White Paper
Intel Information Technology
Computer Manufacturing
Server Virtualization

Evaluating Two- and Four-Socket Server Virtualization Platforms

Intel IT and Intel's Digital Enterprise Group (DEG), End User Platform Integration (EPI), conducted proof of concept (PoC) testing and total cost of ownership (TCO) analysis to compare two- and four-socket servers based on Intel® Xeon® processors for server virtualization. Test scenarios included performance service-level agreements (SLAs), commodity consolidation in which memory is the limiting factor, and data centers with power and cooling constraints or limited LAN and storage area network (SAN) ports. PoC test results indicate that four-socket servers based on Intel Xeon processor 7400 series can offer significant advantages and lower TCO in multiple virtualization deployment scenarios.

Sudip Chahal and Raghu Yeluri, Intel Corporation

September 2008

IT@Intel

Executive Summary

Four-socket servers based on Intel Xeon processor 7400 series can offer significant advantages and lower TCO in multiple virtualization deployment scenarios.

Intel IT and Intel's Digital Enterprise Group (DEG), End User Platform Integration (EPI), conducted proof of concept (PoC) testing and total cost of ownership (TCO) analysis to compare two- and four-socket servers based on Intel® Xeon® processors for a variety of virtualization scenarios. Scenarios included performance service-level agreements (SLAs), commodity consolidation in which memory is the limiting factor, and data centers with power and cooling constraints or limited LAN and storage area network (SAN) ports. Our test results indicate that four-socket servers based on Intel Xeon processor 7400 series can offer significant advantages and lower TCO in multiple virtualization deployment scenarios.

To conduct the PoC, we tested four different servers using the vConsolidate virtualization benchmark suite. Results show that a four-socket server based on Intel Xeon processor X7460 with Intel® 45nm Hi-K process technology and 6 cores delivered 2.16x the throughput of a two-socket server based on Quad-Core Intel Xeon processor X5460, with lower utilization.

Based on our test results, we compared TCO for each server in a variety of data center virtualization scenarios.

- **Performance-centric SLA-focused.** In scenarios focused on performance SLAs, a server based on Intel Xeon processor X7460 could support 26 percent more virtual machines (VMs) for the same TCO than a four-socket server based on Quad-Core Intel Xeon processor X7350.
- **Memory capacity-focused.** In scenarios focused on maximizing consolidation ratios, a server based on Intel Xeon processor X7460 could support 72 percent more VMs for the same TCO than a server based on Quad-Core Intel Xeon processor X5460.
- **Data center-constrained.** In scenarios focused on memory capacity, the four-socket server based on Intel Xeon processor X7460 could deliver more VMs per watt in data centers facing power and cooling constraints.
- **Scalability-focused.** Resource pools using servers based on Intel Xeon processor X7460 could contain about 2.16x as many VMs in performance-centric SLA scenarios and 4x as many VMs in memory capacity-focused scenarios as servers based on Intel Xeon processor 5460; servers based on Intel Xeon processor X7460 would also scale much more consistently when unexpected workload spikes occur.

Our PoC indicates that servers based on Intel Xeon processor 7400 series can offer significant advantages and lower TCO in these and other virtualization deployment scenarios.

Contents

Business Challenge

Like many IT organizations, Intel IT is pursuing server virtualization to reduce costs in areas such as hardware, technical support, and power and cooling.

As part of this strategy, we analyze different server platforms to compare how they help IT meet the compute requirements of the Intel business groups we support. These groups of users have a variety of business requirements, resulting in multiple virtualization deployment total cost of ownership (TCO) scenarios. We have classified these scenarios into three broad groups:

- Performance-centric SLA-focused. These scenarios emphasize the need to meet performance service-level agreements (SLAs) based on response time or application throughput. Examples include end-of-quarter financial processing, trading applications, and engineering design workloads.
- Memory capacity-focused. These scenarios focus on provisioning large numbers of virtual machines (VMs) as cost-effectively as possible; performance is secondary. Typically, this means maximizing consolidation ratios. Examples may include VMs used for development, testing, and demonstration, and lightly loaded production applications.
- VMs per host limited by business policy or IT environment. In these scenarios, external constraints such as risk management concerns or network infrastructure limitations restrict the number of VMs per physical server. As a result, the load on each physical server is relatively low.

Intel IT evaluates server platforms based on Intel® Xeon® processors to determine which delivers the best TCO for each scenario. Our goal is to meet the needs of business groups in the most cost-effective way.

In addition to these three virtualization deployment TCO scenarios, we have identified scenarios in which specific data center or software architectural constraints may influence our virtualization server strategy:

- Data center-constrained. Many data centers are nearing the limit of their power and cooling capacity or have limited LAN and storage area network (SAN) ports. This favors servers that can maximize the number of VMs per watt of power consumption or per port.
- Scalability-focused. These scenarios favor more-scalable systems. One common constraint is imposed by virtualization host management software that limits the number of physical servers within a resource pool or cluster; more-scalable servers provide increased flexibility because each pool can contain a greater number of VMs.

We compare and analyze different server platforms to determine which platforms best meet the requirements of each scenario. To do this, we perform tests to measure each platform's performance, scalability, and power consumption. Based on the results and other factors such as platform and data center costs, we calculate the relative TCO of each platform for specific scenarios.

Proof of Concept

In collaboration with Intel's Digital Enterprise Group (DEG), End User Platform Integration (EPI), Intel IT recently performed proof of concept (PoC) testing of two- and four-socket server platforms based on Intel Xeon processors, including Intel Xeon processor 7400 series based on Intel® 45nm Hi-K process technology and with 6 cores per processor.

We compared this four-socket server with three other server platforms: a four-socket server based on Quad-Core Intel Xeon processor 7300 series, a two-socket server based on Quad-Core Intel Xeon processor 5400 series, and a two-socket server based on Intel Xeon processor 5300 series. To conduct our PoC, we:

- Tested each platform using a virtualization benchmark suite representing typical IT workloads. We compared performance, processor utilization, and power consumption. We derived relative measures of performance and power consumption for use in our TCO analysis.
- Performed an analysis to compare TCO of each platform when supporting each of our three virtualization deployment TCO scenarios: performance-centric SLA-focused, memory capacity-focused, and VMs per host limited by business policy or IT environment. We also compared each platform for data center-constrained and scalability-focused scenarios.

Virtualization Benchmark Tests

We tested four systems:

- A four-socket server based on Intel Xeon processor X7460 (2.66 GHz) with Intel 45nm Hi-K process technology and with 6 cores and 16 MB of cache per processor.
- A four-socket server based on Quad-Core Intel Xeon processor X7350 (2.93 GHz), with 65-nm technology and four cores.
- A two-socket server based on Quad-Core Intel Xeon processor X5460 (3.16 GHz), with 45-nm technology and four cores.
- A two-socket server based on Quad-Core Intel Xeon processor X5365 (3.0 GHz), with 65-nm technology and four cores.

System configurations are shown in Table 1. We used two- and four-socket server configurations with 8 and 32 DIMM slots, respectively. Our configurations included enough RAM to enable us to drive CPU utilization to high levels without encountering memory constraints. We configured the two-socket servers with 16 GB of RAM and the four-socket servers with 32 GB; our testing was not memory-limited on any system.

Architectural differences between the two- and four-socket servers are shown in Figure 1. The four-socket servers included the Intel® 7300 chipset, which provides four dedicated high-speed interconnects (DHSI) to accelerate memory access. The Intel 7300 chipset also integrates a 64-MB Snoop Filter that manages data coherency across

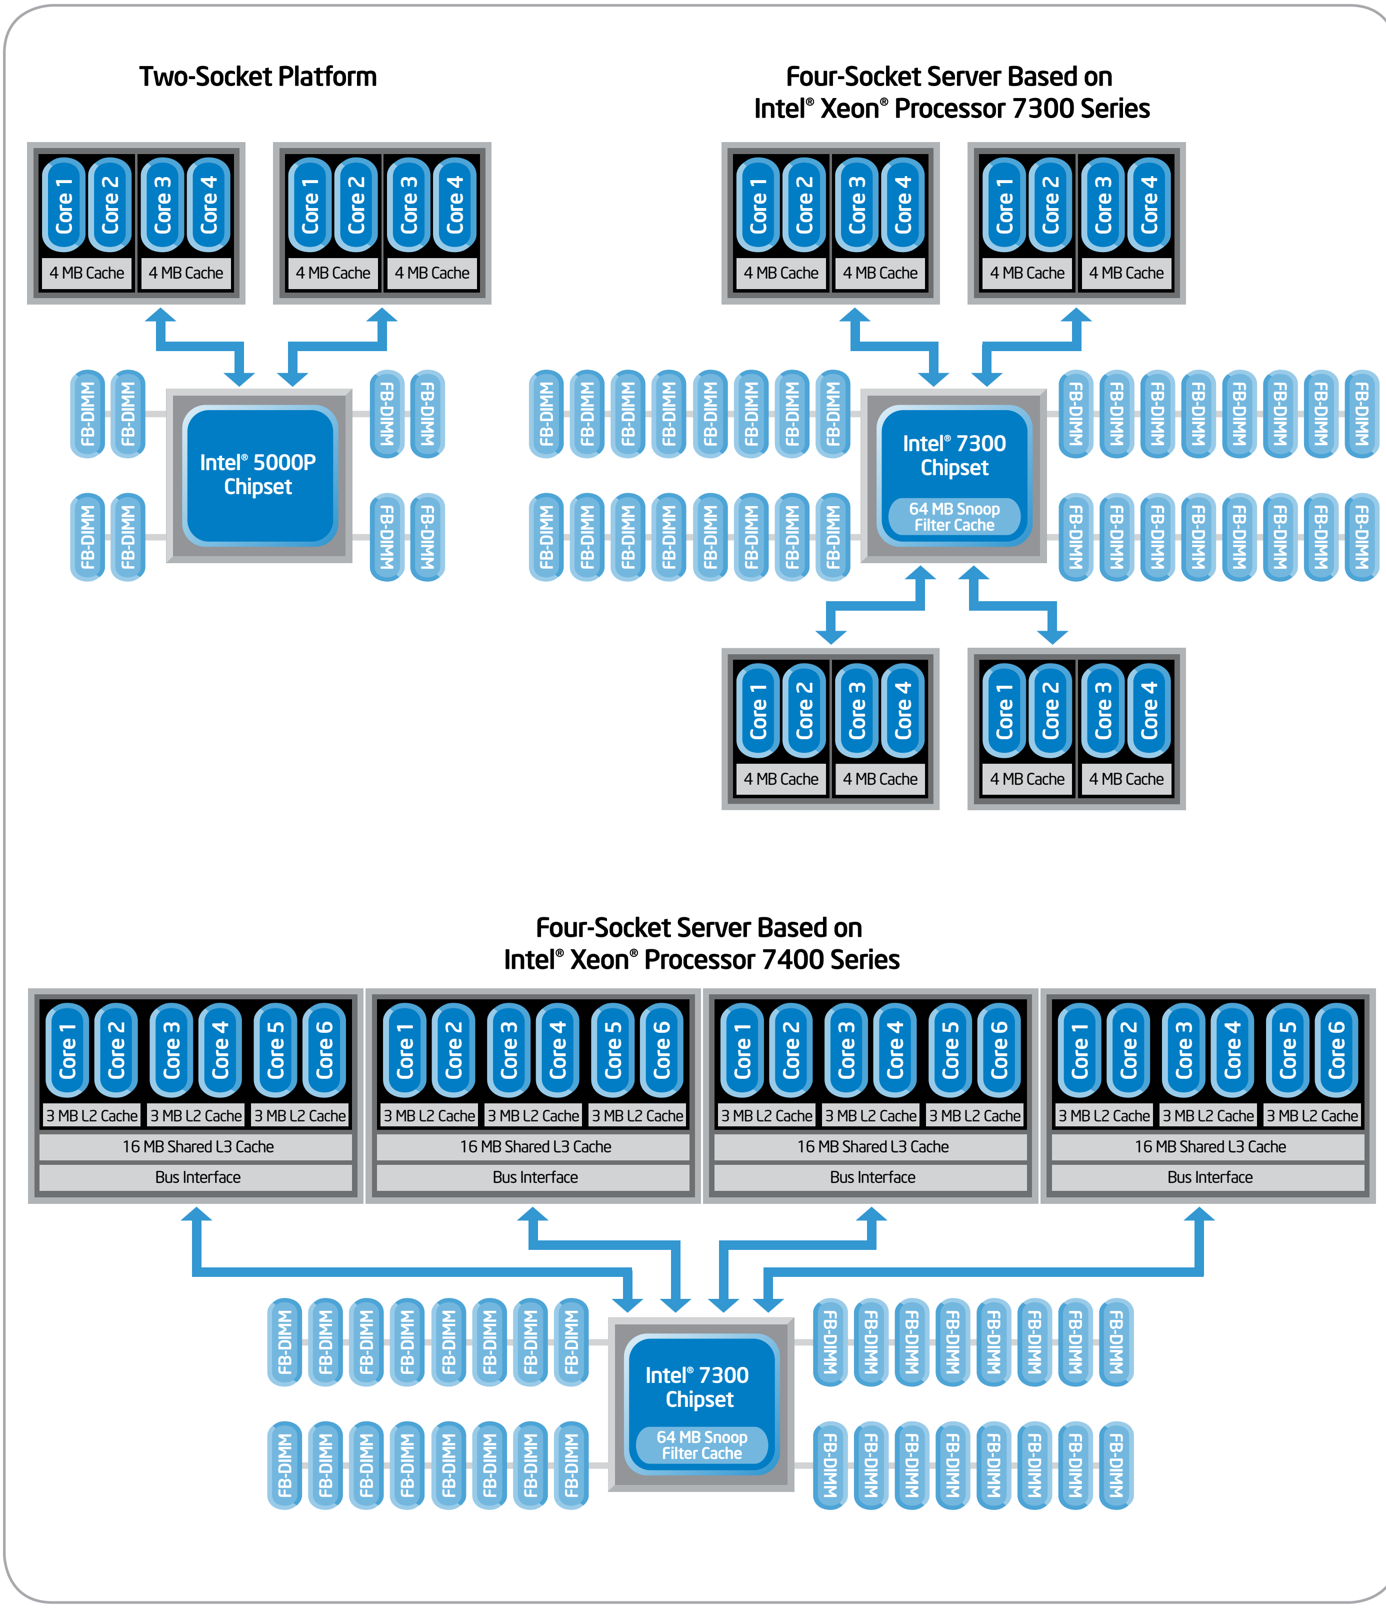

Figure 1. Architecture of two- and four-socket servers based on Intel® Xeon® processors.

Table 1. Proof of Concept Test System Configurations

	Intel® Xeon® Processor X5365	Intel Xeon Processor X5460	Intel Xeon Processor X7350	Intel Xeon Processor X7460
Processor	▪ Intel Xeon processor X5365 (3 GHz) ▪ 2-socket/quad-core ▪ 2 Intel Xeon processor X5365-based CPUs	▪ Intel Xeon processor X5460 (3.16 GHz) ▪ 2-socket/quad-core ▪ 2 Intel Xeon processor X5460-based CPUs	▪ Intel Xeon processor X7350 (2.93 GHz) ▪ 4-socket/quad-core ▪ 4 Intel Xeon processor X7350-based CPUs	▪ Intel Xeon processor X7460 (2.66 GHz) ▪ 4-socket/6 cores ▪ 4 Intel Xeon processor X7460-based CPUs
Chipset	Intel® 5000P Chipset	Intel 5000P Chipset	Intel® 7300 Chipset	Intel 7300 Chipset
Cache	8 MB Intel® Smart Cache	12 MB Intel Smart Cache	8 MB Intel Smart Cache	3x 3 MB L2, 16 MB L3 cache
Front Side Bus Speed	1333 MHz	1333 MHz	1066 MHz	1066 MHz
Memory Capacity/ Speed	▪ 16 GB (8x 2 GB) ▪ FB-DIMM (ECC) 667 MHz	▪ 16 GB (8x 2 GB) ▪ FB-DIMM (ECC) 667 MHz	▪ 32 GB (16x 2 GB) ▪ FB-DIMM (ECC) 667 MHz	▪ 32 GB (16x 2 GB) ▪ FB-DIMM (ECC) 667 MHz

the processors and increases the available memory bandwidth. The two-socket servers included the Intel® 5000P chipset, with a 1333 MHz front side bus (FSB).

Test Environment

For our PoC, we required a benchmark suite that:

- Represented good approximations of IT workloads.
- Produced objective measurements.
- Adequately stressed each of the systems.
- Produced repeatable results.

To meet these needs, we selected the vConsolidate benchmark suite. This tests virtualization performance by simulating real server performance in a typical environment. We used vConsolidate (profile-2) to generate consistent application workloads for the performance analysis.

The test suite runs consolidated stack units (CSUs), each comprising five simultaneously running VMs, including one idle VM, as shown in Table 2. CSUs can be added to increase the system load and CPU utilization.

Our PoC test environment is shown in Figure 2. We connected each tested server to a LAN and to a SAN. Also connected to the LAN were driver systems that generated each CSU workload and controller nodes that managed the CSUs. We created a single set of test VMs and migrated it between test systems using the SAN.

Table 2. vConsolidate Consolidated Stack Unit (CSU)

	Workload	Operating System	Memory	vCPUs
VM 1	Database	64-bit Windows Server 2003*	2.0 GB	2
VM 2	Web	32-bit Windows Server 2003*	1.5 GB	2
VM 3	Mail	32-bit Windows Server 2003	1.5 GB	1
VM 4	Java*	64-bit Windows Server 2003	2.0 GB	2
VM 5	Idle	32-bit Windows Server 2003	0.4 GB	1

vCPU: virtual CPU; **VM:** virtual machine

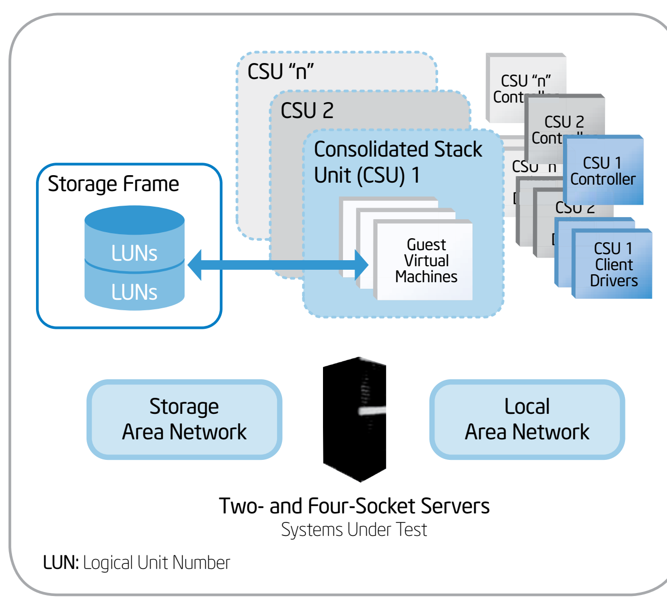

Figure 2. Proof of concept test environment.

When running the benchmark tests, we added CSUs until the systems under test attained high system utilization levels in excess of 80 percent—levels that we would rarely expect to reach in production. This allowed us to more fully exploit the scalability of each platform. If a system has enough headroom, adding CSUs increases the aggregate throughput of the system, which results in a greater vConsolidate score.

We executed three passes of the tests on each system while monitoring CPU utilization and power consumption, and selected the median set of results. For reference, we used the scores from test runs on a reference platform consisting of a two-socket server based on Dual-Core Intel Xeon processor 5160.

We analyzed the capabilities of the four tested servers by comparing performance, performance-per-unit cost, and performance per watt for each system.

vConsolidate Test Results

In our tests, the four-socket server based on Intel Xeon processor X7460 was considerably more scalable than any the other platforms, delivering much greater aggregate throughput at lower utilization levels, as shown in Figure 3.

It required six CSUs to drive utilization to about 90 percent on the server based on Intel Xeon processor X7460. In contrast, the other platforms exceeded 90 percent utilization with two to four CSUs.

Aggregate throughput on the server based on Intel Xeon processor X7460 continued to increase as we increased the number of CSUs to six. As a result, the server achieved a maximum throughput that was 2.16x the maximum throughput of the server based on Intel Xeon processor X5460, with lower utilization.

The maximum throughput of the server based on Intel Xeon processor X7460 was about

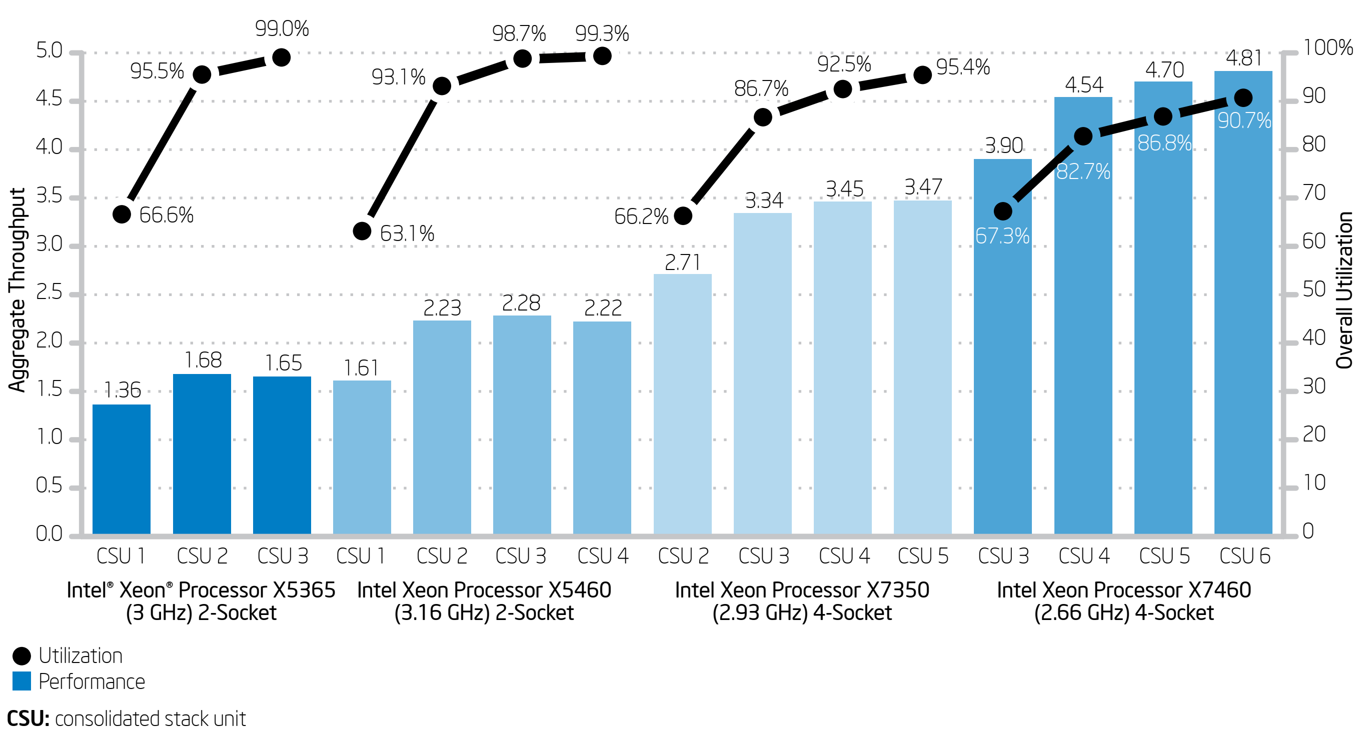

Figure 3. Server aggregate throughput and utilization with increasing workloads. Intel internal measurements, August 2008.

1.39x the throughput of the server based on Intel Xeon processor X7350, with lower utilization.

Maximum server throughput and utilization are compared in Figure 4.

Throughput Per Watt

Processors based on 45-nm technology are more power-efficient than the corresponding processors based on 65-nm technology. The two- and four-socket servers with processors based on 45-nm technology used correspondingly less power when running the tests, as shown in Figure 5. As a result, the server based on the Intel Xeon processor X7460 delivered approximately 46 percent better performance per watt compared with the server based on the Intel Xeon processor X7350 when both were configured with 32 GB RAM.

We then estimated the power consumption of servers configured with more memory for performance-centric SLA-focused scenarios. We assumed servers configured with 4 GB of RAM per core, on the basis that servers with greater throughput would be configured with correspondingly more memory to accommodate a greater number of production workloads. We calculated an adjusted estimate of the power consumption based on 32 GB of RAM on the two-socket servers, 64 GB on the server based on Intel Xeon processor X7350, and 96 GB on the server based on Intel Xeon processor X7460, as shown in Figure 6.

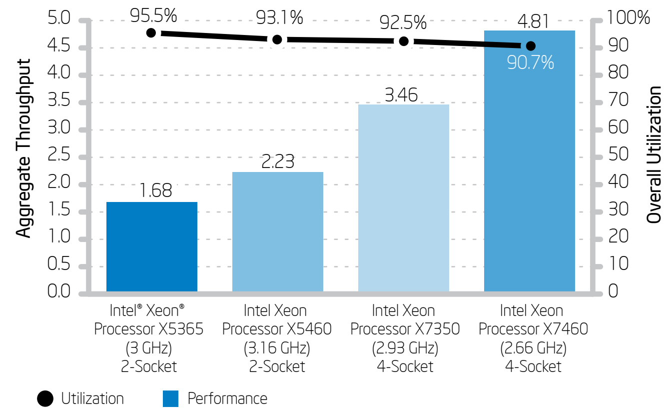

Figure 4. Maximum server throughput and utilization. Intel internal measurements, August 2008.

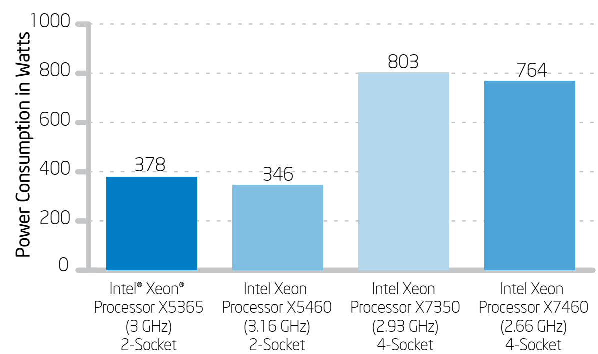

Figure 5. Server relative power consumption when running vConsolidate tests. Intel internal measurements, August 2008.

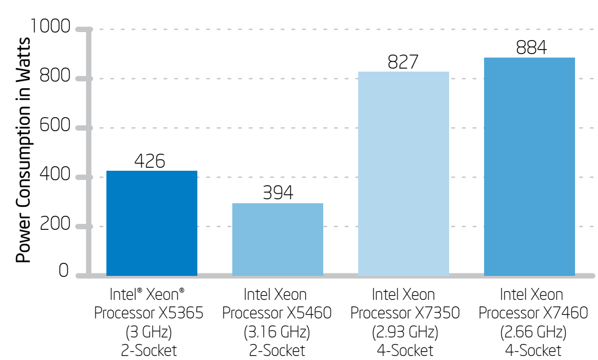

Figure 6. Server relative power consumption adjusted for additional RAM. Estimates based on RAM adjustment assumptions and Intel internal measurements, August 2008.

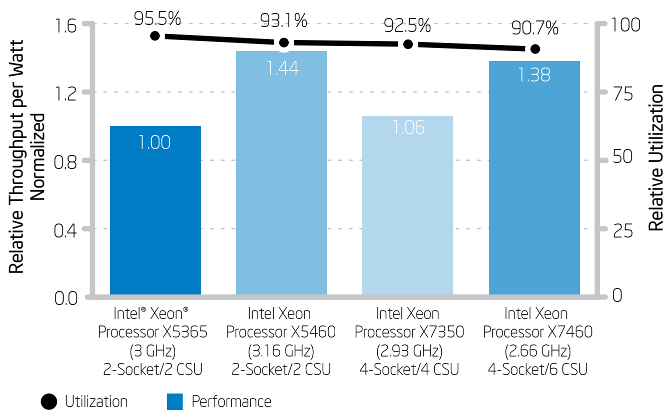

Figure 7. Relative throughput per watt. Intel internal measurements, August 2008.

Table 3. Total Cost of Ownership Assumptions

Category	Assumption
Data Center Physical Plant	▪ Space per rack: 25 square feet ▪ Depreciation cycle: 15 years ▪ Power use: 80 watts per square foot at USD 0.08 per kilowatt-hour ▪ Busy time: 12 hours per day ▪ Cooling power multiplier: 2.0
LAN, Storage Area Network (SAN), and Cabling	▪ Copper and fiber pre-wiring per rack: USD 3,000 over 10 years ▪ Gigabit Ethernet (GbE) LAN port costs per server: USD 300 over 4 years ▪ Fibre Channel (FC) SAN port costs per server: USD 700 over 4 years ▪ LAN/SAN ports per server: 7 GbE (LAN), 2 FC (SAN)
Personnel	▪ USD 100,000 per support employee per year ▪ One support employee per 250 servers (physical server support only, including installation, break, fix, and de-installation). Virtual machine operating system and application support is not included, as it is the same for all alternatives.

We then normalized performance per watt relative to the server based on Intel Xeon processor X5365. Based on the adjusted power consumption, the server based on Intel Xeon processor X7460 with 96 GB of RAM delivered approximately 30 percent better performance per watt compared with the server based on Intel Xeon processor X7350 with 64 GB of RAM, as shown in Figure 7.

TCO Analysis

We used our virtualization test results to calculate TCO for the two- and four-socket servers in each of our deployment scenarios.

In our TCO calculations, we assumed that we would deploy typical rack-mounted two- and four-socket servers configured with cost-effective 4-GB DIMMs. Two-socket servers typically have eight slots, accommodating a maximum of 32 GB of RAM. Four-socket servers typically have 32 slots, accommodating a maximum of 128 GB of RAM.

Our TCO calculations took into account costs in the following areas:

- Hardware (including configured RAM) and software acquisition; we assumed that the server based on Intel Xeon processor X7460 would cost 10 percent more than the similarly configured server based on Intel Xeon processor X7350.
- Depreciation and amortization
- Data center annual depreciation and operating expenses
- Server support personnel
- LAN, SAN, and cabling

Key TCO assumptions are shown in Table 3.

Virtualization Deployment TCO Scenarios

Each of our three distinct virtualization deployment TCO scenarios reflects a unique set of business requirements, as shown in Figure 8. Based on our test results and assumptions, we estimated TCO for each platform when supporting each scenario.

Performance-centric SLA-focused

In these scenarios, application responsiveness or throughput are paramount. Typical examples include end-of-quarter financial processing, trading applications, and engineering design workloads. IT must meet SLAs that specify response times for interactive applications and completion times for batch jobs.

To meet these SLAs, servers must consistently deliver a specific level of performance or throughput. In our TCO analysis, we therefore limited the number of VMs on each server, based on their relative performance in our tests. For example, the server based on Intel Xeon processor X7460 delivered 1.39x the throughput of the server based on Intel Xeon processor X7350, so we assumed that it could support 1.39x as many VMs.

In our TCO analysis, we assumed that servers were configured with 4 GB of RAM per core, as in our performance-per-watt comparison, to reflect the relative servers' relative throughput. We therefore assumed 32 GB of RAM on the two-socket servers, 64 GB on the server based on Intel Xeon processor X7350, and 96 GB on the server based on Intel Xeon processor X7460.

Compared with the server based on Intel Xeon processor X7350, the server based on Intel Xeon processor X7460 supported about 26 percent more VMs for the same TCO, due to its increased throughput, lower power consumption, and greater performance per watt, as shown in Figure 9.

The servers based on 45-nm technology delivered the best TCO for our performance-centric scenarios.

Memory capacity-focused

In these scenarios, the primary objective is to reduce TCO by deploying VMs as cost-effectively as possible. Typically, this means maximizing the number of VMs per system. Performance is usually a secondary consideration. Examples include systems used for testing and development, and production applications with light transaction loads.

Because performance is secondary, the number of VMs per system tends to be constrained by memory capacity rather than processor performance. Therefore, in our TCO analysis, we configured each server with the maximum amount of memory: 32 GB on the two-socket servers and 128 GB on the four-socket servers, using 4-GB DIMMs. We assumed that each server could support as many VMs as the host server physical memory could accommodate.

We found that the four-socket servers delivered the best TCO because of their 4x greater memory capacity, which allowed correspondingly larger consolidation ratios. The four-socket servers

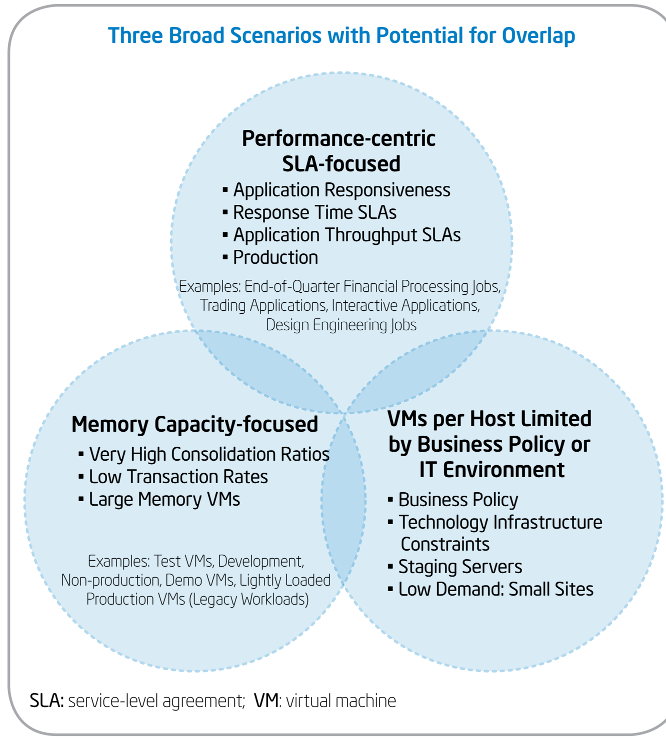

Figure 8. Virtualization deployment total cost of ownership (TCO) scenarios.

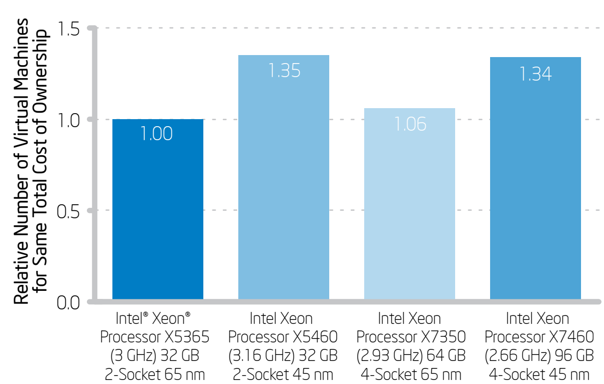

Figure 9. Total cost of ownership (TCO) comparison for performance-centric scenarios. Calculations are based on TCO assumptions as described and Intel internal measurements, August 2008.

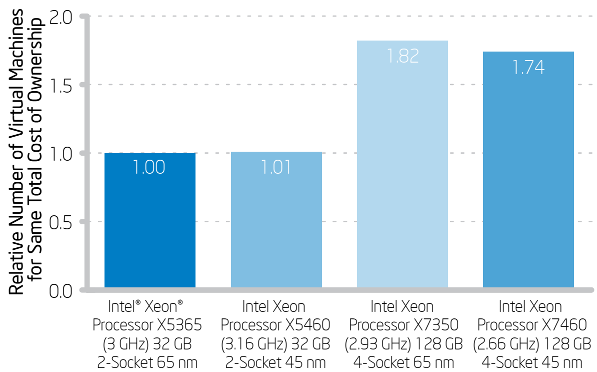

Figure 10. Total cost of ownership (TCO) comparison for memory capacity-focused scenarios. Calculations are based on TCO assumptions as described and Intel internal measurements, August 2008.

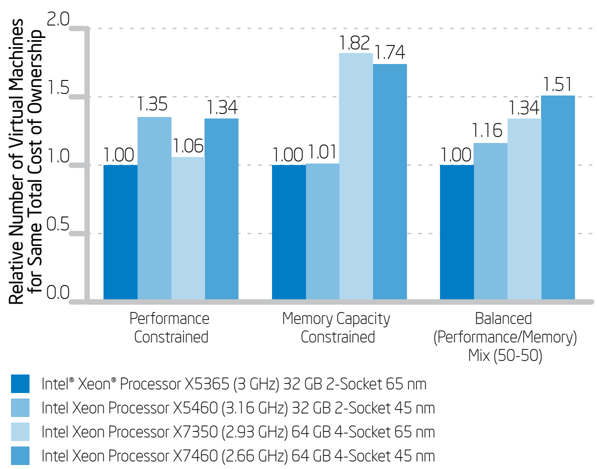

Figure 11. Server virtualization total cost of ownership (TCO) with balanced workloads. Calculations are based on TCO assumptions as described and Intel internal measurements, August 2008.

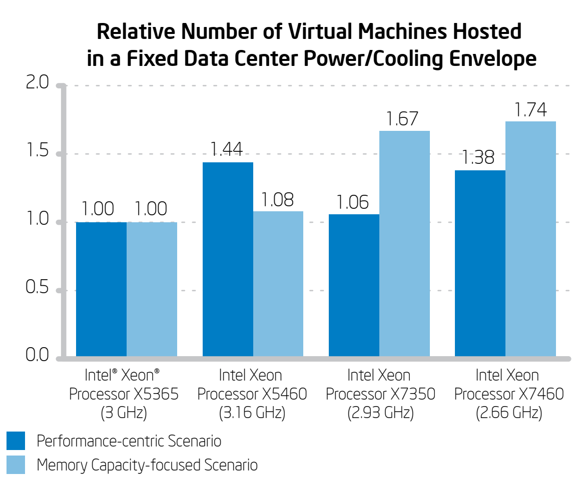

Figure 12. Virtualization server comparison in data centers with power and cooling constraints. Intel internal measurements, August 2008.

based on Intel Xeon processor X7460 could support 72 percent more VMs than the Intel Xeon processor X5460-based server for the same TCO, as shown in Figure 10.

Balance of performance-centric and memory capacity-focused workloads

In situations where a server runs a mix of performance-centric SLA-focused and memory capacity-focused workloads, the server based on Intel Xeon processor X7460 offered a clear TCO advantage, as shown in Figure 11. This suggests that this server would also be a good choice if we do not know in advance which workloads the server will be required to support.

VMs per host limited by business policy or IT environment

In these TCO scenarios, external factors limit the number of VMs on each server. Examples include:

- **Risk management.** To minimize risk, corporate policy may stipulate a limit on the number of VMs per host. This minimizes the business impact if a physical server outage occurs.
- **Backup, network, or other infrastructure constraints.** Backup and restore products may limit the number of VMs per host that can be backed up within an acceptable timeframe.

Other example scenarios include servers deployed at small sites or remote offices, which may have limited processing requirements, or staging servers used for testing new software releases.

Our working assumption was that these deployments do not fully utilize server resources, so two-socket servers would be more cost-effective because of their lower acquisition cost.

Data Center-constrained

In many data centers, the ability to add compute capacity is constrained by power and cooling capacity or by limited LAN and SAN ports. We analyzed which server provided the best fit for each of these common situations.

Power- and cooling-constrained

In data centers facing power and cooling constraints, the servers based on 45-nm processors delivered more VMs per watt for performance-centric scenarios because of their higher performance per watt, as shown in Figure 12. For memory capacity-focused scenarios in these data centers, the four-socket servers delivered more VMs per watt due to their greater memory capacity.

The server based on Intel Xeon processor X7460 was well suited to both of these specialized types of scenario because it combined high performance per watt with large memory capacity. This could enable significant cost avoidance in data centers facing power and cooling constraints by deferring the need for additional data center construction.

Limited LAN or SAN ports

Many data centers are nearing the limit of available LAN or SAN ports. Often, workloads targeted for virtualization are not I/O constrained; in these cases, the server based on Intel Xeon processor X7460 offers significant advantages for both performance-centric and memory capacity-focused scenarios, as shown in Figure 13. Its greater performance per watt and memory capacity enables more VMs to share the same number of limited LAN or SAN ports.

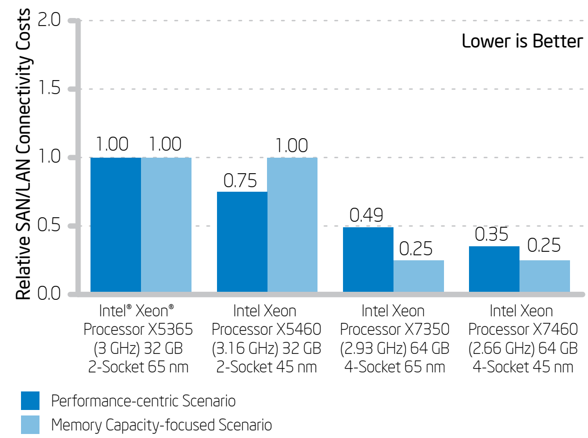

Figure 13. Virtualization server comparison in data centers with limited LAN or storage area network (SAN) ports. Intel internal measurements, August 2008.

Scalability-focused Scenarios

We identified specific scenarios in which features of the virtualization software or the servers themselves favor the use of more-scalable platforms for virtualization.

Large virtualization resource pools

Virtualization host management is typically based on a unit known as a resource pool or cluster; the virtualization management software can use live VM migration to efficiently balance workloads among the servers within the resource pool.

There is a fixed limit on the number of physical servers within a single pool. Because of this, four-socket servers based on Intel Xeon processor X7460 provide significant advantages. Their greater throughput and memory capacity means that the resource pool can include many more VMs, enabling us to use the cluster more efficiently and flexibly. Based on our testing, resource pools using servers based on Intel Xeon processor X7460 could contain about 2.16x as many VMs in performance-centric scenarios, and 4x as many VMs in memory capacity-focused scenarios, compared with using servers based on Intel Xeon processor X5460. This concept is illustrated in Figure 14.

Consistent scalability

In our tests, the server based on Intel Xeon processor X7460 provided much more consistent scaling under increasing load compared with the two-socket servers. There was much less variation in throughput increase between the different workload types as we increased the number of CSUs. This indicates that the server based on Intel Xeon processor X7460 would scale much more consistently when unexpected workload spikes occur.

Memory sharing efficiency

Code and data used by multiple VMs can be stored once and shared among the VMs, saving memory. The hypervisor software we used was able to achieve greater memory sharing efficiency as we increased the number of VMs. When the server based on Intel Xeon processor X7460 was running 10 VMs, each VM used an average of about 1 GB of RAM; when the server was running 30 VMs, each VM used an average of 0.67 GB of RAM, about a third less. This memory sharing efficiency favors the use of larger four-socket servers, which can run more VMs and therefore deliver greater cost savings.

Blade and Specialized Server Considerations

Our TCO analysis is based on mainstream rack-mounted virtualization server configurations. TCO comparisons using blade servers and specialized niche server configurations will depend on the specifics of the server and the target deployment scenario or usage model.

With blades, potential differences include:

- **Memory capacity-focused scenarios.** Depending on the blade design, increasing the physical memory of a four-socket blade may be more expensive than for a four-socket rack-mounted server, resulting in much less difference in TCO between two- and four-socket blade servers in memory capacity-focused scenarios.
- **Data center LAN and SAN constraints.** While four-socket rack-mounted servers offer distinct advantages over two-socket servers in data centers with limited LAN or SAN ports, the extent to which four-socket blade servers offer the same advantage depends on the networking approach used in the blade chassis. With a pass-through networking module, four-socket blade servers continue to offer the same advantage as four-socket rack-mounted servers in supporting more VMs per LAN or SAN port. With a blade switch, the advantage may be minimal.

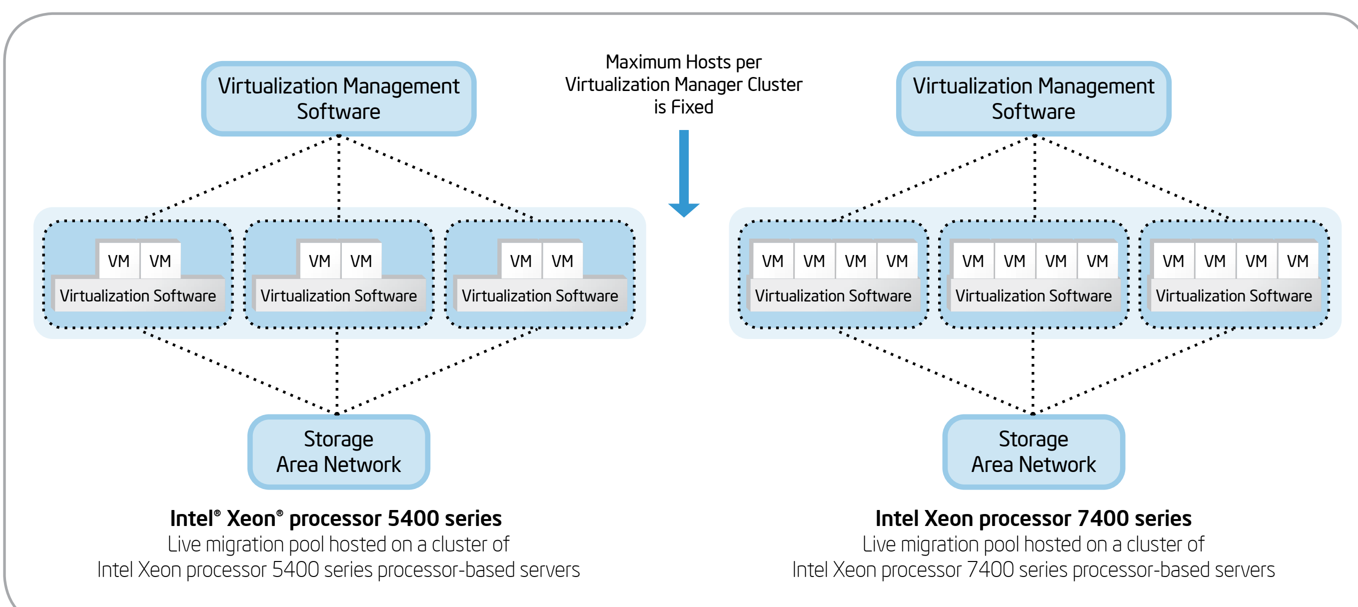

Figure 14. Servers based on Intel® Xeon® processor 7400 series allow larger virtual machine resource pools. Resource pools using servers based on Intel® Xeon® processor X7460 could contain about 2.16x as many VMs in performance-centric scenarios and 4x as many VMs in memory capacity-focused scenarios as servers based on Intel Xeon processor X5460. Intel internal measurements, August 2008.

Conclusion

Four-socket servers based on Intel Xeon processor 7400 series offered significant advantages over other server platforms in almost all the virtualization deployment scenarios that we studied as part of our PoC. The scenarios are summarized in Figure 15.

Our results suggest that servers based on Intel Xeon processor 7400 series would deliver the best TCO in virtualization deployments focused on meeting application performance or throughput SLAs, as well as in deployments where maximizing consolidation ratios is the primary goal. This is due to the server's scalability, power efficiency, and large memory capacity.

These four-socket servers also offer advantages in data centers with power and cooling constraints or limited LAN and SAN ports, by maximizing the number of VMs per watt or per port. Servers based on Intel Xeon processor 7400 series also allow larger VM resource pools, enabling greater efficiency and flexibility in virtualization deployments.

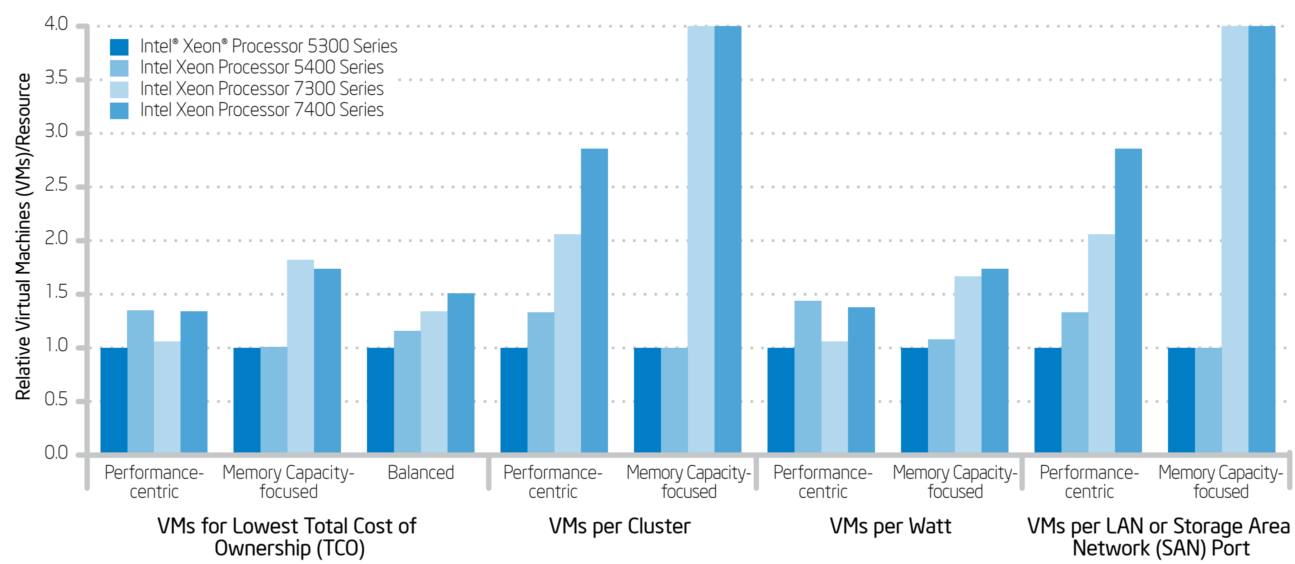

Optimization Goal	Deployment Scenario	Relative Virtual Machines (VMs)/Resource Intel® Xeon® Processor 5300 Series	Intel Xeon Processor 5400 Series	Intel Xeon Processor 7300 Series	Intel Xeon Processor 7400 Series
VMs for Lowest Total Cost of Ownership (TCO)	Performance-centric	1.00	1.35	1.06	1.34
	Memory capacity-focused	1.00	1.01	1.82	1.74
	Balanced (performance and memory)	1.00	1.16	1.34	1.51
	VMs per host limited by business policy or IT environment	High	High		
VMs per Cluster (Resource Pool Scaling)	Performance-centric	1.00	1.33	2.06	2.86
	Memory capacity-focused	1.00	1.00	4.00	4.00
VMs per Watt (Data Center Power-constrained)	Performance-centric	1.00	1.44	1.06	1.38
	Memory capacity-focused	1.00	1.08	1.67	1.74
VMs per LAN or Storage Area Network (SAN) Port (Data Center LAN or SAN Port-constrained)	Performance-centric	1.00	1.33	2.06	2.86
	Memory capacity-focused	1.00	1.00	4.00	4.00

Top Score

Figure 15. Virtualization platform comparison summary. Summary data based on TCO assumptions as described and Intel internal measurements, August 2008.

www.intel.com/IT

Authors

Sudip Chahal is a compute and storage architect with Intel Information Technology.

Raghu Yeluri is an enterprise architect with Intel Digital Enterprise Group, End User Platform Integration.

Acronyms

CSU	consolidated stack unit
DHSI	dedicated high-speed interconnects
FC	Fibre Channel
FSB	front side bus
GbE	gigabit Ethernet
PoC	proof of concept
SAN	storage area network
SLA	service-level agreement
TCO	total cost of ownership
VM	virtual machine

Printed in USA
0908/KAR/KC/PDF

Please Recycle
320069-001US