PRENTICE HALL

Xen

CITRIX

Running Xen

A Hands-On Guide to the Art of Virtualization

Jeanna N. Matthews • Eli M. Dow
Todd Deshane • Wenjin Hu • Jeremy Bongio
Patrick F. Wilbur • Brendan Johnson

Foreword by Ian Pratt, Xen Project Leader
VP Advanced Technology, Citrix Systems

Library of Congress Cataloging-in-Publication Data:

Matthews, Jeanna N.

Running Xen : a hands-on guide to the art of virtualization / Jeanna N. Matthews, Eli M. Dow, Todd Deshane, Wenjin Hu, Jeremy Bongio, Patrick F. Wilbur, Brendan Johnson.

p. cm.

ISBN 0-13-234966-3 (pbk. : alk. paper) 1. Xen (Electronic resource) 2. Virtual computer systems. 3. Computer organization. 4. Parallel processing (Electronic computers) I. Title.

QA76.9.V5M38 2008

005.4'3--dc22

2007052439

ISBN-13: 978-0-132-34966-6
ISBN-10: 0-132-34966-3

Text printed in the United States on recycled paper at Courier Stoughton in Stoughton, Massachusetts.
First printing April 2008

This Book Is Safari Enabled

The Safari® Enabled icon on the cover of your favorite technology book means the book is available through Safari Bookshelf. When you buy this book, you get free access to the online edition for 45 days.

Safari Bookshelf is an electronic reference library that lets you easily search thousands of technical books, find code samples, download chapters, and access technical information whenever and wherever you need it.

To gain 45-day Safari Enabled access to this book:

- Go to http://www.informit.com/onlineedition
- Complete the brief registration form
- Enter the coupon code 9FJZ-X7UK-PMSU-TAKZ-N6SI

If you have difficulty registering on Safari Bookshelf or accessing the online edition, please e-mail customer-service@safaribooksonline.com.

Editor-in-Chief
Mark Taub

Acquisitions Editor
Debra Williams Cauley

Development Editor
Michael Thurston

Managing Editor
Gina Kanouse

Project Editor
Chelsey Marti

Copy Editor
Geneil Breeze

Indexer
Erika Millen

Proofreader
Meg Shaw

Technical Reviewers
Jessie Yu
Ken Hess
Jose Renato Santos
Andrew Warfield
Tom "Spot" Callaway
Simon Crosby
Tom Doeppner
Dan Kuebrich
Jim Owens
Zach Shepherd
Keegan M. Lowenstein
Ryan Kornheisl
Igor Hernandez
Alexander M. Polimeni
Erika Gorczyca
Justin Bennett
Joseph Skufca
Mathew S. McCarrell
Krista Gould
Ron Arenas

Publishing Coordinator
Kim Boedigheimer

Cover Designer
Alan Clements

Senior Compositor
Gloria Schurick

Foreword

The Xen open source hypervisor is changing the world of virtualization. It encourages the broad distribution of a common industry standard hypervisor that runs on a wide range of architectures from super computers to servers to clients to PDAs. By focusing on the hypervisor, the "engine" of virtualization, rather than a specific product embodiment, the Xen open source project enables multiple vendors and the community to combine the common cross platform virtualization features of Xen into exciting new products and service offerings.

To date, the community around the Xen hypervisor has been squarely in the camp of developers and expert users. While the Xen-users mailing list offers a friendly and useful source of advice for those wanting to deploy and manage Xen-based environments, the new user might find herself in need of advice about best practice and step-by-step instructions for the deployment of Xen. *Running Xen: A Hands-on Guide to the Art of Virtualization* speaks directly to this critical need. It provides users with everything they need to know to download, build, deploy, and manage Xen implementations.

To the authors, a set of Xen contributors, practitioners, and researchers, I would like to say thank you on behalf of the broader Xen community for an accessible and immediately useful book. Code might rule, but "know-how" builds the community itself. Clear information, advice, and documentation like this book will allow the Xen project to grow and solidify its user base, to renew its creativity and innovation, to focus itself on a larger set of new virtualization initiatives.

To the readers, I would like to say welcome to the community of Xen users. We look forward to your involvement and contributions! We believe this book will provide you with an excellent introduction to running Xen.

Ian Pratt, Xen Project Leader
VP Advanced Technology, Citrix Systems

Preface

We began using Xen in the fall of 2003 soon after reading the paper "Xen and the Art of Virtualization" published in the Symposium on Operating Systems Principles (SOSP). After attending SOSP and talking to some of the authors, Jeanna Matthews returned excited about Xen. She and her graduate operating systems course at Clarkson University decided to repeat and extend the results reported in that paper. That class included two of the coauthors for this book, Eli Dow (currently at IBM) and Todd Deshane (currently completing his Ph.D.), who were both studying for their Master's degrees at the time. In the process of repeating the results from the 2003 Xen paper, we learned a lot about running Xen—much of it the hard way! Our goal for this book was to write exactly the material we wished was available when we first started using Xen.

In July 2004, we published the paper "Xen and the Art of Repeated Research," describing our experience with Xen and presenting the results we obtained repeating and extending the results. All the authors, in addition to being a part of the Fall 2003 graduate operating systems course, were also members of the Applied Computing Laboratories at Clarkson University, specifically the Clarkson Open Source Institute (COSI) and the Clarkson Internet Teaching Laboratory (ITL). These labs were founded to provide students with hands-on experience with cutting-edge computing technologies and to form a community in which everyone both learns and teaches. Other students in the labs—both graduate and undergraduate—began to use Xen as the basis for both production systems and for research projects. Through the years, we have used Xen as the basis for a number of academic papers as well as the basis of award-winning team projects. In the process, we have learned a lot about running Xen. It is our goal in this book to share this knowledge with you and to make your experience running Xen as smooth and simple as possible.

The book is targeted at individuals and organizations that are deploying Xen systems. It walks the reader through the basics, from installing Xen to using prebuilt guest images. It even tells readers how to experiment with Xen using only a Xen LiveCD. It covers the basics of virtualizations and important elements of all Xen systems like the hypervisor and Domain0. It explains the details of the `xm` commands for managing guest domains. It helps users deploy custom guest images based on operating systems from Linux to Windows. It covers more advanced topics such as device virtualization, network configuration, security, and live migration. We hope you will find it a good mix of introductory and advanced topics that will prove useful from your first Xen deployment experiment to running production Xen systems.

Chapter 1, "Xen—Background and Virtualization Basics," is a quick introduction to virtualization in general and to Xen in particular. Chapter 2, "A Quick Tour with the Xen LiveCD," provides an overview of Xen's functionalities by exploring the Xen LiveCD. Chapter 3, "The Xen Hypervisor," focuses on the hypervisor that is the core of any Xen system and some other trusted components such as Domain0 and `xend`. We build on that common understanding of the Xen hypervisor by concretely showing you how to install and configure your own hard-disk-based Xen installation in Chapter 4, "Hardware Requirements and Installation of Xen Domain0." After you have your own hypervisor installation up and running, this book eases you into using guest images by first showing you how to download and use images available from the Internet in Chapter 5, "Using Prebuilt Guest Images." Chapter 6, "Managing Unprivileged Domains," covers the basics of administering the running DomUs or unprivileged guest domains. You are then guided through the various methods of creating your own custom guest images in Chapter 7, "Populating Guest Images." Now that you have all these guests, Chapter 8, "Storing Guest Images," covers a variety of choices for storing guest images for online use as well as backup and sharing.

The second half of this book delves into more advanced system management topics including device management (Chapter 9, "Device Virtualization and Management"), networking (Chapter 10, "Network Configuration"), security (Chapter 11, "Securing a Xen System"), resource distribution (Chapter 12, "Managing Guest Resources"), and migration (Chapter 13, "Guest Save, Restore and Live Migration"). We conclude with a survey of some of the popular administrative tools available for your Xen systems in Chapter 14, "An Overview of Xen Enterprise Management Tools."

Throughout the book, we include listings illustrating relevant commands and their output. We use the command prompt to indicate where the command should be run.

For example, the following would indicate a command to be run as root on the privileged domain, Domain0:

```
[root@dom0 ]#
```

The following would indicate a command to be run as any user in a regular guest domain:

```
[user@domU]$
```

Watching these command prompts will help you identify which of the many guests in your Xen system should be used for running any given command.

It is our intention to maintain a website with additional information and materials relevant to the book. We have registered the domain, runningxen.com, for this purpose and are working on assembling materials. We invite you to check on our progress and to send questions or suggestions.

Chapter 13

Guest Save, Restore, and Live Migration

In this chapter, we begin by exploring Xen's capability to easily checkpoint the state of a guest domain to disk for quick restoration at a later time. We continue by exploring how Xen makes the migration of guest domains a simple and powerful administrative task. We discuss cold static migration, warm static migration, and live migration of guests, along with the prerequisites and benefits of each.

Representing the State of a Virtual Machine

At the heart of any migration is the ability to fully represent the state of a guest. When a guest virtual machine is completely shut down, this is trivial. An inactive Xen guest is completely defined by its file system image(s), configuration file, and its operating system kernel. Clearly, a guest could be cloned or even moved to another physical machine by making copies of these files. Backup of a guest can be accomplished in this way.

A guest that is active in execution, on the other hand, is a more complicated matter. While a guest is running, saving its state additionally involves creating a snapshot of its memory, device I/O states, open network connections, and the contents of its virtual CPU registers. Xen can save this state information to disk or transfer it over the network, which allows for both backup and migration of VMs.

This idea of saving the state of a running guest is similar to the hibernation feature on many personal computers, which is especially popular among laptop users. In hibernation, a system's state is checkpointed and saved to disk so that the system can park the hard drive heads, power down, and resume its previous state next time it is powered up. Laptop users sometimes rely on this feature to temporarily suspend the state of the machine when moving from place to place, or for conserving battery power when the laptop is not being used. In the case of a virtual machine monitor like Xen, a similar facility can be used to checkpoint states to facilitate rollback in the event a guest fails, or to save the state of a guest past the shutdown of the physical machine on which it is running.

Xen provides a domain save and restore facility to handle the suspension of guest VMs to checkpoint files, operated by the `xm save` and `xm restore` commands. When a guest is saved to disk it is suspended, and its resources are deallocated. As in the case of hibernate, ongoing network connections are not preserved.

With the `xm migrate` command, Xen supports warm static migration (regular migration), where a running guest is temporarily suspended and then relocated to another physical host, as well as live migration, where a guest may be relocated from one host to another seamlessly, without dropping ongoing network connections and with little client perceptible delay. Live migration is

particularly useful when bringing down a physical machine for maintenance. In this case, guests can be relocated to a new physical machine in preparation for the maintenance without a disruption in service. The ability to relocate a guest is also useful for load balancing guests and their resource consumption.

In this chapter, we discuss the uses of `xm save`, `xm restore`, and `xm migrate` in detail.

Basic Guest Domain Save and Restore

Xen makes it possible to suspend a guest domain, save its state to a file, and resume that state later through its domain save and restore facility. Figure 13.1 illustrates the process of saving a guest's state to disk. As with hibernation, when a domain's state is saved on disk, it is suspended, and network connections to and from that guest are interrupted (due to TCP timeouts).

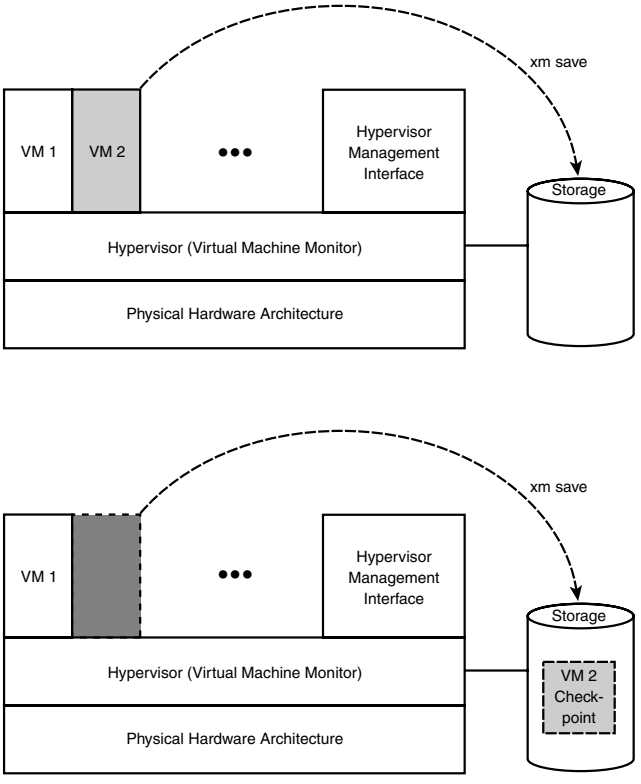

FIGURE 13.1 Xen provides an easy-to-use facility for hibernating guests to checkpoint files, to be restored at a later point in time.

xm save

Saving a guest VM's state first suspends that guest VM and then dumps its state information to disk. The guest VM will not continue to execute until it is restored, much like in the case of hibernation.

Listing 13.1 shows the syntax of this command.

LISTING 13.1 Usage of `xm save`

```
xm save domain filename
```

When performing an `xm save`, an administrator supplies a filename argument to specify where the serialized state of the guest VM should be written. This file is colloquially called a *checkpoint file*. This does not erase any previously saved checkpoints. You are free to save as many distinct checkpoints as you want. Thus it is possible to archive a collection of different running states for a particular guest domain.

Either a guest domain ID number or domain name may be supplied as the `domain` argument. After running this command, Xen suspends the state of the guest-specified domain to the specified file, and the domain no longer executes on the host. Thus it is impossible to save the state of Domain0 in this fashion because the Xen Domain0 must remain operational as the control interface between the administrator and Xen.

In Listing 13.2, we show the output of `xm list` running in Domain0. The output shown illustrates we are running a guest domain named TestGuest that has 64MB of RAM allocated to it.

LISTING 13.2 Domains Running on Domain0 (Before Checkpointing)

```
[root@dom0]# xm list
Name                              ID Mem(MiB) VCPUs State   Time(s)
Domain-0                           0       511     2 r-----   1306.0
TestGuest                          1        63     1 -b----      4.5
[root@dom0]#
```

To save TestGuest's state, the corresponding ID number needs to be noted and passed as an argument to the `xm save` command. In Listing 13.2 we see that the ID that corresponds with TestGuest is 1, and an example suspension using the checkpoint file name `TestGuest.checkpt` is shown in Listing 13.3.

LISTING 13.3 Checkpointing a Guest's State Using `xm save`

```
[root@dom0]# xm save 1 TestGuest.checkpt
[root@dom0]#
```

Note that replacing 1 with TestGuest as an argument to `xm save` also works. We know that the state of TestGuest has successfully finished saving when it is no longer listed as residing on Domain0, which we check by invoking `xm list` as shown in Listing 13.4. The `xm save` command does not return immediately, but instead returns after checkpointing is complete.

LISTING 13.4 Domains Running on Domain0 (After Checkpointing)

```
[root@dom0]# xm list
Name                              ID Mem(MiB) VCPUs State   Time(s)
Domain-0                           0      511     2 r-----   1411.8
[root@dom0]#
```

We can observe that our checkpoint file now exists in the present working directory and determine its size by issuing the `ls -la` command as shown in Listing 13.5. The `ls -lah` command may instead be used to view file sizes that are more human readable.

LISTING 13.5 Checkpoint File

```
[root@dom0]# ls -la
-rwxr-xr-x 1 root root 67266796 Feb  6 03:54 TestGuest.checkpt
[root@dom0]#
```

Now that TestGuest has been suspended, is no longer executing, and has been saved to a state file, we have successfully used Xen's domain save facility. Note that if there is not enough disk space to save the checkpoint file, the guest remains running and the `xm save` command fails with a generic error message.

A checkpoint contains the contents of the entire memory state of the guest. Thus, the time required to save the guest's state is proportional to the amount of memory allocated for the guest. The size of the checkpoint file is also approximately the same as the amount of memory allocated to the guest VM, plus a small amount of extra disk space to store additional state information.

Listing 13.6 shows the output of `xm list` for a system with three sample guests—one with 64 megabytes, one with 128 megabytes, and one with 256 megabytes of RAM allocated. Table 13.1 shows the size of the checkpoint file and the total time taken to save each of these guests. In all three cases, the checkpoint file is slightly larger

than the amount of memory allocated to the guest. It is also clear that the time taken to complete the save grows with the an increase in amount of allocated memory. The actual time taken to save the guest's state varies with the speed of the underlying file system and the hardware used to store the checkpoint file.

LISTING 13.6 Domains Running on Domain0 with Varying Memory Allocations

```
[root@dom0_Host1]# xm list
Name                                ID Mem(MiB) VCPUs State   Time(s)
Domain-0                             0       511     2 r-----   1306.0
Guest64MB                            7        63     1 -b----      3.1
Guest128MB                           8       127     1 -b----      2.8
Guest256MB                           9       255     1 -b----      2.1
[root@dom0_Host1]#
```

TABLE 13.1 Checkpoint File Size and Time Proportions

Actual Guest RAM Allocation (MB)	File Size On Disk (MB)	Time to Save Guest State (sec)
65.5	67.1	0.859
130.8	134.5	2.426
261.9	268.7	4.802

xm restore

Restoring a guest domain from a state file is initiated by the `xm restore` command. Listing 13.7 shows the syntax of this command.

LISTING 13.7 Usage of `xm restore`

```
xm restore filename
```

On the execution of this command, Xen restores the state of the guest located in the specified `filename`. The numerical domain ID of a guest domain is not preserved through save and restore, so no consideration needs to be taken to avoid an ID conflict—a unique ID is automatically assigned when restoring a guest from a checkpoint file.

On Domain0, we currently have a state file for a domain in the present working directory, observed by issuing the `ls` command, as shown in Listing 13.8.

LISTING 13.8 Checkpoint File

```
[root@dom0]# ls
TestGuest.checkpt
[root@dom0]#
```

Note that no guest domains are currently residing on Domain0, by invoking `xm list`, as we do in Listing 13.9.

LISTING 13.9 Domains Running on Domain0 (Before Restoration)

```
[root@dom0]# xm list
Name                              ID Mem(MiB) VCPUs State   Time(s)
Domain-0                           0      511     2 r-----   1702.2
[root@dom0]#
```

To restore the domain from our checkpoint file, we issue the `xm restore` command with its checkpoint's file name as an argument, as in Listing 13.10.

LISTING 13.10 Restoring a Guest's State Using `xm restore`

```
[root@dom0]# xm restore TestGuest.checkpt
[root@dom0]#
```

We know that restoration is complete when we can observe TestGuest residing on Domain0, by invoking `xm list`, shown in Listing 13.11. The `xm restore` command will not return until restoration is complete. Once TestGuest is restored from a state file, it continues executing where it left off at the time it was suspended, though network connections are likely to have timed out.

LISTING 13.11 Domains Running on Domain0 (After Restoration)

```
[root@dom0]# xm list
Name                              ID Mem(MiB) VCPUs State   Time(s)
Domain-0                           0      511     2 r-----   1795.9
TestGuest                          2       63     1 -b----      0.4
[root@dom0]#
```

Xen's domain save and restore facility has several potential uses. For example, the ability to restore from checkpoints may potentially be used in developing a rapid crash recovery procedure, where a guest is restored to a default state much faster than rebooting. Similarly, it is also useful while debugging or testing changes to a system, by allowing an administrator to save quickly restorable checkpoints that can be reverted

to in the event of a failure. Quick installs cannot be performed simply by providing such an image because checkpoint files do not contain a guest's file system contents; instead, quick installs and zero setup installations would require an image of a guest's file system, and optionally a checkpoint file, if it is desired to ship the installation with the guest being in a particular execution state.

Possible Save and Restore Errors

Listing 13.12 shows the error that occurs when there is not enough disk space to store the checkpoint file of a guest at the location of the path specified by the user. You should free up space or specify a different path to store the checkpoint file to fix this problem. If you specify a checkpoint file name that is the same as an existing file name, `xm save` overwrites the existing file without warning.

LISTING 13.12 Error: xm_save failed

```
[root@dom0]# xm save TestGuest TestGuest.checkpt
Error: /usr/lib64/xen/bin/xc_save 46 82 0 0 0 failed
Usage: xc save <Domain> <CheckpointFile>

Save a domain state to restore later.
[root@dom0]#
```

The error shown in Listing 13.13 can occur under several circumstances. This message commonly occurs when the domain contained in the checkpoint file you specified is already running. The message may also occur if the checkpoint file you specified is corrupt or invalid, or in situations where there is not enough RAM on the host system to restore the guest contained in the checkpoint file you specified.

LISTING 13.13 Error: Restore failed

```
[root@dom0]# xm restore TestGuest.checkpt
Error: Restore failed
Usage: xm restore <CheckpointFile>

Restore a domain from a saved state.
[root@dom0]#
```

Listing 13.14 shows a different type of restore error. This occurs when the checkpoint file is inaccessible, which happens if the checkpoint file you specified does not exist, is inaccessible due to permissions, or cannot be read due to device error.

LISTING 13.14 Error Message When `xm restore` Is Unable to Read File

```
[root@dom0]# xm restore TestGuest.checkpt
Error: xm restore: Unable to read file /root/TestGuest.checkpt
[root@dom0]#
```

Types of Guest Relocation

The ability to easily move a guest operating system across the network from one physical host to another can be useful for a number of different administrative tasks such as load balancing or dealing with scheduled maintenance. Xen provides integrated relocation, or migration, support as illustrated in Figure 13.2. It helps manage the process of preparing, transporting, and resuming guests from one nearby host to another.

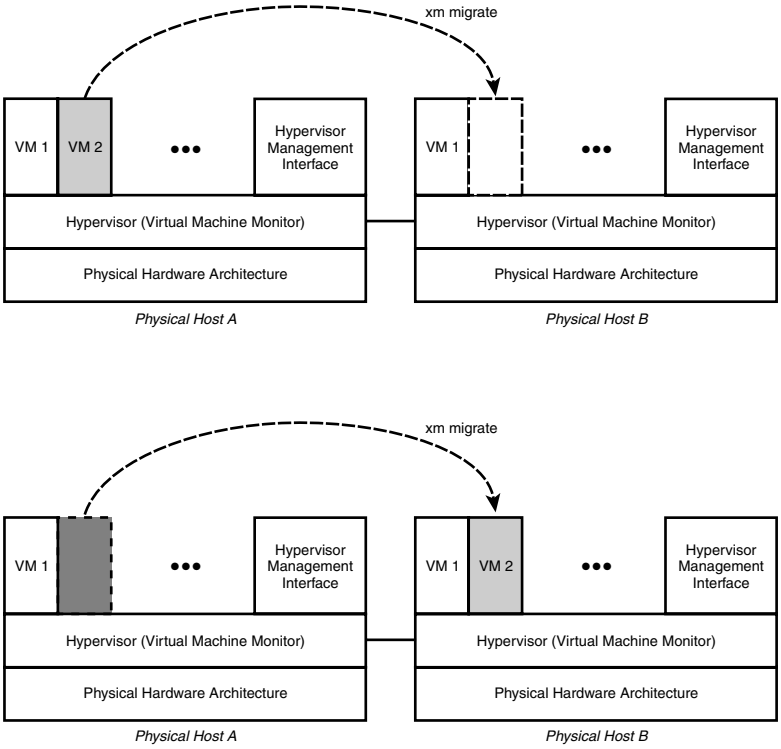

FIGURE 13.2 Guests may be relocated (migrated) between different Xen servers for various administrative purposes.

For comparison purposes, we begin our discussion about guest relocation by introducing the concept of cold static relocation, in which an administrator manually copies all the files that define a guest to a different machine, and executes `xm create` to start the guest at its new location. We then discuss Xen's integrated migration facility, which automates two major paradigms of guest relocation over a network—warm static migration and live migration. For warm static migration (also called regular migration), a guest is suspended on its source host, all relevant state information is transferred to its destination host, and the guest is resumed on its destination host after its state and memory have been safely relocated. This process is effectively the same as checkpointing a guest, manually copying the checkpoint file to another host, and restoring the guest on the new host. Warm static migration does not preserve ongoing network connections and does expose client visible downtime; for live migration, however, a guest is transferred without being suspended, and its services and connections are not only preserved but continue effectively uninterrupted.

Cold Static Relocation

Cold static relocation is accomplished manually without the help of Xen's integrated migration facility. Understanding the elements of this manual process aids in a better understanding of and an appreciation for the `xm migrate` command.

A halted guest may be relocated between two hosts by ensuring that its configuration file appears on and its file systems are available to both hosts. There are two ways of accomplishing this goal. The first occurs when both hosts share underlying storage (network attached storage). The second method involves manually copying the configuration file and file systems from the host to the target hardware. In the latter case, manual copying might take a long time because a guest's file systems might be very large. The transfer of a guest's file systems and configuration file by manual means might occur, for instance, through the use of optical media or FTP/SFTP. A much simpler option is to store the guest's file systems on network-attached storage, which makes copying unnecessary to make the guest's file systems available on both the source and destination hosts.

A running guest might also be relocated using this method, but must first be suspended to a checkpoint file. A guest domain may be checkpointed using the `xm save` command. Once suspended, a guest may be relocated by ensuring its file systems, checkpoint file, and configuration file are accessible on the destination host. When two physical hosts share the same underlying storage, this is equivalent to what `xm migrate` does for our next type of relocation, warm static migration.

Warm static migration could also be performed by manual copying methods if the two physical hosts did not share the storage systems on which all the needed files were stored. Once these three components are available on the desired destination host, the guest may be reactivated using the `xm restore` command.

Be advised that invoking `xm create` for the same guest domain on two different hosts at once has serious ramifications. Although the configuration file and operating system kernel would remain undamaged, multiple guests manipulating the same storage directly would likely lead to file system corruption.

Warm Static (Regular) Migration

Warm static migration, or regular migration, of a guest domain is the combined process of pausing the execution of that guest's processes on its original host, transferring its memory and processes from its origin host to a destination host, and resuming its execution on the destination host. Warm static migration enables a domain to be migrated from one physical host to another, only temporarily pausing its execution, without requiring it to be shut down and restarted. This is illustrated in Figure 13.3. Xen provides fast and simple integrated facility for performing regular migration of guest domains between hosts for load balancing, transferring from old to new hardware, or bringing physical hosts down for maintenance while adopting their guests on other hosts.

Migration requires a guest's memory contents to be transferred, I/O transactions temporarily quiesced, CPU register states transferred, and network connections rerouted and resumed on the destination host. Xen does not currently support the automatic mirroring of a guest's file systems between different hosts, nor does it automatically transfer a guest's file systems, because the act of copying an entire root file system is often too cumbersome to instantaneously transfer during migration. As a result, if a Xen user wants to use Xen's integrated migration support, it is currently necessary to configure guests to access their file systems over network shares that are equally available to both the source and destination host. In regular migration, a guest domain is suspended to easily handle the transfer of guest memory pages that would be continuously changing if the guest was otherwise allowed to continue to run. Suspending the guest also satisfies the need to quiesce I/O.

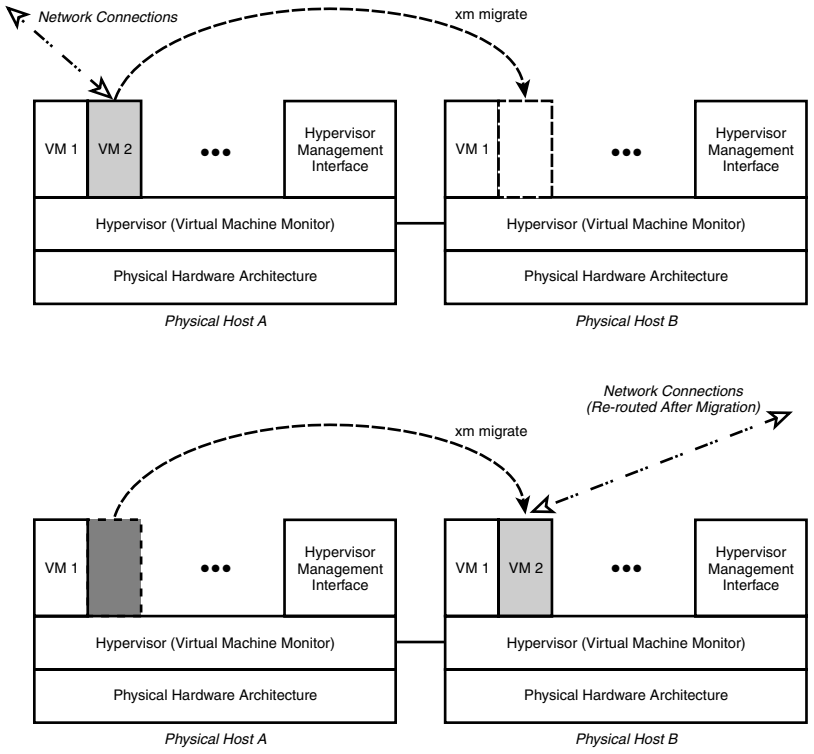

FIGURE 13.3 In warm static migration, a guest domain is temporarily suspended on its source host prior to relocation and resumed on its destination host after its memory contents have been transferred.

Live Migration

The mere ability to migrate a guest domain from one physical host to another is beneficial, but performing migration by temporarily suspending and then restoring a guest's state is not suitable in all applications. This process of migration can incur outages that are perceptible to system users. These sorts of outages are typically on the order of seconds or minutes, depending on the network infrastructure and the amount of memory a guest is allocated. Warm static or regular migration is not applicable for use in situations where in-use guests must be relocated without their services experiencing downtime; instead, in such situations, the ability to migrate a guest while maintaining its current state, operation, and network connections is desired.

Xen's third form of relocation, live migration, enables a domain to be migrated while it is in operation and without the interruption of its services or connections, as illustrated in Figure 13.4. Live migration of a guest is the act of seamlessly moving

its execution to the new physical host, including redirecting established and future network connections away from its original and to its new location.

Live migration is considerably more complicated than regular migration primarily because the state of the guest is changing while it is being copied. This requires an iterative process of copying the state, checking to see what has changed in the meantime, and then copying what has changed.

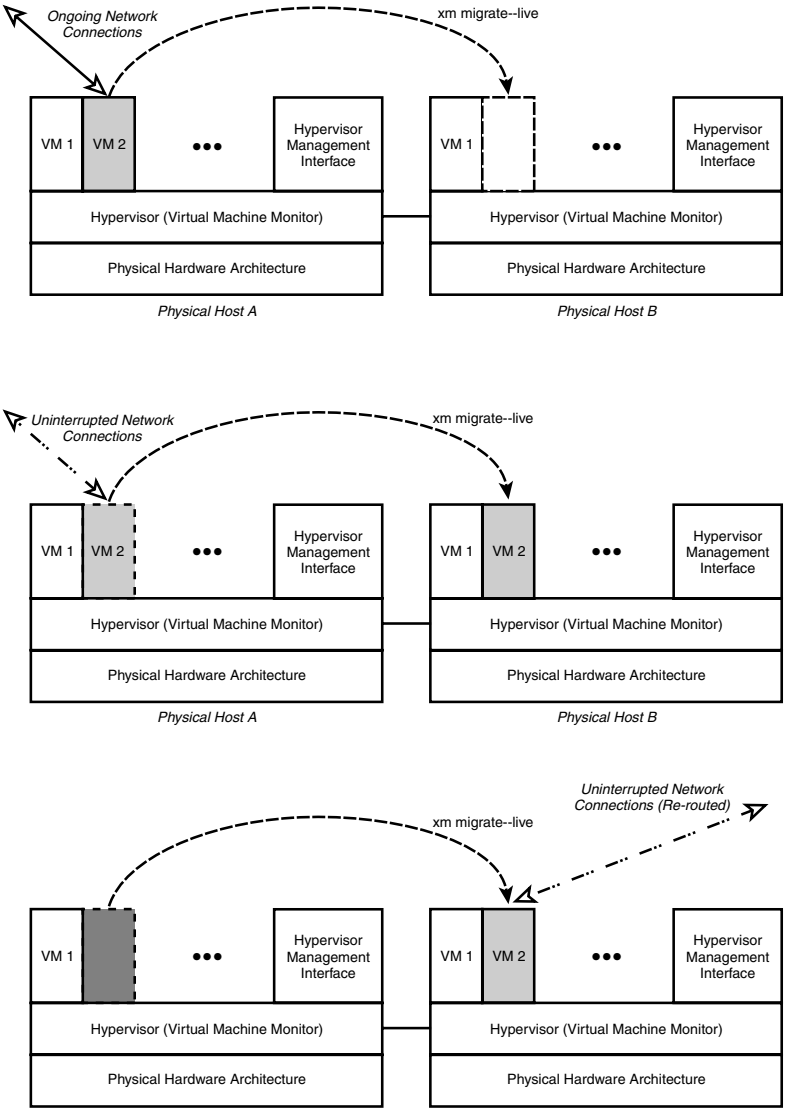

FIGURE 13.4 In live migration, a guest domain continues to execute and remains available throughout relocation, and client connections to its services go uninterrupted.

Making memory that belongs to an active guest available on a different host system is difficult because processes that are executing and memory pages that are swapped to and from disk consistently alter the contents of memory and consistently make attempted copies out-of-date. Like regular migration of a guest, Xen's live migration must also transfer and restore CPU register states, quiesce I/O transactions, and reroute and resume network connectivity on the destination host; however, in live migration, these events must occur in a way and over a small enough duration that there is no perceptible downtime of the guest's services. Additionally, like regular migration, a guest's file systems must be network-accessible to both the source and destination hosts.

The implementation of Xen's live migration involves the novel use of an iterative multipass algorithm that transfers the virtual machine guest memory in successive steps. After the source VM and the destination VM first negotiate to ensure resources are sufficient on the receiving machine, an initial pass over the guest's memory is performed with each page being transferred to the destination. On each successive iteration, only the guest memory that has been dirtied in the interim is sent. This process is executed until either the remaining number of dirty guest pages is sufficiently small enough that the remaining pages can be transmitted quickly or the number of dirty pages remaining to transfer in each pass is not decreasing. At that point, the system is actually quiesced and the final state sent to the new host and the transfer of control to the new physical machine completed. In most cases, this final step can be accomplished so quickly that there is no client perceptible delay and ongoing network connections are preserved.

Preparing for `xm migrate`

To support regular and live migration, a Xen configuration must include the following:

- Two or more physical hosts with `xend` configured to listen for relocation requests
- Hosts that are members of the same layer-2 network
- A storage server that is accessible by guests on both the source and destination hosts, which provides network access to guests' root and all other file systems
- A configuration file for each guest on both the source and destination hosts
- Sufficient resources on the destination host to support the arrival of guests
- The same version of Xen running on both the source and destination hosts

In the subsections that follow, we cover each of these steps in detail.

Configuring xend

To enable relocation (migration) support between hosts we must first edit their xend configuration files. In /etc/xen/xend-config.sxp on each host (or /etc/xen/xend.conf, depending on your system), remove the comment character ("#") from the beginning of each of the following lines if necessary (note that these lines might not immediately follow one another), as illustrated in Listing 13.15.

LISTING 13.15 An Example /etc/xen/xend-config.sxp Prior to Enabling Relocation Support

```
#(xend-relocation-server no)
#(xend-relocation-port 8002)
#(xend-relocation-address '')
#(xend-relocation-hosts-allow '')
```

On the xend-relocation-server line, ensure the value is set to yes. This enables xend to listen for relocation requests from other hosts so that it can receive the relocated guests. These lines should now read as shown in Listing 13.16.

LISTING 13.16 An Example /etc/xen/xend-config.sxp After Enabling Relocation Support

```
(xend-relocation-server yes)
(xend-relocation-port 8002)
(xend-relocation-address '')
(xend-relocation-hosts-allow '')
```

The second and third lines specify the port and address that xend should listen on for incoming migration requests. You may change the xend-relocation-port or leave it the default, port 8002. Leaving the xend-relocation-address field blank between the single quotes configures xend to listen on all addresses and interfaces. The fourth line may be used to restrict access to migration on a particular host to only certain hosts. Leaving this field blank allows all hosts with network access to the relocation port and address on which xend is listening to negotiate migration. Leaving these fields blank suffices for testing purposes but may not provide a desirable level of security in production systems. A better idea is to specify a list of hosts that should be allowed to negotiate migration with a particular host, and/or create a separate, secure network for guest migration. Of course for tighter security, a good idea is to configure both.

It is imperative for the relocation service to be accessible only by trusted systems and for live migration to be performed only across a trusted network. It is important to perform migration across a trusted network because when migrated, the contents

of a guest's memory are transferred from the source to destination hosts in a raw, unencrypted form. The memory contents might contain sensitive data and can be intercepted on a shared network by other hosts capable of scanning the migration traffic. Additionally, allowing unrestricted access to the relocation service on a host could allow an attacker to send fraudulent guest domain data to that host, and possibly hijack or impair system and network resources. Whenever possible, the `xend-relocation-address` should be set on the source and destination hosts to the address of an interface that is connected to an isolated administrative network, where migration can occur in an environment that is not shared with untrusted hosts (or the Internet). Ideally, the transmission medium should be high speed for best results. Having an isolated network for migration and listening for relocation requests only on that network adds natural security to a Xen configuration by physically preventing untrusted hosts from accessing the relocation port or spying on a migrating guest's memory contents. Please see Chapter 11, "Securing a Xen System," for more tips on hardening your Xen configuration.

To have `xend`'s relocation service listen only on one address, specify the address in the `xend` configuration file. For example, if the address 10.10.3.21 is bound to a network interface available to Domain0, the setting illustrated in Listing 13.17 causes `xend`'s relocation service to only listen for relocation requests sent to 10.10.3.21. This configuration feature helps secure a host's relocation service if undesired hosts are blocked from making connections to the specified address, either by firewall rules or by being bound to a network interface connected to a separate administrative network.

LISTING 13.17 An Example of `xend` Bound to a Relocation Address

```
(xend-relocation-address '10.10.3.21')
```

The format for specifying a restricted set of allowed hosts to connect to the relocation service is a list of regular expressions, separated by spaces, that specify hosts that are to be accepted. For instance, the example in Listing 13.18 would cause a particular host to only listen to migration requests from itself and hosts within the our-network.mil domain. Host IP addresses may also be listed in this line. Restricting access to the relocation service by only certain hosts can also be achieved through good firewall practices, either in addition to or instead of setting `xend-relocation-hosts-allow`. Any host that matches any one of the regular expressions listed in this field will be allowed to negotiate migration with the local host.

LISTING 13.18 An Example of Restricting Allowed Source Hosts

```
(xend-relocation-hosts-allow '^localhost$                    ➥
     ^.*\.our-network\.mil$')
```

For our sample configuration, we want our hosts Domain0_Host1 (10.0.0.1) and Domain0_Host2 (10.0.0.2) to accept relocation requests only from themselves and each other, but, unfortunately, our hosts do not have separate real or virtual network interfaces on a separate network just to be used for relocation. Remember that this host setup can be less secure than having a separate network and interface for guest relocation, and can decrease network performance. Our sample xend configuration file for Domain0_Host1, contained in Listing 13.19, shows xend configured only to accept migration requests from its own host and Domain0_Host2.

LISTING 13.19 A Sample xend Configuration File for Domain0_Host1

```
# dom0_Host1 (10.0.0.1)
# Xend Configuration File
#

# = Basic Configuration =
(xend-unix-server yes)
(xend-unix-path /var/lib/xend/xend-socket)


# =*= Relocation Configuration =*=
(xend-relocation-server yes)  # Enable guest domain relocation
(xend-relocation-port 8002)                                  ➥
     # Port xend listens on for relocation requests
(xend-relocation-address '')
     # Interface to listen for reloc requests [ALL]          ➥
(xend-relocation-hosts-allow '^localhost$                    ➥
     ^localhost\\.localdomain$ 10.0.0.2')
     # Hosts that are allowed to                             ➥
     send guests to this host [only 10.0.0.2!]


# = Network Configuration =
(network-script network-bridge)
(vif-script vif-bridge)

# = Resource Configuration =
(dom0-min-mem 256)
(dom0-cpus 0)
```

Likewise, our sample `xend` configuration file for Domain0_Host2, with `xend` configured only to accept migration requests from its host and Domain0_Host1, is contained in Listing 13.20.

LISTING 13.20 A Sample `xend` Configuration File for Domain0_Host2

```
# dom0_Host2 (10.0.0.2)
# Xend Configuration File
#

# = Basic Configuration =
(xend-unix-server yes)
(xend-unix-path /var/lib/xend/xend-socket)


# =*= Relocation Configuration =*=
(xend-relocation-server yes)  # Enable guest domain relocation
(xend-relocation-port 8002)                                    ➥
     # Port xend listens on for relocation requests
(xend-relocation-address '')                                   ➥
     # Interface to listen for reloc requests [ALL]
(xend-relocation-hosts-allow '^localhost$                      ➥
     ^localhost\\.localdomain$ 10.0.0.1')
     # Hosts that are allowed to                               ➥
     send guests to this host [only 10.0.0.1!]


# = Network Configuration =
(network-script network-bridge)
(vif-script vif-bridge)

# = Resource Configuration =
(dom0-min-mem 256)
(dom0-cpus 0)
```

Proximity of Sources and Destinations on the Network

For guest migration to be possible, the source and destination hosts need to be members of the same layer-2 network and the same IP subnet. This is because Xen needs to maintain the same environment for a guest's services before and after migration, including the same IP and MAC addresses. A guest's IP and MAC addresses are

transferred with that guest so that its network services remain accessible to other hosts once its migration completes.

Packet redirection to a new host is generally accomplished through Address Resolution Protocol (ARP), a protocol already familiar with networking hardware. If the source and destination hosts are located on different subnets, connections have to be redirected to the distant destination host in a more complicated way—for example, through the use of tunneling on Domain0, "Mobile-IP," dynamic DNS, or reconnection at the application level. In addition to the administrative complexity these proposed methods would create, they would also cause undesirable effects such as increasing overhead, latency, and even bandwidth usage if tunneling is used. Due to issues such as these, at the time of this publication, Xen does not provide an integrated capability to migrate guests between hosts that are not on the same layer-2 network and IP subnet. Solutions such as IP tunneling in Domain0 must be manually configured as needed. In many scenarios, the need for geographically-isolated source and destination hosts coincides with the need to be able to shift all dependency away from the source location; as such, tunneling is not viable because it depends on the source location to be able to maintain resources at its end of the tunnel.

Network-Accessible Storage

Recall that the migration support for Xen guests requires guest root file systems located on some form of mutually shared storage. Xen does not yet provide a facility for the automatic mirroring of local storage volumes at the Domain0 level, though work is currently being done to explore such possibilities. It is therefore necessary for each guest that is to be migrated to have its file system(s) mapped to network shares because the local devices available to a guest's original host will not be available locally on its destination host following migration.

There are several approaches to making a guest's files network accessible, including NFS, iSCSI, AoE, GNBD, and many others. Services such as iSCSI, ATA-over-Ethernet, and GNBD share access to volume block devices over the network, whereas services such as NFS share access to portions of a file system. See Chapter 8, "Storing Guest Images," for details on configuring a suitable network storage service such as the one mentioned in this section.

Guest Domain Configuration

We configure our guest's virtual block device, storing its root file system to be mapped to an ATA-over-Ethernet (AoE) shared volume. You may choose whichever network

storage service is most convenient for your own needs on your premises. First, we set up the AoE initiator (client) on both of our hosts. Then we create an LVM volume group and volumes on our AoE share to serve as our sample guest's root and swap logical partitions. We define a virtual disk in the guest configuration file that points to our shared block device and also use pygrub as a bootloader.

Both the source and destination Xen hosts need a configuration file for guests that are to be migrated, which will be the same for our hosts because all of our guest's virtual block devices will be accessible identically on both the source and destination hosts. We name this sample guest TestGuest, and Listing 13.21 shows the configuration in `/etc/xen/TestGuest` that is identical on both hosts. If a network block storage service is to be used but LVM is not, configure the `disk` line so that it points directly to the device on Domain0 that corresponds to the appropriate network block device (located in the `/dev/` tree).

LISTING 13.21 A Sample `/etc/xen/TestGuest` with LVM and ATA-over-Ethernet on Both Domain0 Hosts

```
name = "TestGuest"
memory = "64"
disk = [ 'phy:/dev/VolumeGroup0/TestGuest-volume,xvda,w' ]
vif = [ 'mac=00:16:3e:55:9d:b0, bridge=xenbr0' ]
nographic=1
uuid = "cc0029f5-10a1-e6d0-3c92-19b0ea021f21"
bootloader="/usr/bin/pygrub"
vcpus=1
on_reboot   = 'restart'
on_crash    = 'restart'
```

The configuration shown in Listing 13.21 works for network storage services that export entire block devices; however, because NFS exports a file system and not a form of raw access to block devices, configuring a guest to have an NFS-shared root file system is slightly different. The main difference is that the guest configuration file will not define a virtual block device pointing to a shared block device, but instead will have an NFS server configured that stores the guest's root file system. An NFS root on TestGuest may be set up by the configuration in `/etc/xen/TestGuest` on both hosts, shown in Listing 13.22.

LISTING 13.22 A Sample of /etc/xen/TestGuest with NFS on Both Domain0 Hosts

```
name = "TestGuest"
memory = "64"
vif = [ 'mac=00:16:3e:55:9d:b0, bridge=xenbr0' ]
nographic = 1
uuid = "cc0029f5-10a1-e6d0-3c92-19b0ea021f21"
bootloader = "/usr/bin/pygrub"
vcpus = 1
root = "/dev/nfs"
nfs_server = '10.0.0.40'              # Address of our NFS server
nfs_root = '/XenGuestRoots/TestGuest'                          ➥
     # Path on server of TestGuest's root
on_reboot   = 'restart'
on_crash    = 'restart'
```

Version and Physical Resource Requirements

In addition to both the source and destination hosts needing to run the same version of Xen to allow for migration, the destination host also must have sufficient resources available to it in order to support the arrival of a guest. The destination host must have access to memory that is unallocated to either Domain0 or other domains to handle the arriving guest, and the minimum amount of memory needed equals the amount of memory allotted to the guest on the source host, plus an additional 8MB of temporary storage.

Experience with xm migrate

Xen's internal relocation facility supports both warm static migration and live migration. It is available through the xm migrate command.

xm migrate

This command minimally takes the domain ID (or domain name) of the guest that is to be migrated and the destination host to which it is to be migrated to as its first and second arguments, respectively. If a domain with the same numerical domain ID exists on the destination host, migration still occurs, but the guest is assigned a different domain ID on the destination host. Listing 13.23 shows the syntax of the command.

LISTING 13.23 Usage of `xm migrate`

```
xm migrate domain_id destination_host [-l|--live] [-r|--resource rate]
```

`xm migrate` supports two optional arguments: `--live` and `--resource`. The `--live` argument specifies live migration as the form of migration to be performed. If `--live` is not used, regular migration will be the type of migration performed. To reduce network saturation, the optional `--resource` argument, followed by a `rate` in megabits per second, may be used to specify the rate of data transfer during migration. The `--resource` argument is generally unnecessary when using a private network dedicated for migrations to be performed, and it is best to avoid supplying it whenever possible to ensure optimal throughput; however, if a single network or even a single network interface on a host shares both migration and normal guest network traffic, it may be wise to supply this argument to reduce the network saturation so that connections to and from guests are not affected as dramatically.

Xen's relocation facility, as interfaced through the `xm migrate` command, makes guest relocation a simple task. We demonstrate the usage of the `xm migrate` command for performing both regular migration and live migration.

Using `xm migrate` for Warm Static Migration

In this section, we perform a warm static (regular) migration of a guest domain. The Domain0_Host1 system is currently running one guest domain, TestGuest, as seen by invoking `xm list` at a console on Domain0_Host1 and as shown in Listing 13.24.

LISTING 13.24 Domains Running on Domain0_Host1

```
[root@dom0_Host1# xm list
Name                              ID Mem(MiB) VCPUs State   Time(s)
Domain-0                           0      511     2 r-----   5010.6
TestGuest                          1       63     1 -b----     15.3
[root@dom0_Host1]#
```

The guest domain TestGuest is the guest to be relocated using warm static migration, from Domain0_Host1 to Domain0_Host2. TestGuest (10.0.0.5) is currently accessible from a separate workstation on our network, which we reveal using the `ping` command in Listing 13.25.

LISTING 13.25 Demonstrating Remote Accessibility Using `ping`

```
[root@Other_Workstation]# ping TestGuest
PING TestGuest (10.0.0.5) 56(84) bytes of data.
64 bytes from TestGuest 10.0.0.5: icmp_seq=1 ttl=64 time=2.48 ms
64 bytes from TestGuest 10.0.0.5: icmp_seq=1 ttl=64 time=2.31 ms
64 bytes from TestGuest 10.0.0.5: icmp_seq=1 ttl=64 time=2.98 ms
64 bytes from TestGuest 10.0.0.5: icmp_seq=1 ttl=64 time=2.77 ms
[root@Other_Workstation]#
```

Next, to perform a warm static migration of our sample guest from its current host Domain0_Host1 to its destination host Domain0_Host2 (10.0.0.2), we run the `xm migrate` command at a console on Domain0_Host1, illustrated in Listing 13.26.

LISTING 13.26 Performing a Warm Static Migration Using `xm migrate`

```
[root@dom0_Host1]# xm migrate 1 10.0.0.2
[root@dom0_Host1]#
```

After the migration is complete, TestGuest executes on Domain0_Host2. You can check that it has finished successfully and is currently residing on Domain0_Host2 by running `xm list` in a terminal on Domain0_Host2, as illustrated in Listing 13.27.

LISTING 13.27 Domains Running on Domain0_Host2

```
[root@dom0_Host2]# xm list
Name                              ID Mem(MiB) VCPUs State   Time(s)
Domain-0                           0      511     2 r-----    710.1
TestGuest                          4       63     1 -b----     16.2
[root@dom0_Host2]#
```

To demonstrate that the guest is still accessible with its same IP address and hostname after the migration, we repeat our ping from our generic workstation on the network, as shown in Listing 13.28.

LISTING 13.28 Demonstrating Remote Accessibility at Same IP Address after Warm Static Migration

```
[root@Other_Workstation]# ping TestGuest
PING TestGuest (10.0.0.5) 56(84) bytes of data.
64 bytes from TestGuest 10.0.0.5: icmp_seq=1 ttl=64 time=2.99 ms
64 bytes from TestGuest 10.0.0.5: icmp_seq=2 ttl=64 time=2.27 ms
64 bytes from TestGuest 10.0.0.5: icmp_seq=3 ttl=64 time=2.53 ms
64 bytes from TestGuest 10.0.0.5: icmp_seq=4 ttl=64 time=2.43 ms
[root@Other_Workstation]#
```

The domain is still accessible on the network the same way it was prior to migration.

Using `xm migrate` for Live Migration

Now we demonstrate how to perform a live migration of a guest domain. First, let's examine our guest domain, TestGuest, residing on Domain0_Host1. To do so, we invoke `xm list`, as in Listing 13.29.

LISTING 13.29 Domains Running on Domain0_Host1

```
[root@dom0_Host1]# xm list
Name                              ID Mem(MiB) VCPUs State   Time(s)
Domain-0                           0       511     2 r-----   7170.6
TestGuest                          1        63     1 -b----     15.1
[root@dom0_Host1]#
```

Here, we see that TestGuest is residing on Domain0_Host1. The guest domain TestGuest is the guest we will once again migrate, only this time we will perform a live migration. We previously installed the Apache HTTP Web server on our sample guest in preparation for this example, and it is publishing a directory containing a large-sized file. To demonstrate the persistence of connections with our guest during and after live migration, we commence the downloading of a large file from a separate workstation on our network, as shown in Figure 13.5. We observe the status of our download after performing a live migration of our guest to confirm that connections to our guest's services remain uninterrupted throughout the process. We also present a test consisting of constant pings to illustrate the temporary increase in latency during migration when accessing the guest over the network.

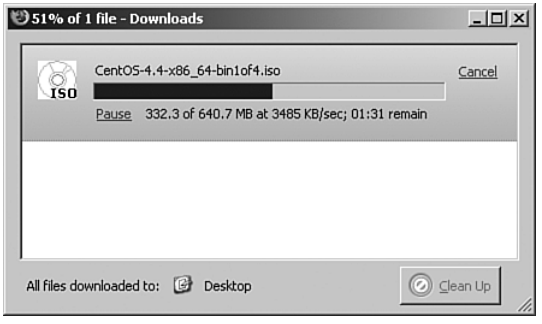

FIGURE 13.5 A connection from another computer on our network is established with our guest domain, and the guest begins serving us a large file.

Next, we request live migration of our sample guest from its current host, Domain0_Host1, to its destination host, Domain0_Host2 (10.0.0.2), to occur through the xm interface. We do this by invoking xm migrate on its source host, Domain0_Host1, as shown in Listing 13.30.

LISTING 13.30 Performing a Live Migration Using xm migrate --live

```
[root@dom0_Host1]# xm migrate --live 1 10.0.0.2
[root@dom0_Host1]#
```

After the migration is complete, TestGuest will reside completely on Domain0_Host2. You can check that it has finished successfully and is currently residing on Domain0_Host2 by running xm list in a terminal on Domain0_Host2, which is illustrated in Listing 13.31.

LISTING 13.31 Domains Running on Domain0_Host2

```
[root@dom0_Host2]# xm list
Name                              ID Mem(MiB) VCPUs State   Time(s)
Domain-0                           0       511     2 r-----   4314.2
TestGuest                          6        63     1 -b----     17.4
[root@dom0_Host2]#
```

To confirm that live migration did not disrupt any connections to our sample guest domain, we observe our download on the client workstation as shown in Figure 13.6.

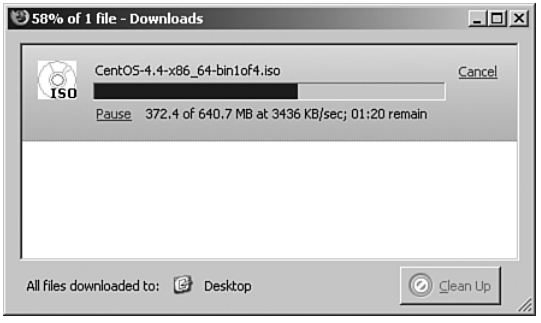

FIGURE 13.6 The connection to our guest domain remains uninterrupted during and after migration.

Although the download rate of the file we were downloading decreased during migration, our connection was not dropped. From the client's perspective, the guest domain remained completely accessible during and after its transfer to Domain0_Host2.

During this example, we ran the ping utility on a separate workstation to illustrate the increase in latency experienced during live migration. We set up this test to perform constant pings before, during, and after TestGuest (10.0.0.5) was live migrated from Domain0_Host1 to Domain0_Host2. Listing 13.32 shows the results of the `ping` command.

LISTING 13.32 Guest Latency During Live Migration

```
[root@Other_Workstation ~]# ping TestGuest
PING TestGuest (10.0.0.5) 56(84) bytes of data.
64 bytes from TestGuest (10.0.0.5): icmp_seq=1 ttl=64 time=2.29 ms
64 bytes from TestGuest (10.0.0.5): icmp_seq=2 ttl=64 time=1.06 ms
64 bytes from TestGuest (10.0.0.5): icmp_seq=3 ttl=64 time=1.07 ms
64 bytes from TestGuest (10.0.0.5): icmp_seq=4 ttl=64 time=1.05 ms
64 bytes from TestGuest (10.0.0.5): icmp_seq=5 ttl=64 time=5.77 ms
64 bytes from TestGuest (10.0.0.5): icmp_seq=7 ttl=64 time=6.13 ms
64 bytes from TestGuest (10.0.0.5): icmp_seq=8 ttl=64 time=4.06 ms
64 bytes from TestGuest (10.0.0.5): icmp_seq=9 ttl=64 time=1.08 ms
64 bytes from TestGuest (10.0.0.5): icmp_seq=10 ttl=64 time=1.09 ms
64 bytes from TestGuest (10.0.0.5): icmp_seq=11 ttl=64 time=1.08 ms
64 bytes from TestGuest (10.0.0.5): icmp_seq=12 ttl=64 time=1.11 ms
[root@Other_Workstation ~]#
```

With Xen's approach for performing live migration, the guest domain is unreachable typically for only 50ms. Due to the brevity of its networking quiescence, the guest domain should have no apparent downtime from the perspective of client computers connected to that guest. Interruption of active connections should not occur because the time the guest domain is unavailable resembles the effects of an otherwise temporary surge in network latency. Even when migrating a computer gaming server, where high latencies often obviously affect application performance in an obvious way to users, this effective latency generally goes unnoticed (as seen in the famous demonstration of the live migration of a Quake game server while players are connected to it and participating in the game). Imperceptibility is what makes the illusion of live migration work.

Possible Migration Errors

Migration, though a fun demo, is more important when the capability is actually needed. In fact, for production systems, migration is something that should not be taken lightly. It is important to consider the level of computing power and other resources available when planning to migrate additional guests to a host. An easy mistake is to saturate all available hosts with guests when load balancing, and then at a later time attempt to migrate guests from one heavily loaded host to another that lacks the available resources to keep up with the demand. It is key to remember that if a host cannot normally support a particular guest due to resource limitations, it is not a viable destination host for the migration of that guest. Thus some planning of your migration infrastructure and contingency migration plans are in order.

In situations where systems are under extreme load, and where a noticeable and unacceptable decrease in performance might result from migrating a guest from one busy host to another, it is wise to have extra guests on the same hosts, and instead try to mix guests with different types of loads in a way that minimizes competition for the same resources of hosts.

The problem shown in Listing 13.33 can be encountered when:

- The destination host does not have enough memory to complete migration and now the guest domain is in a nonfunctioning zombie state.
- The `xm migrate` command is terminated prematurely on the source host.

LISTING 13.33 "Zombie-migrating" Guest Listed

```
[root@dom0_Host1]# xm list
Name                              ID Mem(MiB) VCPUs State   Time(s)
Domain-0                           0      191     2 r-----  49980.3
Zombie-migrating-TestGuest        14      256     1 ---s-d   1421.6
[root@dom0_Host1]#
```

Unfortunately, at the time of this writing, if this situation is encountered, the zombie guest cannot be destroyed, and the host needs to be restarted to alleviate this problem. A possible workaround for the zombie's memory consumption is to use `xm mem-set` to lower the memory allocated to the zombie to the minimum allowed amount (1MB).

Summary

In this chapter, we explored Xen's integrated facilities for guest domain checkpointing and migration that make these tasks trivial and useful for system administration. Xen makes saving and restoring guest domain states to disk possible through the easy-to-use `xm save` and `xm restore` commands. The ability to checkpoint guests makes it possible to pause and resume guest execution at a later date, roll back guest states for faster-than-rebooting crash recovery, and crudely transport guests between hosts.

Xen's integrated support for migration makes relocating guests between distinct physical hosts much more intuitive than manual means of relocating a guest image. Migration is performed with the `xm migrate` command, optionally with the `--live` argument for live migration. Live migration is particularly useful for crash avoidance, load balancing, and bringing down hosts for maintenance without interrupting the operation of important virtual machines.

References and Further Reading

Clark, Christopher et al. "Live Migration of Virtual Machines."
`http://www.cl.cam.ac.uk/research/srg/netos/papers/2005-migration-nsdi-pre.pdf`.

Virijevich, Paul. "Live Migration of Xen Domains." Linux.com.
`http://www.linux.com/articles/55773`.

Xen Users' Manual, Xen v3.0.
`http://www.cl.cam.ac.uk/research/srg/netos/xen/readmes/user/user.html`.

Index

C

D

E

F

H

I

J-K

L

M

N

O

P

S

T

U

V

W

X

Y-Z