b> , 8 chips, 10 cores/chip, 2 threads/core)  Memory: 2 TB (128 x 16 GB, Quad Rank x4 PC3-8500 CL7 ECC DDR3 1066MHz RDIMM)  <b>Red Hat Enterprise Linux Server 6.1 (KVM)</b>  December 2011	<b>Application Server</b> Red Hat Enterprise Linux Server 6.1 (kernel 2.6.32-131.0.15.el6) GlassFish Server Open Source Edition 3.1  <b>Database Server</b> Red Hat Enterprise Linux Server 6.1 (kernel 2.6.32-131.0.15.el6) PostgreSQL 9.0.4  <b>Mail Server</b> Red Hat Enterprise Linux Server 6.1 (kernel 2.6.32-131.0.15.el6) Dovecot 2.0.13  <b>Web Server</b> Red Hat Enterprise Linux Server 6.1 (kernel 2.6.32-131.0.15.el6) Apache 2.2.15  <b>Infrastructure Server</b> Red Hat Enterprise Linux Server 6.1 (kernel 2.6.32-131.0.15.el6) Apache 2.2.15  <b>Idle Server</b> Red Hat Enterprise Linux Server 6.1 (kernel 2.6.32-131.0.15.el6)	92	552	8956 @ 552 VMs
---	--	----	-----	----------------

*Table 3: SPECvirt\_sc2010 Published Results (as of April 2012)*

More information on these benchmark results is available at:  
[http://www.spec.org/virt\\_sc2010/results/](http://www.spec.org/virt_sc2010/results/)



## 2.0 SPECvirt\_sc2010 Benchmark (from [http://www.spec.org/virt\\_sc2010/](http://www.spec.org/virt_sc2010/))

SPECvirt\_sc2010 is designed to be a standard method for measuring a virtualization platform's ability to manage a server consolidation scenario in the datacenter and for comparing performance between virtualized environments. It is intended to measure the performance of the hardware, software, and application layers in a virtualized environment. This includes both hardware and virtualization software and is intended to be run by hardware vendors, virtualization software vendors, application software vendors, academic researchers, and datacenter managers. The benchmark is designed to scale across a wide range of systems and is comprised of a set of component workloads representing common application categories typical of virtualized environments.

Rather than offering a single benchmark workload that attempts to approximate the breadth of consolidated virtualized server characteristics found today, SPECvirt\_sc2010 uses a three-workload benchmark design: a webserver, Java application server, and a mail server workload. The three workloads of which SPECvirt\_sc2010 is composed are derived from SPECweb2005, SPECjAppServer2004, and SPECmail2008. All three workloads drive pre-defined loads against sets of virtualized servers. The SPECvirt\_sc2010 harness running on the client side controls the workloads.

As with all SPEC benchmarks, an extensive set of run rules govern SPECvirt\_sc2010 disclosures to ensure fairness of results. SPECvirt\_sc2010 results are not intended for use in sizing or capacity planning. The benchmark does not address multiple host performance or application virtualization.

### 2.1 Workload Design

The benchmark suite consists of several SPEC workloads that represent applications that industry surveys report to be common targets of virtualization and server consolidation. SPEC modified each of these standard workloads to match a typical server consolidation scenario's resource requirements for CPU, memory, disk I/O, and network utilization for each workload. The SPEC workloads used are:

- **SPECweb2005** - This workload represents a web server, a file server, and an infrastructure server. The SPECweb workload is partitioned into two virtual machines (VMs): a web server and a combined file server and



backend server (BeSim). Specifically, for SPECvirt\_sc2010 the Support workload was modified with modified download file characteristics.

- **SPECjAppserver2004** - This workload represents an application server and backend database server. Specifically, for SPECvirt\_sc2010 the SPECjAppServer2004 was modified such that it creates a dynamic load, the database scale is increased, and the session lengths are decreased.
- **SPECmail2008** - This workload represents a mail server. Specifically, for SPECvirt\_sc2010 the SPECmail IMAP was modified with new transactions.

SPECvirt\_sc2010 created an additional workload called **SPECpoll**. SPECpoll has two purposes: it sends and acknowledges network pings 1) against the idle server in 100% load phase to measure its responsiveness and 2) to all VMs in the 0% load phase (active idle) during power-enabled runs.

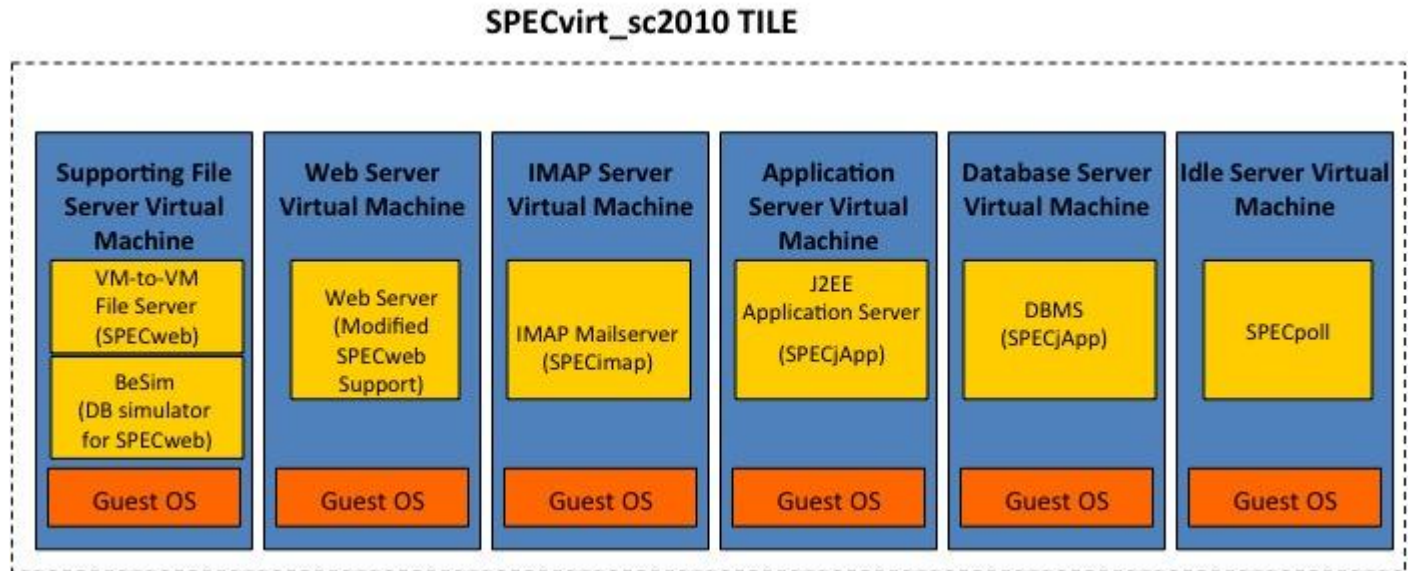
Very lightly-loaded systems are attractive targets when consolidating servers. Even when idle, however, these systems still place resource demands upon the virtualization layer and can impact the performance of other virtual machines.

SPEC researched datacenter workloads and determined suitable load parameters. SPEC refined the test methodology to ensure that the results scale with the capabilities of the system. The benchmark requires significant amounts of memory (RAM), storage, and networking in addition to processors on the SUT. Client systems used for load generation must also be adequately configured to prevent overload. Storage requirements and I/O rates for disk and networks are expected to be non-trivial in all but the smallest configurations. The benchmark does not require that each workload have a maximum number of logical (hardware-wise) processors and is designed to run on a broad range of single host systems.



## 2.2 VMs and Tiles

The benchmark presents an overall workload that achieves the maximum performance of the platform when running one or more sets of Virtual Machines called “tiles.”

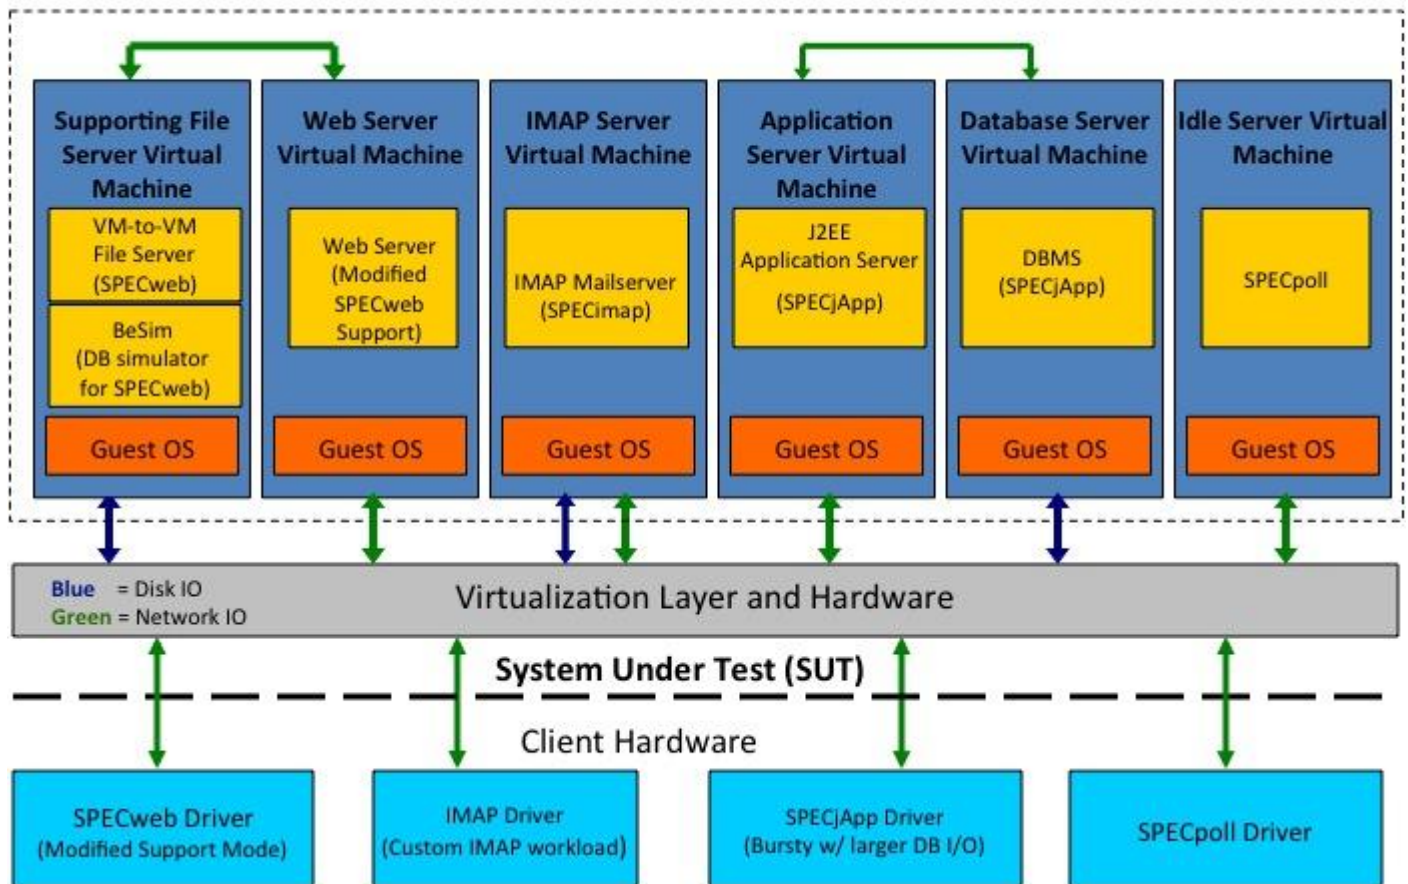


*Figure 2: The Definition to a Tile*

To emulate typical datacenter network use, the webserver and infrastructure server share an internal (private) network connection as do the application server and database server. All VMs use an external (public) network to communicate with each other as well as the clients and controller in the testbed.



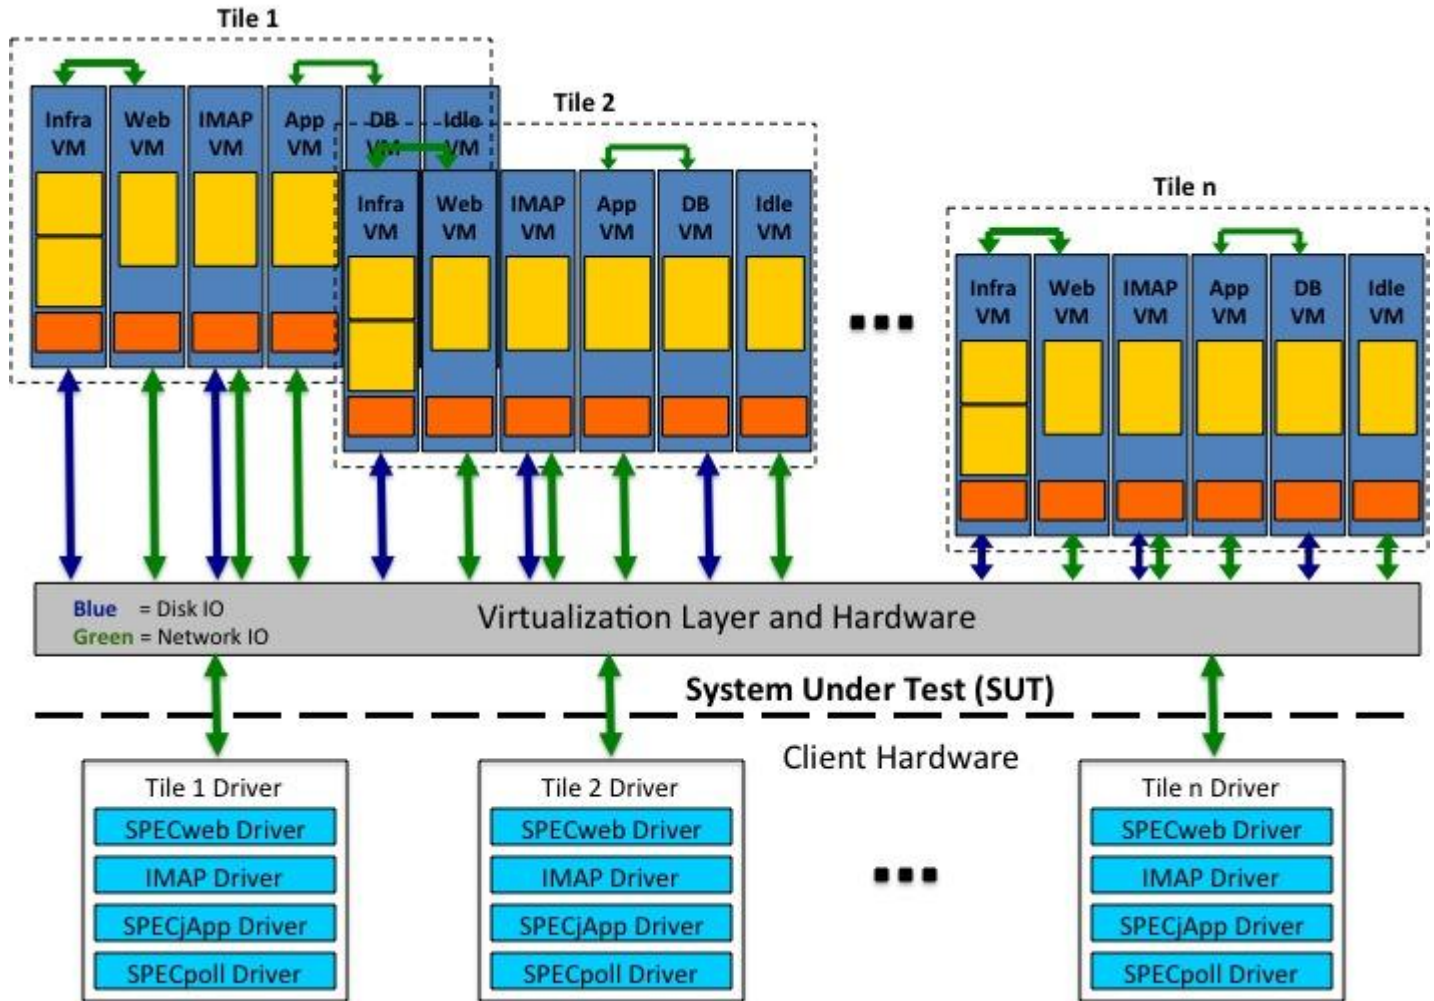
### SPECvirt\_sc2010 TILE



*Figure 3: Interaction between the Tile and Load Driver (Harness)*

Scaling the workload on the SUT consists of running an increasing number of tiles. Peak performance is the point at which the addition of another tile (or fraction) either fails the QoS criteria or fails to improve the overall metric.





*Figure 4: Multi-Tile and Load Driver (Harness) Configuration*

When the SUT does not have sufficient system resources to support the full load of an additional tile, the benchmark offers the use of a fractional load tile. A fractional tile consists of an entire tile with all six VMs but running at a reduced percentage of its full load.

## 2.3 Metrics and Submetrics

The primary metric is the normalized composite of the component submetrics. The benchmark supports three categories of results, each with its own primary metric. Results may be compared only within a given category; however, the benchmarker has the option of submitting results from a given test to one or more categories. The first category is Performance-Only and its metric is SPECvirt\_sc2010 which is expressed as:





SPECvirt\_sc2010 <Overall\_Score> @ <6\*Number\_of\_Tiles> VMs

on the reporting page. The overall score is based upon the following metrics of the three component workloads:

- **Webserver** - requests/second at a given number of simultaneous sessions
- **Mailserver** - the sum of all operations/second at a given number of users
- **JEE Application server** - operations/second (JOPS) at a given injection rate, load factor, and bursty curve (plus additional settings)
- **Idle server** - msec/network ping (not part of the metric calculation)

The overall score is calculated by taking each component workload in each tile and normalizing it against its theoretical maximum for the pre-defined load level. **The three normalized throughput scores for each tile are averaged arithmetically to create a per-tile submetric, and the submetrics for all tiles are added to get the overall performance metric.** The SPECvirt\_sc2010 metric includes reporting this overall metric along with the total number of VMs used (6\* Number\_of \_Tiles).

You can configure one fractional tile to use one-tenth to nine-tenths (at increments of one-tenth) of a tile's normal load level. This allows the benchmarker to saturate the SUT fully and report more granular metrics.

The submetrics must meet the Quality of Service (QoS) criteria adapted from each SPEC standard workload as well as any other validation that the workload requires. The details of the QoS criteria are documented in the Run and Reporting Rules document.

## **2.4 Applications**

The benchmark may use open source or free products as well as commercial products. The benchmark is designed to be open, and the choice of software stack is for the tester to decide. For example, for webserver, any web server software that is HTTP 1.1 compliant can be used. See other sections of this document and the Run and Reporting Rules for more details. Note that variations in implementations may lead to differences in observed performance.

## **2.5 Harness Design**

SPEC developed a test harness driver to coordinate running the component workloads in one or more tiles on the SUT. A command-line-based as well as

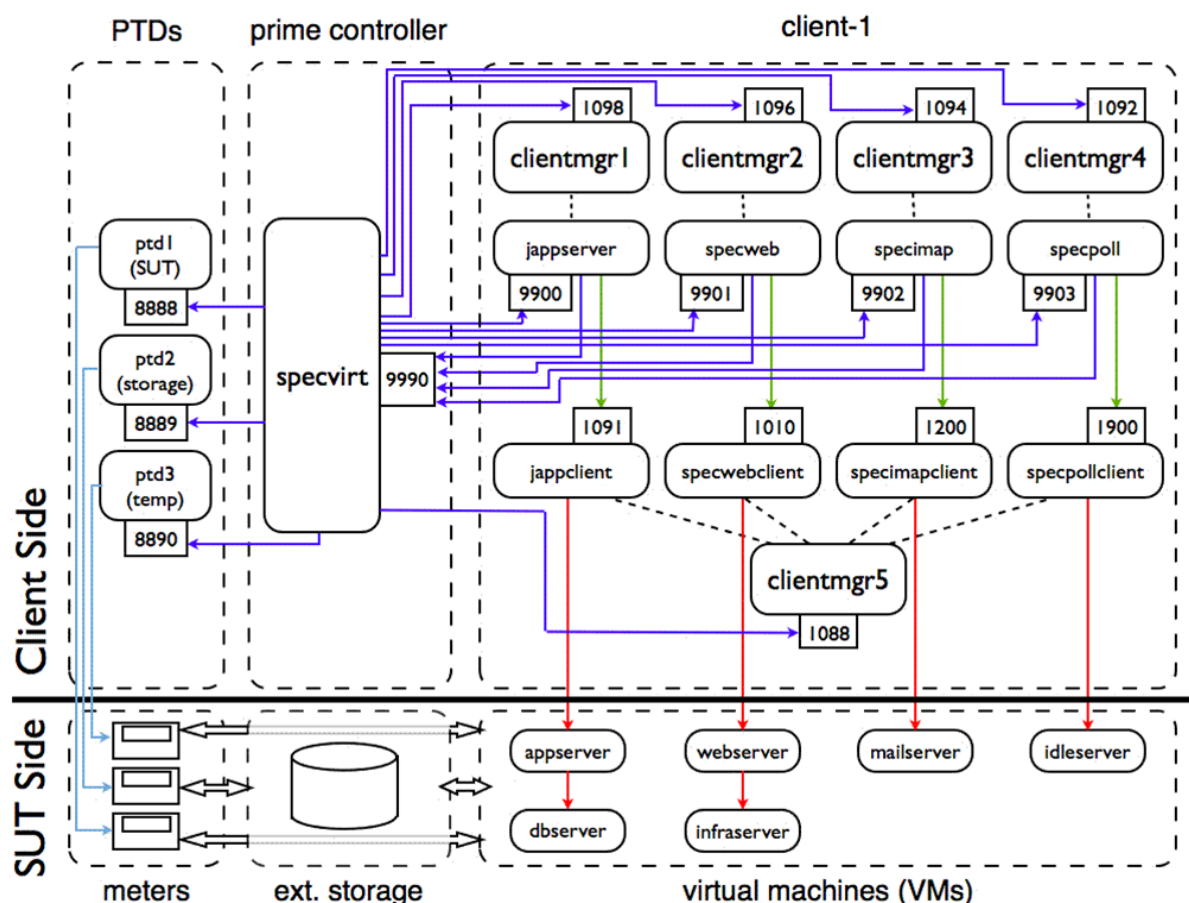


GUI-based front end using Faban allows you to run and monitor the benchmark, collects measurement data as the test runs, post-processes the data at the end of the run, validates the results, and generates the test report.



## 3.0 Running the Benchmark (One Tile Illustration)

This section illustrates the testbed, for a one-tile configuration, including the load driver system (also known as the Client System) which uses a Test Harness to generate and monitor workload requests for the System Under Test (SUT).



*Figure 5: Single-tile, Test-bed Representation*

The above figure represents all components of the test-bed environment and may be a useful visual reference as you work through the remainder of this document. Note, however, that this figure represents a single-tile test configuration. For each additional tile, an additional client box and VM box would be required (and possibly additional external storage, if applicable).

For more details refer to SPECvirt\_sc2010 Client Harness User's Guide at:  
[http://www.spec.org/virt\\_sc2010/docs/SPECvirt\\_Client\\_Harness\\_UserGuide.htm](http://www.spec.org/virt_sc2010/docs/SPECvirt_Client_Harness_UserGuide.htm)  
!



## Appendix A: Fair Use of SPECvirt\_sc2010 Results

Any entity choosing to make public statements using SPECvirt\_sc2010 must follow the OSG Fair Use guidelines [http://www.spec.org/osg/fair\\_use-policy.html](http://www.spec.org/osg/fair_use-policy.html). When public disclosures and competitive comparisons are made using SPECvirt\_sc2010 benchmark results the following benchmark specific rules apply:

1. Results from fully compliant run of the SPECvirt\_sc2010 suite must be used when making competitive comparisons. A fully compliant run consists of a valid run achieved by correctly executing the benchmark workloads on one or more [tiles](#). SPECvirt\_sc2010 supports three categories of results listed below. Cross category comparisons are disallowed, however a submitter has the option of submitting results from the same test run to multiple categories.
  - Performance-only, which produces the SPECvirt\_sc2010 metric,
  - Performance/Power of the Total System Under Test, which produces the SPECvirt\_sc2010\_PPW metric, and
  - Performance/Power of the Server only, which produces the SPECvirt\_sc2010\_ServerPPW metric
2. SPECvirt\_sc2010 uses modified versions of SPECweb2005, SPECjAppServer2004, and SPECmail2008 for its virtualized workloads, as these are established industry-standard workloads. These workloads have been modified to focus on stressing particular aspects of the SUT's resources (CPU, memory, network, disk) typical of server consolidation environments. As such, the modifications are significant enough that comparisons between the the original benchmarks and the versions used in SPECvirt\_sc2010 are not allowed.
3. Primary metrics generated from complete and compliant sets of results for the SPECvirt\_sc2010 workloads are used. For comparisons, if any measured data from the disclosure is used, the primary metrics for the systems being compared must be disclosed in close proximity. "Close proximity" is defined to mean in the same paragraph, in the same font style and size, and either within 100 words or on the same presentation slide.



4. Results must be reviewed and accepted by SPEC prior to public disclosure. The submitter must have a valid SPEC license for this benchmark to submit results. Furthermore, SPEC expects that any public use of results from this benchmark shall follow the [SPEC OSG Fair Use Policy](#) and those specific to this benchmark. In the case where it appears that these guidelines have been violated, SPEC may investigate and request that the offense be corrected or the results resubmitted.
5. Estimates are not allowed.

SPEC expects that the following template be used:

SPEC™, SPECvirt™, and SPECvirt\_sc™ are trademarks of the Standard Performance Evaluation Corp. (SPEC). Competitive numbers shown reflect results published on [www.spec.org](http://www.spec.org) as of <date>. [The comparison presented is based on <basis for comparison>]. For the latest SPECvirt\_sc2010 results visit [http://www.spec.org/osg/virt\\_sc2010](http://www.spec.org/osg/virt_sc2010).

(Note: [...] above required only if selective comparisons are used.)



## Appendix B: References

1. SPECvirt\_sc2010  
[http://www.spec.org/virt\\_sc2010/](http://www.spec.org/virt_sc2010/)
2. SPECvirt\_sc2010 Client Harness User's Guide  
[http://www.spec.org/virt\\_sc2010/docs/SPECvirt\\_Client\\_Harness\\_UserGuide.html](http://www.spec.org/virt_sc2010/docs/SPECvirt_Client_Harness_UserGuide.html)
3. SPECvirt\_sc2010 Run and Reporting Rules  
[http://www.spec.org/virt\\_sc2010/docs/SPECvirt\\_RunRules.html](http://www.spec.org/virt_sc2010/docs/SPECvirt_RunRules.html)
4. SPECvirt\_sc2010 Design Overview  
[http://www.spec.org/virt\\_sc2010/docs/SPECvirt\\_Design\\_Overview.html](http://www.spec.org/virt_sc2010/docs/SPECvirt_Design_Overview.html)
5. SPECvirt\_sc2010 Submitted Results  
[http://www.spec.org/virt\\_sc2010/results/](http://www.spec.org/virt_sc2010/results/)