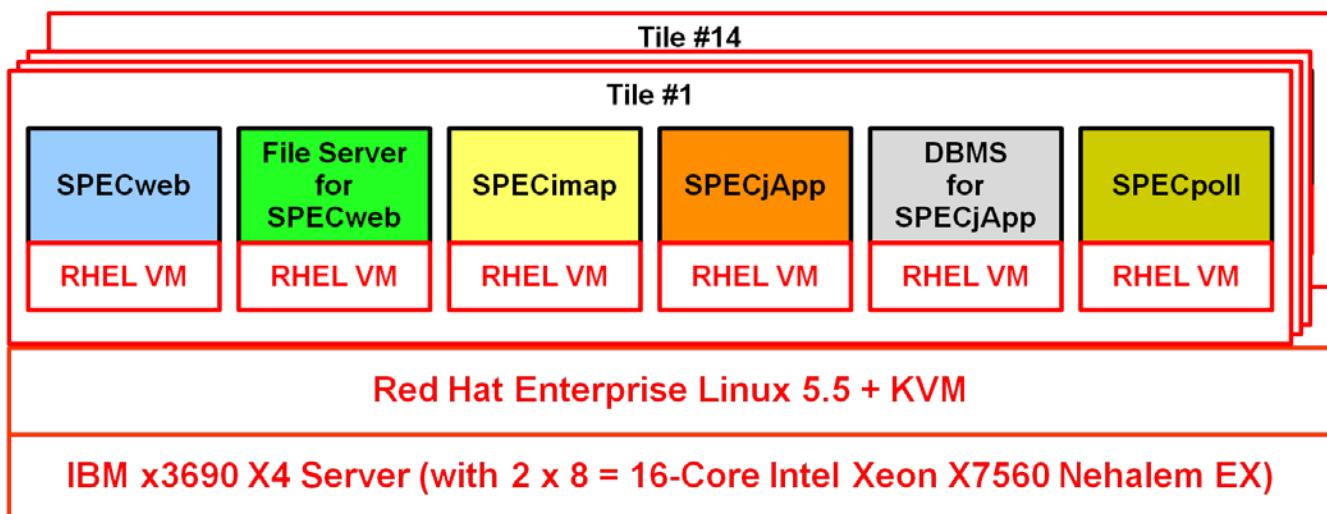




Red Hat Reference Architecture Series

SPECvirt_sc®2010 Benchmark Results using Red Hat® Enterprise Linux® 5.5 and Kernel-based Virtual Machine (KVM)



Version 1.0

November 2010





SPECvirt_sc®2010 Benchmark using Red Hat® Enterprise Linux® 5.5 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® and the benchmark name SPECvirt_sc®, SPECimap®, SPECweb®, SPECjAppServer® are registered trademarks of the Standard Performance Evaluation Corporation (SPEC).

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

© 2010 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 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	5
1.1 RHEL has the best SPECvirt score	5
1.2 All published SPECvirt results used ONLY RHEL guests	6
1.3 SPECvirt Results Summary	7
2.0 SPECvirt_sc2010 Benchmark.....	9
2.1 Workload Design	9
2.2 VMs and Tiles.....	10
2.3 Metrics and Submetrics.....	12
2.4 Applications	13
2.5 Harness Design	13
3.0 Running the Benchmark (One Tile Illustration)	14
4.0 Configuration Used & Tuning – RHEL 5.5 (KVM) on IBMx3650 M3)	15
4.1 Physical Configuration.....	15
4.2 Virtual Configuration.....	17
4.3 Notes, Settings & Tuning.....	20
Physical System Notes.....	20
Storage Notes	20
Virtualization Software Notes.....	20
Web Server VM Notes.....	22
Application Server VM Notes	22
Mail Server VM Notes.....	23
Database Server VM Notes.....	24
Infraserver VM Notes	25
Idle Server VM Notes	26
Client Driver Notes	26
Other Notes	26
5.0 Benchmark Results - RHEL 5.5 (KVM) on IBMx3650 M3)	27
5.1 Performance Summary:.....	27
5.2 Performance Details	28
6. Configuration Used & Tuning - RHEL 5.5 (KVM) on IBMx3690 X5).....	34
6.1 Physical Configuration.....	34
6.2 Virtual Configuration	36
6.3 Notes, Settings & Tuning.....	39
Physical System Notes.....	39
Storage Notes	39
Virtualization Software Notes	39
Web Server VM Notes.....	41
Application Server VM Notes	41
Mail Server VM Notes.....	42
Database Server VM Notes.....	43
Infraserver VM Notes	44
Idle Server VM Notes	45
Client Driver Notes	45



Other Notes	45
7.0 Benchmark Results - RHEL 5.5 (KVM) on IBMx3690 X5).....	46
7.1 Performance Summary.....	46
7.2 Performance Details.....	47
Appendix A: Fair Use of SPECvirt_sc2010 Results.....	54
Appendix B: References	56



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 has the best SPECvirt score

As shown in Figure 1 below, the industry-leading result of **SPECvirt_sc2010 1369 @ 84 VMs** was achieved with **Red Hat Enterprise Linux (RHEL)** using **Kernel-based Virtual Machine (KVM)** using features including the use of:

- SR-IOV (Single Root I/O Virtualization) Web, App, and DB VMs for the 10Gbit Ethernet. (Note that many competitive virtualization technologies do not yet support SR-IOV)
- Hugepages
- NUMA (Non-Uniform Memory Access/Architecture)
- Node binding



Also, in Figure 1 below, the VMware score of 1221 @ 78 SPECvirt_sc2010 (in October 2010) versus the RHEL5 KVM score of virtual 1169 @ 72 SPECvirt_sc2010 (in July 2010) were both run on 2-socket Intel Xeon X5680 Westmere-based servers. However, note that the VMware configuration used 192GB / 1333 MHz of RAM as compared to the RHEL5 KVM configuration with 144GB / 800 MHz of RAM memory in addition to using faster disk drives to achieve a SPECvirt_sc2010 score that was just 4% better. Between July and now, there has been ongoing work on RHEL5 KVM. And RHEL6 is just about to ship where we have made great strides forward in terms of performance and scalability, so expect to see a RHEL6 KVM result in the near future.

1.2 All published SPECvirt results used ONLY RHEL guests

As shown in Table 1 below, ALL SPECvirt_sc2010 results published to date use RHEL as the guest / VM Operating System.

Cleary, Red Hat offers the industry leading platform for virtualization!

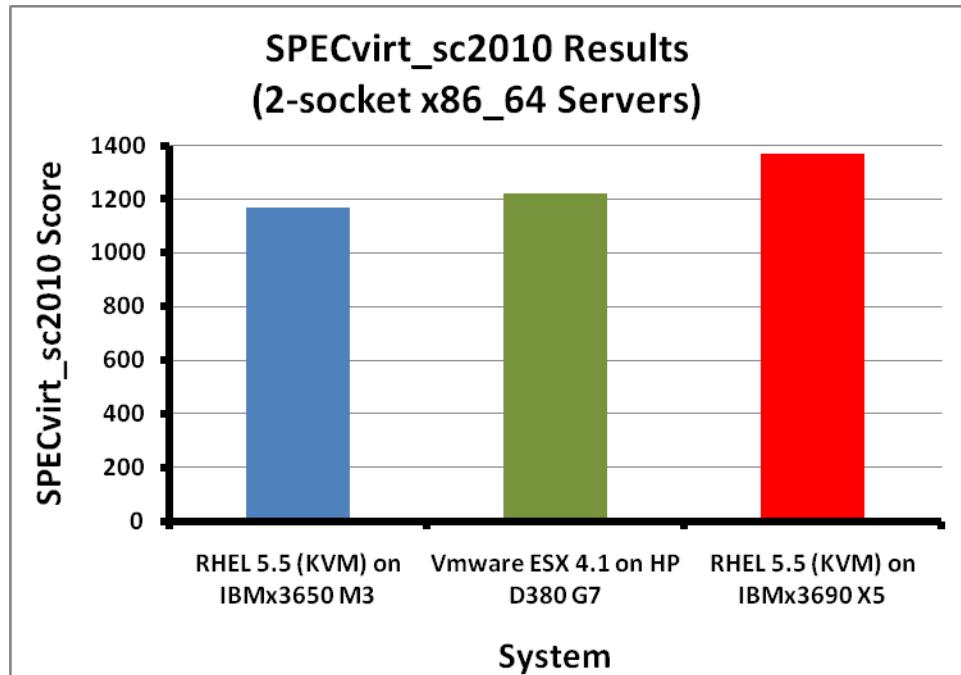


Figure 1: SPECvirt_sc2010 Scores



1.3 SPECvirt Results Summary	VM Software per Tile	SPECvirt_sc2010 Score		
		# Tiles	#VMs (= 6 * # Tiles)	Score
Hewlett-Packard Company ProLiant DL380 G7 Processor: Intel Xeon X5680 (2 chips, 12 cores, 6 cores per chip) Memory: 196608 MB (12 x 16 GB 2Rx4 PC3L-10600R at 1333 MHz) VMware Inc. ESX 4.1	Application Server Red Hat Enterprise Server 5 U5 64-bit (kernel 2.6.18-194.el5) GlassFish Server Open Source Edition 2.1.1 Database Server Red Hat Enterprise Server 5 U5 64-bit (kernel 2.6.18-194.el5) PostgreSQL 8.1.1 Mail Server Red Hat Enterprise Server 5 U5 64-bit (kernel 2.6.18-194.el5) dovecot 2.0.4 Web Server Red Hat Enterprise Server 5 U5 64-bit (kernel 2.6.18-194.el5) Apache 2.2.16 Infrastructure Server Red Hat Enterprise Server 5 U5 64-bit (kernel 2.6.18-194.el5) Idle Server Red Hat Enterprise Server 5 U5 64-bit (kernel 2.6.18-194.el5)	13	78	1221@78
IBM Corporation x3650M3 Processor: Intel X5680 3.33 GHz (2 chips, 12, cores, 6 cores per chip) Memory: 147456 MB SDRAM (18 x 8 GB at 800 Mhz) Red Hat Enterprise Linux 5.5 (KVM)	Application Server Red Hat Enterprise Linux 5.5 (32-bit) WebSphere Application Server 7.0.0.7 Database Server Red Hat Enterprise Linux 5.5 (64-bit) DB2 v9.5 Mail Server Red Hat Enterprise Linux 5.5 (32-bit) Dovecot 1.2.11 Web Server Red Hat Enterprise Linux 5.5	12	72	1169@72



	(64-bit) Apache 2.2.3 Infrastructure Server Red Hat Enterprise Linux 5.5 (32-bit) Apache 2.2.3 Idle Server Red Hat Enterprise Linux 5.5 (32-bit)			
IBM Corporation x3690 X5 Processor: Intel X7560 2.27 GHz (16 cores, 2 chips, 8 cores per chip) Memory: 549755 MB DDR3 (32 x 16 GB at 1066 Mhz) Red Hat: Enterprise Linux 5.5 (KVM)	Application Server Red Hat Enterprise Linux 5.5 (32-bit) WebSphere Application Server 7.0.0.7 Database Server Red Hat Enterprise Linux 5.5 (64-bit) DB2 v9.5 Mail Server Red Hat Enterprise Linux 5.5 (32-bit) Dovecot 1.2.11 Web Server Red Hat Enterprise Linux 5.5 (64-bit) Apache 2.2.3 Infrastructure Server Red Hat Enterprise Linux 5.5 (32-bit) Apache 2.2.3 Idle Server Red Hat Enterprise Linux 5.5 (32-bit)	14	84	1369.23@84

Table 1: SPECvirt_sc2010 Published Results

More information on these benchmark result is available at:
http://www.spec.org/virt_sc2010/results/



2.0 SPECvirt_sc2010 Benchmark

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.”

SPECvirt_sc2010 TILE

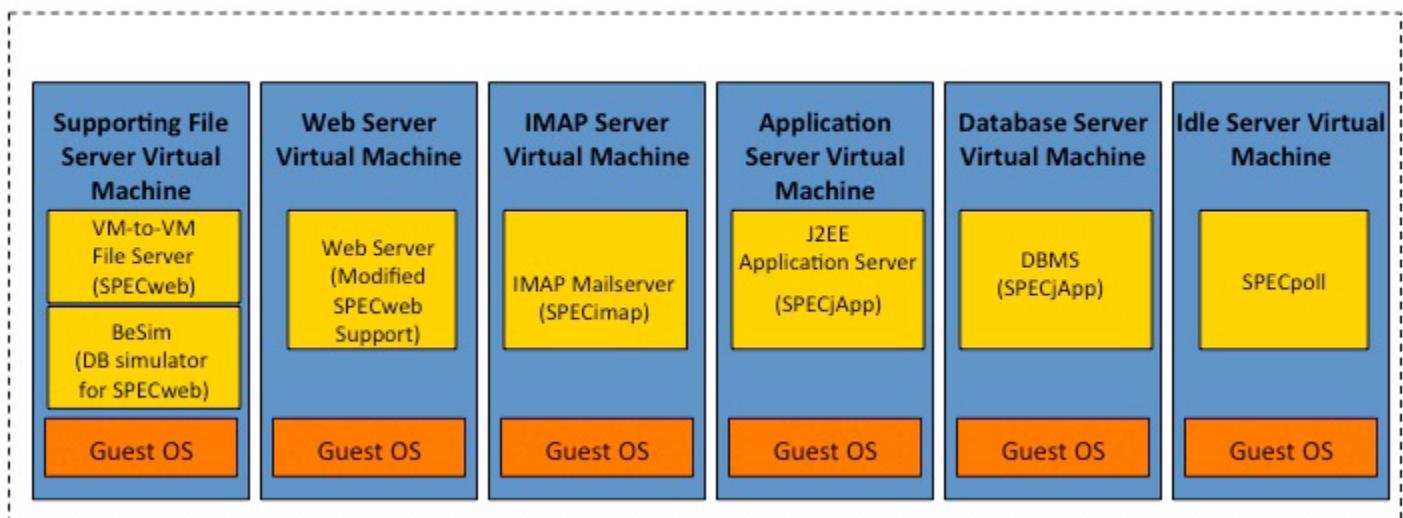


Figure 2: The Definition to a Tile

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

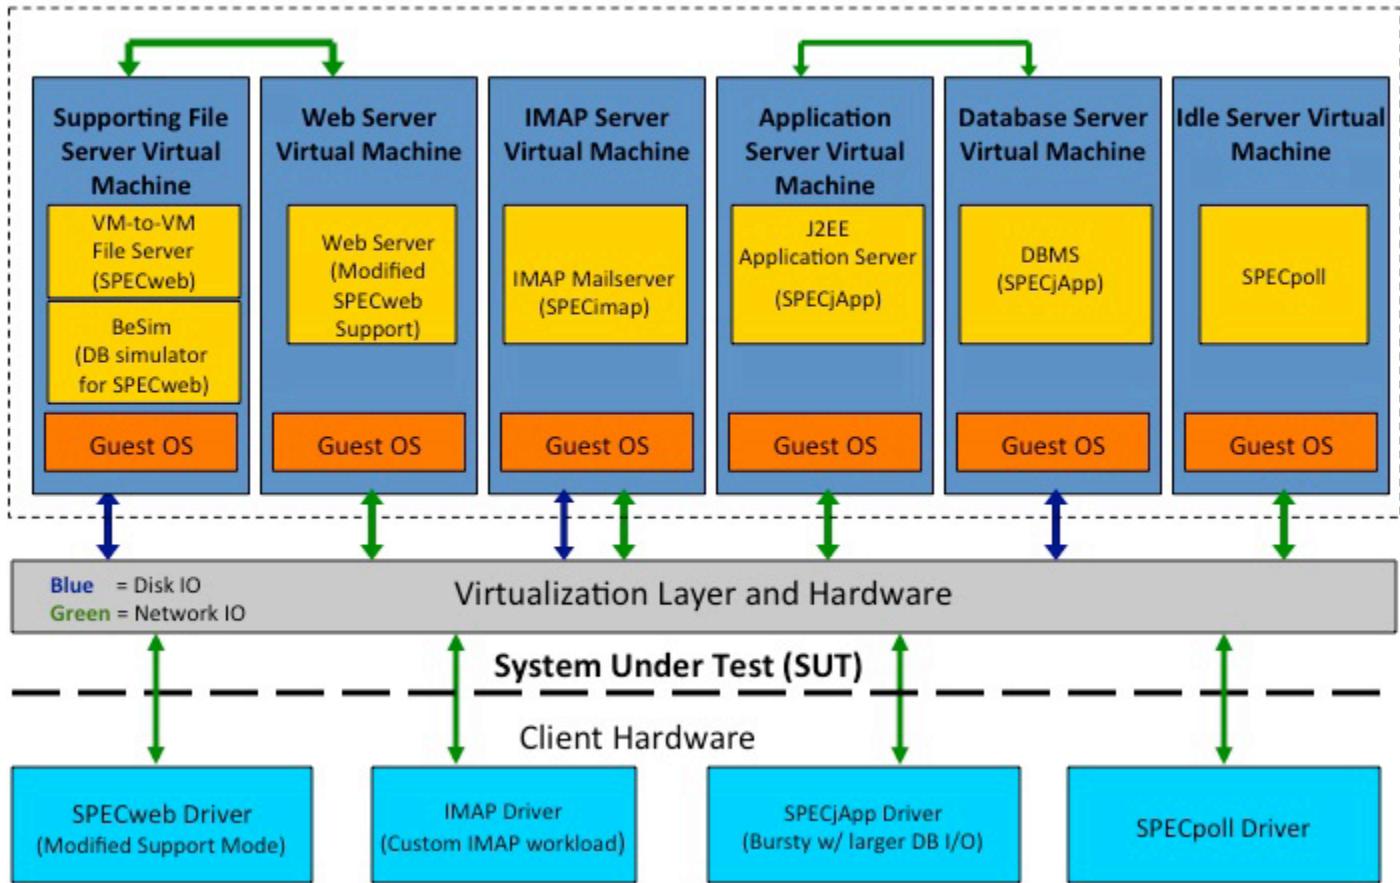


Figure 3: Interaction between the Tile and Load Driver (Harness)

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.

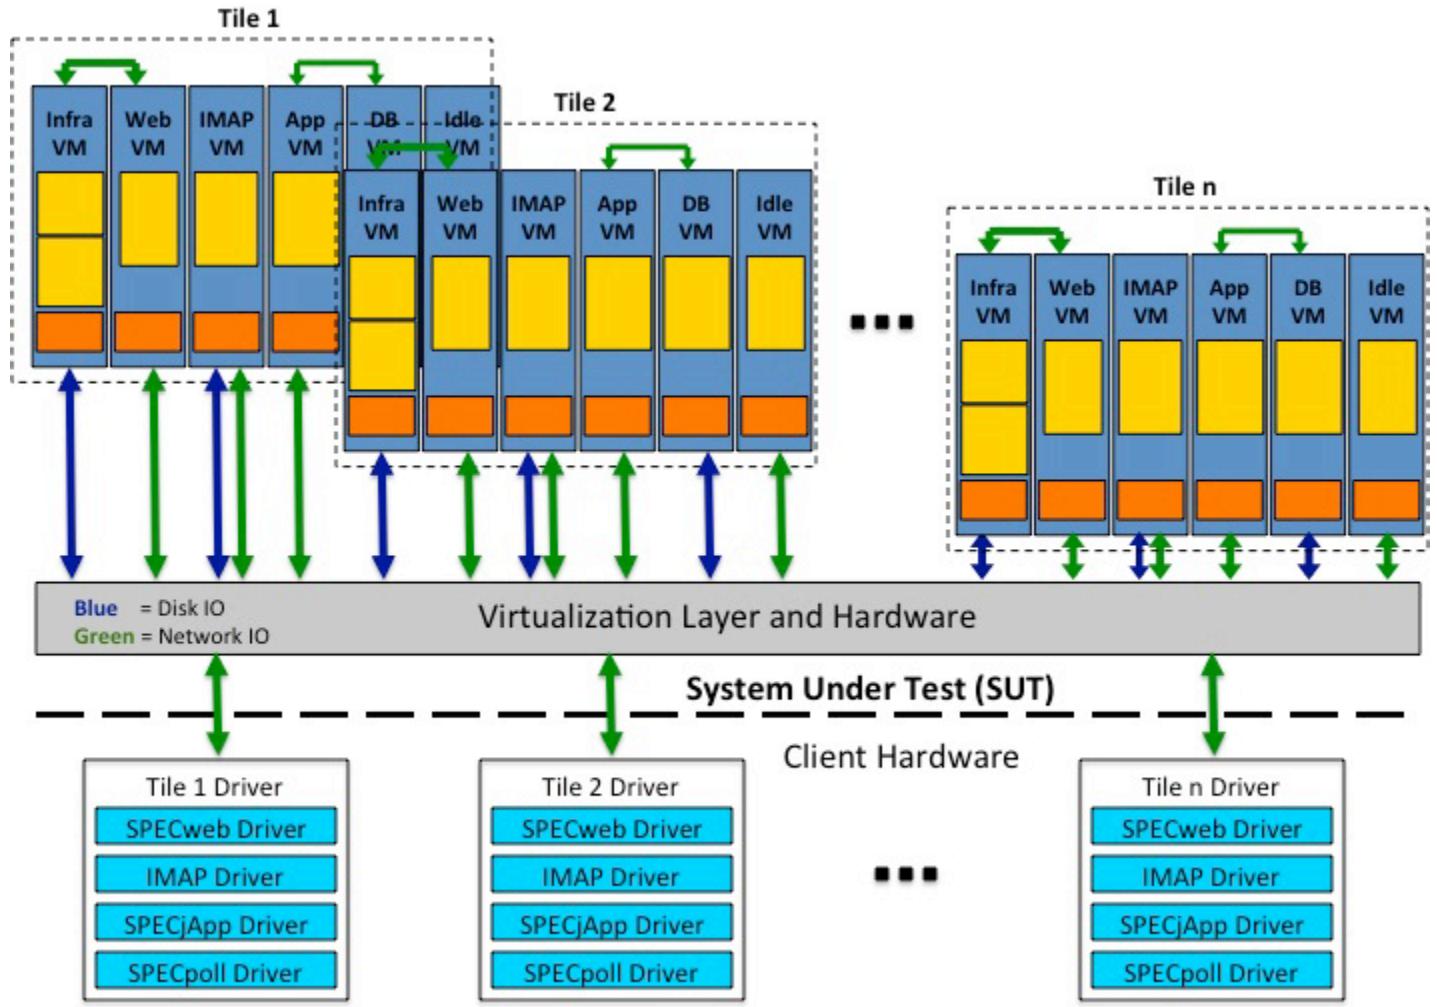


Figure 4: Multi-Tile and Load Driver (Harness) Configuration

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 ($6^* \text{ Number_of_Tiles}$).

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).

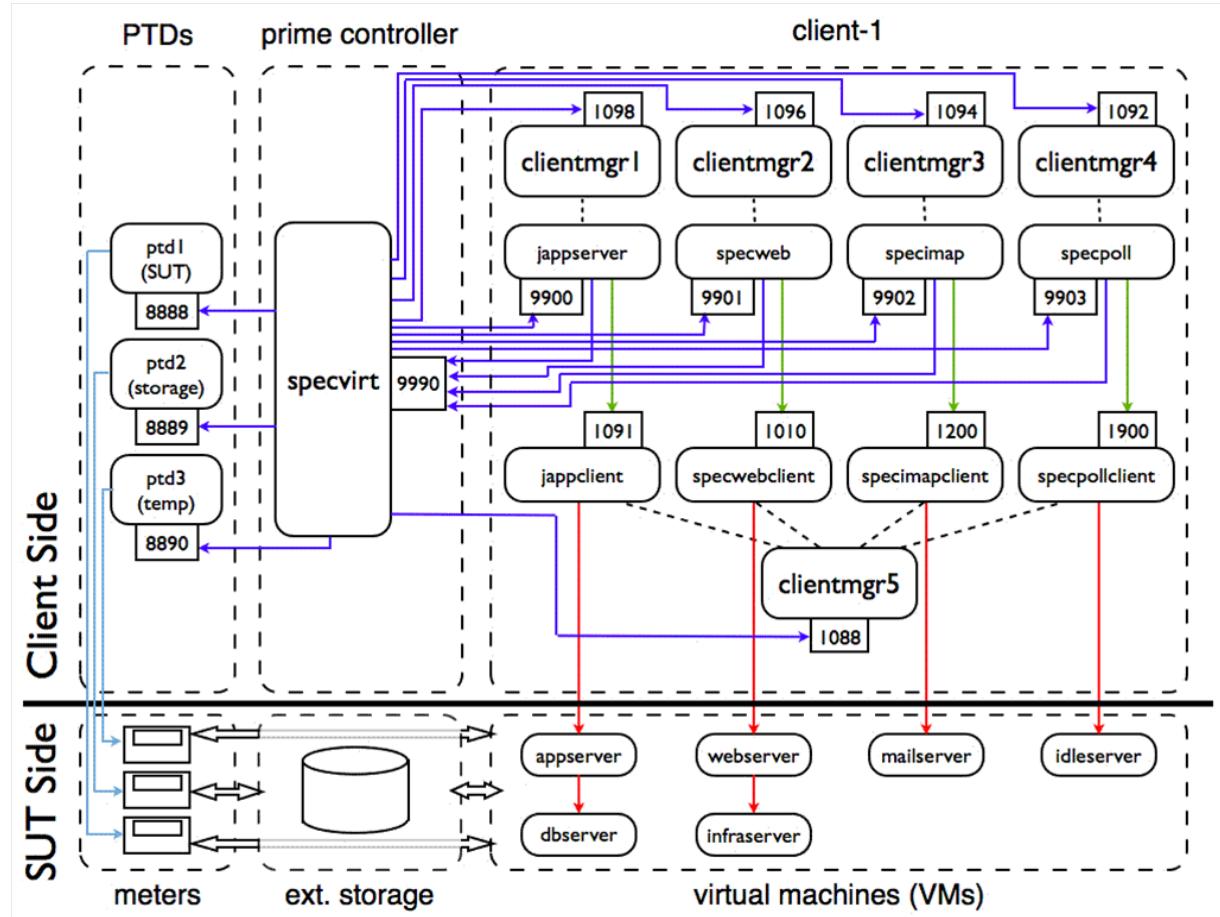


Figure 5: Single-tile, Testbed Representation

The above figure represents all components of the testbed 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.html



4.0 Configuration Used & Tuning – RHEL 5.5 (KVM) on IBMx3650 M3

4.1 Physical Configuration

System Under Test (SUT)	
Virt. Vendor/Product	Red Hat Enterprise Linux 5.5 (KVM)
# of SUTs	1
Server Vendor	IBM Corporation
Server Model	x3650M3
Processor	Intel X5680
Processor Speed (MHz)	3.33 GHz
Processor Cores	12 cores, 2 chips, 6 cores/chip, 2 threads/core
Primary Cache	32KB(I) + 32KB(D) per core
Secondary Cache	256 KB per core
Other Cache	12 MB per chip
Memory	147456 MB SDRAM (18 x 8 GB at 800 Mhz)
Operating System	Red Hat Enterprise Linux 5.5 (2.6.18-194.el5)
File System	ext3
Other Hardware	N/A
Other Software	N/A
SUT Storage	
Storage Controllers	Dual port QLogic Corp. ISP2432-based 4Gb Fibre Channel
Storage Enclosure	2 x IBM DS3400, 6 x IBM DS3000
Disk Description	96 x 73 GB 15k RPM SAS
RAID Level	RAID 5
UPS Required?	No
SUT Network	
Network Controllers	1 x Broadcom NetXtreme II Gigabit Ethernet (dual port, integrated), 3 x Intel Gigabit ET (dual port, PCI-express)
# of Controllers	4 x 2 ports each
# of Networks	8
Network Type	The 8 physical ports are Ethernet, and there are 12 private virtual Ethernet networks, each one for a pair of web & infra VMs. Host access is through the first Ethernet port (eth0).



Network Speed	1 Gb/s
Clients	
Model	x3650M2
# of Clients	5
Processor	Intel E5520
Processor Speed (MHz)	2.27 GHz
# Processors	2
Memory	24 GB SDRAM
Network Controller	Broadcom NetXtreme II (4 ports)
Operating System	Red Hat Enterprise Linux 5.5 (64-bit)
JVM Version	Oracle Java(TM) SE Runtime Environment build 1.6.0_20-b02 (Availability Apr-2010)
Other Hardware	N/A
Other Software	SPECjApp2004 drivers use 1.6.0 IBM J9 2.4 Linux amd64-64 jvmxa6460sr8-20100401_55940 (Availability Mar-2010)
Availability Dates	
SUT Hardware	Apr-2010
Virt. Software	Apr-2010
Other Components	N/A



4.2 Virtual Configuration

VM Configuration Details		Web Server	
# VCPUS	2	Web Server Vendor	Red Hat
VCPU Speed (MHz)	3333 MHz	Web Server Name/Version	Apache 2.2.3
Memory (MB)	5632	Availability Date	Apr-2010
# VNICS	3	Script Vendor	Red Hat
VNIC Description	2 x virt_net, 1 x igbvf	Script Name/Version	PHP 5.1.6
Storage Description	2 x virtio_blk	Script Availability Date	Apr-2010
Virtual Disk Size (MB)	15630	JVM Version	1.6.0 IBM J9 2.4 Linux amd64-64 jvmxa6460sr8-20100401_55940 (Availability Mar-2010)
Datastore Size (MB)	20480	Other Software	Smarty 2.6.26 (Availability Jun-2009), APC 3.0.19 (Availability Aug-2009)
VM OS	Red Hat Enterprise Linux 5.5 (64-bit)	Application Server	
VM OS Availability	Apr-2010	VM Configuration Details	
App Server Configuration Details		VM Configuration Details	
Application Server Vendor	IBM	# VCPUS	2
App. Server Name/Version	WebSphere Application Server 7.0.0.7	VCPU Speed (MHz)	3333 MHz
Availability Date	Jan-2009	Memory (MB)	2560
Emulator Vendor	IBM	# VNICS	2
Emulator Name/Version	WebSphere Application Server 7.0.0.7	VNIC Description	1 x virtio_net, 1 x igbvf
Emulator Availability Date	Jan-2009	Storage Description	2 x virtio_blk
JVM Version	JRE 1.6.0 IBM J9 2.4 Linux x86-32 jvmxi3260sr8-20100401_55940 (Availability Mar-2010)	Virtual Disk Size (MB)	15360
Other Software	N/A	Datastore Size (MB)	20480
VM OS	Red Hat Enterprise Linux 5.5 (32-bit)	VM OS Availability	Apr-2010



Mail Server

VM Configuration Details		Mail Server Configuration Details	
# VCPUS	2	Mail Server Vendor	dovecot.org
VCPU Speed (MHz)	3333 MHz	Mail Server Name/Version	Dovecot 1.2.11
Memory (MB)	512	Availability Date	Mar-2010
# VNICS	1	JVM Version	JRE 1.6.0 IBM J9 2.4 Linux x86-32 jvmxi3260sr8-20100401_55940 (Availability Mar-2010)
VNIC Description	1 x virtio_net	Other Software	N/A
Storage Description	2 x virtio_blk		
Virtual Disk Size (MB)	15360		
Datastore Size (MB)	40960		
VM OS	Red Hat Enterprise Linux 5.5 (32-bit)		
VM OS Availability	Apr-2010		

Database Server

VM Configuration Details		Database Configuration Details	
# VCPUS	2	Database Vendor	IBM
VCPU Speed (MHz)	3333 MHz	Database Name/Version	DB2 v9.5
Memory (MB)	2048	Availability Date	Mar-2010
# VNICS	2	JVM Version	1.6.0 IBM J9 2.4 Linux amd64-64 jvmxa6460sr8-20100401_55940 (Availability Mar-2010)
VNIC Description	1 x virtio_net, 1 x igbvf	Other Software	N/A
Storage Description	2 x virio_blk		
Virtual Disk Size (MB)	15360		
Datastore Size (MB)	20480		
VM OS	Red Hat Enterprise Linux 5.5 (64-bit)		
VM OS Availability	Apr-2010		



Infraserver

VM Configuration Details		Infraserver Configuration Details	
# VCPUS	2	Web Server Vendor	Red Hat
VCPU Speed (MHz)	3333 MHz	Web Server Name/Version	Apache 2.2.3
Memory (MB)	512	Availability Date	Apr-2010
# VNICS	2	Script Vendor	fastcgi.com
VNIC Description	2 x virtio_net	Script Name/Version	FastCGI 2.4.6
Storage Description	2 x virtio_blk	Script Availability Date	Nov-2007
Virtual Disk Size (MB)	15360	JVM Version	JRE 1.6.0 IBM J9 2.4 Linux x86-32 jvmxi3260sr8-20100401_55940 (Availability Mar-2010)
Datastore Size (MB)	61440	Other Software	N/A
VM OS	Red Hat Enterprise Linux 5.5 (32-bit)		
VM OS Availability	Apr-2010		

Idle Server

VM Configuration Details		Idle Server Configuration Details	
# VCPUS	1	JVM Version	JRE 1.6.0 IBM J9 2.4 Linux x86-32 jvmxi3260sr8-20100401_55940 (Availability Mar-2010)
VCPU Speed (MHz)	3333 MHz	Other Software	N/A
Memory (MB)	512		
# VNICS	1		
VNIC Description	1 x virtio_net		
Storage Description	2 x virtio_blk		
Virtual Disk Size (MB)	15360		
Datastore Size (MB)	20480		
VM OS	Red Hat Enterprise Linux 5.5 (32-bit)		
VM OS Availability	Apr-2010		



4.3 Notes, Settings & Tuning

Physical System Notes

- Firmware settings used:
- TurboModeEnable=Enable
- TurboBoost=Traditional
- HardwarePrefetcher=Disable

Storage Notes

- All storage accessed with 2 x 4 Gb Fibre Channel links, through Qlogic 5600 Switch to 2 x DS3400 storage.
- Each DS3400 controller houses 12 73 GB SAS drives and connects to 3 EXP3000 SAS enclosures, each enclosure with 12 73 GB SAS drives.
- Total drive count is 96.
- There are 12 RAID5 arrays, each consisting of 8 drives.
- Each array has one drive defined for one Tile.
- The first array has one extra logical drive defined for the system's boot disk.
- Each drive defined for a Tile is divided into 12 logical volumes (2 per VM) using LVM

Virtualization Software Notes

- The following was run after each boot:
 - echo 69067 >/proc/sys/vm/nr_hugepages
 - mkdir /dev/hugepages
 - mount -t hugetlbfs none /dev/hugepages
 - for i in `seq 1 24`; do echo performance >/sys/devices/system/cpu/cpu\$i/cpufreq/scaling_governor; done
- The following boot options were used:
 - elevator=deadline intel_iommu=on iommu=pt pci=assign-busses



- Network configuration is as follows:
 - Two software bridges are created, br0, br1, and the two Broadcom
 - Ethernet ports are assigned to a bridge (eth0 to br0, eth1 to br1)
 - For every VM, their first virtIO network device is assigned to one of these bridges, even numbered tiles to br0, odd numbered tiles to br1
 - For each tile, there is another bridge created, name t01....tN
 - For all Web and Infra VMs, their second virtIO network device is assigned to a bridge corresponding to the tile number.
 - Kernel module for Intel Gigabit ET Server adapter is loaded with these options to create 36 virtual functions: max_vfs=6,6,6,6,6,6
 - These virtual functions are assigned (one per VM) to the Web, App and DB VMs. When a virtual function is present on a VM, a script is executed to shut down the first virtio_net device and migrate the IP settings to the virtual function device. This is done to facilitate easy IP/dns management when using a mix of VMs with and without virtual functions. The scripts, convert_to_eth[1|2].sh are available in the disclosure archive.
 - On the host, Ethernet devices eth2 - eth7 are the 6 physical functions for the Intel Ethernet ports. These are not used by the host, except to provide physical transport for the virtual functions.
- All VMs are launched with nice utility (-n 19) to adjust their scheduling priority
- The following VMs <VM-type>t<Tile-number> are bound to NUMA node 0 and have a preferred memory allocation node of 0:
 - appt02,appt04,appt06,appt08,appt10,appt12,dbt02,dbt04,dbt06,dbt08,dbt10,dbt12
 - infrat02,infrat04,infrat06,infrat08,infrat10,infrat12,mailt01,mailt02,mailt03
 - mailt04,mailt06,mailt08,mailt10,mailt12,webt02,webt04,webt06,webt08,webt10,webt12
- The following VMs <VM-type>t<Tile-number> are bound to NUMA node 1 and have a preferred memory allocation node of 1:
 - appt01,appt03,appt05,appt07,appt09,appt11,dbt01,dbt03,dbt05,dbt07,dbt09,dbt11
 - infrat01,infrat03,infrat05,infrat07,infrat09,infrat11,mailt05,mailt07,mailt09
 - mailt11,webt01,webt03,webt05,webt07,webt09,webt11
- All VMs' memory but idle VMs are fulfilled with huge (2MB) pages.



Web Server VM Notes

- Disabled the following services via chkconfig:
anacron,atd,bluetooth,cpuspeed,crond,cups,dc_client,dc_server,firstboot,gpm,httpd,ip6tables,iptables,irqbalance,isdn,postgresql,rhnsd,sendmail,smartd,squid,tux,xfs,yum-updatesd
- nofile and nproc set to 102400 in /etc/security/limits.conf
- httpd.conf changes:
 - KeepAlive on
 - MaxKeepAliveRequests 0 (default 100)
 - KeepAliveTimeout 5 (default 15)
 - prefork MPM StartServers=20 (default 8)
 - prefork MPM MinSpareServers=20 (default 5)
 - prefork MPM MaxSpareServers=20 (default 20)
 - prefork MPM ServerLimit=1536 (default 256)
- Boot options: divider=10
 - prefork MPM MaxRequestsPerChild=1000 (default 4000)
 - CustomLog logs/access_log common
- Web VM used NFS to access files from infra VM
- No options were specified in the NFS mount
- File-system /data-disk (ext3) mounted with noatime and data=writeback options

Application Server VM Notes

- Disabled the following services via chkconfig:
anacron,atd,bluetooth,cpuspeed,crond,cups,dc_client,dc_server,firstboot,gpm,httpd,ip6tables,iptables,irqbalance,isdn,postgresql,rhnsd,sendmail,smartd,squid,tux,xfs,yum-updatesd
- Heap size for Application server is 1280 MB, using huge pages
- The following was run after each boot:
 - echo 2818572288 >/proc/sys/kernel/shmmax
 - echo 688128 >/proc/sys/kernel/shmall
- Boot options: divider=10
- File-system /data-disk (ext3) mounted with noatime and data=writeback options



Mail Server VM Notes

- Disabled the following services via chkconfig:
anacron,atd,bluetooth,cupspeed,crond,cups,dc_client,dc_server,firstboot,gpm,httpd,ip6tables,iptables,irqbalance,isdn,postgresql,rhnsd,sendmail,smartd,squid,tux,xfs,yum-updatesd
- dovecot.conf:
 - base_dir = /var/run/dovecot/
 - protocols = imap
 - listen = *:143
 - disable_plaintext_auth = no
 - log_path = /var/log/dovecot
 - ssl = no
 - login_processes_count = 10
 - mail_location = mbox:/mailstore/dovecot/mailldir/%u
 - mail_debug = yes
 - maildir_copy_with_hardlinks = yes
 - mbox_rotate_size = 2048
 - mbox_rotate_min_size = 16
 - mbox_rotate_days = 0
 - protocol lda {
◦ postmaster_address = postmaster@example.com
◦ }
 - auth default {
◦ mechanisms = plain
◦ passdb passwd-file {
◦ args = /etc/passwd.dovecot
◦ }
◦ userdb passwd {
◦ }



user = root

```
}
```

```
dict {
```

```
}
```

```
plugin {
```

```
}
```

- Boot options: divider=10
- File-system /data-disk (ext3) mounted with noatime and data=writeback options

Database Server VM Notes

- Disabled the following services via chkconfig:
anacron,atd,bluetooth,cpuspeed,crond,cups,dc_client,dc_server,firstboot,gpm,httpd,ip6tables,iptables,irqbalance,isdn,postgresql,rhnsd,sendmail,smartd,squid,tux,xfs,yum-updatesd
- Database tuning options:
 - db2 update db cfg for speedb using DATABASE_MEMORY 200000
 - db2set DB2_LARGE_PAGE_MEM=DB; db2set DB2_PINNED_BP=ON
 - alter bufferpool IBMDEFAULTTBP size 160000
 - alter tablespace roll_space NO FILE SYSTEM CACHING
 - alter tablespace C_cust_space NO FILE SYSTEM CACHING
 - alter tablespace C_supp_space NO FILE SYSTEM CACHING
 - alter tablespace C_site_space NO FILE SYSTEM CACHING
 - alter tablespace C_parts_space NO FILE SYSTEM CACHING
 - alter tablespace C_custinv_space NO FILE SYSTEM CACHING
 - alter tablespace M_parts_space NO FILE SYSTEM CACHING
 - alter tablespace M_bom_space NO FILE SYSTEM CACHING
 - alter tablespace M_wo_space NO FILE SYSTEM CACHING
 - alter tablespace M_lo_space NO FILE SYSTEM CACHING
 - alter tablespace M_inv_space NO FILE SYSTEM CACHING
 - alter tablespace O_cust_space NO FILE SYSTEM CACHING



- alter tablespace O_ords_space NO FILE SYSTEM CACHING
- alter tablespace O_ordl_space NO FILE SYSTEM CACHING
- alter tablespace O_item_space NO FILE SYSTEM CACHING
- alter tablespace S_comp_space NO FILE SYSTEM CACHING
- alter tablespace S_sc_space NO FILE SYSTEM CACHING
- alter tablespace S_po_space NO FILE SYSTEM CACHING
- alter tablespace S_site_space NO FILE SYSTEM CACHING
- alter tablespace S_supp_space NO FILE SYSTEM CACHING
- alter tablespace U_space NO FILE SYSTEM CACHING
- update db cfg for specdb using logprimary 64
- reorgchk update statistics on table all
- Boot options: divider=10 hugepages=416
- File-system /data-disk (ext3) mounted with noatime and data=writeback options

Infraserver VM Notes

- Disabled the following services via chkconfig:
anacron,atd,bluetooth,cpuspeed,crond,cups,dc_client,dc_server,firstboot,gpm,httpd,ip6tables,iptables,irqbalance,isdn,postgresql,rhnsd,sendmail,smartd,squid,tux,xfs,yum-updatesd
- nofile and nproc set to 102400 in /etc/security/limits.conf
- httpd.conf changes:
 - Listen port to 81 (default 80)
 - prefork MPM StartServers=100 (default 8)
 - prefork MPM MinSpareServers=100 (default 5)
 - prefork MPM MaxSpareServers=100 (default 20)
 - prefork MPM ServerLimit=1024 (default 256)
 - prefork MPM MaxClients=1024 (default 256)
 - prefork MPM MaxRequestsPerChild=16000 (default 4000)
 - LoadModule fastcgi_module modules/mod_fastcgi.so
 - ServerName besim:80



- CustomLog logs/access_log common
- FastCgiServer /var/www/fcgi-bin/besim_fcgi.fcgi -processes 50
- Boot options: divider=10
- Standard Linux NFS server is enabled to serve files to Web VM
- /etc/exports file contains:
 - /support-download *(rw,no_subtree_check,crossmnt,no_root_squash)
- File-system /data-disk (ext3) mounted with noatime and data=writeback options

Idle Server VM Notes

- Disabled the following services via chkconfig:
anacron,atd,bluetooth,cpuspeed,crond,cups,dc_client,dc_server,firstboot,gpm,httpd,ip6tables,iptables,irqbalance,isdn,postgresql,rhnsd,sendmail,smartd,squid,tux,xfs,yum-updatesd
- Boot options: divider=10
- File-system /data-disk (ext3) mounted with noatime and data=writeback options

Client Driver Notes

- 1 client system is dedicated for prime controller
- 4 remaining clients have 3 VMs each. Each VM is running the prime clients for 1 tile.
- Paravirtual clock is used for client VMs.
- Each VM is assigned 2 vCPUs, 6 GB memory, 1 virtio_net device, and 1 virtio_blk device.
- There are 4 Broadcom 1Gbps Ethernet ports on each physical client. The first is dedicated for host access. The remaining 3 each have a unique software bridge assigned (br1-br3), and 1 VM is assigned per bridge.
- Red Hat Enterprise Linux 5.5 (64-bit, KVM) is used for hosting VMs. All VMs use Red Hat Enterprise Linux 5.5 (64-bit). Each VM uses divider=10 boot option.

Other Notes

- A single Netgear switch, model GSM7352S, is used for all physical Ethernet connections, and a Qlogic 5600 switch for all fibre channel connections.
-



5.0 Benchmark Results - RHEL 5.5 (KVM) on IBMx3650 M3)

5.1 Performance Summary:

Performance							
Tile #	Pct Load	Application Server	Web Server	Mail Server	Idle Server	Per-Tile Score	Overall Score
1	100%	33.47	53.61	88.26	N/A	97.71	1169
2	100%	33.45	52.81	88.08	N/A	97.13	
3	100%	33.40	53.64	88.22	N/A	97.64	
4	100%	33.19	52.78	88.17	N/A	96.90	
5	100%	33.26	53.72	88.22	N/A	97.56	
6	100%	33.41	52.89	88.18	N/A	97.18	
7	100%	33.41	53.65	88.31	N/A	97.69	
8	100%	33.43	53.04	88.24	N/A	97.31	
9	100%	33.46	53.66	88.39	N/A	97.78	
10	100%	33.42	52.89	88.12	N/A	97.16	
11	100%	33.47	53.77	88.08	N/A	97.74	
12	100%	33.44	52.89	88.18	N/A	97.21	

Quality Of Service (QOS)							
Tile #	Pct Load	Application Server	Web Server	Mail Server	Idle Server	Per-Tile Score	Overall Score
1	100%	1.00	1.00	1.00	1.00	1.00	99.54%
2	100%	0.97	0.99	1.00	1.00	0.99	
3	100%	0.99	1.00	1.00	1.00	1.00	
4	100%	0.96	0.99	1.00	1.00	0.99	
5	100%	1.00	1.00	1.00	1.00	1.00	
6	100%	0.98	0.99	1.00	1.00	0.99	
7	100%	0.99	1.00	1.00	1.00	1.00	
8	100%	0.98	0.99	1.00	1.00	0.99	
9	100%	1.00	1.00	1.00	1.00	1.00	
10	100%	0.99	0.99	1.00	1.00	1.00	
11	100%	0.99	1.00	1.00	1.00	1.00	
12	100%	0.99	0.99	1.00	1.00	1.00	



5.2 Performance Details

Tile 1					
Application Server					
Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.35	1.26	7.86	1.75	5
Dealer	20.12	0.14/0.29/0.13	6.62/7.42/6.69	0.40/0.70/0.30	2/2/2
Web Server					
Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.61	384660	731	577	0
Mail Server					
Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.27	102733	102733	0	100.00
Fetch	35.90	258450	258450	0	100.00
Idle Server					
Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec	
Heartbeats	719	7.51	1	901	
Tile 2					
Application Server					
Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.35	1.47	9.18	2.25	5
Dealer	20.10	0.34/0.58/0.35	15.08/26.81/13.36	0.60/1.30/0.50	2/2/2
Web Server					
Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	52.81	374655	4527	1053	0
Mail Server					
Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.18	102089	102089	0	100.00
Fetch	36.08	259784	259784	0	100.00
Idle Server					
Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec	



Heartbeats	719	7.76	1	124
------------	-----	------	---	-----

Tile 3

Application Server

Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.31	1.30	15.34	1.75	5
Dealer	20.09	0.17/0.36/0.15	20.67/22.06/13.47	0.30/0.70/0.30	2/2/2

Web Server

Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.64	385652	480	96	0

Mail Server

Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.12	101670	101670	0	100.00
Fetch	36.00	259182	259182	0	100.00

Idle Server

Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec
Heartbeats	719	6.39	0	234

Tile 4

Application Server

Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.36	1.58	14.40	2.50	5
Dealer	19.83	0.44/0.67/0.47	15.29/27.26/16.69	0.70/1.50/0.60	2/2/2

Web Server

Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	52.78	374540	3642	1820	0

Mail Server

Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.23	102471	102471	0	100.00
Fetch	35.94	258785	258785	0	100.00

Idle Server

Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec
Heartbeats	719	6.43	1	46

Tile 5



Application Server					
Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.27	1.31	14.49	1.75	5
Dealer	19.99	0.15/0.30/0.15	13.13/14.16/13.42	0.30/0.70/0.30	2/2/2
Web Server					
Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.72	386042	463	256	0
Mail Server					
Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.24	102525	102525	0	100.00
Fetch	36.08	259743	259743	0	100.00
Idle Server					
Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec	
Heartbeats	719	6.40	1	90	
Tile 6					
Application Server					
Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.38	1.53	14.86	2.25	5
Dealer	20.03	0.30/0.56/0.28	13.93/15.81/14.05	0.60/1.30/0.50	2/2/2
Web Server					
Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	52.89	376033	3740	1046	0
Mail Server					
Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.13	101753	101753	0	100.00
Fetch	35.99	259113	259113	0	100.00
Idle Server					
Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec	
Heartbeats	719	7.75	1	289	
Tile 7					
Application Server					
Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%



Manufacturing	13.32	1.34	15.34	1.75	5
Dealer	20.09	0.17/0.32/0.17	13.54/14.63/17.17	0.30/0.70/0.30	2/2/2

Web Server

Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.65	385586	540	153	0

Mail Server

Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.09	101424	101424	0	100.00
Fetch	36.10	259901	259901	0	100.00

Idle Server

Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec
Heartbeats	719	7.06	1	144

Tile 8

Application Server

Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.36	1.49	15.69	2.00	5
Dealer	20.07	0.29/0.54/0.28	14.69/23.09/14.49	0.50/1.20/0.50	2/2/2

Web Server

Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.04	377474	3354	1081	0

Mail Server

Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.13	101736	101736	0	100.00
Fetch	36.08	259795	259795	0	100.00

Idle Server

Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec
Heartbeats	719	7.07	1	357

Tile 9

Application Server

Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.33	1.23	4.89	1.50	5
Dealer	20.13	0.13/0.25/0.12	3.66/6.71/2.99	0.30/0.60/0.30	2/2/2



Web Server					
Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.66	385505	528	328	0
Mail Server					
Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.24	102532	102532	0	100.00
Fetch	36.04	259495	259495	0	100.00
Idle Server					
Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec	
Heartbeats	719	8.03	1	978	
Tile 10					
Application Server					
Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.34	1.42	15.28	2.00	5
Dealer	20.08	0.25/0.49/0.25	14.36/15.02/13.54	0.50/1.10/0.40	2/2/2
Web Server					
Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	52.89	376921	3283	615	0
Mail Server					
Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.30	102964	102964	0	100.00
Fetch	35.98	259026	259026	0	100.00
Idle Server					
Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec	
Heartbeats	719	7.69	1	869	
Tile 11					
Application Server					
Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.36	1.31	9.97	1.75	5
Dealer	20.11	0.16/0.30/0.17	12.76/18.26/12.17	0.30/0.70/0.30	2/2/2
Web Server					
Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors



Support	53.77	386585	405	150	0
Mail Server					
Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.14	101799	101799	0	100.00
Fetch	36.11	260024	260024	0	100.00
Idle Server					
Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec	
Heartbeats	719	6.31	1	85	
Tile 12					
Application Server					
Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.36	1.39	9.88	1.75	5
Dealer	20.08	0.24/0.47/0.22	8.49/17.48/7.81	0.50/1.10/0.40	2/2/2
Web Server					
Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	52.89	376175	3676	963	0
Mail Server					
Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.22	102359	102359	0	100.00
Fetch	36.03	259421	259421	0	100.00
Idle Server					
Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec	
Heartbeats	719	5.62	1	150	



6. Configuration Used & Tuning - RHEL 5.5 (KVM) on IBMx3690 X5

6.1 Physical Configuration

System Under Test (SUT)	
Virt. Vendor/Product	Red Hat Enterprise Linux 5.5 (KVM)
# of SUTs	1
Server Vendor	IBM Corporation
Server Model	x3690 X5
Processor	Intel X7560
Processor Speed (MHz)	2.27 GHz
Processor Cores	16 cores, 2 chips, 8 cores/chip, 2 threads/core
Primary Cache	32KB(I) + 32KB(D) per core
Secondary Cache	256 KB per core
Other Cache	24 MB per chip
Memory	549755 MB DDR3 (32 x 16 GB at 1066 Mhz)
Operating System	Red Hat Enterprise Linux 5.5 version 2.6.18-194.el5
File System	ext3
Other Hardware	N/A
Other Software	N/A

SUT Storage	
Storage Controllers	Dual port QLogic Corp. QLE2562 8Gb Fibre Channel IBM ServeRAID MR1015
Storage Enclosure	8 x IBM EXP810 None
Disk Description	112 x 73 GB 15k RPM SAS 2 x 73 GB 15k RPM SAS
RAID Level	RAID 5 RAID 1
UPS Required?	No No



SUT Network	
Network Adapters	1 x Broadcom NetXtreme II Gigabit Ethernet (dual port, integrated) 3 x Intel Gigabit ET (dual port, PCI-express)
SUT Ports Total	8
SUT Ports Used	8
Network Type	Ethernet Ethernet
Network Speed	1 Gb/s 1 Gb/s
Clients	
Model	x3550M2
# of Clients	14
Processor	Intel E5520
Processor Speed (MHz)	2.27 GHz
# Processors	2
Memory	32 GB SDRAM
Network Controller	Broadcom NetXtreme II
Operating System	Red Hat Enterprise Linux 5.5 (64-bit)
JVM Version	Java(TM) SE Runtime Environment (build 1.6.0_20-b02))
Other Hardware	N/A
Other Software	N/A
Availability Dates	
SUT Hardware	Aug-2010
Virt. Software	Apr-2010
Other Components	N/A



6.2 Virtual Configuration

Web Server	
VM Configuration Details	
# VCPUS	2
VCPU Speed (MHz)	2261 MHz
Memory (MB)	26000
# VNICS	3
VNIC Description	2 x virt_net, 1 x igbvf
Storage Description	2 x virtio_blk
Virtual Disk Size (MB)	15630
Datastore Size (MB)	20480
VM OS	Red Hat Enterprise Linux 5.5 (64-bit)
VM OS Availability	April-2010
Web Server Configuration Details	
Web Server Vendor	Red Hat
Web Server Name/Version	Apache 2.2.3
Availability Date	April-2010
Script Vendor	PHP
Script Name/Version	PHP 5.1.6
Script Availability Date	June-1999
JVM Version	1.6.0 IBM J9 2.4 Linux amd64-64 jvmxa6460sr8-20100401_55940
Other Software	APC 3.0.19 Smarty 2.6.26

Application Server	
VM Configuration Details	
# VCPUS	2
VCPU Speed (MHz)	2261 MHz
Memory (MB)	2560
# VNICS	2
VNIC Description	1 x virtio_net, 1 x igbvf
Storage Description	2 x virtio_blk
Virtual Disk Size (MB)	15360
Datastore Size (MB)	20480
VM OS	Red Hat Enterprise Linux 5.5 (32-bit)
VM OS Availability	April-2010
App Server Configuration Details	
Application Server Vendor	IBM
App. Server Name/Version	WebSphere Application Server 7.0.0.7
Availability Date	January-2009
Emulator Vendor	IBM
Emulator Name/Version	WebSphere Application Server 7.0.0.7
Emulator Availability Date	January-2009
JVM Description	vSonic JumpinJVM 1.2.3
JVM Availability	(current)
Other Software	N/A



Mail Server

VM Configuration Details		Mail Server Configuration Details	
# VCPUS	2	Mail Server Vendor	dovecot.org
VCPUs Speed (MHz)	2261 MHz	Mail Server Name/Version	Dovecot 1.2.11
Memory (MB)	512	Availability Date	March-2010
# VNICS	1	JVM Version	1.6.0 IBM J9 2.4 Linux amd64-64 jvmxa6460sr8-20100401_55940
VNIC Description	1 x virtio_net	Other Software	N/A
Storage Description	2 x virtio_blk		
Virtual Disk Size (MB)	15360		
Datastore Size (MB)	40960		
VM OS	Red Hat Enterprise Linux 5.5 (32-bit)		
VM OS Availability	April-2010		

Database Server

VM Configuration Details		Database Configuration Details	
# VCPUS	2	Database Vendor	IBM
VCPUs Speed (MHz)	2261 MHz	Database Name/Version	DB2 v9.5
Memory (MB)	2048	Availability Date	March-2010
# VNICS	2	JVM Version	1.6.0 IBM J9 2.4 Linux amd64-64 jvmxa6460sr8-20100401_55940
VNIC Description	1 x virtio_net, 1 x igbvf	Other Software	N/A
Storage Description	2 x virio_blk		
Virtual Disk Size (MB)	15360		
Datastore Size (MB)	20480		
VM OS	Red Hat Enterprise Linux 5.5 (64-bit)		
VM OS Availability	April-2010		

Infraserver

VM Configuration Details		Infraserver Configuration Details	
# VCPUS	2	Web Server	Red Hat



VCPUs	2261 MHz
Memory (MB)	512
# VNICs	2
VNIC Description	2 x virtio_net
Storage Description	2 x virtio_blk
Virtual Disk Size (MB)	15360
Datastore Size (MB)	61440
VM OS	Red Hat Enterprise Linux 5.5 (32-bit)
VM OS Availability	April-2010

Vendor	
Web Server Name/Version	Apache 2.2.3
Availability Date	April-2010
Script Vendor	fastcgi.com
Script Name/Version	FastCGI 2.4.6
Script Availability Date	November-2007
JVM Version	1.6.0 IBM J9 2.4 Linux amd64-64 jvmxa6460sr8-20100401_55940
Other Software	N/A

Idle Server

VM Configuration Details	
# VCPUS	1
VCPUs	2261 MHz
Memory (MB)	512
# VNICs	1
VNIC Description	1 x virtio_net
Storage Description	2 x virtio_blk
Virtual Disk Size (MB)	15360
Datastore Size (MB)	20480
VM OS	Red Hat Enterprise Linux 5.5 (32-bit)
VM OS Availability	April-2010

Idle Server Configuration Details	
JVM Version	1.6.0 IBM J9 2.4 Linux amd64-64 jvmxa6460sr8-20100401_55940
Other Software	N/A



6.3 Notes, Settings & Tuning

Physical System Notes

- Firmware settings used:
- TurboModeEnable=Enable
- TurboBoost=Traditional
- HardwarePrefetcher=Disable

Storage Notes

- All storage accessed with 2 x 4 Gb Fibre Channel links
- to DS4800 storage. DS4800 is attached
- to 8 EXP3000 SAS enclosures, each enclosure with 16 73 GB SAS drives. Total storage drive count is 114. There are 14 RAID5 arrays, each consisting of 8 drives. Each array has one LUN defined for one Tile. Each drive defined for a Tile is divided into 12 logical volumes

Virtualization Software Notes

- The following was run after each boot:
 - echo 242645 >/proc/sys/vm/nr_hugepages
 - mkdir /dev/hugepages
 - mount -t hugetlbfs none /dev/hugepages
 - for i in `seq 0 31`; do echo performance >/sys/devices/system/cpu/cpu\$i/cpufreq/scaling_governor; done
- The following boot options were used:
 - elevator=deadline intel_iommu=on iommu=pt pci=assign-busses
- Network configuration is as follows:
 - Two software bridges are created, br0, br1, and the two Broadcom
 - Ethernet ports are assigned to a bridge (eth0 to br0, eth1 to br1)
 - For every VM, their first virtIO network device is assigned to one of these bridges, even number tiles to br0, odd numbered tiles to br1



- For each tile, there is another bridge created, name t01....tN
 - For all Web and Infra VMs, their second virtIO network device is to a bridge corresponding to the tile number.
 - Kernel module for Intel Gigabit ET Server adapter is loaded with these options to create 42 virtual functions: max_vfs=7,7,7,7,7,7
 - These virtual functions are assigned (one per VM) to the Web, App and DB VMs. When a virtual function is present on a VM, a script is executed to shut down the first virtio_net device and migrate the IP settings to the virtual function device. This is done to facilitate easy IP/dns management when using a mix of VMs with and without virtual functions. The scripts, convert_to_eth[1|2].sh are available in the disclosure archive.
 - On the host, Ethernet devices eth2 - eth7 are the 6 physical functions for the Intel Ethernet ports. These are not used by the host, except to provide physical transport for the virtual functions.
- All VMs are launched with nice utility (-n 19) to adjust their scheduling priority
 - VMs on even numbered tiles are bound to NUMA node 0 and have a preferred memory allocation node of 0. VMs on odd numbered tiles are bound to NUMA node 1 and have a preferred memory allocation node of 1.
 - All VMs' memory but idle VMs are fulfilled with huge (2MB) pages.
 - All VMs are created with the following command line format (options in [] are for VMs which have PCI device assignment or hugepages):
 - /usr/libexec/qemu-kvm -name \$name -drive file=\$dev1,if=virtio,boot=on,cache=none -drive file=\$dev2,if=virtio,boot=off,cache=none
 - -net nic,model=virtio,vlan=0,macaddr=\$macaddr -net tap,vlan=0,ifname=webt01,script=/etc/kvm/qemu-ifup-br\$bridge_id
 - -net nic,model=virtio,vlan=1,macaddr=\$alt_macaddr -net tap,vlan=1,ifname=webt01b,script=/etc/kvm/qemu-ifup-t\$tile_num
 - [-pcidevice host=\$pci_device] -m \$mem -mem-prealloc [-mem-path /dev/hugepages] -vnc :\$vnc_port -smp \$cpus
 - -monitor telnet::\$monitor_port,server,nowait -serial telnet::\$serial_port,server,nowait -daemonize



Web Server VM Notes

- Disabled the following services via chkconfig:
anacron,atd,bluetooth,cpuspeed,crond,cups,dc_client,dc_server,firstboot,gpm,httpd,ip6tables,iptables,irqbalance,isdn,postgresql,rhnsd,sendmail,smartd,squid,tux,xfs,yum-updatesd
- nofile and nproc set to 102400 in /etc/security/limits.conf
- httpd.conf changes:
 - KeepAlive on
 - MaxKeepAliveRequests 0 (default 100)
 - KeepAliveTimeout 5 (default 15)
 - prefork MPM StartServers=20 (default 8)
 - prefork MPM MinSpareServers=20 (default 5)
 - prefork MPM MaxSpareServers=20 (default 20)
 - prefork MPM ServerLimit=1536 (default 256)
- Boot options: divider=10
 - prefork MPM MaxRequestsPerChild=1000 (default 4000)
 - CustomLog logs/access_log common
- Apache SSL support was disabled
- Web VM used NFS to access files from infra VM
- No options were specified in the NFS mount

Application Server VM Notes

- Disabled the following services via chkconfig:
anacron,atd,bluetooth,cpuspeed,crond,cups,dc_client,dc_server,firstboot,gpm,httpd,ip6tables,iptables,irqbalance,isdn,postgresql,rhnsd,sendmail,smartd,squid,tux,xfs,yum-updatesd
- Heap size for Application server is 1280 MB, using huge pages
- The following was run after each boot:
 - echo 2818572288 >/proc/sys/kernel/shmmax
 - echo 688128 >/proc/sys/kernel/shmall
- Boot options: divider=10



Mail Server VM Notes

- Disabled the following services via chkconfig:
anacron,atd,bluetooth,cpuspeed,crond,cups,dc_client,dc_server,firstboot,gpm,httpd,ip6tables,iptables,irqbalance,isdn,postgresql,rhnsd,sendmail,smartd,squid,tux,xfs,yum-updatesd
- dovecot.conf:
 - base_dir = /var/run/dovecot/
 - protocols = imap
 - listen = *:143
 - disable_plaintext_auth = no
 - log_path = /var/log/dovecot
 - ssl = no
 - login_processes_count = 10
 - mail_location = dbox:/mailstore/dovecot/maildir/%u
 - mail_debug = yes
 - mailldir_copy_with_hardlinks = yes
 - dbox_rotate_size = 2048
 - dbox_rotate_min_size = 16
 - dbox_rotate_days = 0
 - protocol lda {
 postmaster_address = postmaster@example.com
 }
 auth default {
 mechanisms = plain
 passdb passwd-file {
 args = /etc/passwd.dovecot
 }
 userdb passwd {
 }
 }



```
user = root  
}  
dict {  
}  
plugin {  
}
```

- Boot options: divider=10

Database Server VM Notes

- Disabled the following services via chkconfig:
anacron,atd,bluetooth,cpuspeed,crond,cups,dc_client,dc_server,firstboot,gpm,httpd,ip6tables,iptables,irqbalance,isdn,postgresql,rhnsd,sendmail,smartd,squid,tux,xfs,yum-updatesd
- Database tuning options:
 - db2 update db cfg for specdb using DATABASE_MEMORY 200000
 - db2set DB2_LARGE_PAGE_MEM=DB; db2set DB2_PINNED_BP=ON
 - alter bufferpool IBMDEFAULTBP size 160000
 - alter tablespace roll_space NO FILE SYSTEM CACHING
 - alter tablespace C_cust_space NO FILE SYSTEM CACHING
 - alter tablespace C_supp_space NO FILE SYSTEM CACHING
 - alter tablespace C_site_space NO FILE SYSTEM CACHING
 - alter tablespace C_parts_space NO FILE SYSTEM CACHING
 - alter tablespace C_custinv_space NO FILE SYSTEM CACHING
 - alter tablespace M_parts_space NO FILE SYSTEM CACHING
 - alter tablespace M_bom_space NO FILE SYSTEM CACHING
 - alter tablespace M_wo_space NO FILE SYSTEM CACHING
 - alter tablespace M_lo_space NO FILE SYSTEM CACHING
 - alter tablespace M_inv_space NO FILE SYSTEM CACHING
 - alter tablespace O_cust_space NO FILE SYSTEM CACHING
 - alter tablespace O_ords_space NO FILE SYSTEM CACHING



- alter tablespace O_ordl_space NO FILE SYSTEM CACHING
- alter tablespace O_item_space NO FILE SYSTEM CACHING
- alter tablespace S_comp_space NO FILE SYSTEM CACHING
- alter tablespace S_sc_space NO FILE SYSTEM CACHING
- alter tablespace S_po_space NO FILE SYSTEM CACHING
- alter tablespace S_site_space NO FILE SYSTEM CACHING
- alter tablespace S_supp_space NO FILE SYSTEM CACHING
- alter tablespace U_space NO FILE SYSTEM CACHING
- update db cfg for specdb using logprimary 64
- Boot options: divider=10 hugepages=416

Infraserver VM Notes

- Disabled the following services via chkconfig:
anacron,atd,bluetooth,cpuspeed,crond,cups,dc_client,dc_server,firstboot,gpm,httpd,ip6tables,iptables,irqbalance,isdn,postgresql,rhnsd,sendmail,smartd,squid,tux,xfs,yum-updatesd
- nofile and noproc set to 102400 in /etc/security/limits.conf
- httpd.conf changes:
 - Listen port to 81 (default 80)
 - prefork MPM StartServers=100 (default 8)
 - prefork MPM MinSpareServers=100 (default 5)
 - prefork MPM MaxSpareServers=100 (default 20)
 - prefork MPM ServerLimit=1024 (default 256)
 - prefork MPM MaxClients=1024 (default 256)
 - prefork MPM MaxRequestsPerChild=16000 (default 4000)
 - LoadModule fastcgi_module modules/mod_fastcgi.so
 - ServerName besim:80
 - CustomLog logs/access_log common
 - FastCgiServer /var/www/fcgi-bin/besim_fcgi.fcgi -processes 50
- Apache SSL support was disabled



- Boot options: divider=10
- Standard Linux NFS server is enabled to serve files to Web VM
 - /etc/exports file contains:
 - /support-download *(rw,no_subtree_check,crossmnt,no_root_squash)

Idle Server VM Notes

- Disabled the following services via chkconfig:
anacron,atd,bluetooth,cpuspeed,crond,cups,dc_client,dc_server,firstboot,gpm,httpd,ip6tables,iptables,irqbalance,isdn,postgresql,rhnsd,sendmail,smartd,squid,tux,xfs,yum-updatesd
- Boot options: divider=10

Client Driver Notes

- 1 client system is dedicated for prime controller
- 14 IBM x3550 clients are used for the prime clients
- Each client has 2 CPUs, 32 GB memory,
- One Broadcom 1Gbps Ethernet port is used for
- SUT communications.
- Red Hat Enterprise Linux 5.5 64bit is the OS used on the clients.

Other Notes

- A single Cisco Catalyst switch, model 6500-E is used for all physical network connections
- All VMs were running the 2.6.18-194.el5 Linux kernel.
- The Webserver and Database VM were running the RHEL5.5 64bit
- All other VMs were running RHEL5.5 32bit version



7.0 Benchmark Results - RHEL 5.5 (KVM) on IBMx3690 X5)

7.1 Performance Summary

Performance							
Tile #	Pct Load	Application Server	Web Server	Mail Server	Idle Server	Per-Tile Score	Overall Score
1	100%	33.43	53.47	88.67	N/A	97.73	1369
2	100%	33.43	53.64	88.54	N/A	97.79	
3	100%	33.32	53.52	88.73	N/A	97.68	
4	100%	33.47	53.83	88.54	N/A	97.95	
5	100%	33.36	53.51	88.61	N/A	97.67	
6	100%	33.48	53.78	88.78	N/A	98.01	
7	100%	33.26	53.46	88.42	N/A	97.47	
8	100%	33.33	53.74	88.67	N/A	97.81	
9	100%	33.43	53.57	88.66	N/A	97.79	
10	100%	33.40	53.95	88.60	N/A	97.98	
11	100%	33.33	53.70	88.46	N/A	97.70	
12	100%	33.44	53.93	88.67	N/A	98.03	
13	100%	33.43	53.42	88.57	N/A	97.67	
14	100%	33.50	53.71	88.65	N/A	97.94	

Quality Of Service (QOS)							
Tile #	Pct Load	Application Server	Web Server	Mail Server	Idle Server	Per-Tile Score	Overall Score
1	100%	1.00	0.99	1.00	1.00	1.00	99.82%
2	100%	1.00	1.00	1.00	1.00	1.00	
3	100%	0.99	0.99	1.00	1.00	1.00	
4	100%	1.00	1.00	1.00	1.00	1.00	
5	100%	1.00	0.99	1.00	1.00	1.00	
6	100%	1.00	1.00	1.00	1.00	1.00	
7	100%	1.00	0.99	1.00	1.00	1.00	
8	100%	1.00	1.00	1.00	1.00	1.00	
9	100%	0.99	0.99	1.00	1.00	1.00	
10	100%	0.99	1.00	1.00	1.00	1.00	
11	100%	0.99	1.00	1.00	1.00	1.00	



12	100%	1.00	1.00	1.00	1.00	1.00	
13	100%	1.00	0.99	1.00	1.00	1.00	
14	100%	1.00	1.00	1.00	1.00	1.00	

7.2 Performance Details

Tile 1					
Application Server					
Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.34	1.23	14.99	1.50	5
Dealer	20.08	0.12/0.24/0.11	13.15/14.93/13.22	0.30/0.60/0.20	2/2/2
Web Server					
Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.47	382674	1681	612	0
Mail Server					
Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.18	102083	102083	0	100.00
Fetch	36.12	260057	260057	0	100.00
Idle Server					
Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec	
Heartbeats	719	3.56	1	159	
Tile 2					
Application Server					
Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.36	1.19	14.61	1.50	5
Dealer	20.07	0.12/0.22/0.12	14.64/17.55/13.61	0.20/0.50/0.20	2/2/2
Web Server					
Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.64	384510	875	810	0
Mail Server					
Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.22	102368	102368	0	100.00



Fetch	36.23	260891	260891	0	100.00
Idle Server					
Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec	
Heartbeats	719	4.29	1	504	
Tile 3					
Application Server					
Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.30	1.28	15.73	1.50	5
Dealer	20.02	0.21/0.28/0.16	19.98/26.74/22.82	0.20/0.50/0.20	2/2/2
Web Server					
Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.52	381831	1942	1606	0
Mail Server					
Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.29	102872	102872	0	100.00
Fetch	36.12	260067	260067	0	100.00
Idle Server					
Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec	
Heartbeats	719	3.66	1	145	
Tile 4					
Application Server					
Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.37	1.16	14.30	1.50	5
Dealer	20.10	0.10/0.18/0.09	13.12/13.55/13.12	0.20/0.40/0.20	2/2/2
Web Server					
Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.83	385939	917	729	0
Mail Server					
Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.24	102555	102555	0	100.00
Fetch	36.10	259929	259929	0	100.00
Idle Server					



Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec
Heartbeats	719	3.98	1	340

Tile 5

Application Server					
Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.36	1.22	14.11	1.50	5
Dealer	20.01	0.11/0.23/0.10	13.07/13.38/13.27	0.20/0.50/0.20	2/2/2

Web Server

Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.51	382307	1430	1555	0

Mail Server

Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.17	102034	102034	0	100.00
Fetch	36.13	260119	260119	0	100.00

Idle Server

Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec
Heartbeats	719	4.06	1	426

Tile 6

Application Server					
Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.33	1.24	19.58	1.50	5
Dealer	20.15	0.12/0.21/0.11	17.25/17.78/16.75	0.20/0.40/0.20	2/2/2

Web Server

Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.78	385360	798	1091	0

Mail Server

Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.23	102482	102482	0	100.00
Fetch	36.18	260507	260507	0	100.00

Idle Server

Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec
Heartbeats	719	4.14	1	404



Tile 7

Application Server

Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.36	1.23	19.63	1.50	5
Dealer	19.90	0.13/0.23/0.15	25.94/30.13/25.38	0.20/0.50/0.20	2/2/2

Web Server

Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.46	382251	1765	927	0

Mail Server

Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.18	102110	102110	0	100.00
Fetch	36.10	259933	259933	0	100.00

Idle Server

Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec
Heartbeats	719	3.80	1	305

Tile 8

Application Server

Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.27	1.19	15.20	1.50	5
Dealer	20.07	0.11/0.19/0.10	13.55/15.20/13.91	0.20/0.40/0.20	2/2/2

Web Server

Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.74	385849	605	449	0

Mail Server

Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.26	102694	102694	0	100.00
Fetch	36.12	260095	260095	0	100.00

Idle Server

Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec
Heartbeats	719	3.68	1	90

Tile 9

Application Server



Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.31	1.29	17.34	1.50	5
Dealer	20.12	0.16/0.28/0.16	14.75/20.43/23.66	0.30/0.60/0.20	2/2/2

Web Server

Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.57	382530	2034	1126	0

Mail Server

Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.29	102906	102906	0	100.00
Fetch	36.15	260293	260293	0	100.00

Idle Server

Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec
Heartbeats	719	3.47	1	17

Tile 10

Application Server

Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.29	1.22	20.57	1.50	5
Dealer	20.11	0.12/0.23/0.12	17.54/18.19/15.07	0.20/0.50/0.20	2/2/2

Web Server

Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.95	387527	466	438	0

Mail Server

Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.24	102555	102555	0	100.00
Fetch	36.23	260826	260826	0	100.00

Idle Server

Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec
Heartbeats	719	3.61	1	154

Tile 11

Application Server

Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.34	1.27	18.48	1.50	5



Dealer	19.98	0.20/0.26/0.14	23.93/28.79/17.53	0.20/0.50/0.20	2/2/2
--------	-------	----------------	-------------------	----------------	-------

Web Server

Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.70	384700	1300	631	0

Mail Server

Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.27	102738	102738	0	100.00
Fetch	36.06	259617	259617	0	100.00

Idle Server

Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec
Heartbeats	719	5.33	1	1325

Tile 12

Application Server

Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.30	1.21	14.64	1.50	5
Dealer	20.14	0.09/0.18/0.09	13.10/13.84/13.35	0.20/0.40/0.20	2/2/2

Web Server

Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.93	386481	830	1003	0

Mail Server

Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.27	102720	102720	0	100.00
Fetch	36.07	259727	259727	0	100.00

Idle Server

Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec
Heartbeats	719	3.96	1	415

Tile 13

Application Server

Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.32	1.20	17.59	1.50	5
Dealer	20.11	0.11/0.22/0.10	15.19/18.21/13.64	0.20/0.50/0.20	2/2/2

Web Server



Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.42	380853	2763	1016	0

Mail Server

Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.26	102680	102679	1	100.00
Fetch	36.10	259929	259929	0	100.00

Idle Server

Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec
Heartbeats	719	4.07	1	637

Tile 14

Application Server

Req. Type	Req/sec	Avg Resp. Time	Max Resp. Time	90th%	Required 90th%
Manufacturing	13.38	1.13	6.51	1.25	5
Dealer	20.12	0.07/0.16/0.07	2.00/7.86/1.24	0.20/0.40/0.20	2/2/2

Web Server

Web Wkload	Req/sec	Good	Tolerable	Fail	Valid. Errors
Support	53.71	385264	628	818	0

Mail Server

Req. Type	Req/sec	Total Count	Pass Count	Fail Count	Pass Pct
Append	14.30	102993	102993	0	100.00
Fetch	36.20	260621	260621	0	100.00

Idle Server

Req. Type	Total	Avg. Resp. Msec	Min Resp. Msec	Max Resp. Msec
Heartbeats	719	3.56	0	57



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. 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.
 - o Performance-only, which produces the SPECvirt_sc2010 metric,
 - o Performance/Power of the Total System Under Test, which produces the SPECvirt_sc2010_PPW metric, and
 - o 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 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.

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



Appendix B: References

1. SPECvirt_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
3. SPECvirt_sc2010 Run and Reporting Rules
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
5. SPECvirt_sc2010 Submitted Results
http://www.spec.org/virt_sc2010/results/