

A Superior Hardware Platform for Server Virtualization

Improving Data Center Flexibility, Performance and TCO with Intel® Virtualization Technology

Technology Brief

Intel® Virtualization Technology


Server Virtualization

Server virtualization is helping IT organizations improve data center productivity in fundamental ways. It lets you consolidate multiple operating systems and applications per physical server and deploy new applications in minutes. It also lets you move running applications from one server to another without downtime for flexible workload management, high availability and planned or unplanned maintenance. The benefits in utilization, energy savings, manageability, service levels and cost models can be dramatic.

Realizing the full value of virtualization requires servers that are built to handle the heavy demands of a virtualized and consolidated computing environment. Servers of just a few years ago were designed to host a single operating system. Successful virtualization with these systems requires software that can emulate a complete hardware environment for every guest operating system. This is a compute-intensive process that introduces significant performance overhead. It can slow application response times, limit scalability, and create complexity that can impact reliability and security.

“Choosing the right hardware platform for server virtualization is just as important as choosing the right virtualization software.”¹

– Ken Cayton, IDC Analyst



Intel® Virtualization Technology

- For Intel® Xeon® processors: Intel® VT-x
- For Intel® Itanium® processors: Intel® VT-i



Intel® Virtualization Technology for Directed I/O

(Intel® VT-d)



Intel® Virtualization Technology for Connectivity

(Intel® VT-c)

- Virtual Machine Device Queues (VMDq)
- PCI-SIG Single-Root I/O Virtualization (SR-IOV)
- Intel® I/O Acceleration Technology (Intel® I/OAT)

Intel integrates hardware assists for virtualization into all key server components to help IT organizations consolidate more applications and heavier workloads on each server, and to improve flexibility, reliability, and TCO.

Superior Virtualization through Comprehensive Hardware Support

Intel® Virtualization Technology (Intel® VT) addresses these challenges at the silicon level, by providing comprehensive hardware assists that boost virtualization software performance, improve application response times and provide greater reliability, security and flexibility. These integrated hardware assists accelerate fundamental virtualization processes throughout the platform to reduce latencies and avoid potential bottlenecks. They also reduce the demands placed on the virtualization software, so more processor cycles are available for running business applications. As a result, you can consolidate more applications and heavier workloads per server to get better value from your server and software investments.

Intel Virtualization Technology consists of three technology suites that function together to improve virtualization performance in every part of the server platform, including:

- **The processor:** Intel® Virtualization Technology in Intel® Xeon® processors² (Intel® VT-x) and Intel® Itanium® processors (Intel® VT-i) provides enhanced virtualization across the full range of 32-bit and 64-bit applications.
- **The chipset:** Intel® Virtualization Technology for Directed I/O (Intel® VT-d).
- **I/O devices:** Intel® Virtualization Technology for Connectivity (Intel® VT-c).

Intel works with VMware, Microsoft, the Xen open source community, Parallels and many other virtualization software vendors to help ensure that these technologies are broadly supported in today's and tomorrow's solutions, so they deliver high value while being completely transparent to IT organizations and end-users. The functionality of your virtualization solutions is unchanged. Your virtual servers are simply more responsive, more scalable and more reliable.

Intel® VT-x

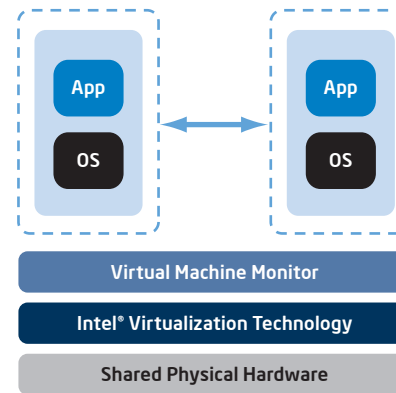
Better Virtualization Support in Intel® Processors

The primary implementation of Intel VT-x helps to improve the fundamental flexibility and robustness of software-based virtualization solutions.³ It reduces VMM interventions by eliminating the need for the VMM to listen, trap and execute certain instructions on behalf of the guest OS as is required in software-only virtualization. It also provides hardware support for transferring platform control between the VMM and guest OSs, so when VMM intervention is required, handoffs are faster, more reliable and more secure.⁴

More recent feature enhancements to Intel VT-x add valuable new capabilities:

- **Intel® VT FlexPriority:** When a processor is performing a task, it often receives requests or “interrupts” from other devices or applications that need attention. To minimize the impact on performance, a special register in the processor (the APIC Task Priority Register, or TPR) monitors the priority of tasks, so only interrupts that have a higher priority than the currently running task receive immediate attention. Intel FlexPriority creates a virtual copy of the TPR which can be read and in some cases changed by guest OSs without VMM intervention. This can deliver major performance improvements for 32-bit OSs that make frequent use of the TPR. (For example, it can improve performance by as much as 35 percent for applications running on Windows Server 2000⁵)

- **Intel® VT FlexMigration:** One of the key benefits of virtualization is the ability to migrate running applications from one physical server to another without downtime. Intel VT FlexMigration is designed to enable seamless migrations among current and future Intel processor-based servers, even though newer systems may include enhanced instruction sets. With this technology, hypervisors can establish a consistent set of instructions across all servers in the migration pool, enabling seamless migration of workloads. The result is a more flexible and unified pool of server resources that functions seamlessly across multiple hardware generations⁶



With Intel® VT-x, fewer VMM interventions are required and control can be passed between a guest OS and the VMM more quickly, reliably and securely.

Intel® Virtualization Technology for Directed I/O (Intel® VT-d)

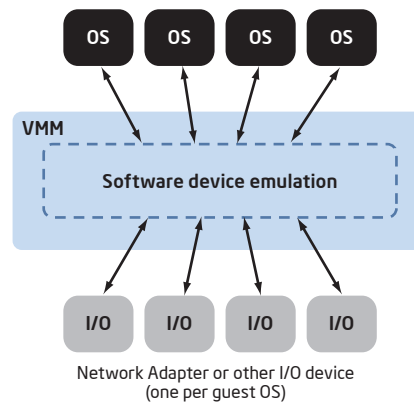
Better Virtualization Support in Intel® Chipsets

As more guest OSs are consolidated per server, the movement of data into and out of the system (I/O traffic) increases and becomes more complex. Without hardware assistance, the VMM is directly involved in every I/O transaction. This not only slows down data movement, but also increases the load on server processors due to the higher VMM activity. It's as if every shopper in a busy shopping mall had to enter or exit the mall through a single door and get directions only from the mall manager. This would not only slow down customers, but would also prevent the manager from attending to other pressing issues.

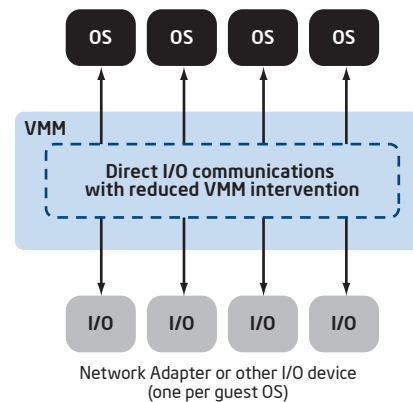
Intel VT-d speeds data movement and eliminates much of the performance overhead by reducing the need for VMM involvement in managing I/O traffic. It accomplishes this by enabling the VMM to securely assign specific I/O devices to specific guest OSs. Each device is given a dedicated area in system memory that can be accessed only by the device and by its assigned guest OS.

Once the initial assignments are made, data can travel directly between a guest OS and its assigned devices. I/O traffic flows more quickly and the reduced VMM activity decreases the load on the server processors. Security and availability are also improved, since I/O data intended for a specific device or guest OS cannot be accessed by any other hardware or guest software component.

Sharing of I/O devices without VT-d



Sharing of I/O devices with VT-d



With Intel® VT-d, the VMM can establish direct links between guest OSs and their assigned I/O devices, so traffic flows more quickly and there is less need for VMM intervention.


Intel® Virtualization Technology for Connectivity (Intel® VT-c)

Better Virtualization Support in Intel® I/O Devices

Intel VT-c further enhances server I/O solutions by integrating extensive hardware assists into the I/O devices that are used to connect your servers to your data center network, storage infrastructure and other external devices. In essence, this collection of technologies functions much like a post office that sorts an enormous variety of incoming letters, packages and envelopes and delivers them to their respective destinations. By performing these functions in dedicated network silicon, Intel VT-c speeds delivery and reduces the load on the VMM and server processors.

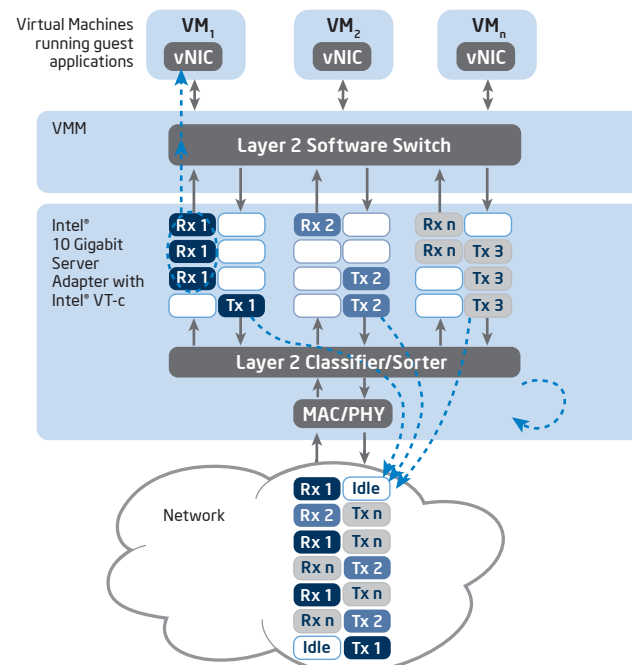
Intel VT-c includes three key technologies, which are now supported in all Intel® 10 Gigabit Server Adapters and selected Intel® Gigabit Server Adapters.

▪ **Virtual Machine Device Queues (VMDq):** In a traditional server virtualization environment, the VMM has to sort and deliver every individual data packet to its assigned virtual machine. This can consume a lot of processor cycles. With VMDq, this sorting function is performed by dedicated hardware in Intel Server Adapters. All the VMM has to do is route the presorted packet groups to the appropriate guest OSs. I/O latency is reduced and the processor has more cycles available for business applications.



Intel® VT-c can more than double I/O throughput, so more applications can be consolidated per server with fewer I/O bottlenecks. (The results shown are based on internal tests performed by Intel and VMware.?)

- **Intel® I/O Acceleration Technology (Intel® I/OAT):** This suite of technologies accelerates I/O communications across the platform through a combination of network, chipset and software enhancements.
- **PCI-SIG Single-Root I/O Virtualization (SR-IOV):** This industry-standard specification has been ratified by the Peripheral Component Interconnect Special Interest Group (PCI-SIG) forum. As discussed in the previous section, Intel VT-d enables a direct communication channel between a guest OS and an I/O device. SR-IOV extends this by enabling multiple direct communication channels for each I/O device. For example, each of ten guest OSs could be assigned a protected, dedicated, 1 Gb/sec link to the corporate network through a single Intel 10 Gigabit Server Adapter. These direct communications links bypass the VMM to enable faster I/O performance with less load on the server processors.



Intel® VT-c offloads network I/O management to dedicated network silicon that helps to accelerate throughput and reduce the load on the VMM and server processors.

More and Better Support to Come

Intel VT is a multi-generational roadmap of increasingly powerful enhancements to Intel processors, chipsets and I/O devices. New generations will continue to build on previous advances to provide more secure, scalable, reliable, flexible, and responsive virtualization solutions. These technologies are fully integrated, thoroughly tested and widely supported by leading virtualization software solutions. They provide IT organizations with a proven, industry-leading foundation for optimizing the value of their server and virtualization investments.

For the latest information about Intel Virtualization Technology, visit the Intel Web site at:
www.intel.com/technology/virtualization/

For detailed information about Intel VT-d, visit: www.intel.com/technology/virtualization/

For detailed information about Intel VT-c, visit: www.intel.com/go/vtc

¹ Source: Choosing the Right Hardware for Server Virtualization, an IDC white paper sponsored by Intel, Doc # 211622, April 2008. <http://www.intel.com/business/technologies/IDCchoosingvirthardware.pdf>

² Intel® VT-x supports both 32-bit and 64-bit Intel® Xeon® processor-based solutions (Intel® 64 and IA-32).

³ Intel® VT-x is included in Intel® Xeon® processors; VT-i is included in Intel Itanium processors and delivers comparable benefits.

⁴ Based on Intel tests, transition times have decreased about 6 times from 2004 (when Intel processors did not support VT-x) to 2008 (with the latest VT-x enhancements).

⁵ Intel tests demonstrate 35 percent performance gains for Windows Server® 2000 and 2003 SP1 versions running as guest operating systems.

⁶ Intel® VT FlexMigration supports live VM migration across all Intel® Core™ microarchitecture-based servers. It is included in the new Intel® Xeon® processor 5400 series, and provides backward compatibility for live VM migration with current dual-core Intel® Core™ microarchitecture products (Xeon 5100 and Xeon 3000) and forward compatibility with future dual- and multi-core processors. Contact your preferred VMM vendor for support requirements.

⁷ For more information, see the Intelligent Queueing Technologies for Virtualization, An Intel-VMware perspective: Enhanced Performance in Virtualized Servers. http://download.intel.com/network/connectivity/products/whitepapers/Intel-VMware_VMDq_wp_May08.pdf

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web site at www.intel.com.

Copyright © 2008 Intel Corporation. All rights reserved. Intel, the Intel logo, Xeon, and Itanium are trademarks of Intel Corporation in the U.S. and other countries.

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

