

**Understand why virtualization
is so important**



Virtualization

FOR DUMMIES®

Sun and AMD Special Edition

**A Reference
for the
Rest of Us!®**

FREE eTips at dummies.com®

**Find out how
Sun and AMD
technologies work
together!**



**Bernard Golden
Clark Scheffy**

Virtualization
FOR
DUMMIES®
SUN AND AMD SPECIAL EDITION

Bernard Golden and Clark Scheffy



WILEY

Wiley Publishing, Inc.

Virtualization For Dummies[®], Sun and AMD Special Edition

Published by
Wiley Publishing, Inc.
111 River Street
Hoboken, NJ 07030-5774

Copyright © 2008 by Wiley Publishing, Inc., Indianapolis, Indiana

Published by Wiley Publishing, Inc., Indianapolis, Indiana

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Legal Department, Wiley Publishing, Inc., 10475 Crosspoint Blvd., Indianapolis, IN 46256, (317) 572-3447, fax (317) 572-4355, or online at www.wiley.com/go/permissions.

Trademarks: Wiley, the Wiley Publishing logo, For Dummies, the Dummies Man logo, A Reference for the Rest of Us!, The Dummies Way, Dummies Daily, The Fun and Easy Way, Dummies.com, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. AMD, the AMD Arrow logo, AMD Opteron, AMD Virtualization, AMD-V, and combinations thereof are registered trademarks of Advanced Micro Devices, Inc. HyperTransport is a licensed trademark of the HyperTransport Technology Consortium. Sun, the Sun logo, Solaris, StorageTek, Sun Fire, Sun xVM Ops Center, and Sun Ray are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

All other trademarks are the property of their respective owners. Wiley Publishing, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, please contact our Customer Care Department within the U.S. at 800-762-2974, outside the U.S. at 317-572-3993, or fax 317-572-4002. For details on how to create a custom *For Dummies* book for your business or organization, contact bizdev@wiley.com. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN: 978-0-470-29264-8

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

Project Editor: Jennifer Bingham

Editorial Manager: Rev Mengle

Sun Contributors: Laura Davenport, Julianne Raedeke, Tina Shibue, Vijay Sarathay, Henry Kennedy, Michael Nalls, Amy Ahlers, Lisa Clark, Mark Thacker, Cindy Gallie, Allison Michlig, Frances Sun

AMD Contributors: Valeri Reeder, Steve Demski, Jeff Jones



WILEY

Introduction



Virtualization is the latest in a long line of technical innovations designed to increase the level of system abstraction and enable IT users to harness ever-increasing levels of computer performance.

At its simplest level, *virtualization* allows you, virtually and cost-effectively, to have two or more computers, running two or more completely different environments, on one piece of hardware. For example, with virtualization, you can have both a Linux machine and a Windows machine on one system. Alternatively, you could host a Windows 95 desktop and a Windows XP desktop on one workstation.

In slightly more technical terms, virtualization essentially decouples users and applications from the specific hardware characteristics of the systems they use to perform computational tasks. This technology promises to usher in an entirely new wave of hardware and software innovation. For example, and among other benefits, virtualization is designed to simplify system upgrades (and in some cases may eliminate the need for such upgrades), by allowing users to capture the state of a virtual machine (VM), and then transport that state in its entirety from an old to a new host system.

Virtualization is also designed to enable a generation of more energy-efficient computing. Processor, memory, and storage resources that today must be delivered in fixed amounts determined by real hardware system configurations will be delivered with finer granularity via dynamically tuned VMs.

About This Book

Virtualization For Dummies, Sun and AMD Special Edition explains how virtualization works and how it can benefit your organization. The book covers the kinds of issues virtualization can address and how it addresses them.

Icons Used in This Book

In the margins of this book, you find several helpful little icons that can make your journey a little easier:



This icon flags information that you should pay attention to.



This icon lets you know that the accompanying text explains some technical information in detail. You don't need to know this stuff to get what you need from the book, but it may be interesting.



A Tip icon lets you know some practical information that can really help you out is on the way. These tips can help save you time, effort, or money.

Chapter 1

Wrapping Your Head around Virtualization

.....

In This Chapter

- ▶ Understanding virtualization
 - ▶ Looking at the different types
 - ▶ Examining hardware and virtualization
-

It seems like everywhere you go these days, someone is talking about virtualization. Technical magazines trumpet the technology on their covers. Virtualization sessions are featured prominently at technology conferences. And, predictably enough, technology vendors are describing how *their* product is the latest word in virtualization.

Why Virtualization Is Hot, Hot, Hot

What's all the shouting about? Why is virtualization the sensation of the season? This section goes over four reasons virtualization is so important.

Trend #1: Underutilized hardware

Today, many data centers have machines running at only 10 or 15 percent of total processing capacity. In other words, 85 or 90 percent of the machine's power is unused. However, a lightly loaded machine still takes up room and draws electricity, so the operating cost of today's underutilized machine can be nearly the same as if it were running flat-out.

It doesn't take a rocket scientist to recognize that this situation is a waste of computing resources. And, guess what? With the steady improvement in performance characteristics of computer hardware, next year's machine will have twice as much spare capacity as this year's (and so on, for the foreseeable future). Obviously, there ought to be a better way to match computing capacity with load. And that's what virtualization does — by enabling a single piece of hardware to seamlessly support multiple systems. By applying virtualization, organizations can raise their hardware utilization rates dramatically, thereby making much more efficient use of corporate capital.

So, the first trend that is causing virtualization to be a mainstream concern is the unending growth of computing power brought to us by the friendly folks of the chip industry.

Trend #2: Data centers run out of space

The business world has undergone an enormous transformation over the past 20 years. Business process after business process has been captured in software and automated, moving from paper to electrons.

The rise of the Internet has exponentially increased this transformation. Companies want to communicate with customers and partners in real-time, using the worldwide connectivity of the Internet. Naturally, this has accelerated the move to computerized business processes.

The net effect of all this is that huge numbers of servers have been put into use over the past decade, which is causing a real estate problem for companies: They're running out of space in their data centers. And, by the way, that explosion of data calls for new methods of data storage. These methods go by the common moniker of *storage virtualization*, which, as you may guess, means making it possible for storage to be handled independently of any particular piece of hardware.

Virtualization, by offering the ability to host multiple guest systems on a single physical server, allows organizations to

reclaim data center territory, thereby avoiding the expense of building out more data center space. This is an enormous benefit of virtualization, because data centers can cost in the tens of millions of dollars to construct.

Trend #3: Green initiatives demand better energy efficiency

Power costs used to rank somewhere below what brand of soda to keep in the vending machines in most company's strategic thinking. Companies could assume that electrical power was cheap and endlessly available.

The assumption regarding availability of reliable power was challenged during the California power scares of a few years ago. Although later evidence caused re-evaluation about whether there was a true power shortage, the events caused companies to consider whether they should look for ways to be less power dependent.

Furthermore, the impact of the green revolution has meant that companies are increasingly looking for ways to reduce the amount of energy they consume — and one of the places they look first is their data center.

To show the level of concern about the amount of energy being consumed in data centers, consider these facts:

- ✔ A study commissioned by AMD and performed by a scientist from the Lawrence Berkeley National Laboratory showed that the amount of energy consumed by data centers in the U.S. doubled between 2000 and 2005. Furthermore, energy consumption is expected to increase another 40 percent by the end of the decade. Current energy consumption by data center servers and associated cooling costs represents 1.2 percent of the total energy consumed in the U.S.
- ✔ Based, in part, on the results of this study, the United States Environmental Protection Agency (EPA) has convened a working group to establish standards for server energy consumption and plans to establish a new "Energy Star" rating for energy efficient servers.

The cost of running computers, coupled with the fact that many of the machines filling up data centers are running at low utilization rates, means that virtualization's ability to reduce the total number of physical servers can significantly reduce the overall cost of energy for companies.

Trend #4: System administration costs mount

Computers don't operate all on their own. Every server requires care and feeding by system administrators. Common system administration tasks include: monitoring hardware status; replacing defective hardware components; installing operating system (OS) and application software; installing OS and application patches; monitoring critical server resources like memory and disk use; and backing up server data to other storage mediums for security and redundancy purposes.

As you can imagine, these tasks are pretty labor intensive. System administrators — the people who keep the machines humming — don't come cheap. And, unlike programmers, system administrators are usually co-located with the servers, because they need to access the physical hardware.

As part of an effort to rein in operations cost increases, virtualization offers the opportunity to reduce overall system administration costs by reducing the overall number of machines that need to be taken care of. Although many of the tasks associated with system administration (OS and application patching, doing backups) continue even in a virtualized environment, some of them disappear as physical servers are migrated to virtual instances. Overall, virtualization can reduce system administration requirements drastically, making virtualization an excellent option to address the increasing cost of operations personnel.

Sorting Out the Types of Virtualization

Now that you have a rough idea of virtualization and why it's an important development, what are your options regarding

it? In other words, what are some common applications of the technology?

A number of common uses for virtualization exist, all centered around the concept that virtualization represents an abstraction from physical resources. In fact, there are enough kinds of virtualization to make it a bit confusing to sort out how you might apply it in your organization. The two most common types of virtualization applied in the data center are server virtualization and storage virtualization. Within each main type there are different approaches or “flavors,” each of which has its benefits and drawbacks.

Server virtualization

There are three main types of server virtualization: operating system virtualization; hardware emulation; and paravirtualization, a relatively new concept designed to deliver a lighter weight (in terms of application size), higher performance approach to virtualization.

Operating system virtualization

Operating system (OS) virtualization (sometimes called *containers*) runs on top of an existing host operating system and provides a set of libraries that applications interact with, giving an application the illusion that it is (or they are, if there are multiple applications) running on a machine dedicated to its use. The key thing to understand is that, from the application’s execution perspective, it sees and interacts only with those applications running within its virtual OS, and interacts with its virtual OS as though it has sole control of the resources of the virtual OS. Crucially, it can’t see the applications or the OS resources located in another virtual OS.

This approach to virtualization is extremely useful if you want to offer a similar set of operating system functionalities to a number of different user populations while using only a single machine. This is an ideal approach for Web hosting companies: They use container virtualization to allow a hosted Web site to “believe” it has complete control of a machine, while in fact each hosted Web site shares the machine with many other Web sites, each of which is provided its own container. Operating system virtualization imposes little overhead for the virtualization capability, thereby ensuring most of the

machine's resources are available to the applications running in the containers.

There are some limitations to operating system virtualization, though. First and foremost, this approach typically limits operating system choice. Containerization usually means that the containers offer the same operating system as the host OS and even be consistent in terms of version number and patch level. As you can imagine, this can cause problems if you want to run different applications in the containers, since applications are often certified for only a certain OS version and patch level. Consequently, operating system virtualization is best suited for homogenous configurations — for those arrangements operating system virtualization is an excellent choice.

Companies offering operating system virtualization include Sun (as part of their Solaris™ operating system) and SWSOft, which offers the commercial product Virtuozzo as well as sponsoring the open source operating system virtualization project called OpenVZ.

Hardware emulation

In *hardware emulation*, the virtualization software (usually referred to as a *hypervisor*) presents an emulated hardware environment that guest operating systems operate upon. This emulated hardware environment is typically referred to as a *virtual machine monitor* or *VMM*.

The VMM provides a standardized hardware environment that the guest OS resides on and interacts with. Because the guest OS and the VMM form a consistent package, that package can be migrated from one machine to another, even though the physical machines the packages run upon may differ. The hypervisor, which resides between the VMM and the physical hardware, translates the calls from the VMM to the specific resources of the physical machine.

This approach to virtualization means that applications run in a truly isolated guest OS, with one or more guest OSs running, one per VMM. The VMMs all reside on the virtualization hypervisor. Not only does this approach support multiple OSs, it can support dissimilar OSs, differing in minor ways (for example, version and patch level) or in major ways (for example, completely different OSs like Windows and Linux can be run simultaneously in hardware emulation virtualization software).

Common applications for hardware emulation are software development and quality assurance, because it allows a number of different OSs to be run simultaneously, thereby facilitating parallel development or testing of software in a number of different operating system environments. Hardware emulation is also used in server consolidation, where a number of operating system/application environments are moved from separate physical servers to a single physical server running virtualization software.

There are a couple of drawbacks to hardware emulation, however. One is that the virtualization software hurts performance, which is to say that applications often run somewhat slower on virtualized systems than if they were run on unvirtualized systems.

Another drawback to hardware emulation is that the virtualization software presents a standardized hardware interface (the VMM) to the guest operating system. The hypervisor provides an interface to the VMM and then translates that into calls to the actual physical resources on the machine. This means that the hypervisor must contain the interfaces to the resources of the machine; these resources are referred to as *device drivers*. If you've ever installed new hardware in a PC, you know that you often have to install a device driver into the operating system so that the new hardware and the operating system can communicate.

The device driver issue for hardware emulation is that the hypervisor contains the device drivers and there is no way for new device drivers to be installed by users (unlike on your typical PC). Consequently, if a machine has hardware resources the hypervisor has no driver for, the virtualization software can't be run on that machine. This can cause problems, especially for organizations that want to take advantage of new hardware developments.

Companies offering hardware emulation virtualization software include VMware (in two versions, VMware Server and ESX Server) and Microsoft, which offers a product called Virtual Server. VMware supports x86 servers only, with an emphasis on the Microsoft OS. Microsoft's Virtual Server is anticipated to be replaced by Hyper-V, which is included as a

component of Microsoft Windows Server 2008. Another possibility is Xen, a hypervisor-based open source alternative.

Paravirtualization

Paravirtualization is the name for another approach to server virtualization. In this approach, rather than emulate a complete hardware environment, the virtualization software is a thin layer that multiplexes access by guest operating systems to the underlying physical machine resources.

There are two advantages to this approach. First, it imposes less performance overhead because it is a very small amount of code. Hardware emulation, you'll recall, inserts an entire hardware emulation layer between the guest operating system and the physical hardware. By contrast, paravirtualization's thin software layer acts more like a traffic cop, allowing one guest OS access to the physical resources of the hardware while stopping all other guest OSs from accessing the same resources at the same time.

The second advantage of the paravirtualization approach compared to hardware emulation is that paravirtualization does not limit you to the device drivers contained in the virtualization software; in fact, paravirtualization does not include any device drivers at all. Instead, it uses the device drivers contained in one of the guest operating systems, referred to as the *privileged guest*. Without going into too much detail about this architecture here, suffice it to say that this is a benefit, since it enables organizations to take advantage of all the capabilities of the hardware in the server, rather than being limited to hardware for which drivers are available in the virtualization software as in hardware emulation virtualization.

It might seem that paravirtualization would be the way to go. However, there has been one significant drawback to this approach to virtualization: Because it is lightweight and multiplexes access to the underlying hardware, paravirtualization requires that the guest operating systems be modified in order to interact with the paravirtualization interfaces. This can only be accomplished by having access to the source code of the guest operating system. This access is possible for open source operating systems like Solaris and Linux,

and is only possible for Microsoft operating systems with Microsoft source code access. The good news is that Quad-Core AMD Opteron™ processors featured within Sun x64 systems provide functionality that enables unmodified operating systems to be hosted by a paravirtualized hypervisor. Consequently, this drawback to paravirtualization will diminish as servers with these new chips take their place in production infrastructures.

One example of paravirtualization is a relatively new open source product called Xen, which is sponsored by a commercial company called XenSource. Xen is included in the recent Linux distributions from Red Hat and Novell, as well as being available for many community Linux distributions like Debian and Ubuntu. XenSource itself sells Xen-based products as well. Another possibility is Virtual Iron, a Xen-based solution.

Storage virtualization

The amount of data organizations are creating and storing is exploding. Due to the increasing shift of business processes to Web-based digital applications, every company is being inundated with data.

This explosion of data is causing problems for many of them. First, from a sheer storage capacity, many applications generate more data than can be stored physically on a single server. Second, many applications, particularly Internet-based ones, have multiple machines that need to access the same data. Having all of the data sitting on one machine can create a bottleneck, not to mention presenting risk from the situation where many machines might be made inoperable if a single machine containing all the application's data crashes. Finally, the explosion of machines mentioned earlier in the chapter causes backup problems; in other words, trying to create safe copies of data is a Herculean task when there are hundreds or even thousands of machines that need data backup.

For these reasons, data has moved into virtualization as well. Companies use centralized storage (virtualized storage) as a way of avoiding data access problems. Furthermore, moving to centralized data storage can help IT organizations reduce costs and improve data management efficiency.

Virtualization Makes Hardware More Important

Even though virtualization is a software technology, it has the effect of making hardware more important. This is because removing lots of servers and migrating their operating systems to virtual machines makes the remaining servers that support all those virtual machines even more important.

Although IT organizations could allow individual servers to go down in a “one application, one server” environment because each system failure would inconvenience a single user population, virtualization is very different. Each server supports multiple virtual machines and multiple user populations. If a virtualization host goes down, it may affect many applications and all of those applications’ users.

The importance of hardware is only going to increase as new, virtualization-ready hardware comes to market. There are significant virtualization capabilities being cooked up by hardware manufacturers, so don’t overlook the role of hardware in your virtualization infrastructure.

Consequently, as you move forward with virtualization software, the perhaps unexpected effect is that your hardware environment becomes more important.

Chapter 2

Understanding AMD Virtualization™ Technology

In This Chapter

- ▶ Looking at operating system state
 - ▶ Managing memory with AMD Virtualization (AMD-V™) Technology
 - ▶ Explaining AMD-V
-

Although you may consider a computer as just one of those boring pizza boxes (the term stems from the resemblance that 1U rack-mount servers have to the box that pizzas are delivered in, although we tend to believe that it's more reflective of the fact that pizza and technology are inextricably intertwined in the lives of true geeks), in fact a computer combines a number of different resources to enable the automated processing of data.

Four of these resources are crucial to virtualization:

- ✓ **Processor:** The central processing unit (CPU) is what turns random information into organized data. CPUs manipulate strings of characters, add and subtract numbers, and arrange for information to flow in and out of the system. As you'll remember from the previous chapter, virtualization enables a single physical computer to support multiple virtual guest systems. The ability to coordinate processor access by the separate guest systems is one of the main challenges of virtualization, particularly since the x86 processor was never really designed to support multiple guests.

- ✓ **Memory:** A computer contains physical memory to store the data that the processor manipulates. Similar to the processor, memory must be carefully managed to enable multiple guests to share a single set of physical memory without allowing separate guest systems to overwrite one another's data. And, as you might have guessed, x86 memory was not designed with multiple guest access in mind.
- ✓ **Network:** Today's computers are, by default, social; they communicate with one another as well as sending and receiving data from the great cloud that is the Internet. While data flows back and forth on the physical network card within a virtualized system, it's critical to ensure that each virtual system receives the appropriate network traffic.
- ✓ **Storage:** The fourth critical resource that is affected by virtualization is storage — data residing in a place that it can be retrieved from. If you've ever installed a hard drive in your own computer, you've managed storage! To repeat the refrain, each virtual guest system must have its own data storage and the virtualization software must keep each guest system's storage isolated.

Managing Operating System State

State is a term used within computing to (pardon the pun) state the obvious: at each moment in time, the operating system has a number of pieces of data that reflect its current condition. For example, if you are writing a document, the state reflects the sentences you've written, the file location on the disk where the document is stored, the individual values each system resource contains, and so on. The value that every resource has at a given moment of time is described as the operating system's state. Each of those values is stored in system memory — when you make a change to an individual character in your word processing document, a tiny bit of memory is changed to reflect the new state of your document.

In a virtualized world, where multiple guest operating systems share a single set of system resources, the ability to save one guest's state and restore another guest's state is vital.

Because each guest system needs to have its state in system memory so that it may operate, the virtualization hypervisor needs to be clever (and fast!) enough to swap state in and out of system memory so that each guest can share the system resources without trampling on one another's state.



State can be difficult to understand — it sort of reminds co-author Bernard of when he learned calculus — it didn't make any sense at all until the day it finally did — and then he realized the beauty of the mathematics that Isaac Newton put together. Think of state as a hotel room. Each guest brings his or her possessions to the room. Those possessions reflect the guest's state at that moment — a suitcase, certain clothing, toiletries, perhaps a briefcase containing a computer, paper notebook, and a book or two. When one guest leaves, another moves into the hotel room and brings his or her possessions. The hotel room is a virtual dwelling that is shared by all guests. Now, imagine if the guests all shared the room simultaneously, with each getting the room for a five minute stretch. You can see the challenge of unpacking and packing every five minutes — well, that's what virtualization accomplishes for operating systems. Virtualization enables operating systems to pack and unpack their state — except it happens thousands of times per second! You can see that the ability to save and restore state — to manage system memory which is where state resides — is *the* key capability for virtualization software.

Virtualizing Memory

Long before computer scientists came up with the notion of virtualizing an entire system, architects had already invented techniques to virtualize memory management. The Atlas computer at the University of Manchester was the first system to incorporate virtual memory technology. Virtual memory technology lets a system with a limited amount of physical memory look much larger to application software. To create this illusion, the OS stores the full memory image of the application and its data on the system's hard drive, and transfers required pieces of this image into the system's DRAM memory as the program executes.

To translate the virtual addresses seen by each application into physical DRAM memory addresses, the system relies on a map (known as a *page table*) that contains references linking chunks of virtual memory to real memory. Contemporary x86 processors include hardware features known as translation look-aside buffers (TLBs) that cache the translation references for recently accessed chunks of memory, thus speeding up the process. TLBs play a role in almost all memory references, so the manner in which they perform their translations can play a significant role in determining overall system performance.

Architects soon learned that TLB design can seriously impact multitasking systems operations. Most tasks in such systems have unique page tables. This forces the operating system to reset (or, more colorfully, “flush”) the TLB each time it switches from one task to another. Then, as the new task executes, its page table entries fill up the TLB, at least until the next task switch. This constant flushing and reloading can really eat into performance, especially if each task runs for only a few milliseconds before the next switch.

To mitigate the impact of task switching, architects added a task ID field to each TLB entry. This allows the system to retain the mapping information of multiple tasks in the TLB while switching between tasks, because it only uses the entries for the task actually executing at any point. This in turn eliminates the need for performance-inhibiting TLB flushes. At least until virtualization entered the scene.

Allocating memory the old way

Because the guest OS running on a virtual machine (VM) is unaware of other guests, it can only assign unique task IDs within its own environment. Thus, multiple VMs can have tasks with the same ID, confusing the TLB and making a real mess. There’s a simple solution to this problem — the hypervisor merely flushes the TLB every time it switches from one VM to another. This forces the tasks executing in the next VM to reload the TLB with its own page table entries. Unfortunately, this approach seriously impacts virtual system performance, giving architects everywhere déjà vu.

A better way of allocating memory

AMD's CPU architects had a better idea. They merely added a new, VM-specific tag called an address space identifier (ASID) to the TLBs in the AMD Opteron™ processors with AMD Virtualization technology. This concept is known as a tagged translation look-aside buffer, or tagged TLB.

Each VM has a unique ASID value, known only to the hypervisor and the TLB hardware. The ASID is invisible to the guest OS, thus eliminating the need to modify the guest, preserving the virtual illusion and avoiding any performance degradation. Figure 2-1 illustrates the tagged TLB concept. Figure 2-2 shows how competing processors handle memory in a virtualized environment.




Figure 2-1: Tagged translation look-aside buffer.




Figure 2-2: The non-AMD-based virtual machine's memory (shown on the right side of Figure 2-1) must be flushed every switch, which can slow performance.

Solving one problem . . . creating another

Although tagging the TLB allows a system to support multiple VMs without the need to flush the TLB when switching between tasks, it doesn't solve one sticky problem with managing virtualized memory. Remember, we're trying to fool the OS into thinking that it's in control of a real system, when in actuality all we've done is provide it access to a virtualized system with virtualized physical memory (also known as Guest physical memory). We still need to map from that Guest physical memory to the actual physical memory plugged into the system (also called Host physical memory). The way we do this in many cases is with a technique called *Shadow Page Tables*.

Without going into too many techie details, Shadow Page Tables are what connect Virtual Machines (VMs) to the actual hardware from a memory standpoint. The TLB translates virtual memory to Guest physical memory (which is still virtual), Shadow Page Tables translate Guest physical memory to Host physical memory. Although this process may sound fairly straightforward, there's a whole bunch of complex software required to manage all of this — and that extra management can have the negative side effect of slowing down virtualized applications. So what do smart CPU engineers do when they come across a complex software problem? They try to solve that problem quicker and more efficiently in hardware.

AMD's virtualization extensions

If you recall back from earlier in the chapter, CPU architects created a TLB to store page tables, which are maps that translate from virtual to physical memory. Modern CPU architects, being the smart people they are, have developed something called Virtualized Page Tables that translate all the way from virtual memory to Guest physical memory to Host physical memory. But unlike Shadow Page Tables, which perform the same function in software, Virtualized Page Tables are built directly into the CPU. Not stopping there, modern CPU architects also created a guest TLB, which is where these Virtualized Page Table translations are stored.

The best part of all of this memory-management hocus focus is that virtualized applications can now achieve near-native performance and responsiveness.

AMD Opteron™ Processor: The Green Chip

One of the big reasons companies are moving to virtualization is to reduce energy consumption. Big data centers can use incredible amounts of power, and anything that helps reduce power needs is welcome.

The new generation of chips from AMD is even more environment-friendly than before. Compared to the previous generation of chips, Third Generation AMD Opteron processors (also known as Quad-Core AMD Opteron processors) can save significant amounts of energy by cleverly adjusting power consumption according to processing loads.

Although it may seem that the power savings for one chip would be relatively unimportant, keep in mind that today's data centers can contain thousands of machines. Even a virtualized data center, where many physical machines have been converted to guest virtual machines, can contain hundreds of physical servers — and the energy savings of AMD Opteron processors with AMD-V can be enormous when spread across that number of servers.

By moving to Quad-Core AMD Opteron processors with AMD-V in a virtualized environment, users can save energy in two ways:

- ✓ Each machine uses less energy by dynamically adjusting to processing demands. This means less overall energy consumption (and saving money on energy costs).
- ✓ Because each machine uses less energy, it throws off less heat (chips generate heat as they process information). The reduced heat means less air conditioning is needed in the data center, further reducing energy consumption (and saving even more money on energy costs!).

So not only does AMD-V make your virtualization run more effectively, it makes your data center run more efficiently.

But AMD is not resting on its laurels. It has even more virtualization enhancements included in Quad-Core AMD Opteron processors, which are described in the next chapter.

Chapter 3

Looking into AMD's Virtualization Initiatives

In This Chapter

- ▶ Looking at AMD Virtualization™ technology with Rapid Virtualization Indexing
- ▶ Examining I/O virtualization

You may think that AMD would be content with its current advances designed to support virtualization. However, performance is an ongoing pursuit. The better the performance, the more responsive virtual guests are.

Even more important is the fact that as individual virtual machines improve, more of them can be supported on a given piece of hardware; in other words, if virtual machine performance goes up, more virtual machines can be squeezed onto a server, thereby achieving higher virtual machine density.



Virtual machine density refers to the ratio of virtual machines to physical machines. The higher the number of virtual machines that can be supported on a physical system, the higher the virtual machine density. Put another way, the higher the density, the lower the number of physical machines required to run an organization's virtual systems. And, since one of the main goals of virtualization is to get rid of physical servers, high density is very desirable, indeed.

AMD Manages Memory

One of the most important tasks for a hypervisor is memory management. The ability to keep track of the memory for individual processes within a virtual machine, and, perhaps more important, to ensure that each virtual machine's overall memory is managed is critical for virtualization. The latter task may be referred to as keeping track of virtual machine *state* — the settings of all critical system variables at each moment in time.

Keeping track of all those different bits of memory is critical to ensure that one virtual machine does not modify another's memory; after all, if a system's memory is compromised, none of its data can be trusted.

However, it's not enough that a hypervisor keeps accurate track of all the memory settings of the various guest virtual machines. Accuracy must be matched with speed, because if the hypervisor takes too long to swap memory in and out, the performance of the virtualization solution will be unacceptable (or, to use a fancy technical term, the achievable virtual machine density will be unacceptably low). It's not an overstatement to say that efficiently managing memory is the key determinant of virtualization performance.

In the early days of virtualization, all of the memory management was done by the hypervisor software. While the creators of the virtualization software were extremely smart, it's undeniable that software always runs slower than hardware. Therefore, figuring out how to move memory management into hardware would significantly improve performance and raise VM density.

In the previous chapter, we discuss the hardware advances in AMD Opteron™ processors with AMD-V™ technology for managing memory:

- ✓ The use of address space identifier (ASID) to segregate each VM's translation look-aside buffer (TLB).
- ✓ The use of Shadow Page Tables that enable a virtual guest's virtual memory to be mapped through to the physical memory of the underlying hardware.

In Quad-Core AMD Opteron processors, AMD implemented a further hardware optimization to memory management called Rapid Virtualization Indexing.

To understand what Rapid Virtualization Indexing accomplishes, it's important to understand the flow of memory in a virtualized environment:

1. The virtual machine operating system has its own virtual memory that enables the system to “pretend” its total available memory is larger than is really available. Page tables swap memory back and forth onto disk to enable this.
2. The virtual machine has its actual memory that this virtual memory is swapped into and out of as needed. In a virtualized environment, this actual memory is managed by the hypervisor and is, in fact, also virtual.
3. The hypervisor itself manages a pool of memory that may be larger than the physical memory available on the underlying server. It has pages that it swaps back and forth to disk to support this virtual memory.
4. Finally (at last!) there is the physical memory on the hardware system, which is where actual processing occurs.

This flow means that there are four steps between a virtual machine starting an instruction that requires memory access and the access itself.

AMD Virtualization technology with Rapid Virtualization Indexing removes one of those steps. It provides the ability to map the virtual machine's physical memory (remember, this is actually virtual, since it is provided by the hypervisor) directly to the actual physical memory of the underlying hardware. Please see Figure 3-1 for an illustration of AMD's Rapid Virtualization Indexing.

In essence, when the virtual machine first comes online, the hypervisor sets up the Shadow Page Tables that the VM requires, creating the mapping necessary for the four-step virtualization process to occur. It then creates a direct connection between the “physical” memory of the virtual machine and the actual physical memory of the underlying hardware. This enables subsequent memory accesses by the virtual machine to bypass the hypervisor virtual memory.

Hypervisor translates a page in guest virtual address space to machine physical space through a two-level translation

- First, map guest virtual address to guest physical address
- Then, map guest physical address to machine physical address




Figure 3-1: Design of AMD's Rapid Virtualization Indexing.

Cutting out a step increases the performance of the virtualized system. The performance increase is somewhat dependant upon the type of workload the virtual machine is executing; memory-intensive applications see more performance improvement than applications that are not heavily dependent on memory access.

Overall, however, reducing the number of steps required to manipulate memory, and moving the memory access functions into hardware has the net effect of improving performance and virtual machine density.

Extending Virtualization to Devices

The previous chapter described the big four resources of computing: processor, memory, network, and storage. And, if you've read that chapter and the material at the beginning of

this chapter, it's clear that AMD has done a great job in optimizing virtualization performance for processors and memory.

That still leaves network and storage. But never fear: AMD plans to implement improvements to the way network and storage interaction is handled to improve performance for these resources in a virtualized environment.

The changes for these resources go by the name I/O Memory Management Unit, shortened to the acronym IOMMU.

To understand these changes, it's important to recognize that in order for information to flow back and forth to these input/output devices (this is where the IO in IOMMU comes from), it first must be moved to memory that is accessible by the I/O device. Each I/O device connected to a computer has its own specific location in the system memory. The operating system knows just where that memory is, so when it wants to send data, say, across the network, it transfers data from the processor's memory to the I/O device's memory, where the I/O device can access it and send it on its merry way. When data returns, the I/O device puts it into its assigned memory location, where the OS grabs it and transfers it into the processor's memory, where it can be, well, processed.

This gets complicated in a virtualized environment. The guest operating system writes I/O bound data where it *thinks* the I/O device can grab it. The virtualization hypervisor cleverly intercepts the attempt by the guest OS to write to physical memory and maps it through its own memory and then on to the actual physical memory that the I/O device is attached to.

The hypervisor must keep track of and map all the different guests' virtual I/O memory locations and constantly swap the virtual representations of the guests' I/O memory locations into the actual physical I/O memory. As should be pretty clear, this requires really, really smart hypervisor software so that all of this I/O interaction can be kept straight. After all, you wouldn't want your CRM system to be reading data from your DNS (Domain Name Service) system, would you?



Of course, one wouldn't use the term "kept straight" in a complex computer science topic like virtualization — it sounds so . . . casual. The official term for keeping data straight is *data integrity*, which sounds much more formal and

well-thought-out. But data integrity just refers to the fact that data must be associated with the resource that's using the data. In this case, it refers to the fact that data must be reliably communicated only with the virtual machine that has sent it or is waiting to receive it. If the hypervisor fails to keep the data associations straight, it demonstrates poor data integrity.

As should also be clear, with I/O memory mapping, besides requiring careful management to ensure data integrity, the hypervisor must perform this task efficiently — very efficiently. This is particularly important because I/O is critical for overall system performance — after all, no computer operates without accessing data on a hard drive and in today's computing world applications typically interact with users or other systems across a network. So I/O performance is an important area that really needs optimization to ensure acceptable performance, and, of course, good virtual machine density.

AMD has taken the lead in moving functions originally performed by the hypervisor in software into hardware, and IOMMU is no different. In fact, IOMMU takes advantage of some of the same architectural approaches we've already seen in Shadow Page Tables.

In essence, IOMMU subdivides the I/O memory associated with an I/O device and allows a hypervisor to create dedicated subsections of the memory that may be assigned to virtual machines (see Figure 3-2). In this way, each virtual machine has a section of memory dedicated to its I/O use, which means the hypervisor can set up the original dedicated connection, and then let the virtual machine communicate directly with the I/O device without needing to be involved. This reduces the software processing overhead and improves performance.

Subdividing the memory assigned to an I/O device presents a challenge: how to keep track of each of the subdivisions, ensure that the I/O device places the appropriate data into the right memory subdivision, and also ensure that only the right virtual machine accesses that subdivision to get its assigned data.




Figure 3-2: How IOMMU works.

How does IOMMU meet this challenge? Remember our old friend the translation look-aside buffer (TLB)? Well, IOMMU takes advantage of TLBs to isolate the various subdivisions of the I/O memory. Furthermore, it uses a Domain ID to keep track of the assignments between particular memory areas and the guest virtual machines on the system.

By implementing IOMMU, a hypervisor can reduce the number of steps required to get data from a virtual machine out to a physical I/O device — and every step that can be bypassed means fewer software instructions that need to be executed, higher performance virtualization, and, ultimately, higher virtual machine density.

IOMMU is planned for future releases of AMD chips which will be Fourth-Generation AMD Opteron processors. This improvement requires virtualization hypervisors to be modified in order to take advantage of the hardware changes, so when Fourth-Generation AMD Opteron processors ship look for updated hypervisors that implement support for IOMMU. This also requires a chipset (the silicon chip that connects CPUs with I/O devices) that implements an IOMMU. AMD is planning an upcoming platform that will include a chipset that implements an IOMMU and also supports Fourth-Generation AMD Opteron processors.

With AMD-V, Rapid Virtualization Indexing, and IOMMU, AMD has addressed all four critical computing resources: processor, memory, storage, and network. By shifting virtualization functions from software to hardware, overall performance is improved, making virtualization even more capable and providing better virtualization capability to end users.

Chapter 4

Making Servers Virtualization-Ready

.....

In This Chapter

- ▶ Looking at commodity servers
 - ▶ Exploring servers that address computing's four key resources
 - ▶ Examining a beautiful relationship: Sun and AMD
-

The previous chapters described all the ways AMD is improving its chips to better support virtualization: AMD-V™, Rapid Virtualization Indexing, and IOMMU.

However, most of us don't just buy system resources. We buy entire computers and rely on the manufacturer to design and build servers that contain the key computing resources: processor, memory, storage, and network.

From a user perspective, it's vital to get the right servers to host virtualization, because, as Chapter 1 points out, virtualization makes hardware more important.

The New Role of Commodity Servers in a Virtualized World

Hardware is more important in a virtualized world because the nature of what servers do is changing. In a non-virtualized data center, the rule of thumb is “one application, one server.” This refers to the fact that most IT organizations tend to segregate applications, assigning each to its own physical server. This simplifies system accounting (every application has its

own hardware, making cost assignment much easier) and ensures that applications do not interfere with one another. With the vastly lower price of servers, it's been financially possible to support this "one application, one server model," although the proliferation of machines has caused its own problems as described in Chapter 1: overcrowded data centers stuffed with underutilized machines, each of which takes a full ration of power and cooling.

Virtualization, by contrast, breaks the "one application, one server" mold. Rather than support just one system on a server, virtualization supports many systems on a server. And that makes the hardware more important, because more is riding on the availability of each piece of hardware.

With Virtualization, Servers Need More of the Four Key System Resources

When a server is used to host a number of virtual machines, it is faced with much higher levels of demand for system resources than would be presented by a single operating system running a single application.

Obviously, with more virtual machines running on the server, there will be more demand for processing. Even with two processors, virtualization can outstrip the processing capability of a traditional commodity server.

Also, with more virtual machines on the server, there will be far higher storage and network traffic as each virtual machine transmits and receives as much data as would be demanded by a single operating system performing in the old "one application, one server" model. Furthermore, because virtualization makes the robustness of hardware more important, most IT organizations seek to avoid so-called Single Point of Failure (SPOF) situations by implementing redundant resources in their servers: multiple network cards, multiple storage cards, extra memory, and multiple processors — all doubled or even tripled in an effort to avoid a situation where a number of virtual machines (and user populations) can be stalled due to the failure of a single hardware resource.

Finally, the lack of higher amounts of memory can severely impact virtual machine performance. The available system memory must be shared among all the virtual machines, not to mention the memory used by the hypervisor itself.

Although 4 GB of memory may be sufficient to support a single operating system and application, it can limit server responsiveness in an environment in which the server is asked to support 5, 10, or even 20 virtual machines. And processor advancements require more memory today and will require even more in the future. In fact, the single biggest bottleneck experienced by IT organizations when they implement virtualization is inadequate memory, because too little memory forces additional page swapping, thereby impacting system performance.

To put the matter another way, the availability of system resources directly affects the achievable virtual machine density for a given server: the more resources, the higher the achievable density. And, because one of the main motivations for IT organizations to move to virtualization is to reduce the number of physical servers in their data centers and thereby increase the overall virtual machine density, it's obvious that resource availability is the critical determining factor in your virtual machine density level.

Consequently, the system design goals of traditional commodity servers, perfectly adequate for the "one application, one server" environment, may no longer be sufficient for a virtualized data center.

Fortunately, Sun Microsystems created a new generation of servers based on AMD Opteron processors that marry the virtualization improvements of AMD Virtualization technology with innovative system designs from Sun that address the resource requirements of virtualization.

Sun Microsystems' Servers Based on AMD Opteron™ Processors

Clearly, servers intended for a virtualized environment should utilize the latest generation of chips that are optimized for virtualization.

This means that these servers should use AMD Opteron™ processors as their processing foundation. Sun was an early proponent of AMD Opteron processor-based systems, recognizing that AMD Opteron processors with Direct Connect Architecture were a breakthrough design.

The advancements represented by AMD Opteron processors with AMD-V enable Sun servers to offer high performance, implemented by the processor and memory improvements outlined in Chapters 2 and 3.

Furthermore, by using AMD Opteron processors, Sun is able to offer exceptional energy efficiency because of the low power consumption characteristics of AMD Opteron processors.

Sun provides an entire line of AMD Opteron processor-based systems beginning with dual processor machines and scaling on upwards. And, remember, those are multi-core processors, so each chip represents multiple processing units. Today's Sun machines offer dual-core processors and with the new Quad-Core AMD Opteron processors, Sun will offer quad-core systems.

The world of virtualization demands servers that offer much larger amounts of system resources. Sun has taken a leadership position in this arena with its Sun Fire™ x64 servers and Sun Blade Modular Systems. These servers are designed from the ground up to support virtualization.

The most powerful Sun x64 rackmount server, the Sun Fire X4600 M2 Server (see Figure 4-1), offers the following resource capabilities:

- ✓ Up to eight Quad-Core AMD Opteron processors in a single chassis
- ✓ Up to 256 GB of memory to ensure adequate amounts of memory for even the most demanding virtualization environments
- ✓ Four gigabit ethernet ports to allow multiple network connections and avoid network SPOF
- ✓ Eight expansion slots to allow multiple storage connections and avoid storage SPOF




Figure 4-1: Sun Fire X4600 M2 Server.

The Sun Blade 8000 Modular System (see Figure 4-2), among the most powerful blade platforms on the market, offers the following resource capabilities:

- ✓ Runs up to four Quad-Core AMD Opteron processors per server module, with up to 20 server modules in a single rack
- ✓ Up to 128GB of memory per server module, 2.56TB per rack
- ✓ Up to 192Gbps I/O throughput per blade

The innovative Sun Blade 6000 Modular System enables virtualization deployment with few servers — helping reduce costs, power consumption, and complexity.

Furthermore, because of the Sun x64 line's innovative design and the incorporation of AMD Opteron processors, these systems can consume much less energy than comparable commodity x86-based systems, thereby making them perfect for green data center initiatives.

Using the new generation of servers

To provide a concrete example of how the new generation of servers delivers real virtualization capability, here's an example. Ventyx (formerly known as NewEnergy, a Siemens subsidiary) is an energy consulting and software company. It faced a common problem: running out of space in its data center, coupled with skyrocketing power and air conditioning costs.

Using a Sun Fire X4200 server and virtualization software from VMware,

Ventyx was able to consolidate 18 1U servers onto a single 4U machine with a 30 percent heat savings. Plus, additional capacity remains on the machine, providing headroom for future growth in computing needs.

As this example shows, the capability of the new generation of hardware is impressive indeed.

For more information, visit www.sun.com/customers/servers/newenergy.xml.




Figure 4-2: Sun Blade 8000 Modular system.

Chapter 5

Managing Virtualization End-to-End

In This Chapter

- ▶ Examining Sun server virtualization
 - ▶ Looking at Sun desktop virtualization
 - ▶ Perusing Sun storage virtualization
 - ▶ Combining virtualization and management
-

You might be forgiven if you believe that virtualization is a server-only technology. Most of the attention in the industry has focused on the area of server consolidation.

But the fact is that virtualization applies throughout the enterprise. This chapter will describe how Sun can help you virtualize all aspects of your computing infrastructure.

Sun Virtualizes the Entire Enterprise

Although the data center has proven a fruitful area for server virtualization, servers only represent a portion of enterprise computing. In fact, the typical data flow of a transaction goes through all these layers of the enterprise computing topology:

- ✓ **Interactions at the client level:** Someone, somewhere decides to do some work, whether create a spreadsheet, write a document, enter a transaction, or retrieve some data. All of this requires a client platform for the individual to interact with. Client platforms (for instance, a PC)

are some of the most underutilized computing platforms in the enterprise — sporadically used during the work day, and then sometimes left running during the 12 to 16 hours outside of working hours.

- ✓ **Processing at the server level:** After an individual kicks off a computing process, data typically flows to a server-based application, where it must be operated upon, or processed. This has been the primary area of focus for virtualization heretofore.
- ✓ **Stored and retrieved at the storage level:** It's called data processing because the key element is data — information. Without a reliable way to store and retrieve data, all of the other components in the enterprise computing environment are useless. A further complication is that most storage today is isolated in islands of direct-attached storage; that is, hard drives inside of servers, where the data can't be conveniently shared among servers and applications.

In order for an enterprise to fully take advantage of virtualization, it's vital that it look to all of these areas to assess how virtualization might be applied.

It might surprise you that virtualization has a role to play beyond the data center, but in fact the concept of virtualization — the abstraction of logical resources from physical ones — may be applied to clients, servers, and storage. Each of these layers can benefit from having virtualization applied to the current mode of operation.

Sun has created virtualization solutions for all layers of the enterprise, and offers options for every customer to realize the complete benefits of virtualization.

Tying Virtualization Together End-to-End

Sun provides a complete range of products to offer virtualization for all elements of the IT infrastructure — from client through server all the way to storage. Sun addresses all the layers of the IT infrastructure, offering solutions that not only

virtualize each layer but also tie the entire infrastructure together into one unified virtualized environment.

Sun provides the ability for IT organizations to take advantage of hardware and software improvements to achieve the goals of virtualization: better hardware utilization, reducing data center sprawl, green initiatives through lower energy consumption, and reduced administrative costs through reducing the number of machines that must be managed.

Sun Server Virtualization

In the previous chapter, we discuss AMD Opteron™ processor-based Sun servers and the processing and energy advantages they provide. However, the benefits of Sun's server virtualization initiatives go well beyond the hardware capability of AMD Opteron processor-based machines.

Sun's advanced hardware designs enable individual servers to be divided into separate domains, isolated from one another to ensure no interference between one domain and another. Each domain can contain one or more virtual machines. Domain isolation provides assurance to customers that different systems have no way of interfering with one another's integrity (there's that integrity word again!).

Sun also supports several different types of server virtualization:

- ✓ **Operating System (OS) virtualization:** As discussed earlier, OS Virtualization is a form of virtualization in which the native OS exports libraries so that applications have the "illusion" that they are operating in separate OSs. Solaris™ Containers, an integral part of Sun's Solaris™ OS, isolates software applications and services using flexible, software-defined boundaries, allowing many private execution environments to be created within a single instance of the OS.
- ✓ **Hard partitions:** This is the capability that enables a single system to be broken into separate domains to ensure isolation.

✓ **Virtual machines:** This is the type of virtualization that most people think of when they hear the term virtualization. Each virtual machine contains a completely separate operating system, each with its own application or applications. The isolation between virtual machines is complete, with the hypervisor ensuring that virtual machines can't access one another's applications or data. Both VMware ESX and Sun™ xVM Server run as the primary application on a dedicated system, with guest operating systems running on top of them. Sun™ xVM VirtualBox provides developers a way to create multiple guest OSs on top of their existing laptop or workstation. Developers can get started quickly, then move their virtual machine images onto a production server running VMware ESX or Sun xVM Server.

By providing a complete range of server virtualization solutions, Sun ensures that its customers can apply the type of virtualization best suited for their environment and needs.

Sun Desktop Virtualization

The desktop can be one of the most wasteful areas of the enterprise in terms of capacity utilization and energy use.

Sun provides the Sun Virtual Desktop Solution with a typical power consumption of only 4 watts, which can represent a substantial improvement over a traditional desktop solution.

The Virtual Desktop Solution applies the following technologies:

✓ **The physical desktop device:** This can be a traditional standalone PC or a SunRay virtual display client. Unlike the traditional standalone clients, in the Virtual Desktop Solution the client doesn't carry the storage of applications and no client configuration administration is necessary. The desktop device is used to display data and interact with the user, but all processing takes place on the backend server.

- ✓ **Sun Virtual Desktop Infrastructure Software:** Many desktop environments can be hosted on a single server, with Sun VDI Software providing the bridge allowing users to access their desktop environments from traditional PCs and Mac OS X computers, as well as thin clients from Sun and other vendors. Each virtual desktop functions as though it were running directly on the user's computer, but critical data is kept in the data center where it can be more easily managed and be less susceptible to loss or theft.
- ✓ **Sun Servers:** The powerful, energy-efficient Sun servers run each client instance, offloading processing from the physical desktop device onto a backend server. This enables lower-spec client machines to have their life-cycles extended, helping make use of corporate capital more efficient.
- ✓ **Sun Storage:** With the physical desktop device no longer used as a repository of data, data storage can be centralized, which is not only more efficient, but can also raise utilization rates of data devices.
- ✓ **VMware's Virtual Desktop Infrastructure:** This software runs each client instance as a virtual machine on the backend server, keeping all configuration and administration in the data center instead of on the office floor.

The move to desktop virtualization is less well-established than server virtualization, but it is rapidly growing in importance.

Desktop virtualization holds the potential of enormous financial rewards because huge amounts of IT dollars are tied up in desktop system administration — installing and reinstalling operating systems, keeping them patched, ensuring the data on the machines is backed up, and keeping antivirus and anti-spam software current.

By implementing Sun's desktop virtualization solution, organizations can achieve significant savings and reduce IT burdens in a time of stretched human resources.

Sun Storage Virtualization

The move to server and client virtualization means a change in how data is stored. Traditionally, data has been associated

directly with the machine generating the data. This type of storage, known as Direct-Attached Storage (DAS), is convenient to implement but problematic in a virtualized environment.

Every virtual machine requires its own data; while managing the data for one machine was not difficult in a pre-virtualized environment, attempting to manage local storage for 10 or 20 virtual machines located on a single physical server becomes a logistical nightmare.

Beyond the difficulty of ensuring sufficient DAS storage to support many virtual machines, a further difficulty is raised if an IT organization wishes to move to more advanced forms of virtualization, such as virtual machine migration or server pooling, in which individual virtual machines are automatically placed on one of a number of physical hosts, based on where the virtualization management software decides the virtual machine should be placed.

In these more advanced forms of virtualization, having a virtual machine's storage locked to a single server is unworkable — while local storage may be perfect if a virtual machine never moves from its original location, there is a strong likelihood that a machine migrated to another location may not be able to access its data located on the original machine.

Of course, in server pooling, it is uncertain where a virtual machine will be instantiated; since the virtual machine can be placed on any physical server in the server pool, there can be no concept of local storage.

Fortunately, storage itself is now being virtualized. In a virtualized environment, storage is moved off local systems and into a remote storage environment, where it can be accessed by virtual machines, no matter where they are located.

Storage virtualization pays benefits beyond making storage available to virtual machines independent of their location.

By migrating data to a specialized environment, storage virtualization can accomplish the following:

- ✓ **Better utilization of storage resources:** Local storage may be inefficiently allocated — one system may have its drives completely filled, while another has huge amounts of unused capacity. By moving all storage into a central location, storage virtualization can ensure that every virtual machine has adequate storage capacity while not wasting money on excess capacity.
- ✓ **Easier expansion of storage resources:** IT organizations seem to have an insatiable hunger for storage. Trying to increase capacity in a DAS environment is administratively complex, not to mention the potential to outstrip capacity on individual machines. This is particularly likely to happen in a virtualized environment where many virtual machines share the DAS storage of a single server.
- ✓ **More efficient management of storage resources:** By moving data from individual servers to a centralized location, it is easier to manage the data resources of the IT organization; moreover, critical data management tasks like backup are easier to track and perform when data resides in a single location rather than being spread throughout the data center.

Sun provides a full range of networked storage options that can take advantage of server virtualization technologies. In addition, Sun offers products specifically designed to virtualize storage assets — disk or tape — independent of the virtualization scheme deployed on the server side.

Tying the Virtual Environment Together: Virtualization Management

For IT organizations, virtualization presents a blessing . . . and a curse. All the benefits of virtualization bring their own challenge: complexity. The ease of instantiating new virtual machines, migrating them from one server to another, accessing virtualized storage — well, it's a long way from “one application, one server” with that one application talking to on-board storage.



All this complexity cries out for more powerful management capabilities, which Sun delivers with its Sun xVM Ops Center, a next-generation management tool designed to ease the administrative challenge of managing a complex heterogeneous environment of different operating systems comprised of both physical and virtual servers, all of which must be provisioned, updated, managed, and kept track of. Sun has designed Sun xVM Ops Center to enable IT organizations to meet the management challenges that virtualization presents.

Key capabilities of Sun xVM Ops Center include:

- ✓ **Discover:** Sun xVM Ops Center can identify all the servers in your network, whether physical or virtual, even if they're powered off. Because one of the major challenges of virtualization is virtual machine sprawl, brought on by the ease of virtual machine instantiation, this capability provides the ability to track all of the assets present in a data center infrastructure.
- ✓ **Provision:** Sun xVM Ops Center makes it easy to remotely install virtual machine operating systems, packages and RPMs, as well as firmware. This reduces the need to physically log on to guest virtual machines as well as offering centralized control of the provisioning process.
- ✓ **Update:** In a complex, heterogeneous environment, one of the biggest challenges is ensuring that all operating system instances are kept up-to-date with versions, patches, and bugfixes. This is particularly important with regard to security-related fixes. Sun xVM Ops Center provides patch management functionality to ensure that all Red Hat, SUSE, and Solaris operating system instances are kept up-to-date and secure.
- ✓ **Manage:** Server sprawl means that many more systems need to be managed in terms of user management, disk utilization, system performance, and so on. Sun xVM Ops Center provides the ability to manage all systems remotely from a centralized location, thereby easing the burden of system management.
- ✓ **Report:** Sun xVM Ops Center enables IT compliance tracking by providing a compliance auditing solution. Using Sun xVM's Ops Center audit reporting capability ensures that IT organizations can fulfill their audit requirements quickly and completely.

Sun xVM Ops Center allows organizations to gain the full benefit of virtualization. By providing a fully-rounded management capability, Sun xVM Ops Center offers IT organizations the opportunity to efficiently use virtualization as a tool to better meet business requirements while reducing the administrative burden of a heterogeneous environment.

Professional Services

Sun offers a complete set of professional services to assist customers with their virtualization initiatives. Working with Sun's Professional Services organization can help companies to:

- ✔ Meet mandates to reduce operating costs (including utility spending, maintenance, administration, and management) while improving service levels
- ✔ Build a forward looking, next generation data center that is environmentally responsible and uses best-in-class virtualization technology to both plan and execute
- ✔ Get expertise and support in developing a roadmap for change that provides a rapid return on investment and aligns with their business drivers

Chapter 6

Ten Steps to Virtualization Success

.....

If you've read the rest of the book, you're probably raring to go with your virtualization project. To cut to the chase, here are ten key things for you to do when you get started.

- ✓ **Don't wait for "all the kinks to be worked out."** Start small and experiment. You find out more by doing than you do by reading or even talking to others whose organization has implemented virtualization. Don't wait for "all the kinks to be worked out," because there are going to be kinks for the foreseeable future. You need to get on board today.
- ✓ **Don't skimp on training.** Because virtualization is a new technology, you can't count on people already knowing how to use it. It's critical that you understand that there will be a period of learning as employees get up to speed on the new infrastructure. Don't compound the challenge by failing to educate employees on how to use and manage the new software.
- ✓ **Don't imagine that virtualization is static.** Not only will your business conditions dictate that you continually evaluate how well your virtualization infrastructure meets current business realities, but virtualization itself is constantly changing. This means that your state-of-the-art virtualization solution implemented 18 months ago may need to be examined in light of new virtualization developments.
- ✓ **Don't overlook a business case.** In these times of short rations for IT organizations, there's no surer way to get your project shot down than by ignoring the business case for it. On the other hand, there's no surer way to ensure your project gains executive support and sails through the approval process than by demonstrating the impressive financial benefits available by moving forward with the project.

- ✔ **Don't overlook the importance of hardware.** Virtualization is software that enables other software resources to take better advantage of underlying hardware. But don't imagine that the hardware itself has no effect on virtualization. Far from it. The type and capability of the hardware you use to host your virtualization solution can dramatically impact the virtualization density you achieve, as well as the performance levels available for your virtual machines.
- ✔ **Examine your administrative processes.** Virtualization can reduce the administrative burden by managing huge numbers of machines, not to mention simplifying tasks like backup. Examine your administrative processes to determine what tasks can be reduced or replaced by more virtualization-appropriate ones.
- ✔ **Look at the virtualization possibilities throughout your enterprise.** While every movie has its stars, it's often the case that great performances are given by actors who get less attention. This is the case with virtualization as well. Keep storage and client virtualization in mind as you move forward with your virtualization initiatives.
- ✔ **Find a management solution that incorporates virtualization.** Take a look at the management tools your hardware providers have available. Many of them have been extended to incorporate virtualization management into the existing software and hardware management already present. An integrated management tool can help reduce your burden and keep your administrative personnel happier. The Sun™ xVM family helps address virtualization and management of both physical and virtual, multiplatform Linux, Windows, and Solaris™ environments.
- ✔ **Collaborate with leading virtualization providers.** Most of the focus in the world of virtualization has been on the hypervisor providers: VMware, Xen, and Microsoft. Certainly a capable hypervisor is a prerequisite for a successful virtualization project. However, don't overlook the ability of hardware providers to make your virtualization initiative more successful. By relying on virtualization-enabled processors from AMD and virtualization-focused hardware from Sun, you can be assured that the hardware supporting your virtualization project will be robust and high-performing.
- ✔ **Don't forget to have a project party.** Last, but not least, be sure to celebrate your virtualization success.



**Get more from your
IT investments!**

Harness ever-increasing levels of computer performance

It seems like everywhere you go these days, someone is talking about virtualization. Technical magazines trumpet the technology on their covers. Virtualization sessions are featured prominently at technology conferences. And, predictably enough, technology vendors are describing how *their* product is the latest word in virtualization. This book helps you understand how virtualization works and whether it's right for you. It also discusses Sun and AMD technological offerings, how they work together, and how they can benefit your business.

**THE
DUMMIES
WAY**

Explanations in plain English
"Get in, get out" information
Icons and other navigational aids
Top ten lists
A dash of humor and fun

Discover how to:

*Save energy, time, and
money*

*Allocate memory where
it's needed*

Improve scalability

*Understand the different
types of virtualization*

*Achieve virtualization
success*

Get smart!
@ www.dummies.com

- ✓ Find listings of all our books
- ✓ Choose from many different subject categories
- ✓ Sign up for eTips at etips.dummies.com

ISBN: 978-0-470-29264-8
AMD tracking number: 44911-A
Sun tracking number: GNOT14021-0
Book not for resale

For Dummies®
A Branded Imprint of

