

## WHITE PAPER

---

# Optimizing Hardware for x86 Server Virtualization

---

Sponsored by: Intel

---

Jean S. Bozman

Gary P. Chen

August 2009

## EXECUTIVE SUMMARY

Currently, virtualization is one of the most talked about new technologies in IT infrastructure. The ability to virtualize servers and to reclaim excess capacity has caught the interest of datacenter managers for many reasons, including more efficient processing and better resource utilization, reduced power/cooling costs, and improved management capabilities. The server virtualization marketplace has been evolving rapidly, and IDC has seen customer adoption of x86 server virtualization mature rapidly — with the average number of virtual machines (VMs) per physical server increasing, the types of workloads trending toward enterprise applications, and the hypervisor being leveraged for extended use cases such as high availability and disaster recovery. This changing marketplace has manifested itself with new requirements for virtualization products as follows:

- ☒ **Support for hypervisor performance.** Hardware virtualization acceleration, improved memory performance, and increased VM density are a must for maintaining service-level agreements (SLAs) in the enterprise.
- ☒ **Support for virtualized I/O.** Reduced I/O latency, resource contention, and increased throughput are also required to maintain SLAs.
- ☒ **Support for enhanced security.** Hardware enforcement of VM isolation, hardware protection of I/O streams and data, and support for energy efficiency in the datacenter allow for a balanced system approach, offering a high ratio of performance per watt.

Broadly, IDC looks at server virtualization platforms as comprising four general components:

- ☒ **Virtualization platform.** This includes both the core software and hardware platform. The software is built on the core hypervisor, basic resource controls, and application programming interfaces (APIs). Competitive differentiation includes number of sockets, number of processors in a VM, number of guests supported by the license, and operating system (OS) support. Competitive hardware differentiation includes CPU architecture, number and type of processors supported, chipset features, memory capacity, and hardware-enhanced virtualization feature set.

- ☒ **Virtual machine management (VMM).** This includes the host-level management as well as management across virtualized servers and datacenters. Today the differentiator between vendors is whether they offer virtualization management and at what scale.
- ☒ **Virtual machine infrastructure (VMI).** These are the value-added features that drive most customer purchasing decisions today and represent the area with the biggest gaps between vendors. They include live migration, automatic restart, and workload balancing of virtual machines across hosts.
- ☒ **Virtualization solutions.** These represent the bundling of the aforementioned technologies with some enabling workflows and process automation capabilities to meet specific business needs. These solutions are just now emerging and likely represent an important future competitive front. This category includes solutions such as VDI or disaster recovery.

---

## SITUATION OVERVIEW

---

### **Virtualization: A Mainstream Technology**

Virtualization in computing is a broad term that refers to the abstraction of computer resources. It includes making a single physical resource (such as a server, an operating system, an application, or a storage device) appear to function as multiple logical resources, or it can include making multiple physical resources (such as storage devices or servers) appear as a single logical resource. Early virtualization technology, as used in mainframes or host servers, carved up a single, large computing resource into partitions running multiple workloads or jobs, as they were frequently called. This was accomplished using both hardware and software implementations that were tightly coupled in a proprietary, vendor-specific environment. Today, in x86 platforms, technologies in processors and hypervisors are leveraged as industry-standard components that can be incorporated into products from many server OEMs worldwide.

---

### **Virtualization Moving into the Mainstream on x86 Server Platforms**

Software for implementing "virtual machines" (hypervisors) and virtual server infrastructure is being rapidly adopted on x86 server platforms. Today's state-of-the-art technology for x86 virtualization is also progressing in terms of use cases, technology, and maturity.

The first phase of customer adoption of virtualization is really a continuation of a trend in the industry that began back in 2000. Predominantly, this phase involved IT simplification. Following the economic downturn of 2001–2003, customers recognized that there was a need for datacenter consolidation, physical server consolidation, and asset inventory.

Physical server consolidation began to merge with virtualization, and customers began to do legacy rehosting of nonsupported operating environments, such as Microsoft Windows NT4, to get the benefits of the new hardware power that was available to them. Prior to this wave of consolidation, many Windows NT servers were running in a standalone, dedicated mode, in which one operating system was dedicated to one application or database, due to limited physical resources — such as memory or I/O.

Early on, x86 virtualization focused on testing and development environments, allowing these environments to be isolated within VMs, and allowing many programmers to share a single physical server while viewing their own virtual programming space on that server. In recent years, adoption of virtualization has been rising rapidly for production applications — and IDC demand-side, customer-based research shows that the average of two to four VMs per physical server in 2005–2006 grew to an average of eight or more VMs per physical server in 2008. Some intensive users are deploying 10 or more VMs per physical server, depending on the types of workloads and the amount of server resources required to run them.

The current phase of virtualization adoption is leveraging live migration of VMs — for planned upgrades and repairs — and also leveraging VM migration for purposes of high availability and disaster recovery, when coupled with storage solutions such as data replication.

These live migration scenarios for VMs join a spectrum of HA solutions — which also include traditional high-availability software to protect workloads running inside VMs — and fault-tolerant solutions that "mirror" the state of applications across the network.

In summary, the combination of multicore processors and virtualization is making it possible to consolidate workloads that have been running on underutilized servers. IDC believes that we will continue to see a diversification in the types of reasons why customers deploy virtualization software — whether it is around planned migrations or business-level, mission-critical high availability. The next phase of customer adoption will likely be centered on the concept of automation or a utility computing environment whereby services are delivered based on policies and are moved from a fixed-cost model to a variable-cost model within the organization. While the first phase of virtualization focused on capex, the next phase of virtualization will focus on operational cost reduction.

---

## **The Benefits of Server Virtualization**

Virtualization offers a myriad of benefits to an enterprise. The capex savings from server consolidation is an initial benefit from virtualization that is now well understood by most customers. The market is evolving to leverage virtualization for more than just consolidation, utilizing hypervisors as the foundational technology for many other uses cases that drive down opex costs. High availability, fault tolerance, disaster recovery, and workload balancing are several top use cases expected to gain significant traction in 2009 and 2010. While these features have been available in the datacenter for some time, they can be achieved less expensively and more easily via virtualization in many cases. The next generation of

virtualization brings in advanced management tools for greater automation and orchestration of the datacenter to reduce operational costs and improve service levels. These tools offer the following benefits:

- ☒ The ability to rapidly save, copy, and provision a virtual machine that enables zero-downtime maintenance and supports new "go live" initiatives
- ☒ A dynamic sharing of idle resources across server platforms, resulting in improved performance and use while eliminating stovepipes
- ☒ Higher technology standardization and currency (up-to-date systems), resulting in lower operations and maintenance costs
- ☒ Seamless failover when a virtual server component fails, resulting in higher system availability
- ☒ Reduced complexity, resulting in improved logical and physical disaster recovery
- ☒ Easy scalability for applications by dynamically moving applications to a server with more resources or by spinning up new instances
- ☒ Resource optimization to drive great efficiency in the datacenter by automatically balancing workloads, resource pooling, and creating multiple virtualization tiers
- ☒ Cloud computing with the creation of internal clouds and enabling the ability to interface to external clouds

### ***Server Virtualization Usage and Adoption***

More than 50% of today's virtual servers support production workloads, and 60% of all VM spending is for business processing and decision support workloads. IDC forecasts that by 2010, customer spending on server virtualization will amount to some \$20 billion, a 68% increase over five years.

---

## **Virtualization Platform Requirements**

The platform requirements for virtualization in support of multiple applications in the business environment involve a mix of software and hardware considerations.

### ***Software Requirements for Server Virtualization***

In terms of the core virtualization platform, today there are more similarities between all of the software vendors than differences. All have the ability to manage processor, memory, network, and disk resources. All have support for at least 16 processor cores in a host and either at least two or four virtual CPUs per VM, which covers the vast majority of x86 systems. All support both Microsoft Windows and Linux operating environments, and some offer support for the Solaris Unix environment as well.

The four intangibles that can differentiate the virtualization software offerings are channel to market, performance, reliability, and advanced management. The operating system vendors, until recently, have had an easier route to market because customers are/would be getting virtualization capabilities "free" with the purchase of

the base operating system. Over time, this advantage has been diminished as the hypervisor vendors have gone directly to the server OEMs to integrate virtualization directly with the hardware.

With this virtualization model, customers can partition the server into an unlimited number of virtual machines and, more important, get this capability at a much lower price than if they had purchased a standalone hypervisor.

Performance and reliability become greater considerations as enterprises virtualize more of their infrastructure and more of their most mission-critical, performance-sensitive applications. By virtualizing their servers, customers are making a greater percentage of their datacenters "mission critical," and reliability becomes essential. Here incumbents such as VMware have the advantage of time. They have been in the market the longest and have a strong installed base that is vocal about the reliability of the company's platform. The other vendors are still very much demonstrating the robustness of their platforms and are challenged to create the early customer buzz to get the word out and influence others.

In terms of performance, the rest of the market has become much more competitive with VMware, with new hypervisor entries from Microsoft, Citrix, and others. One area of differentiation that remains is I/O performance, a key aspect when virtualizing certain types of workloads such as databases. As the virtualization vendors seek to expand their footprint and virtualize the remaining applications previously thought to be poor candidates, performance will be a focus of customer purchase decisions.

The greatest differentiator for the virtualization market will be advanced management. As customers mature in their virtualization implementations, they will look to move beyond simple consolidation. Hypervisors will be leveraged for high availability, disaster recovery, load balancing, resource optimization, automation, and orchestration, all of which require sophisticated management tools. In addition, as the install base of virtualization gets to a certain point, specialized virtualization management becomes a prerequisite.

In terms of the hypervisor platform, the differences between all of the instantiations of the technology are diminishing rapidly — and reliability, performance, ease of acquisition, ease of deployment, and most importantly, management functionality matter more than the "speeds and feeds." This is why IDC believes there is a coming battle for real estate in the virtualization market.

### ***Hardware Requirements for Server Virtualization***

While the purpose of virtualization software is to effectively divide and use the hardware to run multiple applications and increase the usage of the computing resources, it is limited to the capabilities of the system or systems on which it is running. If one looks at the major components that virtualization software controls — CPU, memory, network, and disk resources — it becomes easy to see the importance of having the right balance of resources.

As system loads increase from between 5% and 20% to between 40% and 50% or more, there is a greater demand for virtualization-optimized systems with higher levels of processor performance, more processor cores, increased memory capacity and speeds, expanded I/O capabilities, lower latency, higher-throughput network bandwidth, and accessibility to networked storage resources. Virtualization not only increases the workload but also segments the workload into multiple pieces that increase the demand for more resources running at the same time. As virtualization deployment becomes more complex and hosts increasingly mission-critical applications, customers are requiring servers that contain these virtualization-optimized attributes. Today, about two-thirds of all new virtualization licenses are deployed on new server hardware, an indication that the market clearly demands a new class of server for its virtualization projects.

---

## Choosing the Right Hardware Platform

Choosing the right hardware platform for server virtualization is just as important as choosing the right virtualization software. To make the proper choice, one must consider the following:

- ☒ **Performance.** Increased workloads require more system performance to maintain desired service levels and application response times. More performance and throughput can be achieved by using new processor architectures with multiple cores per processor and systems designed to utilize this additional performance and functionality. In addition to the increase in performance needed to run more workloads in native mode is the added resource requirement for running the VMs. The hardware assist features in processors and chipsets play a key role in minimizing hypervisor overhead, especially in the area of I/O, which has been a performance limiter in the past.
- ☒ **Energy efficiency and space.** Power and cooling are most understood when cost or physical constraints affect an end user's ability to function normally; for example, when costs exceed budget or power requirements exceed available resources. Floor space can be a huge issue when an expanding IT footprint results in building out datacenter or IT space. Virtualization and consolidation help reduce hardware footprint and energy utilization. More energy-efficient systems further reduce power requirements, both saving power and lowering operating costs. Systems with reduced power requirements and or higher performance per watt aid IT in controlling energy and space requirements.
- ☒ **TCO/ROI.** Since the mass adoption of x86 servers, systems management costs have grown significantly. The increasing operating costs have resulted in diverting resources and capital from initiatives aimed at driving innovation and increasing the value of IT. Using a virtualized environment has the potential to increase system utilization, lower power and cooling requirements, decrease space requirements, and simplify operations, reducing system management requirements.
- ☒ **Optimized platform features.** Component and systems vendors are increasingly adding features that are optimized for virtualization. For example, both Intel and AMD have added virtualization capabilities to their processors.

These capabilities both ease the design of robust virtualization software and reduce the performance overhead typically seen by applications running in a virtualized environment. They can also provide security features such as I/O isolation and crucial features for mission-critical applications such as QoS to ensure certain VMs get priority.

- ☒ **Industry collaboration.** Virtualization on x86 is best characterized by looking at the industrywide collaboration that has occurred. Because virtualization involves multiple levels of hardware and software, hardware component and systems vendors, operating system and application vendors, and virtualization vendors have collaborated to create the first working solutions. Now the industry players are optimizing for virtualization on their new platforms.
- ☒ **Compatibility.** Maintaining compatibility across platform generations is a key attribute for future systems. Moving VMs from one server to another requires architectural considerations so that IT does not create islands of virtual resources but actually maximizes flexibility within its infrastructure. New platform features now allow live migration of VMs to a wider range of hardware.

## **FUTURE OUTLOOK: x86 PLATFORMS AND VIRTUALIZATION ADOPTION**

Virtualization brings a set of capabilities to the x86 platform that were originally developed in mainframes and then added to RISC-based machines. These capabilities have now become available on mainstream servers.

The worldwide server market has undergone significant technology shifts over the past two years, which have led to marked changes in customer buying behavior. In recent years, both multicore technology and server virtualization capabilities were introduced to the x86 server market, and they are now well established in IT organizations in the United States and throughout the rest of the world.

---

### **Impact of Virtualization and Multicore**

Each of these technologies is impactful to the market in its own right. However, the use of multicore technology in conjunction with server virtualization tools has a compounding impact on server configurations and accelerates the ability of IT organizations to exploit the benefits of multicore technology as follows:

- ☒ The introduction of multicore processing and virtualization into the x86 server segment will significantly disrupt the worldwide server industry. Multicore will increase the effective processing capacity in the market through relatively fewer systems.

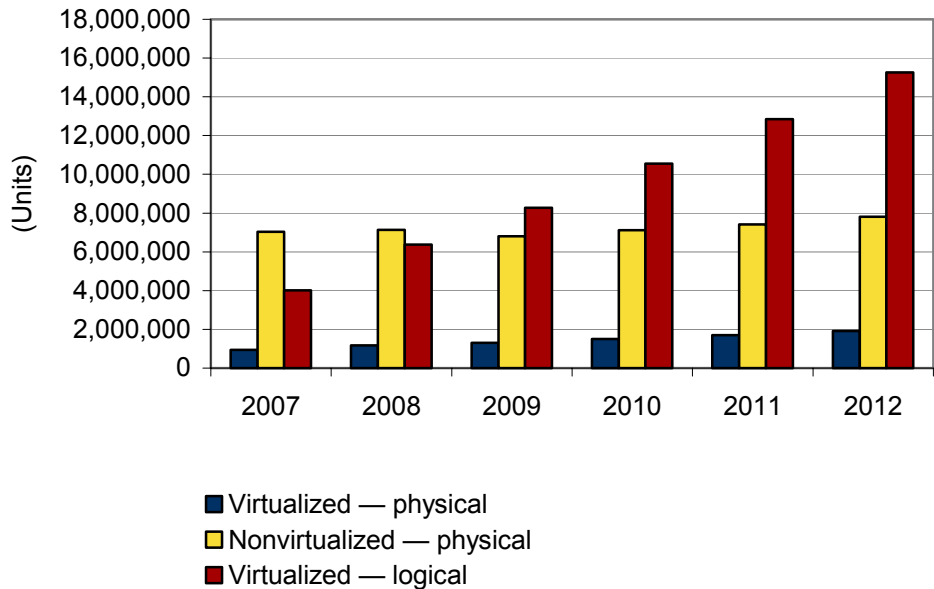
Virtualization enables customers to raise system utilization rates, consolidating servers within their IT infrastructures. Based on IDC's 2008 virtualization study, overall x86 utilization, which was recorded at 36% before virtualization, is currently shown at 56% with virtualization. The planned increase with virtualization will take utilization to 67% over the next two years. This projection is in line with growth prospects for virtualized servers.

- ☒ Customers are buying richer configurations in servers with multicore processors, larger memory, and more redundancy to increase VM density and provide better insurance against hardware failure, which becomes more disruptive due to consolidation.
- ☒ While most customers are using virtual machines as part of their production and test environments, more than half use VMs as part of their enterprisewide high-availability strategies and another 47% are using VMs in their disaster recovery environments.

Figure 1 shows the rapid growth of both virtualized servers and the logical server units that run on top of the physical server hardware. While the share of physical virtualized servers is projected to remain relatively flat at about 8% from 2007 through 2012, the share of virtualized logical units (or virtual servers running on the virtualized server platforms) is expected to increase from 33% of total logical units in 2007 at a 31% CAGR to reach 61% in 2012. This compares with a projected CAGR of 2% for physical server units. The number of logical servers generated on virtualized servers will surpass the number of nonvirtualized physical server units in 2009.

**FIGURE 1**

Virtualized Server Versus Nonvirtualized Server Forecast, 2007–2012

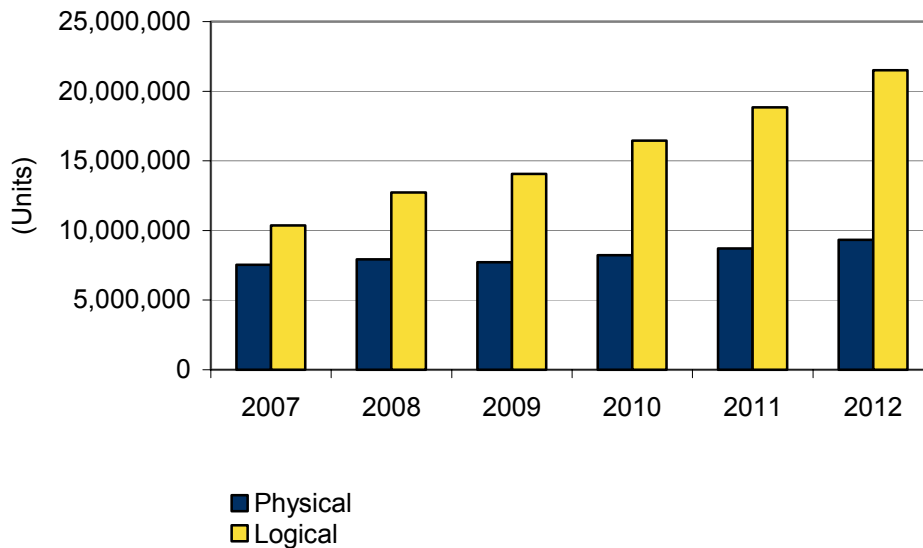


Source: IDC, 2009

Figure 2 compares the projected growth of total physical server units with the projected growth of total logical server units (the virtual servers running on top of the physical servers that are virtualized). Physical servers are expected to increase at a 6% CAGR from 2007 to 2012, while logical servers (running on top of the physical servers) are expected to increase at an 18% CAGR.

**FIGURE 2**

Physical Server Versus Logical Server Forecast, 2007–2012



Source: IDC, 2009

### **Platform Choice: System Scaling for Virtualization and Increased Utilization**

IDC conducts ongoing research to look at how systems have been deployed. IDC has seen virtualization drive more robust systems, increasing the number of processors and the amount of memory, I/O, and disk capacity.

IDC believes this change is driven by the higher workload activity running on each system. This generally occurs through consolidation of multiple applications or increasing individual application workloads. For example, a system that was running at 10% of capacity may have been upgraded to run at 50% of capacity. Maintaining a balanced system with adequate resources may require additional processors, memory, I/O, and disk storage capacity. The increase may occur in only one of the mentioned areas, or it may occur in several or all of the areas.

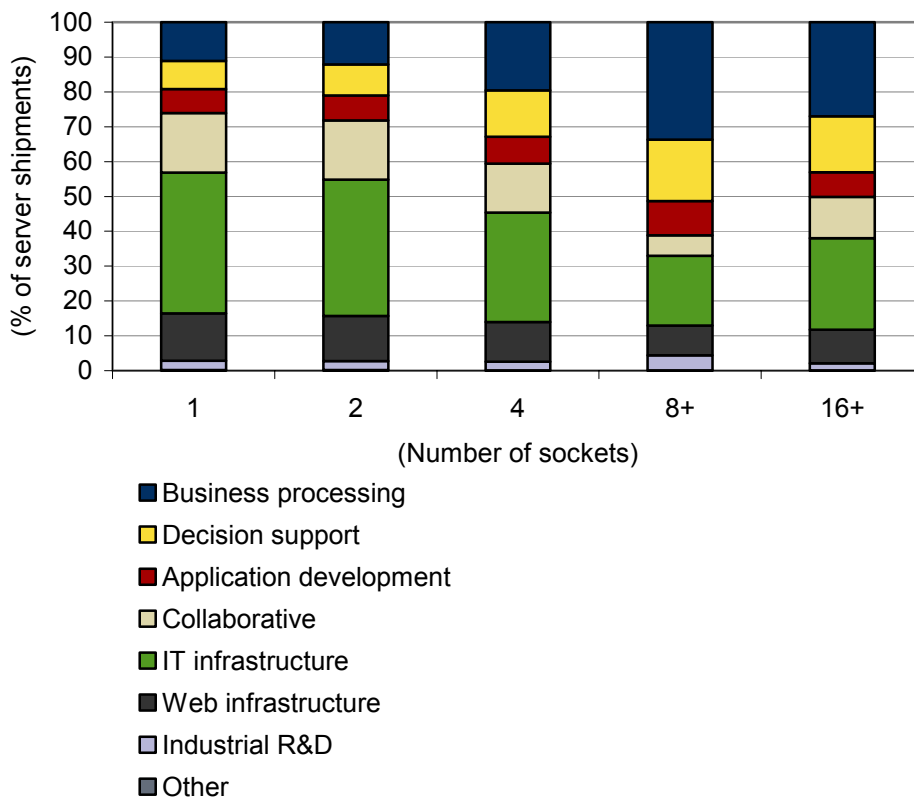
The change is also driven by increased compute capability provided by multicore processors and the need to add more memory to fully utilize this added compute resource. Virtualization becomes a key tool on the software side to leverage this added capability as well.

In short, IDC is seeing end users buy more fully configured systems to balance the amount of work that can be done and efficiently utilize the systems' capabilities. Virtualization has become a major tool to maximize the utilization of these resources. It is also driving the purchase of larger systems, even with new, more powerful systems coming on the market. As user adoption of virtualization increases, denser computing and support for more demanding enterprise applications result.

System size and configuration are driven by the size, type, and number of workloads running on a system or set of systems (see Figure 3).

**FIGURE 3**

Server Shipment Share by Socket Number and by Workload, 2008



Source: IDC, 2009

IDC Workloads data is demand-side (customer-based) research data that clearly shows the "profile" of workloads running on servers. This annual study of more than 1,000 IT sites takes a "snapshot" of all workload types and shows how they are run on a variety of server platforms and operating environments. IDC's sample of IT sites studied crosses all major vertical industry categories (e.g., financial services, government, manufacturing, retail) and all business sizes, including small (fewer than 100 employees), medium sized (100 to 1,000 employees), large (more than 1,000 employees), and very large (more than 10,000 employees). IDC Workloads data has

several large segments, as shown in Figure 3 (workloads are color-coded in the figure's legend), including business processing (e.g., OLTP, line-of-business [LOB] applications such as ERP and CRM); collaborative (e.g., email and groupware); decision support (including database analysis); technical (e.g., scientific and engineering workloads and market simulations); Web infrastructure (e.g., Web serving; proxy, caching); IT infrastructure (e.g., network support, file/print); and industrial R&D, among others. Database workloads underlie several of the named workloads, including business processing and decision support.

Systems with four or more sockets have a higher percentage of business processing and decision support workloads than one- and two-socket systems. More than two-thirds of today's virtual servers support production workloads, and 38% of all VM spending is for business processing workloads. Revenues for decision support, Web infrastructure, and IT infrastructure are increasing at a rate that exceeds that of other workloads. Virtualization tends to drive larger systems with richer configurations, and these workloads generally run on larger systems.

---

## **Current Server Virtualization Trends**

For much of the industry, the most exciting story around virtualization software has focused mostly on server consolidation and reducing physical footprints and the resulting hardware acquisition cost reductions, resulting in lower power and cooling costs.

During 2009, watch for customers to begin leveling their virtualization road maps from a relatively simplistic technology play to a more comprehensive and longer-term end-to-end adoption road map that will impact servers, clients, storage, and even network architectures over the next decade. Simply stated, virtualization begins a process that will extend this technology phenomenon far beyond its current general role as a tool for server consolidation.

Among the bigger-picture benefits CIOs and IT managers will demand are an integrated approach to business continuity, fault-resistant (and -tolerant) and high-availability architecture designs, a reduction in unplanned downtime, and vastly increased business agility that allows companies to become better able to address emerging market opportunities and to do so more quickly than less agile competitors.

Live migration, or the ability to move virtual servers to alternate computing resources while running, was popularized by VMware with its VMotion product. Used most often for planned downtime scenarios, applying upgrades and security patches, or repairing the physical server on which the VM resides, live migration is being embraced by customers as a useful and important element of improving high availability for their virtualized x86 environments. IDC believes that live migration will be further improved over time by the reduction in any "gap" in processing time so that workloads will be moved from resource to resource in a seamless way, with no perceptible interruption for end users accessing the migrating virtual machine.

Live migration will become a key building block of the next-generation dynamic datacenter, which will automatically optimize resources and be able to react in real time to the needs of the business. Server virtualization will become more tightly integrated with network, I/O, and storage virtualization, creating a fully virtualized

datacenter that is governed by automated policies to achieve unprecedented levels of efficiency and service. Virtualization will enable an internal cloud delivery model and allow for easier interfacing with external clouds.

---

## **Virtualized Platform Characteristics: A Review of Intel Features**

On March 30, 2009, Intel introduced the Nehalem EP quad-core processors, based on the Nehalem microarchitecture and formally named Intel Xeon processor 5500 series. These processors include an integrated memory controller and a new Intel QuickPath Interconnect (QPI) link that replaces the front-side bus (FSB) for on-processor I/O.

As a result of these and other design changes, including improved support for hyperthreading (HT) of software applications, Intel said that the Xeon 5500 series processors are in the range of two to three times faster than the Xeon 5400 series processors that preceded them. They also have energy-efficiency features that result in power/cooling savings of 40% to 50%, compared with Xeon 5400 series processors, and improved support for virtualization, via the Intel Virtualization Technology (Intel VT).

### ***Intel VT***

Intel VT, which is found in several separable components, improves connections to the physical system's I/O and access to networked devices. Intel VT connects the requests from the virtual machine monitors to the physical devices, forging a link between logical servers (virtual machines issuing requests) and the virtualized systems' physical server resources, which are directly attached to the outboard storage and network resources. Intel VT simplifies management of onboard resources — and provides a connection to outboard resources, including storage and network links, at near-native speed. This is important because virtualized I/O is a key factor in performance, especially for data-intensive workloads such as databases and enterprise LOB applications.

### ***Intel VT for Processors (Intel VT-x)***

Overall, Intel Virtualization Technology for processors is called Intel VT-x technology, which allows the operating system inside the VM to run at its original privilege level, even though there are abstraction layers involved in virtualization. In addition, VT-x includes a variety of hardware assists that reduce virtualization overhead to increase VM performance. The latest additions to Intel VT-x include three key elements: FlexMigration, FlexPriority, and Extended Page Tables (EPT).

### **Intel VT FlexMigration**

This feature provides compatibility with servers based on earlier Xeon processors, dating back to the Xeon 5100 (Woodcrest) and the Xeon 5500. In addition, the compatibility pool includes servers based on Xeon 7300 and Xeon 7400 processors. This ability to move VMs between Xeon-based x86 servers of various ages is especially important in sites with large pools of physical x86 server resources. In this way, VMs can be seamlessly moved from one server to another, within the installed base, for proactive repair or software upgrade purposes. Equally important is the support for VM migration for purposes of availability during planned downtime — and for purposes of preparing for disaster recovery scenarios.

### **Intel VT FlexPriority**

This feature optimizes performance in response to system "interrupts" that come from other devices or applications that need attention. FlexPriority virtualizes a register, called the APIC Task Priority Register (TPR), that monitors and manages the priority of tasks. This capability is especially important for Windows 32-bit workloads, which make frequent use of the TPR; the 32-bit workloads, many of them carried forward from older x86 processors, can run on the Xeon processors, which support both 32-bit and 64-bit workloads.

### **Extended Page Tables**

The Extended Page Tables feature increases virtualization performance by providing a hardware page table that does the work of translating between the guest address of a VM and the physical address. This increases memory performance and reduces hypervisor overhead, especially for applications that are memory intensive, such as databases. It is a new enhancement to VT-x with the release of the Xeon 5500 processors.

### ***Intel VT for Connectivity and Directed I/O (Intel VT-c and Intel VT-d)***

As they are delivered into the marketplace, several Intel virtualization features are seen by OEMs and customers as supporting virtualization technology, including Intel Virtualization Technology for Connectivity (Intel VT-c) for virtualized network connections and Intel Virtualization Technology for Directed I/O (Intel VT-d), which efficiently routes I/O requests to specific hardware devices, such as memory.

- ☒ **Intel VT-c.** This technology provides hardware assistance to network devices via Virtual Machine Device Queues (VMDq) technology, which reduces the compute overhead of sorting network packets to the appropriate VM and Virtual Machine Direct Connect (VMDc) technology, which allows the network device to be partitioned into multiple virtual devices and allows capacity and priority to be allocated to these partitions.
  
- ☒ **Intel VT-d.** This technology is instantiated in the chipset to ensure reliability and protection by means of device isolation. It works by directly assigning devices, such as memory addresses, to specific I/O streams, thereby reducing the I/O overheads in the virtual machine monitoring software. By doing so, VT-d reduces system overhead associated with virtualization — and provides improved system performance.

The Xeon 5500 processor's integrated memory controller allows users to increase the density of VMs running on each physical server, reducing latency that would impact the overall performance of the virtualized computing environment. IDC demand-side, customer-based research shows that the average number of VMs per physical server is increasing, over time, as adoption of virtualization in the x86 server space continues to rise. With VM densities having risen to an average of eight or more in 2008, customers will need to ensure performance for the workloads running with VMs and managed by software hypervisors that are running on the physical hardware platform.

## Optimizing Hardware for Virtualization

The software on x86 platforms generally runs regardless of system size, so the focus here is on some differing capabilities in the hardware system. Size matters — larger systems run larger workloads and can support more VMs. This gives IT the ability to run larger and more workloads and also more flexibility to manage its systems to meet changing needs and scalability. Hardware is being optimized to run virtualization better and to provide more control and easier management of VMs and workloads. Details of the Intel technology supporting virtualization include the following elements:

- ☒ **VT-x.** Baseline virtualization capability in the processor simplifies software and enables mixed 32- and 64-bit operating environments.
  - ☐ **VT FlexMigration.** This technology eases migration between multigenerational platforms (x86 processors made in different generations of Intel technology), allowing VMs to be moved more freely among a "pool" of processors in a virtualized IT infrastructure.
  - ☐ **VT FlexPriority.** Interrupt virtualization reduces interrupt overhead.
  - ☐ **Extended Page Tables.** Memory/page table virtualization reduces memory/page table overhead and is available on Xeon 5500 processors.
- ☒ **VT-c.** This technology supports network connections at near-native link speeds with two main features, as follows:
  - ☐ **VMDq.** Optimization in the NIC reduces server overhead, improves LAN performance in a virtualized environment, and reduces hypervisor overhead associated with network processing by offloading it to the I/O silicon (the network processor on the NIC).
  - ☐ **VMDc.** This technology enables VMs to access the network directly using the PCI-SIG SR-IOV standard, which provides near-native network performance, improving VM scalability and enabling flexibility and mobility.
- ☒ **VT-d.** I/O virtualization provides reliability and protection through device isolation, and performance through direct assignment of devices, which results in reducing the I/O overheads in the virtual machine monitor.
- ☒ **Virtual Processor ID (VPID).** This technology assigns a different VPID to each virtual processor to tag translations in the translation lookaside buffer (TLB). This prevents the need to flush the TLB on each VM entry and exit, improving performance.

### ***Processor Architecture Enhancements***

The Intel Xeon 5500 processors, based on Intel's Nehalem microarchitecture, provide improved support for virtualization capabilities. Improved memory use and improved I/O speeds — along with overall improved energy efficiency (compared with previous Xeon processors through the Xeon 5400 series) — will be leveraged to improve performance of workloads running on a virtualized x86 server.

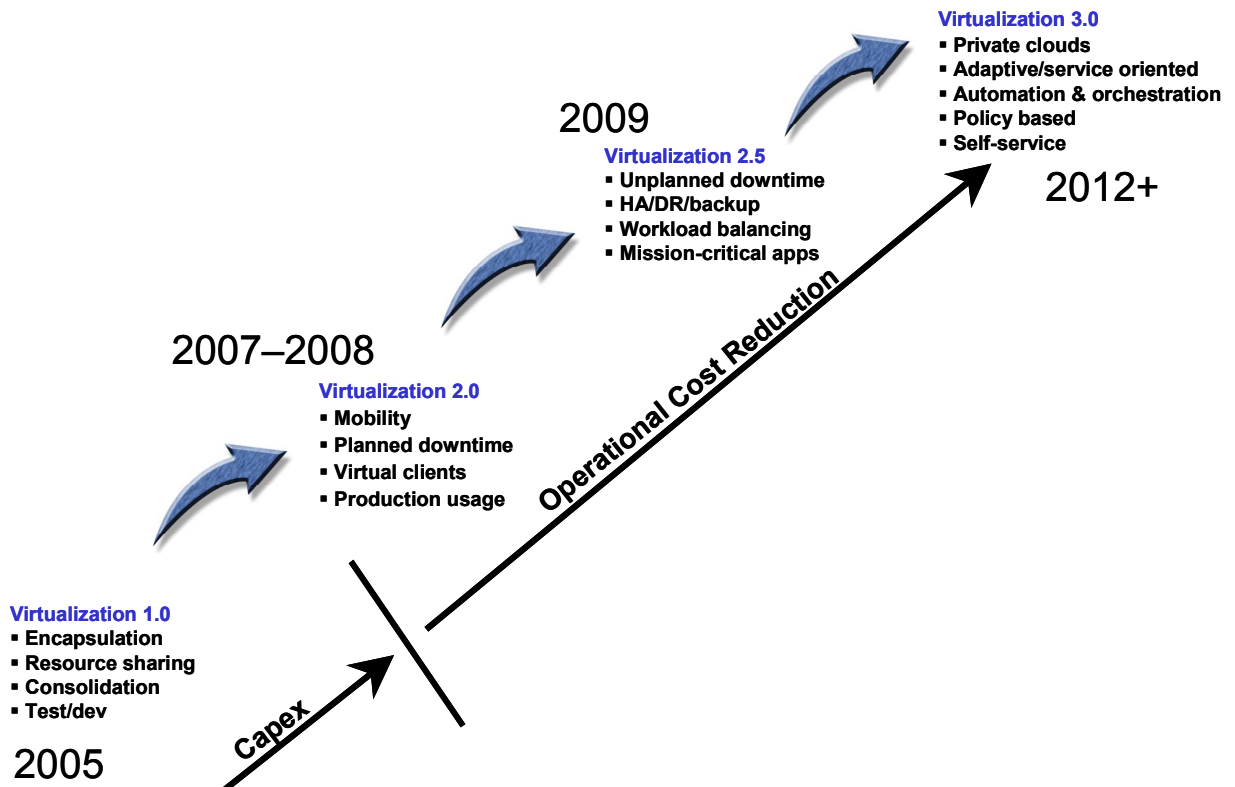
The key to these added features is a platform that runs virtualized workloads close to the speed of native workloads running on the physical server machine. In short, these features are designed to eliminate overhead and to increase efficiency. When combined with Xeon 7400 MP processors, the VT series of features has allowed Intel to produce numbers on virtualization benchmarks that show current generation systems outperforming the Intel Xeon 7100 dual-core processors by more than 150%. From another perspective, a four-socket system will support 2.5 times the number of VMs that a two-socket system can support. This illustrates the scalability gained by a four-socket Xeon-based system.

## The Future of Virtualization

As the market moves toward Virtualization 2.5 and 3.0 (see Figure 4), higher-level capabilities such as high availability, disaster recovery, workload balancing, and automation are gaining more attention. IDC is seeing more companies move into or at least toward 3.0 capabilities and requirements. The previously mentioned hardware capabilities are becoming a baseline requirement.

**FIGURE 4**

### Virtualization Milestones



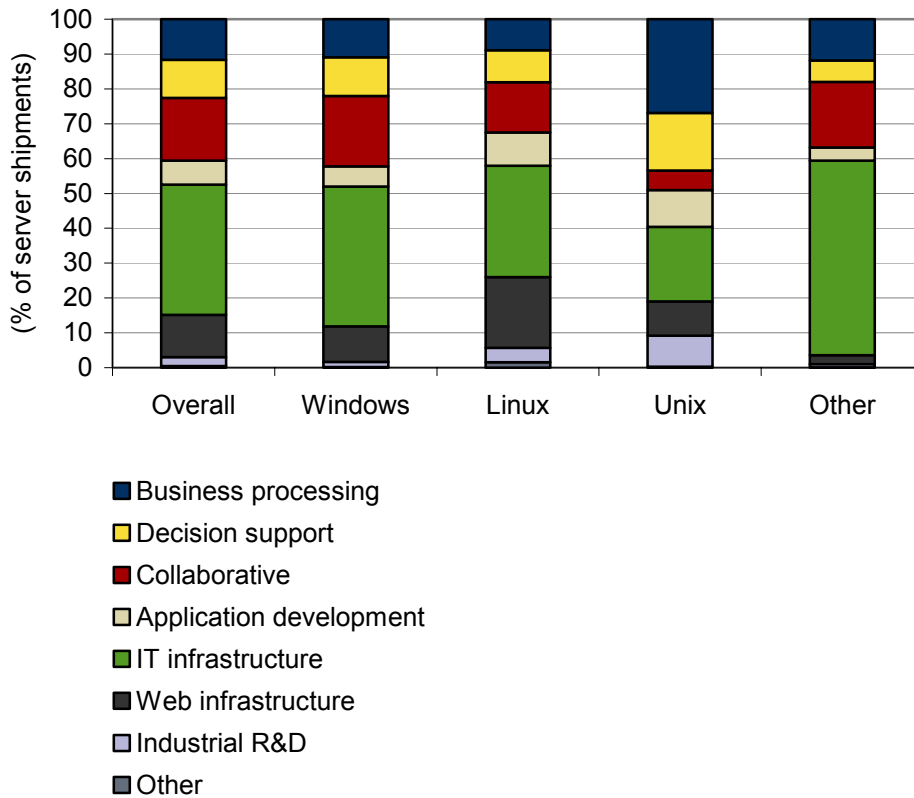
Source: IDC, 2009

## Workloads That Benefit from Virtualization

In IDC's annual survey of end users regarding workloads, virtualization was used to support all workload segments. IT infrastructure is the most popular workload, but highly critical business applications, such as business processing and decision support, make up more than 20% of the workloads (see Figure 5). This number increases for Unix and other platforms. Web applications are more likely to be virtualized on Linux platforms relative to other operating environments.

**FIGURE 5**

Server Workloads by Operating Environment, 2008



Source: IDC, 2009

## CHALLENGES/OPPORTUNITIES

While virtualization is now considered to be a mainstream technology in servers, it is still a mystery to many IT shops that are using unvirtualized servers — or that are considering an upgrade from servers instated three to five years ago, particularly in the midmarket, where IT staff resources are more limited than those in large enterprises. Intel, VMware, and others engaged in the market still need to educate end users and provide guidance on adoption strategies. In addition, those who have done server consolidation projects and are moving on to extended virtualization use cases and ultimately the next-generation dynamic datacenter will need proof points that virtualization can serve as the foundation for these scenarios and that new levels of efficiency are achievable.

The mix of workloads and the usage model for running those applications are different for every IT shop. Standardizing environments and operations is on the to-do lists of most IT leaders. This is both a challenge and an opportunity. It is a challenge to create a working plan and execute on that plan. It is an opportunity to increase the market for x86-based systems as a foundation for standardization in virtualized environments.

Virtualization is a hot topic in the IT industry and in customer sites. It is a growing market where leadership platforms and technologies are valued. Virtualization is accelerating purchases of new, multicore servers as platforms for workload consolidation — and it is driving transitions to newer systems. Intel needs to continue educating the market on its capability and ability to lead with innovations that will help to enhance the market and improve virtualization performance.

## CONCLUSION

Virtualization technology is on a roll, driven by adoption on x86-based servers. While the software market ecosystem around virtualization is heating up, hardware vendors are leveraging that momentum and creating platforms that are optimized for virtualization. Importantly, virtualization is a springboard for cloud computing services — in which workloads can be easily provisioned from a "pool" of virtualized resources and scaled up rapidly, depending on requests for data services. IDC expects cloud computing services to be provided both inside enterprises (private clouds) and outside enterprises (public clouds) — and even to be a mix of both (hybrid cloud services).

IT shops are the beneficiary of this rapid change, which impacts operational costs at a time of many economic challenges worldwide. This market change occurred at the same time that multicore processors arrived, creating even more need and opportunity for server virtualization solutions that could use onboard system resources more efficiently than when dedicated servers (one physical server, one operating system) predominated in the datacenter. In addition, the rapid buildup of x86-based systems in the datacenter over the past decade created a need to drive more efficient use of systems and a need to control energy and space consumption.

The future of the datacenter clearly lies in virtualization. Customers are demanding higher VM density, better performance for their most sensitive applications, and new

architectures to accommodate the next-generation dynamic datacenter. Virtualization is driving new hardware designs that are optimized for these unique and challenging needs. To fully realize all that virtualization brings today and will bring in the future, customers will require new and innovative system designs.

---

### **Copyright Notice**

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2009 IDC. Reproduction without written permission is completely forbidden.