

White Paper

DATA CENTER

www.novell.com

Harness the Power of Virtualization for Server Consolidation

Xen-based Virtualization with SUSE® Linux Enterprise on AMD Processors
with AMD Virtualization

AMD 

Novell.

Harness the Power of Virtualization for Server Consolidation

By consolidating multiple server environments onto a single server, you can harness more of your data center's unused computing power and get up to eight times more work from the same amount of resources.

The average server workload in an enterprise environment ranges from five to 40 percent, leaving at least 60 percent of the available capacity unused (based on October 2005 Novell® customer interviews). Still, the majority of today's data centers run a single application on a single server, creating server sprawl and resulting in more money spent on hardware and the resources required to manage it. To significantly reduce these costs, data centers need to take better advantage of the high processing power of today's server CPUs by employing server virtualization solutions that reduce sprawl and maximize workloads.

Virtualization lets you reduce your total number of servers by giving you the ability to migrate multiple environments on different physical machines to individual virtual machines hosted on a single server. These virtual machines each run their own independent and self-contained operating systems and server applications.

By consolidating multiple server environments onto a single server, you can harness more of your data center's unused computing power and get up to eight times more work from the same amount of resources. Virtualization also provides the added benefit of dynamic provisioning—moving applications and systems in one virtual machine from one server to another as needed. By harnessing the server-consolidation power of virtualization, your organization can enjoy the following:

- *Increased resource efficiency.* When you allow multiple applications to coexist on one physical server, you harness more of a server's capacity.
- *Reduced server provisioning time,* since virtualization can reduce server provisioning time by up to 50 percent. This enables you to be more responsive to customer needs.
- *Improved troubleshooting.* Because instances of virtual machines can be remotely managed at any time from anywhere on the network, you reduce the need for physical support visits.
- *Increased flexibility and scalability.* Because new virtual machines are easy to deploy, your IT team can respond more rapidly and flexibly to business requirements for new IT resources.
- *Increased productivity from being able to run multiple operating systems on a single computer.* Your organization can do more in-depth development and testing, as well as take on a greater number of additional projects.
- *High availability and non-stop maintenance.* Through temporarily migrating virtual machines and their hosted applications to other physical servers for planned hardware and software maintenance, you'll have zero application downtime.
- *Tighter application security by isolating and minimizing vulnerabilities through the creation of virtual machines for each application.* Protection can extend to the application, libraries, services and operating system, as each is protected in its own container.
- *More control over configuration change management since unapproved changes made to a running application in production can be protected through the ability to rollback servers to a warehoused image.*
- *Decreased hardware costs through reducing the number of physical servers that need to be maintained, physical space, power and cooling requirements.* Virtualization can allow for the deployment of 12 or more virtual machines on each physical server for a 12:1 consolidation ratio.

- *Scalability on demand. Automatically deploying virtual-machine images to other existing hardware adds more processing power at peak demand times.*

The SUSE® Linux Enterprise 10 platform is the first Linux* distribution to fully integrate of the Xen 3.0 hypervisor as a core component of the operating system. Available from Novell, SUSE Linux Enterprise 10 offers data centers the first enterprise-ready Linux virtualization solution on the market. Xen is the industry's fastest and most secure open source infrastructure virtualization software technology. It has been endorsed and adopted by more than 20 of the industry's major vendors, including AMD, whose next-generation AMD Opteron processor incorporates AMD Virtualization (AMD-V). Xen utilizes specific extensions that enable data centers to provide hardware-based protection/isolation of individual virtual machines and to achieve higher levels of performance when running fully virtualized servers.

The Value of Virtualization

You want to enjoy the benefits of server consolidation by moving all of your existing server applications to a single server with a new, high-performance processor. However, you're afraid that if one crashes, it will bring all the other applications down as well. Or maybe your concerns are security-related. You want to keep critical applications isolated to protect them from being infected with malware or viruses from the server's other applications. Or perhaps your legacy applications require older—or a variety of—operating systems, preventing you from running them on the latest hardware or in the same server environment.

Virtualization addresses all of these issues, facilitating data-center server consolidation. In a virtualized environment, a single server hosts multiple virtual machines. In essence, each of these virtual machines “thinks” it is running on its own physical hardware.

In reality, it's running on a piece of software known as the hypervisor, which presents each machine with a virtualized view of its own native hardware and operating system. The hypervisor lies on top of the physical layer of the server. It runs at the most privileged hardware-protection ring and has the responsibility to allocate resources for each virtual machine.

With virtualization, machines are isolated from each other. They act and operate as distinct and separate server environments. If one virtual machines crashes, it does not affect the other virtual machines or the applications running on them. This also prevents interaction among virtual machines; malware or viruses infecting one virtual machine won't spread to others despite their being on the same host server.

Since the hypervisor presents each virtual machine with a distinct virtualized view of its own native hardware and operating system, virtual machines are not required to run the same operating system. On the same physical server, one virtual machine might run Linux, another NetWare® and a third Windows NT*. Consequently, you can consolidate your various (and even incompatible) server environments onto a single physical server.

The Xen Hypervisor and SUSE Linux Enterprise 10

Based on an open source project hosted by the University of Cambridge, Xen is currently the best-performing and most secure hypervisor in the industry—and it typically has 10 times less overhead than competitive proprietary offerings. As an open source technology, Xen provides a number of benefits over proprietary solutions, including improved functionality, better performance and greater extendibility. Unlike proprietary hypervisors, which rely on software-only virtualization, Xen is the industry's first supported software base to take advantage of

Available from Novell, SUSE Linux Enterprise offers data centers the first enterprise-ready Linux virtualization solution on the market.

Both Novell and its partners have worked with the open source community to ensure that the Xen hypervisor meets the business needs of its mutual customers that want to leverage the consolidation benefits of server virtualization.

As one of the major contributors to the Xen open source project, Novell has integrated the Xen hypervisor as a core function of its SUSE Linux Enterprise 10 platform, making Novell the first major commercial vendor to bring an enterprise-ready Linux virtualization solution to market.

hardware-based virtualization enhancements. This allows running unmodified guest operating systems on the next-generation AMD processors with AMD-V.

A number of Novell strategic alliance partners have worked with the open source community and Novell to help ensure that the Xen hypervisor meets the business needs of its mutual customers that want to leverage the consolidation benefits of server virtualization. As one of several major contributors to the Xen open source project, Novell has integrated the Xen hypervisor as a core function of its SUSE Linux Enterprise 10 platform, making Novell the first major commercial vendor to bring an enterprise-ready Linux virtualization solution to market. The Xen hypervisor runs under SUSE Linux Enterprise 10 as the control operating system for managing the virtual-server environment.

Paravirtualization

Xen's unique performance benefits accrue from its use of paravirtualization. With paravirtualization, the operating system running inside of a virtual machine (known as a guest operating system) is modified to run on top of a hypervisor. A para- (or partially) virtualized operating system instance is aware that it is running in a virtualized state and has been fine-tuned for optimal performance in that environment.

Paravirtualization allows the hypervisor to avoid hard-to-virtualize processor instructions by replacing them with procedure calls that

provide that functionality. A paravirtualized operating system loads and runs virtual drivers that are capable of interacting with Xen to access resources on the host virtual server. In other words, it does not require complete emulation of computer devices.

For example, a virtual machine with a paravirtualized operating system wouldn't require an emulated graphics card, eliminating the need for the hypervisor to emulate its video data. Consequently, a paravirtualized operating system requires less overhead management, resulting in improved performance.

The initial release of SUSE Linux Enterprise Server 10 includes support for para-virtualized instances of SUSE Linux Enterprise Server 10. In late 2006, Novell will add support for a paravirtualized version of SUSE Linux Enterprise Server 9 running as a guest operating system on a SUSE Linux Enterprise Server 10 host. Novell plans to add support in 2007 for the following operating systems that will be modified to run as paravirtualized guest operating systems on top of the Xen hypervisor in a SUSE Linux Enterprise Server 10 environment:

- *NetWare 6.5 SP3 (delivered as part of the next version of Novell Open Enterprise Server)*
- *Red Hat* Enterprise Linux 4 and 5*
- *Solaris* x86*

Full Virtualization

One of the additional advantages that the Xen hypervisor in SUSE Linux Enterprise 10 provides is that it supports—in addition to paravirtualized operating systems—full virtualization on CPUs that have been designed specifically for virtualization. (Examples include the next-generation AMD processors with AMD-V.) A fully virtualized operating system is one that has not been modified specifically to run in a virtual environment, so it is unaware that it is being virtualized. As a result, the hypervisor traps and emulates

every I/O and hardware instruction that is deemed privileged by the hypervisor.

Typically, the overhead occurring from these trapping and emulation operations would have a significant impact on performance. However, the AMD processors with AMD-V have been designed specifically for virtualization. The Xen hypervisor interacts with the virtualization extensions in the AMD processors not only to improve performance and efficiency, but also to provide hardware-based isolation between these unmodified guest operating systems running on a virtualization server.

Even with these enhancements, fully virtualized guest operating systems will typically not be able to achieve the same levels of performance as paravirtualized guest operating systems. However, since it is unlikely that all operating systems of the past—or even the future—will be modified for virtualization, the main benefit of full virtualization comes from its ability to host legacy operating systems that have not been paravirtualized. The ability to host these legacy operating systems in a virtualized environment is critical to a data center's server-consolidation efforts. This feature is mandatory for virtualizing proprietary operating systems, including those from Microsoft*.

The following platforms will operate as unmodified guests on the Xen hypervisor on a SUSE Linux Enterprise 10 virtualization host server when using the AMD processors with AMD-V:

- *Microsoft Windows* operating systems, including Windows NT, 2000, 2003, XP and Vista*
- *SUSE Linux Enterprise Server 8 and 9*
- *Solaris x86*

Additionally, unmodified guests in an AMD-V virtual machine may exist alongside a virtual machine using a paravirtual guest.

Graphical Management Interface

SUSE Linux Enterprise Server 10 offers the first set of integrated management tools for the Xen hypervisor through the YaST graphical interface management module. Within the graphic-based YaST tool, you can create, start, stop and manage the guest virtual machines running on your host SUSE Linux Enterprise Server. The YaST management tool automatically detects whether or not your server's hardware supports both paravirtualization and full virtualization. If your hardware supports both, it will let you choose whether to run the virtual machine using full virtualization or paravirtualization.

AMD Processors with AMD Virtualization

The original introduction of the AMD Opteron™ processor unveiled a processor that extended the x86 instruction set to support both 32-bit and 64-bit applications. In addition, its overall architecture design was progressive enough to become an ideal foundation for upcoming computing technologies, including virtualization. Its Direct Connect Architecture connects multiple processors, the memory controller and the I/O directly to the CPU to increase processing power, increase bandwidth, and optimize memory performance. Moreover, the increased system efficiency and performance that the Dual-Core AMD Opteron processor provides for multithreaded and multitasking applications are even more crucial on a virtual server hosting multiple virtual machines and their guest operating systems.

With the next-generation AMD Opteron processors, AMD extended the Direct Connect Architecture further with hardware virtualization instructions that reduce software emulation overhead in full-virtualization environments. These extensions establish the framework required to create fully virtualized guest-mode virtual machines, to assign memory to guest contexts, to let an x86 operating system run in guest mode and to keep individual guest-mode virtual machines isolated from each other.

Within the graphic-based YaST tool, you can create, start, stop and manage the guest virtual machines running on your host SUSE Linux Enterprise Server.

The Xen hypervisor in SUSE Linux Enterprise 10 has been designed to work seamlessly with these extensions to enable data centers to create, start, stop and manage fully virtualized guest mode virtual machines.

Just as SUSE Linux Enterprise was one of the first Linux distributions to support AMD's 64-bit and dual core capabilities, it is also the first major Linux distribution to support the virtualization extensions in the AMD processors with AMD-V. The Xen hypervisor in SUSE Linux Enterprise 10 has been designed to work seamlessly with these extensions to enable data centers to create, start, stop and manage fully virtualized guest mode virtual machines.

Guest Mode

One of the biggest challenges in a virtualized environment is to be able to truly separate one virtual machine from another in memory access. The nature of the x86 world assumes that the operating system is always running at the most privileged layer of code and that there is only one operating system running. x86 wasn't built with the idea that you could have multiple operating systems that were all assigning memory or doing I/O. The AMD processors with AMD-V address this issue with its new Guest Mode.

Guest Mode enables the Xen hypervisor and the SUSE Linux Enterprise Server 10 host operating system to run in ring 0 while giving the hypervisor the ability to create virtual machines running in Guest Mode on rings 1, 2 or 3. It also employs certain techniques to enable these modes to run isolated from each other. But the primary ability of Guest Mode is allowing operating systems to run as guests without requiring modification for virtualization. In other words, it's Guest Mode that enables full virtualization of legacy operating systems. While Guest Mode is mostly associated with full virtualization, it is also capable of hosting a paravirtual operating system with the same benefits.

The following additional virtualization-specific extensions to the AMD processors with AMD-V help establish and maintain the Guest Mode:

- *Virtual Machine Control Block*
- *VMRUN*
- *Tagged TLB*
- *Paged Real Mode*
- *Intercept-based Virtualization*

Virtual Machine Control Block

The Virtual Machine Control Block (VMCB) is a new data structure in the AMD processors with AMD-V that describes a virtual machine guest. The VMCB contains a list of instructions or events in a guest that need to be intercepted. It indicates whether a specific guest has virtual interruptions or if it can access true interruptions. It also indicates whether a guest can write directly to a page table. Moreover, it contains various control bits that specify the execution environment of the guest, indicating any special actions that need to be taken before running guest code. The VMCB also stores the CPU state of all the guest operating systems running as virtual machines on a server.

VMRUN

While there are 9 new instructions in AMD Virtualization, the VMRUN instruction is the cornerstone. VMRUN is the instruction in the AMD processors with AMD-V that the Xen hypervisor calls to start a guest mode. The guest is loaded when issuing the VMRUN instruction with the characteristics described by the VMCB structure.

Tagged TLB

The Tagged TLB extension allows the Xen hypervisor to assign a specific address space identifier to a guest in the memory page table. These enable the hypervisor to distinguish between different host and guest address spaces, as well as allow the switching of a new process in a guest without having to flush the TLBs from memory. It also allows the hypervisor to efficiently decide when TLBs do need to be flushed from memory.

Paged Real Mode

Real mode is essentially how DOS runs; the operating system has complete access control of the physical machine. Paged Real Mode in the AMD processors with AMD-V allows a guest virtual machine to behave as if it were running in real mode. To do this, the Xen hypervisor sets up a shadow page table that maps the guest's physical memory to the appropriate host physical addresses. It does this in a way that makes the guest operating system "believe" that it has linear access to memory when it is in fact being paged. As a result, this functionality allows any application or code that expects to be in real mode to run on top of a fully virtualized guest operating system in paged real mode.

Intercept-based Virtualization

Intercept-based virtualization enables the hypervisor to control why a guest stops running. This allows the hypervisor to fine tune how it handles interrupts and I/O based on a guest operating system's workload for optimum efficiency or based on privileges. These characteristics are described for a guest in the VMCB.

Reaping the Benefits of Virtualization from Novell and AMD

With the integration of the Xen hypervisor in SUSE Linux Enterprise 10, Novell provides the first enterprise-ready Linux virtualization product to reach the market. As data center customers look to virtualization for server consolidation, the Novell offering gives them out-of-the-box ability to create virtual machines running paravirtualized guest operating systems for optimal performance. It also allows

As data center customers look to virtualization for server consolidation, the Novell offering gives them out-of-the-box ability to create virtual machines running paravirtualized guest operating systems for optimal performance.

them to leverage next-generation AMD processors with AMD-V to fully virtualize legacy operating systems at faster speeds than those provided by the original, native processors.

Through the cooperation of Novell, AMD and the open source community, SUSE Linux Enterprise 10, AMD processors with AMD-V, and the Xen hypervisor enable enterprises to easily leverage the power of server virtualization. As a result, enterprises can:

- *Decrease hardware costs*
- *Increase resource-usage efficiency*
- *Reduce the time and effort required to provision servers*
- *Improve server availability*
- *Increase productivity*
- *Improve control over change-configuration management*
- *Tighten application security*

Most of all, they can increase their ability to rapidly and flexibly respond to ever-changing business requirements. For more information on the benefits of virtualization, refer to the *Virtualization in the Data Center* Novell white paper found at: www.novell.com/rc/docrepository/public/37/RC-21582/26254.pdf

Through the cooperation of Novell, AMD and the open source community, SUSE Linux Enterprise 10, AMD processors with AMD-V, and the Xen hypervisor enable enterprises to easily leverage the power of server virtualization.

www.novell.com



Contact your local Novell
Solutions Provider, or call
Novell at:

1 800 714 3400 U.S./Canada
1 801 861 1349 Worldwide
1 801 861 8473 Facsimile

Novell, Inc.
404 Wyman Street
Waltham, MA 02451 USA



Novell.