

Making Everything Easier!™

2nd Sun and AMD Special Edition

Virtualization

FOR DUMMIES®

Learn to:

- Save energy, time, and money
- Allocate memory where it's needed
- Improve scalability
- Understand the different types of virtualization



Bernard Golden



Virtualization
FOR
DUMMIES®
2ND SUN AND AMD SPECIAL EDITION

by Bernard Golden



WILEY

Wiley Publishing, Inc.

Virtualization For Dummies®, 2nd Sun and AMD Special Edition

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

Copyright © 2009 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 Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://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.com, Making Everything Easier, 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 Hyper Transport 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 877-762-2974, 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: 978-0-470-53269-0

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

Acquisitions, Editorial, and Media Development

Project Editor: Jennifer Bingham

Editorial Manager: Rev Mingle

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: Jeff Jones, Adrienne Miller, and Tim Mueting



Table of Contents

.....

Introduction	1
About This Book	1
Icons Used in This Book.....	2
Chapter 1: Wrapping Your Head around Virtualization . . .	3
Why Virtualization Is Hot, Hot, Hot	3
Sorting Out the Types of Virtualization	6
Virtualization Makes Hardware More Important	10
Chapter 2: Understanding AMD-Virtualization™ Technology	11
Managing Operating System State	12
Virtualizing Memory	13
AMD Opteron™ Processor: Moving toward Green	17
Chapter 3: Looking into AMD's Virtualization Initiatives	19
AMD Manages Memory	20
Extending Virtualization to Devices	23
Chapter 4: Making Servers Virtualization-Ready	27
The New Role of Commodity Servers in a Virtualized World	27
With Virtualization, Servers Need More of the Four Key System Resources	28
Sun Microsystems' Servers Based on AMD Opteron™ Processors	30
Chapter 5: Managing Virtualization End-to-End	35
Sun Virtualizes the Entire Enterprise	35
Tying Virtualization Together End-to-End	36
Sun Server Virtualization	37
Sun Desktop Virtualization	38
Sun Storage Virtualization	39
Tying the Virtual Environment Together: Virtualization Management	42
Sun Services	44
Chapter 6: Ten Steps to Virtualization Success	45

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 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 operating system and a Microsoft Windows operating system on one server. 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 is likely to usher in an entirely new wave of hardware and software innovation. For example, and among other benefits, virtualization can 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 transporting 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, 2nd 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.

This book was sponsored by and written in collaboration with Sun and AMD.

Icons Used in This Book

In the margins of this book, you find several helpful 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, data centers may 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 still be high.

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 may have up to 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 by enabling a single piece of hardware to seamlessly support multiple systems, that's what virtualization does. By applying virtualization, organizations can raise their hardware utilization rates dramatically, thereby helping to use corporate capital much more efficiently.

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 too 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. But many servers require a large amount of space. This can cause a real estate problem for companies: They're running out of room. Making the space problem more difficult is the fact that, with an explosion of new data comes the need to store that data someplace.

Virtualization, by offering the ability to host multiple guest systems on a single physical server, allows organizations to reclaim data center territory, thereby helping to avoid the

expense of building out more data centers. This can be an enormous benefit of virtualization, because data centers can cost in the tens of millions of dollars to construct.

Trend #3: Green initiatives seek better energy efficiency

Years ago, power costs factored into strategic thinking at about the same level as what brand of soda to keep in the vending machines. Companies could assume that electrical power would be 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 of whether there was a true power shortage, the events caused many 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.

Recently, companies have turned this focus to their data centers. 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 a large number of servers, coupled with the fact that many of the machines filling up data centers are running at low utilization rates, means that a company that can reduce the total number of physical servers may significantly reduce its overall energy costs. This is another way in which virtualization can help.

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 can be 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.

Virtualization can help rein in operations cost increases by reducing the overall number of machines that need to be taken care of. Although many of the tasks associated with system administration (for instance, OS and application patching, doing backups, and so on) 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, thereby making virtualization an excellent option to help 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 benefits and drawbacks.

Server virtualization

Server virtualization enables you to consolidate many different types of workloads and operating systems onto virtual environments all running on a single hardware platform. Virtual servers or virtual machines are independent operating environments that use virtual resources. One of the most common approaches to virtualization is to use hypervisor technology. Today’s hypervisors provide the greatest level of flexibility in how virtual resources are defined and managed and have become the primary choice for server virtualization. Hypervisors use a thin layer of code in software to achieve fine-grained, dynamic resource sharing.

There are two types of hypervisors. Type 1 hypervisors run directly on the system hardware, whereas Type 2 hypervisors run on a host operating system that provides virtualization services such as I/O device support and memory management. Virtualization solutions that use a Type 2 hypervisor are also referred to as operating system (OS) virtualization, and in some environments are called containers.

Type 1 hypervisor solutions

You will find that Type 1 hypervisors are typically the preferred approach for server consolidation because they can achieve higher virtualization efficiency by dealing directly with the hardware. Type 1 hypervisors provide higher performance and efficiency and use hardware assisted virtualization technology like AMD-V™, now found in today’s multi-core AMD Opteron™ processors.

Type 1 hypervisors use a thin layer of code to provide resource sharing within a single hardware platform. Another way to look at it is that the hypervisor provides a standard emulated hardware environment that the guest OS, sometimes referred to as the virtual machine (VM), resides on and interacts with. The VM encapsulates the guest operating systems and the application into a single entity that provides isolation from the underlying hardware. It is because of this encapsulation that the VM can be migrated from one physical machine to another without any service interruption.

Not only does this approach support running multiple VMs, it can support multiple VMs running different types and/or versions of operating systems (for example, completely different OSs like Windows and Linux can be run simultaneously on the same physical server).

Type 1 hypervisor solutions are often used for server consolidation to achieve higher levels of resource utilization. Software development and quality assurance environments will also benefit greatly from this type of virtualization due to the ability to allow a number of different operating systems to be run simultaneously. This can facilitate parallel development or testing of software in a number of different operating system environments in a quicker, and potentially more efficient, manner.

Early Type 1 hypervisor solutions experienced some performance degradation due to the overhead of the virtualization software, but with the introduction of hardware-assisted virtualization into today's multi-core processors and platform improvements, such as expanded memory and I/O capabilities, this performance degradation has been minimized. In fact, there is evidence of some native applications running better in a virtual environment due to these improvements.

Type 1 hypervisor solutions include ESX Server from VMware, Hyper-V from Microsoft, XenServer from Citrix, and xVM from Sun. Another possibility is Xen, an open source alternative.

Operating system virtualization — type 2 hypervisor

Type 2 hypervisors run on a host operating system that provides virtualization services such as I/O device support and memory management. These services give 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.

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 must offer the same operating system as the host OS and even be consistent in terms of version number and patch level. As you may imagine, this can cause problems if you want to run different applications in the containers, because 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 Parallels, which offers the commercial product Virtuozzo as well as sponsoring the open source operating system virtualization project called OpenVZ.

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, companies are being inundated with data.

This explosion of data can cause problems. 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 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, as mentioned earlier in the chapter, the increase in numbers of machines 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.

For example, in an organization that has a “one application, one server” environment, if a single server went down, this would inconvenience a single user population. However, virtualization is very different. If each server supports multiple virtual machines, then a server that goes down would inconvenience multiple user populations. Thus, the hardware can take on greater importance in a virtual environment.

This shift in importance can be seen through the introduction of virtualization-ready hardware such as AMD Virtualization™ technology found in all AMD multi-core processors. Other hardware manufacturers have introduced similar enhancements to support virtualization implementations, so don't overlook the role of hardware in your virtualization infrastructure.

Chapter 2

Understanding AMD-Virtualization™ Technology

In This Chapter

- ▶ Looking at operating system state
 - ▶ Managing memory with AMD Virtualization Technology
 - ▶ Explaining AMD's technology
-

Although you may consider a computer as just one of those boring pizza boxes (the term arose 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.



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 in quick succession, 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, allowing the system TLB to be subdivided into sections, each dedicated to a single task, each with its own task ID. 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

In the virtualized world, multiple virtual machines (each with its own OS) run on top of the hypervisor. Each OS assigns its own task IDs, which can lead to confusion — because the hypervisor has a single TLB and each OS pushes memory with its own task IDs into the TLB — raising the potential that two separate OSs might assign identical task IDs to the system TLB. This would lead to confusion, possible data corruption, and security issues. The simple solution to this problem is to have the hypervisor flush 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 replicates the same issue that led to TLBs in the first place — performance-sapping memory copies.

AMD-V: Better virtualization through better allocation

Clearly, a more efficient approach to managing hardware resources would help overall virtualization performance. To that end, AMD introduced AMD Virtualization Technology (AMD-V) with its Dual-Core AMD Opteron™ processors in 2006. AMD-V is designed to improve virtualization security and performance with processor-based hardware enhancements. Today, AMD-V is comprised of two types of technology:

- ✔ **AMD-Vc**, which offers a hardware-assisted CPU and memory improvements. AMD-Vc is designed to help virtualization software run more efficiently and hide the complexity of hardware infrastructure.
- ✔ **AMD-Vi**, which offers hardware-assisted chipset enhancements designed to improve I/O operations (for instance, disk and network data transfer).

AMD-V also leverages Direct Connect Architecture with an integrated memory controller to provide fast and efficient memory management.

So how did this play out with respect to that pesky memory management? AMD chip architects added a new, VM-specific tag called an *address space identifier* (ASID) to the TLBs in the AMD Opteron processors. 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.

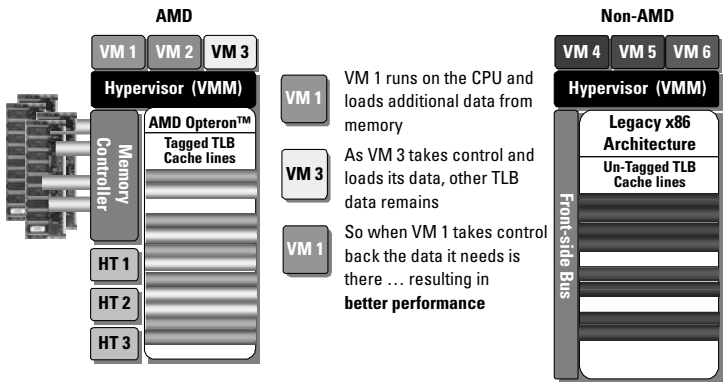


Figure 2-1: Tagged translation look-aside buffer.

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, you're trying to fool the OS into thinking that it's in control of a real system, when in actuality all you've done is provide it access to a virtualized system with virtualized physical memory (also known as guest physical memory). You 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 to do this, in many cases, is with a technique called shadow page tables.

Shadow page tables connect virtual machines (VMs) to the physical memory managed by the hypervisor. The TLB translates virtual memory to guest physical memory (which, in a virtualized environment, is still a virtual representation); shadow page tables translate guest physical memory to host physical memory. Although conceptually simple — a straightforward mapping via the shadow page table keeps the relationship straight between the virtualized guest physical memory and the hypervisor-managed system physical memory — the overhead of the mapping can impose a performance hit on

guest systems. 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 rapid virtualization indexing (RVI) that translates all the way from virtual memory to guest physical memory to host physical memory. But unlike shadow page tables, which perform these translations in software, RVI performs the memory translation in the CPU. Not stopping with RVI, modern CPU architects also created a guest TLB, which is where these nested page table translations are stored.

The net effect of these hardware extensions is that performance-sapping software translations via shadow page tables are no longer required. The mapping between the virtualized guest physical memory and hypervisor-managed system physical memory is performed in hardware, which can improve overall performance and reduce system overhead.

AMD Opteron™ Processor: Moving toward Green

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 is welcome.

The new generation of chips from AMD is even more environmentally-friendly than their predecessors. Compared to the previous generation of chips, 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 insignificant, 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:

- ✔ By dynamically adjusting to processing demands, each machine can tailor the amount of energy needed to the specific job. Thus, users can save on overall energy consumption by using energy on an as-needed basis.
- ✔ Using less energy means less heat in the chip (chips generate heat as they process information). By reducing the heat generated by the machine, users will need less air conditioning in the data center, further reducing energy consumption (and saving even more money on energy costs!).

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

But AMD isn't resting on its laurels. It has even more virtualization enhancements included in Quad-Core AMD Opteron processors (see Chapter 3 for more info on this).

Chapter 3

Looking into AMD's Virtualization Initiatives

In This Chapter

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

You may think that throwing raw power at a virtualization solution is the key to great performance. However, the optimal solution for virtualized environments is balancing raw performance and energy efficiency — also known as performance-per-watt.

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.

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 importantly, 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 technical lingo, 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 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. This led to replacement of shadow paging by nested page tables, which moved page use tracking from software in the hypervisor to hardware in the processor, thereby improving performance.

In Quad-Core AMD Opteron processors, AMD implemented a hardware optimization to memory management called Rapid virtualization indexing (RVI).

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. (This is called guest virtual memory.)
2. The virtual machine has 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. (This is called guest physical memory.)
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 is called host physical memory.)

RVI is the process of translating these virtual memory addresses to physical memory addresses. Virtualization introduces another level of translation — that of the guest (or virtual machine) virtual address to the guest physical address which is then translated to the host physical address. See Figure 3-1.

As discussed earlier, software-only virtualization solutions used a technique called shadow paging. Shadow paging uses stored information about the physical location of guest memory in shadow page tables and managed these structures in the hypervisor. Under shadow paging, the hypervisor intercepts guest page table updates to keep the shadow page tables in synch with the guest page tables.

Rapid virtualization indexing (RVI) now provides a hardware mechanism to determine the physical location of guest memory by walking an extra level of page tables — called nested page tables.

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

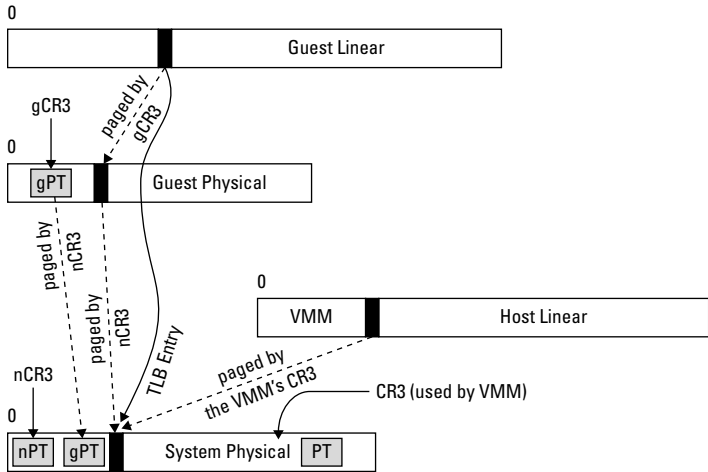


Figure 3-1: Design of AMD's Rapid Virtualization Indexing.

In a native system, the operating system maintains a mapping of logical page numbers to physical page numbers in page table structures. When a logical address is accessed, the hardware walks these page tables to determine the corresponding physical address. For faster memory access, the x86 hardware caches the most recent mappings in its translation lookaside buffer (TLB).

In a virtualized system, the guest operating system maintains page tables just like in a native system, but the hypervisor maintains an additional mapping of the guest physical page numbers to host physical page numbers in the shadow page table structures.

RVI enables the hypervisor to maintain this guest-physical to host-physical mapping in a second level of page tables, called nested page tables in hardware.

When a logical address is accessed, the hardware walks the guest page tables as in the case of native execution, but for

every guest physical page accessed during the guest page table walk, the hardware also walks the nested page tables to determine the corresponding host physical page address. This eliminates the need to maintain shadow page tables and synchronize them with the guest page tables.

Overall, this can help reduce the overhead of accessing memory, thereby increasing system performance.

The extent of this performance increase is somewhat dependant on 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. But by and large, moving the memory access functions into hardware can improve both performance and virtual machine density.

Extending Virtualization to Devices

AMD has implemented improvements to the way network and storage interaction is handled to improve performance for these resources in a virtualized environment.

AMD Virtualization technology for I/O virtualization is called ADM-Vi and refers to the hardware assists for I/O Virtualization in AMD Chipsets. AMD-Vi uses an I/O Memory Management Unit or an IOMMU to control how a device accesses memory. The IOMMU also provides a mechanism to isolate memory accesses by devices.

To understand the capabilities of AMD-Vi, 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 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. By leveraging AMD-Vi, systems ensure data integrity ("keep data straight") as well as enforce security, since the functionality isolates individual VM I/O from one another. One VM can't access the I/O memory locations of another VM.

With I/O memory mapping, the hypervisor must perform this task efficiently — very efficiently. Doing so 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 can reduce the software processing overhead and improve 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.

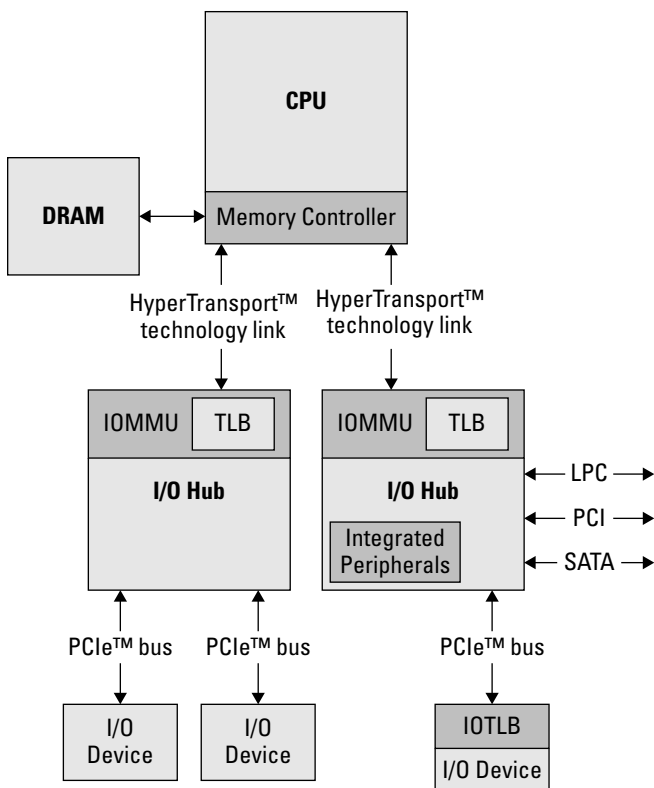


Figure 3-2: How IOMMU works.

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 the fourth generation of AMD Opteron processors. This improvement requires virtualization hypervisors to be modified in order to take advantage of the hardware changes, so when future generations of 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.

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
-

Most people don't just buy system resources. They 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 — virtualization makes hardware more important.

The New Role of Commodity Servers in a Virtualized World

Hardware takes on greater importance in a virtualized world because of the changing nature of servers. 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 don't 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: 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 supporting 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 is 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 is far higher storage and network traffic because 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 4GB 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 likely 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 use the latest generation of chips that are optimized for virtualization.

AMD has continued to optimize its line of Opteron processors for virtualized environments, and Sun, as an early proponent of AMD Opteron processor-based systems, recognizes that AMD Opteron processors with Direct Connect Architecture have 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 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. Sun x64 systems offer dual-core, quad-core, and six-core AMD Opteron processors.

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 X4640 Server (see Figure 4-1), offers the following resource capabilities:

- ✓ Up to eight Six-Core AMD Opteron processors in a single chassis
- ✓ Up to 512GB 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 PCIe and PCI expansion slots to allow multiple storage connections and avoid storage SPOF

The Sun Fire X4640's current VMMark score is 29.11 (www.vmware.com/products/vmmark/results.html).




Figure 4-1: Sun Fire X4640 Server.

Using the new generation of servers

In 2006, the Las Vegas Valley Water District IT team decided to replicate its Oracle 9i database into different versions for use in multiple projects. Each development group needed continuous access to the latest data in order to integrate and test it with their separate applications. After evaluating its processing, database, and storage requirements, the district turned to Sun for a data storage and replication solution.

The district selected the Sun StorageTek 9990 system for its scalability, reliability, and ability to provide virtualization for both its own internal storage and other tiers of storage across a Fibre Channel SAN. The district chose the Sun Fire E490 and T2000 servers based on

price-performance, as well as Solaris Containers to help run multiple applications and database copies on the same server. The district also valued the ability of the T2000 servers to run multithreaded applications with up to three times the throughput and one third the power and space requirements of competing systems.

Following the successful rollout of the database solution, the district migrated its Microsoft Windows VMware environment from eight X86-based servers to four Sun Fire X4600 servers with AMD Opteron processors. The district estimates cost-performance improvements and reductions in electric power and cooling costs.

The Sun Blade 6000 Modular System (see Figure 4-2), among the most powerful blade platforms on the market, offers the following resource capabilities:

- ✔ Runs up to four Six-Core AMD Opteron processors per server module, with up to 40 server modules in a single rack
- ✔ Up to 256GB of memory per server module, 10TB per rack
- ✔ Up to 192Gbps I/O throughput per blade




Figure 4-2: Sun Blade 6000 Modular system.

The innovative Sun Blade 6000 Modular System enables virtualization deployment with fewer servers, which helps 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 less energy than comparable commodity x86-based systems, thereby making them a great choice for green data center initiatives.

Chapter 5

Managing Virtualization End-to-End

In This Chapter

- ▶ Examining Sun server virtualization
 - ▶ Looking at Sun desktop virtualization
 - ▶ Delving into Sun Storage for server and desktop virtualization
-

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 describes 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 on or processed. This has been the primary area of focus for virtualization so far.
- ✔ **Storage in virtualization:** The scope of data expands dramatically in a virtualized environment where diverse operating systems are involved. The role of storage here is to not only provide storage for virtualized application data, but also to work in tandem with the physical server to improve the performance of the applications that rely on stored data to process the data flow.

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 can help IT organizations take advantage of hardware and software improvements to achieve the goals of virtualization: better hardware utilization, reduced 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

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:** 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 operating systems. 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. 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 Microsoft Hyper-V run as the primary application on a dedicated

system, with guest operating systems running on top of them. Sun™ 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 Microsoft Hyper-V.

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 as low as 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 that allows 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 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

lifecycles extended, helping make use of corporate capital more efficient.

- ✔ **Sun Storage:** Typically, data and applications for Sun VDI desktops are stored on networked storage systems, such as the Sun Storage 7000, which is used as a centralized repository for documents, images, and spreadsheets. The Sun 7000 series provides a wide choice of storage solutions to connect to Sun VDI desktops, preferably using NFS because it is simple, extremely fast, and can be an inexpensive alternative to traditional storage such as Storage Area Networks (SANs). For an explanation of these advantages see the section “Sun Storage Virtualization,” later in this chapter.
- ✔ **VMware’s Virtual Desktop Infrastructure:** This software runs each client instance as a virtual machine on the back-end 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/desktop virtualization means carefully planning how storage is organized and managed. There are different ways to achieve a good storage infrastructure to support server and desktop virtualization with each having their own advantages and trade-offs:

- ✔ **Storage Area Network (SAN):** SANs are typically deployed in large data centers for big ERP applications and use expensive fiber channel connections and switches.

Many companies utilize their SAN for storage in server virtualization projects because it is what they are familiar with. SANs provide high performance storage at a premium price, but the additional cost, complexity, and manageability issues can overcome the advantages of this type of storage.

- ✔ **iSCSI SAN:** This storage protocol is used typically in virtualization deployments to boot and initialize virtual machines. iSCSI SANs may suffer performance scalability issues that prevent them from being used throughout a virtualization architecture. Companies should balance Fiber Channel SANs and NAS with the complexity of managing iSCSI protocols in medium to large storage deployments.
- ✔ **Network File System (NFS):** NFS is the preferred storage protocol for virtualization due to simplicity, scalability, security, and relatively low cost of deployment and management. From a hardware perspective NFS is “plug and play,” using all of the network transport media elements that are ubiquitous in today’s data center.

The Network File System, invented by Sun in 1984, offers a simple method to manage data because it operates as a file-based network storage system. NFS-based files are easy to create, grow, copy, back up, and replicate. NFS has matured in the data center and customers have used it in various operating systems for a variety of applications including Windows. Sun continues its investment in NFS (such as in the Sun Storage 7000) and has provided Flash SSD integration to provide significant performance improvements to virtual machines and their applications. Sun recommends Sun 7000 with NFS because of the value it adds by simplifying the virtual environment that otherwise can easily get complicated and misinterpreted by IT staff.

The Sun Storage 7000 supports server and client virtualization for both iSCSI and NFS protocols and has key features that make managing them much easier. DTrace Analytics provides observability into VM storage workloads and a new innovation, called the Hybrid Storage Pool, accelerates performance significantly. The Hybrid Storage Pool combines Flash SSD, high capacity, lower energy hard disk drives, and the ZFS file-system to deliver high performance throughput at a reasonable cost.

The Sun Storage 7000 with NFS offers following advantages to accomplish a better virtualization deployment:

- ✔ **High availability of data in virtualization:** The Sun Storage 7000 supports continuous availability in virtualization systems. It provides a no-single-point-of-failure solution by clustering two storage systems together. Additional reliability is provided by a system of self healing for CPU, memory, and other key elements of the storage architecture. The ZFS filesystem provides superior data integrity features that respond to situations where issues with data integrity are detected.
- ✔ **Observability of VM I/O performance:** When applications and desktops are virtualized, they add an element of complexity that sometimes makes it more difficult to diagnose problems than prior to virtualization, especially when problems arise. To understand and respond to crashes and performance problems you need detailed information about what's going on inside the storage system. The Sun Storage 7000 feature called DTrace Analytics provides visibility into virtual machine performance and storage workloads and helps improve performance by identifying bottlenecks — thereby enabling rapid fixes.
- ✔ **Performance boost for VM data access:** Efforts to maintain and improve performance extend the life of existing physical servers, saving valuable future capital outlays. The Sun Storage 7000 with Hybrid Storage Pool technologies helps improve performance by accelerating I/O throughput using Flash SSD technology in concert with intelligent caching algorithms built in to the ZFS filesystem to stage files that are being referenced heavily and by accelerating writes to disk. This can improve performance for virtual machines.
- ✔ **Thin-provisioning of VM data for storage savings:** Being efficient in allocating storage space can equate to considerable savings on capital expenses. So one of the important features of a storage system is thin provisioning, which avoids unnecessary disk space consumption. In a virtualization environment, it is important to have features that support and extend existing storage capacity. The Sun Storage 7000's thin provisioning features extend the useful capacity of storage because thin provisioning

enables VMs to allocate and use only the amount of storage that applications actually consume.

✓ **Data services for VM data management, movement, and migration:** With virtualization comes the challenge of manageability. Point in time snapshots provide an instantaneous copy of a virtual machine's data and are useful for backup and cloning of VMs. These features improve overall data center productivity because they can be utilized for application development, migrating a virtual machine's data to another system, creating new virtual machines, or simply as a backup from which an instantaneous restore can be accomplished.

The Sun Storage 7000 series are versatile, high-performance storage engines for server and client virtualization infrastructures. They work in concert with server virtualization technologies to improve simplicity, performance, and data availability.

Tying the Virtual Environment Together: Virtualization Management

For IT organizations, virtualization can be a blessing . . . and a curse. The benefits of virtualization bring with them the challenge of 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:

- ✔ **Discovering:** 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 the assets present in a data center infrastructure.
- ✔ **Provisioning:** Sun xVM Ops Center makes it easy to remotely install virtual machine operating systems, packages, and RPMs, as well as firmware. This helps reduce the need to physically log on to guest virtual machines as well as offering centralized control of the provisioning process.
- ✔ **Updating:** 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 bug fixes. 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.
- ✔ **Managing:** 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 helping to ease the burden of system management.
- ✔ **Reporting:** 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 helping reduce the administrative burden of a heterogeneous environment.

Sun Services

Sun offers a complete set of services to assist customers with their virtualization initiatives. Sun's Professional Service organization can help companies make decisions on a virtualization strategy that makes sense for their IT organization, create the business case, and develop a roadmap that starts small and builds over time with the success of each project. Some of the benefits can include:

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

Sun can also help keep your staff abreast of new software and hardware technologies through training. Sun Learning Services offers an extensive list of courses. A beginner can start with a virtualization course on concepts and datacenter trends. Once decisions are made about specific technologies, courses can be found on point products such as xVM Ops Center, Solaris Containers, LDOMs, Secure Global Desktop Software, virtual disk, and virtual tape solutions.

Once the technology has been implemented, users can decide whether to manage it internally using tools such as xVM Ops Center or to look to Sun for a selective outsourcing strategy with Sun Managed Services.

Sun Managed Services can help organizations achieve and sustain the benefits of a virtualized datacenter improvement project. Managed Services offers a portfolio of complex IT infrastructure monitoring and management capabilities, which can help customers transition to new technologies rapidly and confidently.

Chapter 6

Ten Steps to Virtualization Success

In This Chapter

- ▶ Getting on board today
 - ▶ Selling a business case
 - ▶ Collaborating with the right vendors
-

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, 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 help 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, multi-platform 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. Hardware that was designed with compatibility in mind, such as AMD processors and virtualization-focused hardware from Sun, can help ensure that 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.

