



# Intel® Virtualization Technology: A Primer

by Andrew Binstock, and updated by Matt Gillespie

April 17th, 2006

## Introduction

Intel Virtualization Technology is a set of silicon-based features available from Intel® server, desktop, and mobile processors that complement software-based virtualization technologies to add greater manageability, security, and hardware utilization to the enterprise.

In order to understand Intel Virtualization Technology, it is worthwhile to begin with a description of virtualization generally, which is to say, virtualization that is accomplished by hardware measures, software measures, or a combination of the two. In general, then, virtualization is a technique by which hardware resources can be abstracted, divided, and allocated among multiple virtual partitions on a single machine. Each of these virtual partitions contains its own operating-system environment. Virtualization can be accomplished by a variety of techniques such as simulation, emulation, and hardware or software partitioning of the resources. An ideal virtualization solution, from a user's perspective, should offer sufficient isolation between different virtual machines and uncompromised performance of each virtual machine, as well as availability, reliability, and security of the entire platform.

One of the many abilities that virtualization provides is to run multiple operating systems simultaneously on a single hardware platform, which is fast becoming an important solution to numerous problems that confront Information Management. This Primer examines the benefits provided by virtualization such as reliability, security, and manageability for information technology (IT) managers and increased productivity for developers, and it explains how virtualization solutions are enhanced by Intel Virtualization Technology. It provides the technical background that implementers and architects need to determine how to realize the full opportunity presented by Intel Virtualization Technology.

## Why Is Virtualization Important?

There are many benefits to virtualization, but the key reasons for its adoption at IT sites is that it delivers better use of resources and greater manageability of systems. Virtualization delivers a wide variety of benefits:

- **Consolidation and reduced Total Cost of Ownership (TCO):** Many applications run on underutilized, stand-alone servers because they cannot be migrated to new platforms or consolidated onto a single platform. The primary obstacles to this migration and consolidation are factors such as the application's requirement to run on a dedicated system for reliability and supportability, or else that it uses an operating system not used elsewhere at the site. Often, such applications are legacy solutions that use out-of-date operating systems and cannot be easily upgraded to run on newer systems. Such
-

considerations often effectively block migration to consolidated platforms. By using virtualization, one hardware platform can run multiple instances of various operating systems, each running in its own, isolated space, called a virtual machine (VM), from which it shares access to hardware resources.

TCO is lowered as consolidation reduces the cost of maintaining systems, and better management of space and power is possible within the data center. Enterprises must also consider the software licensing costs that may arise from virtualization, however. In cases where additional instances of operating systems and application software must be licensed, there can be additional costs, depending on how the solution is designed. As vendors work through these licensing issues, network architects and administrators need to keep up on policy changes.

- **Manageability:** Servers that use virtualization can accept migration of VMs from other systems, as workload requirements require. In fact, workload-balancing agents can perform this migration automatically. The Xen project uses workload migration on servers and clients, and similar technology exists in commercial products, such as VMotion\* from VMware ([www.vmware.com](http://www.vmware.com)\*).
  - **Security and fault isolation:** Virtualization can provide security through isolation between different VMs. By running applications in separate instances, one VM cannot be corrupted by a virus or malware resident in another VM. Any damage caused by malicious software is contained to the specific session in which it is running. Likewise, if the application should hang due to a bug or error, it locks up only its own session. This VM can then be restarted without affecting any of the other VMs running on the system. Furthermore, security schemes such as different VMs for different user logins can be provided.
  - **Reliability and Availability:** Virtualization solutions can provide reliable failover schemes and system backup on a single physical system.
  - **Improved productivity:** Upgrading and updating software is one of the common yet critical tasks that every IT manager performs. Sites that want to test upgrades of software packages on the same hardware on which it will be deployed can do so easily by running a new VM session on the platform. Virtualization can increase productivity for developers and quality assurance personnel. Developers and software testers working on cross-platform projects find it practical to run sessions with target operating systems on their development workstation. This enables them to verify cross-platform functionality. Developers who work on creating and testing kernel components can reduce downtime and turn-around time by performing their testing on an isolated VM. In the absence of a reliable virtualization solution, an erroneous kernel component might result in reinstall or rebuilding of the OS.
-

As processors become increasingly powerful with multi-core architectures and hardware multithreading such as Hyper-Threading Technology (HT Technology), virtualization solutions would be sought to address under-utilized, standalone server and legacy application situations.

Virtualization solutions could allocate dedicated core(s) to different virtual machines, while providing the aforementioned benefits. For example, in a multi-core processor platform, one core could be dedicated for a VM that is expected to take over in case of a failover, thus providing reliability with uncompromised performance while providing ease of manageability and consolidation benefits.

Another key aspect to the value proposition associated with Intel Virtualization Technology is its complementary nature with other emerging Intel® processor features. For example, [Intel® Active Management Technology](#) (Intel® AMT) allows network administrators to remotely discover, heal, and protect hardware, even if it is powered off or has a corrupted operating system. [LaGrande Technology](#) enables applications to run within their own protected space, helping to guard against software-based attacks and to protect the confidentiality and integrity of data stored or created on the client PC. [Intel® I/O Acceleration Technology](#) (Intel® I/OAT) increases the speed of networking and I/O operations. Together with these capabilities, Intel Virtualization Technology adds distinct value to the customer.

## What Is the Problem?

Software-only virtualization solutions—known as virtual machine monitors (VMMs)—handle all virtualization of the system. Because the VMM must create the perception that the hosted OS is communicating directly with the hardware, it must resort to a little magic. Two approaches are used:

**Paravirtualization:** This technique requires changes to the source code of the OS, especially the kernel, so that it can be run on the specific VMM. This approach is akin to the mainframe approach, in which custom OS extensions are closely matched to the hardware. Paravirtualization, of course, will not work with off-the-shelf operating systems.

**Binary translation:** The VMM makes changes to the binaries of the operating system as they are loaded into the VM. This approach is common in commercial products and has the singular limitation that only specific versions of the OS can be loaded, as each new release of the OS requires proof, testing, and possibly upgrades to the VMM software.

The problem with both solutions is that the software cannot work in concert with the underlying hardware, and so it must use complex schemes to emulate certain hardware features to the software, and it must fool the hosting operating system into thinking that the VM is just another application. VMMs also face other technical challenges.

---

## **Current Challenges for Virtualization Software**

### *Use of Private Memory for VMM Use Only*

In order to store system information, VMMs must use private blocks of memory that only they can access. The problem is how to allocate this memory in such a way that the guest OS will not access it (either inadvertently or on purpose). The principal solution is for the VMM to intercept accesses to these memory areas and to emulate the expected result of the initial access. This cumbersome process is required by hardware that does not support hardware-based virtualization. On processors with Intel Virtualization Technology, however, certain memory pages for use by the VMM can be made accessible only from software—such as the VMM—that has the highest level of privilege, as granted by the processor. This step makes these areas inaccessible—and most importantly—invisible to all other software.

### *Use of VMM Interrupt Handling*

Interrupts—events that require immediate system attention—must be handled by the VMM. The problem is that operating systems have the ability to prevent delivery of interrupts. This mechanism is used to block interruptions of certain activities that must be completed without interference from an external event. VMMs can manage the flow of interrupts to guest operating systems, but to do so they must monitor the attempts to mask and unmask (that is, block and allow) these interrupts. Some operating systems make heavy use of this feature, which causes significant performance penalties on the VMM.

## **What Is the Solution?**

Intel Virtualization Technology provides robust hardware support for virtualization that addresses the problems of software-only solutions and gives much-needed support to VMM vendors. It enables VMMs to run off-the-shelf operating systems and applications without recourse to binary translation or paravirtualization. This capability greatly facilitates the deployment of VMMs and provides greater reliability and manageability of guest operating systems and applications.

### **How Does Intel Virtualization Technology Work?**

VMMs must do two things well. They must completely emulate the hardware environment to the point that the hosted OS cannot tell it does not own the entire hardware platform, and they must handle all unusual circumstances that can arise either in the OS (such as hardware malfunctions) or the application (software errors). Both tasks must be performed with high levels of reliability and low performance overhead.

---

Hardware that does not support hardware-based virtualization makes it difficult for VMMs to meet these goals, because traditional processors were designed primarily to run a single instance of a single operating system. As a result, VMMs face a number of challenges that are addressed by Intel Virtualization Technology. Let's look at these in greater detail.

### **Privilege Levels**

All modern processors and operating systems implement the concept of privilege levels, which define what actions can be performed by specific processes. Intel® architecture provides four levels of privilege, called rings, that are numbered 0-3. The highest level, 0, is used by the operating system; the lowest level, 3, is employed by applications. For various reasons, levels 1 and 2 are rarely, if ever, used. Only operating systems running in ring 0 have unrestricted access to the hardware. By limiting this ring to use by a single OS, the processor enables the OS to have complete knowledge of the state of the hardware.

For the VMM to work properly, it needs to run at ring 0 and create the illusion to the guest OS that the guest OS is running in ring 0. However, since the VMM is itself running in ring 0, no guest OS can run at this privilege level. In fact, today they typically run at ring 1—a technique known as “ring deprivileging.” This practice creates enormous difficulties for the VMM, which must constantly monitor the activities of the VMs to trap hardware accesses and certain system calls, executing them itself and emulating the results.

Intel Virtualization Technology solves this problem by creating two classes of rings: the privileged “root” ring—referred to as ring 0P—for use by the VMM, and the deprivileged “non-root” ring—ring 0D—for the operating systems. In this way, the VMM can function as the fundamental layer and all OSs can run above it with the necessary benefits of ring 0. By use of this approach, hosted OSs and applications run within their expected ring levels and are unaware of the VMM—each hosted OS thinks it owns the entire machine.

### **How Does Intel Virtualization Differ Between IA-32 and Itanium® Architectures?**

Intel Virtualization Technology is similar on the surface between IA-32 and Itanium architectures, but because the architectures themselves are so different, separate specifications govern the implementation of the technology on these two platforms. The use of the technology will be similar if not identical to most business users, but a summary of the differences is available from the May, 2005, cover feature of *Computer* magazine, [Intel Virtualization Technology](#).

The IA-32 version of Intel Virtualization Technology is referred to as VT-x, and documentation on VT-x can be found in [“Intel Virtualization](#)

---

Technology Specification for IA-32 processors (VT-x).” The Intel Itanium architecture version is referred to as VT-i, and documentation on VT-i can be found in “[Intel Virtualization Technology Specification for the Intel Itanium Architecture \(VT-i\)](#).”

## What Must I Change and What Will I Get with Intel Virtualization Technology?

Interestingly, IT sites need to do nothing to their applications to leverage Intel Virtualization Technology. In fact, this is the whole goal of Intel Virtualization Technology: to run any application without modification in a VM. Intel Virtualization Technology is primarily oriented toward engineers at vendors of VMMs who can exploit the features to deliver better virtualization software. By making use of Intel Virtualization Technology, the new VMM products will be:

- **Robust:** VMMs will no longer need to use paravirtualization or binary translation. This means that they will be able to run off-the-shelf OSs and applications without any special steps.
- **Enhanced:** Intel Virtualization Technology enables VMMs to run 64-bit guest OSs—a first on IA x86 processors.
- **Reliable:** Due to the hardware support, VMMs can now be smaller, less complex, and more efficient. This improves reliability and availability and reduces the potential for software conflicts.
- **Secure:** The use of hardware transitions in the VMM strengthens the isolation of VMs and further prevents corruption of one VM from affecting others on the same system.

Few technologies deliver so much benefit without requiring sites to change or upgrade their software. To make use of Intel Virtualization Technology, simply make sure to include it when specifying your virtualization solutions.

## A Growing Ecosystem

Virtualization software is available today from a number of providers, giving Intel® architecture-based servers capabilities that were previously available only on mainframes. Examples of the solutions available in this ecosystem include the following:

- VMWare(EMC): ESX Server\*, VMWare Server\*, VMWare Player\*, and VirtualCenter\*
  - Microsoft: Virtual Server\* and Virtual PC\*
  - Xen opensource community: Xen
  - Virtual Iron: Virtual Iron\*
-

- SW Soft: Virtuozzo\*
- Parallels: Parallels Workstation\*

Intel is actively working with software vendors to help this ecosystem develop. Intel and VMware are collaborating on several market acceleration and education initiatives, including educating and bringing the value of virtualization directly to IT managers through direct engagements and targeted materials. To help educate the market and drive virtualization ubiquity, the companies are also investing in the development of a comprehensive virtualization starter kit containing fully featured products and the supporting resources required for new customers to start using virtualization.

Intel and Microsoft have also joined together to extend Intel Virtualization Technology to include support for mapping I/O devices to virtual machines on servers with a new specification called Intel® Virtualization Technology for Directed I/O (Intel® VT-d). Part of the Intel VT family of technologies, Intel VT-d helps improve the reliability, flexibility and performance of I/O in a virtualized environment. Microsoft has collaborated with Intel on development of the specification to help ensure it provides optimal functionality for users.

Intel platforms supporting Intel Virtualization Technology started shipping in 2005 for desktop and 2006 for mobile platforms, as well as Intel® Xeon® processor-based servers and workstations. Intel® Itanium® processor-based servers supporting Intel Virtualization Technology will start shipping later in 2006.

## **Learn More about Intel Virtualization Technology**

For further research into Intel Virtualization Technology, the following resources provide a good starting point:

[Software Developer FAQ: Intel® Virtualization Technology](#) gives an overview of issues and resources for developers and decision-makers

- [Intel® Virtualization Technology Web site](#) provides in-depth information about specifications and technologies
- [Virtualization for the Desktop PC](#) includes measures that businesses can employ to improve security and productivity using Intel Virtualization Technology
- [Virtualization for Servers](#) provides guidance about how to benefit from Intel Virtualization Technology on enterprise servers

More in-depth information is available from the following articles and white papers:

---

- [Virtualization: Bringing Flexibility and New Capabilities to Computing Platforms](#) shows how Intel Virtualization Technology is transforming IT
- [Enhanced Virtualization on Intel® Architecture-Based Servers](#) demonstrates how virtualization technology helps derive better value from IT investments
- [Improving IT Management with Multi-Port NICs and a Virtual Infrastructure](#) provides provisioning advice for virtualization in the enterprise

## About the Author



Andrew Binstock is the principal analyst at Pacific Data Works, LLC. He was previously a senior technology analyst at PricewaterhouseCoopers, and earlier editor in chief of *UNIX Review and C Gazette*. He is the lead author of "Practical Algorithms for Programmers," from Addison-Wesley Longman, which is currently in its 12th printing and in use at more than 30 computer-science departments in the United States.



Copyright © 2006 Intel Corporation. All rights reserved. BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino logo, Chips, Core Inside, Dialogic, EtherExpress, ETOX, FlashFile, i386, i486, i960, iCOMP, InstantIP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Core, Intel Inside, Intel Inside logo, Intel. Leap ahead., Intel. Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Viiv, Intel XScale, IPLink, Itanium, Itanium Inside, MCS, MMX, MMX logo,

Optimizer logo, OverDrive, Paragon, PDCharm, Pentium, Pentium II Xeon, Pentium III Xeon, Performance at Your Command, Pentium Inside, skool, Sound Mark, The Computer Inside., The Journey Inside, VTune, Xeon, Xeon Inside and Xircom are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

\* Other names and brands may be claimed as the property of others.