

## WHITE PAPER

---

# Virtualization on Itanium: Intel Enables Choice and Flexibility for Customers

Sponsored by: Intel

---

John Humphreys

Ken Cayton

November 2007

## EXECUTIVE SUMMARY

As virtualization is quickly becoming an important technology across all parts of the IT environment, the technology is rapidly being incorporated in storage, networks, and client environments. By far the most visible adoption of virtualization technology is happening in servers — from the largest Unix server to the smallest volume system. Virtualization software breaks the link between a given software-driven application or service and the hardware resources of the underlying system and encapsulates and isolates that stack or service.

Server virtualization has a long history, starting in mainframes in the 1960s. For well over a decade, high-end RISC-based servers with Unix operating systems (OSs) have utilized virtualization technologies. The Itanium mix of hardware and software virtualization capabilities is taking its lead from these platforms to provide similar capabilities (physical partitioning) and extensive firmware and software support in a platform for mission-critical applications at volume economics. The inherent capabilities of virtualization on Itanium are intended to help customers move beyond simple server consolidation and leverage virtualization capabilities for high availability, disaster recovery, as well as security and physical isolation purposes for both business-critical and mission-critical application environments.

Intel is partnering with a host of virtualization providers to deliver customers as much choice and flexibility as possible on which to run their mission-critical applications. Key virtualization providers include HP, Hitachi, Fujitsu, NEC, Red Hat, and SWsoft. By partnering so broadly, Intel believes it is delivering on the customer need for choice in selecting the right tool for the job. Because virtualization is such a key and important focal means for IT to return value to the organization, it is critical to have a deep tool box.

This paper discusses the key attributes for virtualization adoption and highlights the offerings of the core group of Intel partners leveraging Itanium processor solutions with Intel for a variety of virtualization solutions.

## SITUATION OVERVIEW

---

### **Virtualization on Itanium: A Mainstream Technology**

Virtualization in computing is a broad term that refers to the abstraction of computer resources. This includes making a single physical resource (such as a server, an operating system, an application, or 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 used in mainframes was primarily used to partition a single large machine/resource into multiple machines 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 environment.

Software for implementing "virtual machines" (hypervisors) and virtual server infrastructure is technology that is now spreading rapidly on x86 platforms. x86 virtualization is also progressing in terms of use cases, technology, and maturity.

Virtualization on Itanium is designed for a more comprehensive solution than the typical virtualization play seen on x86. The focus is on extending to more efficient levels of resource management for mission-critical applications. Virtualization on Itanium-based platforms can be accomplished through physical partitioning, soft partitioning, and a combination of both techniques.

The mainframe-class ability to physically partition a system — processors, memory and I/O channels — (allowing multiple environments to share one physical server to give greater flexibility in managing workloads) has been effectively absent in Windows server-class systems. Physical partitioning technology typically consists of a combination of hardware and firmware, sometimes with some level of control software, that creates multiple physical hardware configurations. Each of these partitions is provisioned with a dedicated operating system and layered infrastructure and application software, and this configuration appears to the software environment to be a dedicated hardware resource. Physical partitioning on Itanium is widely available today.

Soft partitioning through a hypervisor enhances resource management providing finer-grain, more dynamic resource control, thereby raising enterprise IT's efficiency in meeting peak demands.

The combination of physical partitions and soft partitions is offered to provide maximum control, flexibility, and high-availability capabilities. The section on the Itanium ecosystem highlights several examples of how OEMs have implemented various virtualization capabilities.

Intel Virtualization Technology (VT) now available on Itanium platforms replaces lower-level hypervisor functions, which results in expanded guest OS support for existing virtualization solutions and enables an architectural foundation for new solutions. Intel VT is a complementary technology to virtualization software products that enhances today's virtualization solutions and lays the foundation for future platform virtualization. Intel VT provides hardware assists to the virtualization software, reducing its size and complexity and enabling lower-cost, more efficient, and more powerful virtualization solutions.

Intel VT present on Itanium processors (VT-i) enables a new privilege space where the virtual machine management software can operate. It reduces the size and complexity of the virtual machine management software, improving its efficiency and enabling greater functionality.

Intel VT-i provides the foundation for widely deploying virtualization solutions across a broad set of customer applications and production workload environments by working to address the following challenges associated with software-only virtualization:

☒ **Overhead and performance.** Software-only virtual machine monitors (VMMs) introduce processing overheads to "emulate" server resources. These overheads occur in such areas as IO operations, memory management (paging), and simulating "privileged" CPU instructions. Intel VT-i, by supporting these functions with hardware virtualization, helps to reduce virtual machine overheads and hence expands the portfolio of applications and workloads suitable for virtualization.

☒ **Less complexity.** Software virtualization solutions today mean that the OS has to go through the VMM when communicating with the underlying hardware because the OS is running in the address space and privilege level where applications are normally run. With software-only solutions, enabling an OS to run as a guest in a virtualized environment is achieved via binary translation and patching of the OS itself, to "trick" the OS into believing it is running on a "bare machine." Translation can be done dynamically, at runtime, or statically, in advance (known as paravirtualization). The translation approach means that every time the OS is updated (e.g., new release, service pack, patch), the VMM must also be patched to maintain support.

Hardware-assisted virtualization eliminates the need for binary translation or patching by providing a new architecture that allows the OS to run as an unmodified guest.

☒ **Reliability.** Many users worry that software translation or patching reduces the reliability of the overall solution. Intel VT addresses these concerns via hardware support for privilege ring expansion with resulting simplification of the hypervisor. Privilege ring expansion means that the VMM runs in a new, higher privilege ring, thus allowing the guest OS to run in its native privilege ring (ring 0).

- ☒ **Breadth of support.** Intel VT broadens the number of OSs that can be supported, compared with today's current software-only solutions. Intel VT enables direct support of unmodified, legacy operating systems. The aggressive ramp of Intel 64-bit Xeon solutions means we will see an increased need for VMs to support 64-bit guest OSs. Today's solutions do not support this capability, but solutions with Intel VT will. In addition, VT-based solutions support a full range of legacy OSs (multiple versions of Linux and Windows).
- ☒ **Flexibility.** A key goal of Intel VT is to make VMM software independent of specific OS software in order to remove the necessity of constantly updating VMs to keep up with OS changes and patches.
- ☒ **Need for platform reliability.** As customers go from largely "one server, one application" environments today to many applications being hosted on a virtualized server, there is tremendous need for reliable hardware. To mitigate the risks associated with having many "eggs in one basket," users must search for the platform with the most comprehensive reliability, availability, and serviceability (RAS) features.

Overall, Intel VT-i hardware-assisted virtualization, combined with ongoing support from virtualization vendors, will provide enhanced virtual environments.

---

## **Why Virtualize Systems? The Business Case**

At the highest level, IDC has found two major benefits associated with the virtualization of system resources: cost benefits and greater flexibility in managing systems. Because of the benefits, in some cases, virtualized systems are becoming the standard deployment platform for all new applications in customer environments.

In terms of cost benefits, we have found that virtualization impacts both capital costs and operational costs. Partitioning allows users to increase the server utilization of the current systems in their environments. Data suggests that system utilization averages between 5% and 25% in most situations, which means that 75% to 95% of the capacity is unused. This information, combined with the fact that many x86 customers still assign one server per application, underscores the reality that many customers have dramatically overprovisioned their hardware resources. Virtualization is a means to increase utilization by consolidating multiple applications on a single host so that capital costs are used efficiently. This is one of the major reasons that most customers initially started to virtualize their servers.

From an operational perspective, virtualization is impacting costs in a few ways. By separating the application from the underlying hardware, IT administrators are able to more easily manage, provision, restart, and migrate applications on a shared pool of server hardware. IT is finding that virtualization is a means to speed change in the datacenter. Companies using virtualization technologies report that the time to deliver a new server within their organizations has been reduced from days or weeks to just a few hours or, in some cases, only minutes. This ability provides IT with a means to be more responsive to business requirements.

Customers also report massive increases in server-to-administrator ratios, which significantly drive down the cost of managing large IT infrastructures. This ability to grow infrastructures without adding staff members can lead to competitive advantages for businesses.

In addition, the reduction in physical hosts associated with better utilization of resources can lead to dramatic savings in electricity and cooling costs. Not only is this savings recurring, but the reduction in power consumption can prolong datacenter life, thus helping organizations avoid having to expand or build new facilities.

---

## Server Virtualization Use Cases

Because of the benefits of virtualization, including the capital and operational cost benefits and the flexibility and ease of managing applications, users are developing an increasingly broad set of usage scenarios for virtualization technologies in their environments. These major scenarios are as follows:

- ☒ **Server consolidation.** The ability to run multiple production applications while sharing physical hardware resources leads to reduced hardware costs (and depending on the solution, potentially lower software costs), easier administration, and simplified manageability. Not only does virtualization help to reduce server footprints, but it also can greatly reduce power consumption, cooling demands, and cabling and free up space in overfilled datacenters.
- ☒ **Test and development.** A popular use is to run pilot solutions in different virtual partitions. This approach takes advantage of all the benefits of server consolidation (reducing the number of "nonproductive" servers) and allows increased agility through accelerated application rollout. New applications can be validated inside partitions before being rolled out across the datacenter. Additionally, developers can easily compile and test their applications under multiple operating systems using a single server, without heavy hardware investments.
- ☒ **Hosting legacy applications.** Virtualization allows users to run applications along with older operating systems that are no longer supported and that can no longer run on today's servers. The business case for legacy applications is that customers benefit from the increased performance of newer hardware without the cost of having to port the application off unsupported legacy operating environments.
- ☒ **Hardware migrations and upgrades.** Virtualization allows a staged migration of applications onto new hardware. By validating the solutions in virtual partitions, IT shops reduce the possibility of disruption of service due to migration. When migrating to a new architecture, staging the migration in virtualized servers allows robust validation before a widespread rollout. The ROI with migration comes from the level of automation and reduced IT staff time needed to move applications across platforms.

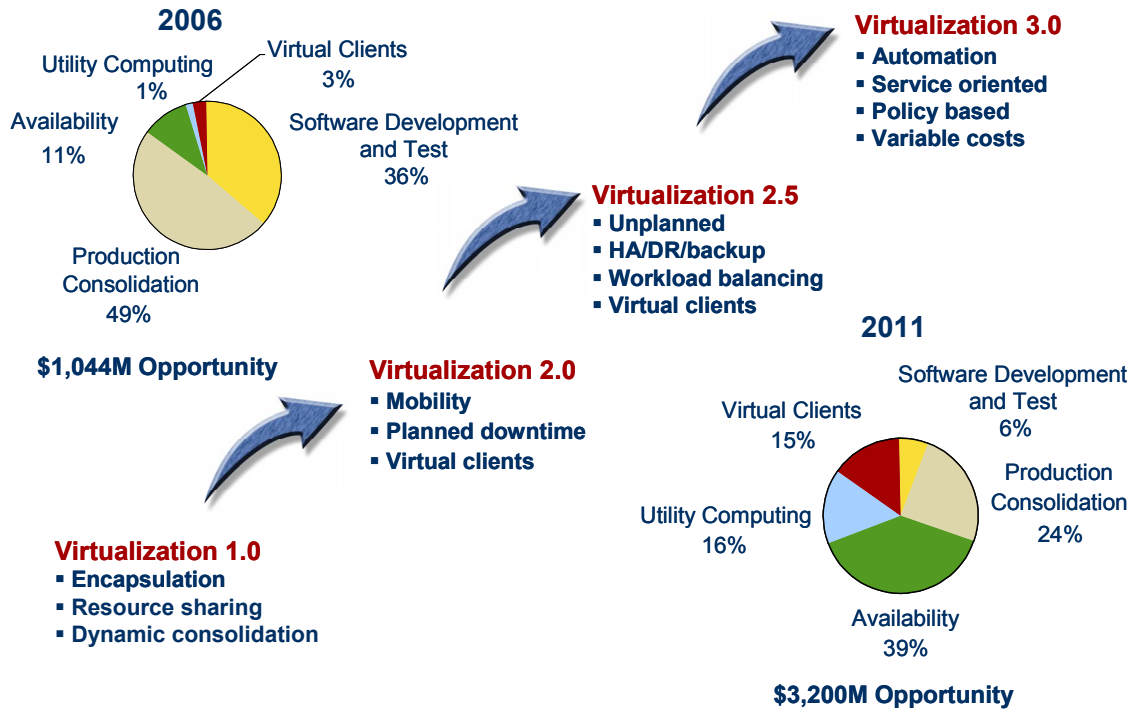
- ☒ **Business continuity.** Most IT shops today deploy some form of failover, usually involving replicated servers. Because many failures are associated with software, users are recognizing that replicating servers in virtual partitions on the same machine provides the benefits of higher application availability from built-in virtualization failover at a reduced hardware cost.
- ☒ **Capacity planning.** Virtual partitions can be sized and resized as required. This capability offers great flexibility to IT shops, which can configure large compute resources when required (e.g., end of the quarter), then scale down as needed. By scaling compute resources as needed, IT shops can optimize the overall utilization of the server. In the future, the provisioning of compute resources will become even more dynamic, providing the ability to add and remove resources as required.
- ☒ **Resource management.** Resource management refers to policy-based tools that monitor utilization and move virtual machines as needed to balance peak capacity with headroom. The benefit of load balancing is that users can more effectively pool their infrastructures and share a centralized managed pool across multiple applications. This ability to balance loads and provision applications dynamically helps further reduce hardware costs and increase service levels and application availability.

This ever-expanding portfolio of usage for virtualization is what makes the technology so compelling to end users and, in turn, is why Intel, its partners, and leading virtualization vendors are working together to create a robust platform that encompasses the hardware, software, and management layers of the solution. This collaborative approach is critical if the full potential of the technology is to be realized. When choosing a virtualization solution, customers must examine not just the cost but also their business needs, the application, the operating environment, the level of isolation, and the flexibility and performance they require. Only through a thorough examination of these criteria can users strike the right balance and find the best virtualization solution for their specific needs.

The move from what IDC has termed Virtualization 1.0 to Virtualization 2.0 is predicated on the mobility of the virtual environment — specifically being able to migrate a live running virtual environment from one physical host to another. This ability is a key enabler of the new use cases that help to take virtualization beyond a tool for simple server consolidations (see Figure 1).

**FIGURE 1**

Virtualization Milestones



Source: IDC, 2007

**Virtualization Solutions: An Itanium Ecosystem**

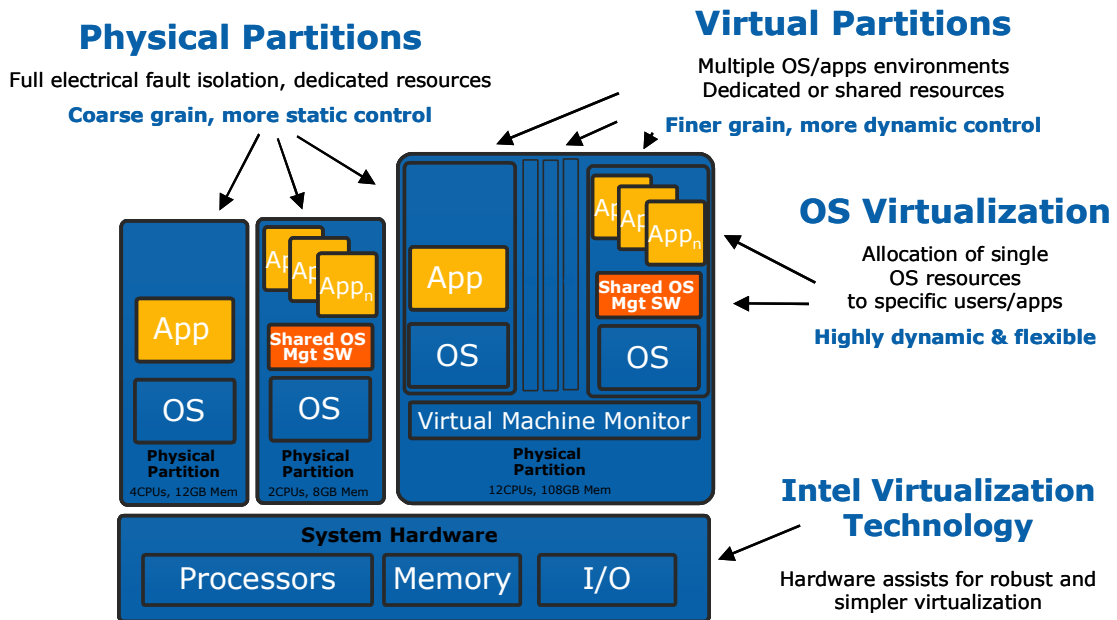
In essence, virtualization is a hardware resource-sharing strategy. The core advantage of implementing server virtualization is the ability to drive higher utilizations on physical servers by sharing the use of hardware resources across several workloads, moving away from the "one application per server" approach that often characterized the use of nonvirtualized servers in the past. Other advantages are the ability to isolate one workload from another, in terms of memory, data, and hard drive contents, avoiding such problems as driver conflicts when two applications are running on the same server, and the ability to provide better security and minimize compliance issues while sharing resources across multiple workloads.

As the operational and business benefits delivered by virtualization have been realized, implementing virtualization has become a major initiative for organizations. There is growing recognition among customers of the need for diversity in how they virtualize or decouple the application stack from the underlying hardware. This need for diversity stems directly from the application because different applications have different requirements in terms of isolation, flexibility, and performance. Having the right virtualization tool for the job is as critical as choosing a hardware platform or vendor.

Itanium system vendors, operating systems vendors (OSVs), and independent software vendors (ISVs) have responded by working to optimize the choice and capabilities of virtualization solutions on Itanium (see Figure 2) to meet the needs of mission-critical IT customers as detailed in the following sections.

**FIGURE 2**

Itanium Virtualization Options



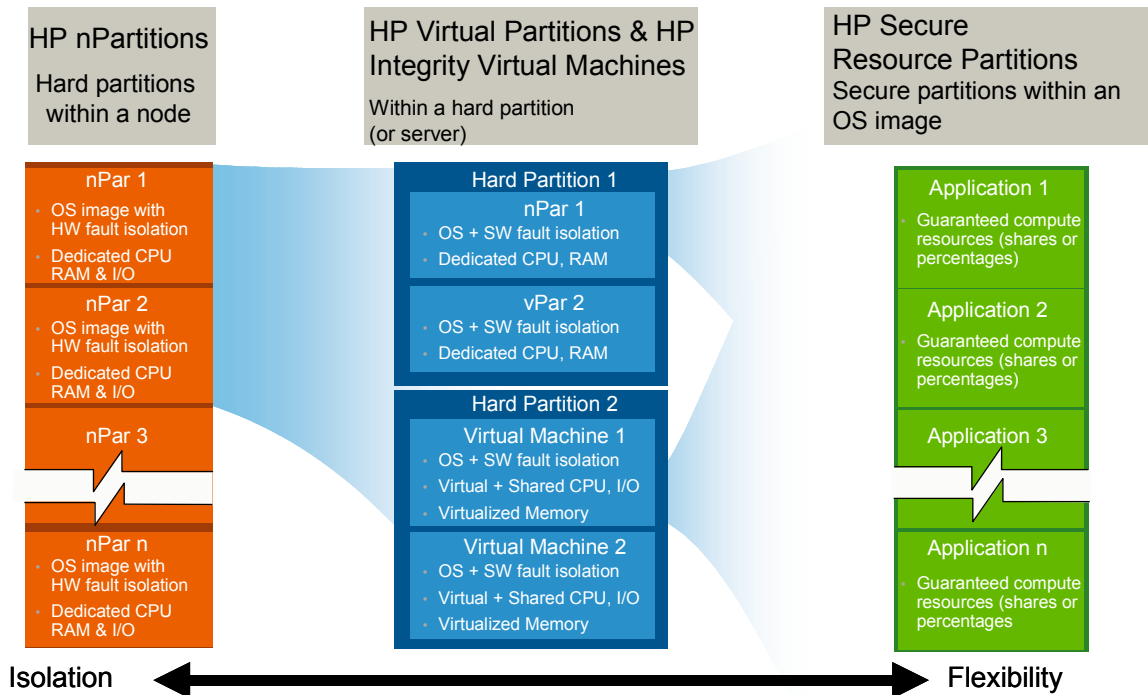
Source: Intel, 2007

**The Platform Approach: Virtualization from HP**

HP provides what it terms a "continuum" of partitioning and virtualization options for customers on Integrity systems (see Figure 3). This continuum ranges from physical partitioning at one end to resource balancing at the other — with virtual partitioning in between. The intent of this spectrum of offerings is to enable the customer to mix, match, and/or nest each of these virtualization solutions to devise the best solution for the workload and environment.

**FIGURE 3**

HP's Virtual Server Continuum



Source: HP, 2007

Within the virtual partitioning space, HP has two virtualization solutions. The first is designed for HP-UX and is referred to as Virtual Partitions. Virtual Partitions allow multiple instances of HP-UX to run on a single system or nPartition. This capability provides finer granularity than nPartitions through the dynamic reallocation of CPUs and memory between Virtual Partitions.

The Virtual Partitions product is available on HP Integrity platforms and is implemented by assigning dedicated hardware components such as CPU, memory, and I/O to each partition. Rather than boot HP-UX directly off the hardware, the vPar Monitor is booted. The vPar Monitor is configured so that it can identify all the partitions as well as the physical components that belong to each partition. The vPar Monitor loads each Virtual Partition, which then boots its OS. When the OS queries the kernel to determine what hardware components are available, the kernel responds with only the subset of the resources that are allocated to that partition. An advantage to this model is that the Virtual Partition is able to communicate directly with the hardware resources, resulting in a very low overhead.

Virtual Partitions can adjust the amount of CPUs and memory allocated to each partition while they are running. This allows CPUs and memory to be moved between Virtual Partitions by deallocating them from one Virtual Partition and subsequently allocating those same CPUs to another Virtual Partition.

The other solution in HP's partitioning continuum is called HP Integrity Virtual Machines, which is a fully virtualized environment for running applications. The VM Host runs on any HP Integrity system or nPartition, which means that a cell-based system is not needed to use Integrity Virtual Machines. Virtual machines, which present themselves to the VM Host as physical servers, are run on top of the VM Host. When an OS is installed on a virtual machine, the virtual machine becomes a guest. In an Integrity Virtual Machines environment, all of the resources in a guest are virtualized. The physical CPUs, memory, and I/O devices are managed by the VM Host. What the OS inside the virtual machine sees is a virtual resource that is mapped on top of the physical devices in the system. This allows the physical resources to be shared by multiple guests.

The virtualization that is provided by Integrity Virtual Machines is so complete that the operating systems running inside the virtual machines run without modification. This means that the same Integrity version of each OS (HP-UX, Windows, and Linux) that runs on a standalone system or nPartition will run on a virtual machine.

Because of this level of isolation, software running inside a virtual machine cannot determine that it is not a physical system. All of the resources that are presented are virtual, and the virtual machine might be sharing those resources with multiple other virtual machines. As a result, the resource allocation to virtual machines is very dynamic by default, but it is generally transparent to applications running inside the virtual machine.

HP has developed a robust set of virtualization technology that brings together the flexibility of hypervisor-based virtualization with the reliability of its long-standing hardware and software features and capabilities.

### ***The Embedded Approach: Hitachi Virtage in BladeSymphony***

Intel and Hitachi have worked closely together to drive virtualization into the hardware. Virtage is integrated in the firmware, enabling customers to choose which blades to virtualize. Like mainframes, Virtage employs "direct execution." As a result, Virtage-based systems can offer a tremendous performance enhancement as more guest operations can be directly executed, thereby minimizing host intervention.

The work by Intel and Hitachi is manifested in the BladeSymphony 1000 and includes Virtage embedded virtualization technology. The platform will enable customers to deploy Intel Dual-Socket, Quad-Core Xeon and/or Intel Dual-Core Itanium-based server blades within the same chassis. The Virtage virtualization technology can partition physical blades into multiple isolated logical partitions (LPARs). A single server module can be configured with up to 16 LPARs, and each of these environments can run its own independent software stack.

In addition to the virtualization benefits of BladeSymphony, customers gain centralized management capability, high-performance I/O, and sophisticated RAS features. In line with mission-critical expectations, BladeSymphony 1000 features a very modular design to maximize flexibility. System elements are redundant and hot-swappable so that the system can be easily expanded without downtime or unnecessary disruption to service levels. Because the software is embedded, Hitachi can deliver a strong virtualization platform and yet stay very competitive from a solution cost perspective.

## The Hardware Approach

### Dynamic Partitioning from NEC

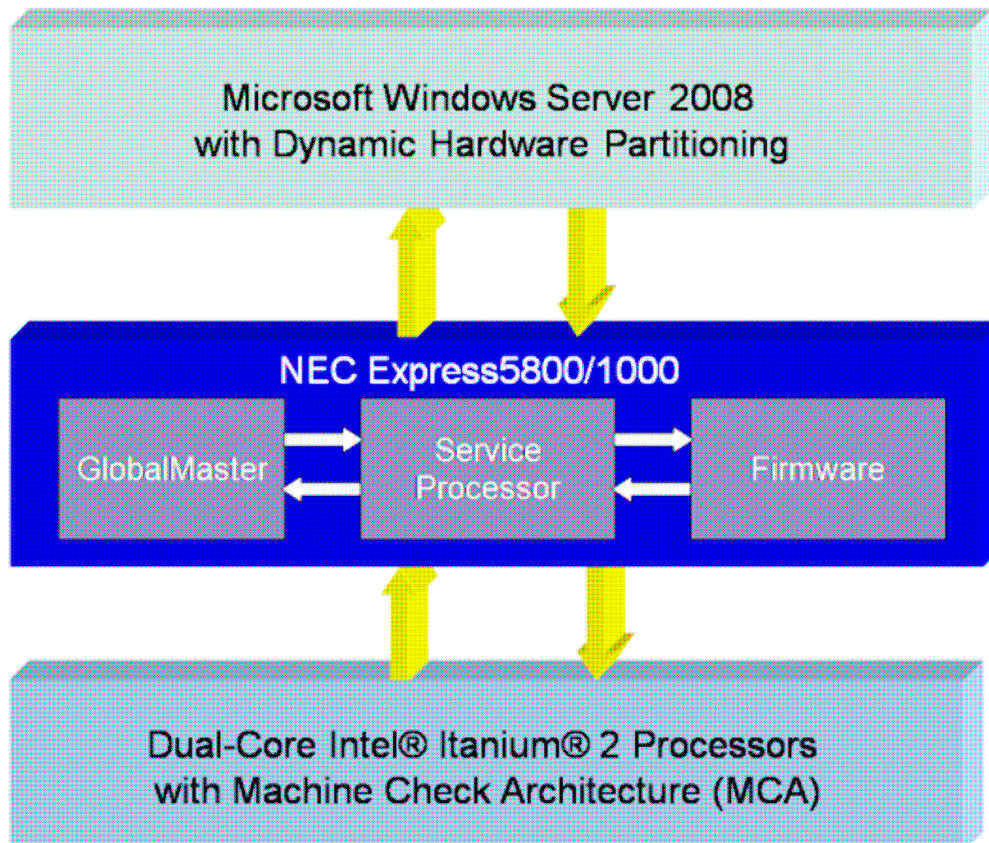
The NEC Express5800/1000 Series includes the capability for hardware partitioning of the processors, memory, and I/O within the server into physically isolated partitions, each running its own operating system environment.

The NEC Express5800/1000 Series of servers utilizes the Intel Itanium Machine Check Architecture (for error handling) and adds another layer of intelligence through a service processor and related software called the GlobalMaster (managed via either a GUI or a CLI). The GlobalMaster also communicates with Microsoft Windows Server 2008 to convey to the operating system that a hot add (of processors, memory, or I/O) or hot replace (of processors and memory) is taking place.

NEC calls the combination of hardware, firmware, and operating system technologies Dynamic Hardware Partitioning (see Figure 4). The combination of these technologies allows the automation of a hot add or hot replace of processors and their associated memory, based on thresholds and policies, to a running OS partitioned on the server, with no disruption to applications on the critical OS.

**FIGURE 4**

NEC Dynamic Hardware Partitioning



Source: NEC, 2007

## **Dynamic Partitioning from Fujitsu**

Fujitsu has developed virtual machine capability for its PRIMEQUEST mission-critical IA server environment. The functions are enabled by virtualization software included in Red Hat Enterprise Linux (RHEL) 5 and Fujitsu's support for that software. This includes Fujitsu proprietary value-added software for performance enhancement, improved reliability, and installation support. The PRIMEQUEST virtual machine function is available with all servers, starting with the PRIMEQUEST 500 series.

The PRIMEQUEST virtual machine offers concurrent operation of multiple operating systems (Linux and Windows) with up to 60 virtual machines or "guests" per host server. The Fujitsu virtualization software also allows users flexible and detailed allocation of system resources to each operating system. Each virtual machine can expand or contract dynamically to handle real-time fluctuations in demand, and allocation of system resources including CPU, memory, and I/O devices is possible in quite smaller increments. This greater granularity makes it much easier to change resource allocations.

The unified operation management of both physical and virtual machine environments is another advantage. It is achieved by combining the PRIMEQUEST virtual machine function with Fujitsu's Systemwalker integrated operation management software product set. By combining the virtual machine function with Systemwalker Operation Manager, users can automate start-up, termination, and resource allocation changes to virtual machines and guest OSs based on a defined schedule.

Beyond developing virtual machine-type virtualization for PRIMEQUEST, Fujitsu also offers physical partitioning capabilities, which include Physical PARTitioning (PPAR) and eXtended PARTitioning (XPAR), on the Itanium-based system. Both partitioning technologies enable the splitting of the physical hardware at hardware component boundaries for those customers that want, need, or require hard partitions.

## ***The Container Approach: Operating System Virtualization from SWsoft***

The Virtuozzo product from SWsoft offers customers yet another approach and alternative when it comes to virtualization. Instead of virtualizing below the operating system, SWsoft chooses to virtualize on top of the operating system stack. This capability, which the company terms "container technology," enables customers not only to consolidate systems but also to consolidate operating systems instances. The ability to securely run multiple applications and operating system images on a single server helps customers reduce server sprawl, thereby lowering power and cooling requirements while also minimizing the OS images, which in turn reduces moves, adds, and changes associated with OS deployment and patches. This approach to virtualization can complement what others are doing in terms of physical partitioning, hypervisor virtualization, and resource management.

Because of the performance of Virtuozzo for Web-based workloads and the high density of virtual environments, the majority of adoption has occurred among service providers and companies offering software as a service. This situation appears to be poised to change. As SWsoft has added more "enterprise-class" features such as live migration, IDC believes that enterprise customers are beginning to take note.

Additionally, while hypervisor virtualization enables server consolidations, it does nothing to help manage OS sprawl. For companies with hundreds or thousands of images, OS deployment and patching is an issue as well. By consolidating operating systems, customers have the ability to significantly reduce the operational costs associated with software moves, adds, and changes.

The requirements of enterprise customers, the challenges associated with virtual machine sprawl, and the drive for increased virtual environment density are reasons that SWsoft has chosen to make Itanium a key platform to support.

### ***The Open Source Approach: Xen Hypervisor Virtualization from Red Hat***

The Itanium ecosystem is also working with the open source community to incorporate virtual machine-type virtualization into the portfolio of options available to customers. Red Hat has announced that in the upcoming release of Red Hat Enterprise Linux (RHEL) version 5.1, integrated software virtualization using Xen will be supported and optimized for Itanium 2-based systems.

Xen has emerged as the open source alternative for virtualization leveraging virtual machines. With RHEL 5.1, the base hypervisor (which is freely available to all) will be incorporated into RHEL 5 and paired with management and monitoring tools to streamline deployment for customers. As customers subscribe to RHEL 5.1, they will be able to leverage the incorporated Xen hypervisor to freely partition and manage their virtual instances.

Red Hat is an early supporter of running virtual machines on the Itanium platform because it enables customers to put more and more applications on a single system. In turn, as the system becomes even more critical because of the consolidated applications, it requires the reliability, availability, and serviceability as well as the balance of an enterprise-class system. The company believes that supporting Itanium makes it easier and safer for customers to take virtualization deeper into the datacenter.

Additionally, Red Hat supports running paravirtualized operating systems. By taking advantage of the capabilities inherent in VT-i, customers can receive added performance benefits on top of the RAS and balance associated with Itanium systems.

## **CHALLENGES/OPPORTUNITIES**

As virtualization continues to gain momentum in the marketplace and is projected to grow rapidly over the next five years, the challenge for Intel and many of the server OEMs is to position their platforms as the right one for virtualization of mission-critical applications.

For Intel and the system OEMs, the essence of this challenge is to have their platforms recognized as having the reliability, availability, and serviceability (RAS) characteristics to support mission-critical applications. Virtualization can concentrate risk as more applications are placed on fewer systems. As such, enterprise features

and design become increasingly important. Clearly articulating the RAS of the Itanium systems and underlying processor has been and will continue to be an important first step in the battle for share of virtualized hardware.

In addition to demonstrating the hardware benefits, Intel will need to show a rich ecosystem of vendors and virtualization options in the marketplace. The vendors detailed in this paper show that the company has significant momentum and is offering customers choice as to how they virtualize — which is encouraging. These efforts to port virtualization software to Itanium still need to continue to deliver even more choice for virtualizing on Itanium systems.

Intel and other vendors must continue to engage customers in industrywide educational discussions that articulate how to map IT requirements to the different virtualization, operating environment, and server platform choices available as well as what virtualization tools work best with specific applications or for certain business solutions. Having a broad set of choices is only the first step in full coverage of customer requirements.

## **CONCLUSION**

The virtualization solutions discussed in this white paper are specifically designed to work on Itanium and are focused on meeting the needs of mission-critical environments in which RISC and mainframe systems dominate. The breadth in virtualization development on Itanium not only should give customers comfort that Itanium momentum continues to build but also should provide customers with more flexibility. With all these different virtualization solutions enabled to run on Itanium, customers can weigh the benefits and challenges of each solution relative to their specific business requirements to make the best choice on how they virtualize their mission-critical environments as they move to more open solutions.

---

### **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 2007 IDC. Reproduction without written permission is completely forbidden.