

VISIT US AT

www.syngress.com

Syngress is committed to publishing high-quality books for IT Professionals and delivering those books in media and formats that fit the demands of our customers. We are also committed to extending the utility of the book you purchase via additional materials available from our Web site.

SOLUTIONS WEB SITE

To register your book, visit www.syngress.com/solutions. Once registered, you can access our solutions@syngress.com Web pages. There you may find an assortment of value-added features such as free e-books related to the topic of this book, URLs of related Web sites, FAQs from the book, corrections, and any updates from the author(s).

ULTIMATE CDs

Our Ultimate CD product line offers our readers budget-conscious compilations of some of our best-selling backlist titles in Adobe PDF form. These CDs are the perfect way to extend your reference library on key topics pertaining to your area of expertise, including Cisco Engineering, Microsoft Windows System Administration, CyberCrime Investigation, Open Source Security, and Firewall Configuration, to name a few.

DOWNLOADABLE E-BOOKS

For readers who can't wait for hard copy, we offer most of our titles in downloadable Adobe PDF form. These e-books are often available weeks before hard copies, and are priced affordably.

SYNGRESS OUTLET

Our outlet store at syngress.com features overstocked, out-of-print, or slightly hurt books at significant savings.

SITE LICENSING

Syngress has a well-established program for site licensing our e-books onto servers in corporations, educational institutions, and large organizations. Contact us at sales@syngress.com for more information.

CUSTOM PUBLISHING

Many organizations welcome the ability to combine parts of multiple Syngress books, as well as their own content, into a single volume for their own internal use. Contact us at sales@syngress.com for more information.

SYNGRESS®

VIRTUALIZATION WITH XEN™

INCLUDING

XenEnterprise,™
XenServer,™ and
XenExpress™

David E. Williams Technical Editor

Juan Garcia

Foreword by Simon Crosby

CTO and Founder, XenSource



Elsevier, Inc., the author(s), and any person or firm involved in the writing, editing, or production (collectively “Makers”) of this book (“the Work”) do not guarantee or warrant the results to be obtained from the Work.

There is no guarantee of any kind, expressed or implied, regarding the Work or its contents. The Work is sold AS IS and WITHOUT WARRANTY. You may have other legal rights, which vary from state to state.

In no event will Makers be liable to you for damages, including any loss of profits, lost savings, or other incidental or consequential damages arising out from the Work or its contents. Because some states do not allow the exclusion or limitation of liability for consequential or incidental damages, the above limitation may not apply to you.

You should always use reasonable care, including backup and other appropriate precautions, when working with computers, networks, data, and files.

Syngress Media®, Syngress®, “Career Advancement Through Skill Enhancement®,” “Ask the Author UPDATE®,” and “Hack Proofing®,” are registered trademarks of Elsevier, Inc. “Syngress: The Definition of a Serious Security Library”™, “Mission Critical™,” and “The Only Way to Stop a Hacker is to Think Like One™” are trademarks of Elsevier, Inc. Brands and product names mentioned in this book are trademarks or service marks of their respective companies.

KEY	SERIAL NUMBER
001	HJIRTCV764
002	PO9873D5FG
003	829KM8NJH2
004	BPOQ48722D
005	CVPLQ6WQ23
006	VBP965T5T5
007	HJJJ863WD3E
008	2987GVTWMK
009	629MP5SDJT
010	IMWQ295T6T

PUBLISHED BY
Syngress Publishing, Inc.
Elsevier, Inc.
30 Corporate Drive
Burlington, MA 01803

Virtualization with Xen™: Including XenEnterprise™, XenServer™, and XenExpress™

Copyright © 2007 by Elsevier, Inc. All rights reserved. Printed in the United States of America. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher, with the exception that the program listings may be entered, stored, and executed in a computer system, but they may not be reproduced for publication.

Printed in the United States of America
1 2 3 4 5 6 7 8 9 0
ISBN: 978-1-59749-167-9

Publisher: Amorette Pedersen
Acquisitions Editor: Andrew Williams
Technical Editor: David E. Williams
Cover Designer: Michael Kavish

Project Manager: Gary Byrne
Page Layout and Art: Patricia Lupien
Copy Editors: Audrey Doyle, Mike McGee
Indexer: Julie Kawabata

For information on rights, translations, and bulk sales, contact Matt Pedersen, Commercial Sales Director and Rights, at Elsevier; email m.pedersen@elsevier.com.



Acknowledgments

Syngress would like to thank Simon Crosby, Ian Pratt, Roger Klorese, Aimee Francioni, Leo Zarkhin, and the rest of the team at XenSource for your support and guidance throughout the development of this project. We could not have done it without you.



Lead Author

David E. Williams is a principal at Williams & Garcia, LLC, a consulting practice based in Atlanta, GA, specializing in effective enterprise infrastructure solutions. He specializes in the delivery of advanced solutions for x86 and x64 environments. Because David focuses on cost containment and reduction of complexity, virtualization technologies have played a key role in his recommended solutions and infrastructure designs. David has held several IT leadership positions in various organizations, and his responsibilities have included the operations and strategy of Windows, open systems, mainframe, storage, database, and data center technologies and services. He has also served as a senior architect and an advisory engineer for Fortune 1000 organizations, providing strategic direction on technology infrastructures for new enterprise-level projects.

David studied Music Engineering Technology at the University of Miami, and he holds MCSE+I, MCDBA, VCP, and CCNA certifications. When not obsessed with corporate infrastructures, he spends his time with his wife and three children.



Contributors

Kris Buytaert is Founder and CTO of X-Tend. He is a longtime Linux, Security, and Open Source consultant. He has consulting and development experience with multiple enterprise-level clients and government agencies. In addition to his high-level technical experience, he is also a team leader who likes to deliver his projects on time. He is a contributor to the Linux Documentation Project and author of various technical publications. Kris is a Red Hat Certified Engineer and is currently the maintainer of the openMosix HOWTO Web site. Kris is also a frequent speaker at Linux and OpenSource conferences. He is currently focusing on Linux clustering (both HA and HPC), virtualization, and large infrastructure management.

Juan R. Garcia is a Principal Consultant at Williams & Garcia, LLC. He provides strategic and technical consulting in legacy systems migrations, enterprise architecture, disaster recover planning, and enterprise IT resource consolidation to Williams & Garcia's customers. He specializes in open systems (UNIX/Linux), virtualization technologies (VMware, Xen, and AIX 5L), storage solutions, and RDMBS technologies. Juan's previous positions include Solutions Architect for Bellsouth, Senior Enterprise Architect for John H. Harland Co., and Technical Manager for Sun Professional Services.

Rami Rosen (B.Sc, Computer Science, Technion—Israel High Institute of Technology) is working as a Linux and Open Solaris kernel programmer accompanying advanced networking and security projects. His background includes positions in Ethernet switching and Avionic operating system start-ups. His specialities include virtualization technologies and kernel networking internals. His articles are occasionally published in the *Linux Journal* and the lwn.net Web site.



Foreword Contributor

Simon Crosby is an industry evangelist for the Xen™ open source hypervisor and CTO of XenSource Inc. In this position, he is responsible for XenEnterprise R&D, technology leadership, and product management. He also maintains a close affiliation to the Xen project run by Ian Pratt, the founder of XenSource. Prior to XenSource, Simon was a principal engineer at Intel, where he led strategic research in distributed autonomic computing, platform security, and trust. Before joining Intel, Simon founded CPlane Inc., a network optimization software vendor, and held a variety of executive roles while there, including president & CEO, chairman, and CTO. Prior to his position at CPlane, Simon was a tenured faculty member at the University of Cambridge, U.K., where he led research on network performance and control, and multimedia operating systems. He is an

author of more than 35 research papers and patents on a number of data center and networking topics, including security, network and server virtualization, and resource optimization and performance. Simon is a frequent speaker on the topic of enterprise-grade virtualization with open source technologies, and he has most recently been a presenter at such well-known industry events as LinuxWorld, Interop, and the Server Blade Summit.

Contents

Foreword	xix
Chapter 1 An Introduction to Virtualization	1
Introduction	2
What Is Virtualization?	2
The History of Virtualization	3
The Atlas Computer	3
The M44/44X Project	4
CP/CMS	4
Other Time-Sharing Projects	5
Virtualization Explosion of the 1990s and Early 2000s	6
The Answer: Virtualization Is... ..	8
Why Virtualize?	9
Decentralization versus Centralization	9
True Tangible Benefits	13
Consolidation	15
Reliability	17
Security	18
How Does Virtualization Work?	19
OS Relationships with the CPU Architecture	20
The Virtual Machine Monitor and Ring-0 Presentation	22
The VMM Role Explored	23
The Popek and Goldberg Requirements	23
The Challenge: VMMs for the x86 Architecture ...	25
Types of Virtualization	26
Server Virtualization	26
Storage Virtualization	29
Network Virtualization	30
Application Virtualization	31
Common Use Cases for Virtualization	32
Technology Refresh	32
Business Continuity and Disaster Recovery	34
Proof of Concept Deployments	35
Virtual Desktops	35

- Rapid Development, Test Lab,
and Software Configuration Management 36
- Summary 38
- Solutions Fast Track 38
- Frequently Asked Questions 41
- Chapter 2 Introducing Xen. 43**
- Introduction 44
- What Is Xen? 44
 - Features of Xen 46
 - The XenServer Product Family 47
- Xen’s Virtualization Model Explored 50
 - Architecture Overview 50
 - Processor Architecture 51
 - Paravirtualization with Xen 52
 - Xen Domains 54
- CPU Virtualization 58
 - Exceptions 59
 - CPU Scheduling 60
 - Time 62
- Memory Virtualization 63
 - Memory Allocation 64
 - Page Tables and Segmentation 66
 - Virtual Address Translation 69
- I/O Virtualization 72
 - Device I/O Rings 73
 - Event Channels 76
 - Virtual I/O Devices and Split Device Drivers 76
 - Network I/O 78
 - Block I/O 79
 - Trusted Platform Module and Other Devices 79
 - Driver Domains 79
 - Software and Hardware IOMMUs 81
 - SWIOTLB 81
 - Grant Tables 82
- The Xenstore 82
- Summary 87
- Solutions Fast Track 87
- Frequently Asked Questions 92

Chapter 3 Deploying Xen: Demystifying the Installation	95
Introduction	96
Determining Which Xen to Choose	96
System Requirements	97
Thinking Before You Start	99
Installing Xen on a Free Linux Distribution	101
Fedora Core 6	101
VirtManager	114
Installing Windows XP	121
Installing the XenServer Product Family	125
What Is XenServer?	125
XenServer Requirements	126
Getting and Installing XenServer	127
Installing the Host	127
Client Installation	135
Installing an Initial Virtual Machine on XenServer . .	138
Other Xen Installation Methods	142
Using the XenSource Binaries and LVM	142
Configuring Xen	146
Getting Xen on Your Network	149
Summary	153
Solutions FastTrack	153
Frequently Asked Questions	156
Chapter 4 The Administrator Console and Other Native Tools	159
Introduction	160
Native Xen Command-Line Tools	161
The xe Command-Line Interface	161
Installing and Cloning XenVMs	162
Starting Up, Shutting Down, Rebooting, Suspending, and Resuming XenVMs . .	162
Shutting Down and Rebooting XenHosts	163
Query Options for XenHosts	163
XenServer Administrator Console	163
System Requirements for the Administrator Console . .	164
Installing the Administrator Console	164

- Installing the Administrator Console on Windows (XP/2000/2003)164
- Installing the Administrator Console on Linux169
- Using the Administrator Console170
 - Working with Hosts171
 - Deploying and Configuring XenVMs173
- Summary177
- Solutions Fast Track177
- Frequently Asked Questions178

Chapter 5 Managing Xen with Third-Party Management Tools 179

- Introduction180
- Clusters openQRM180
 - Xen Management with openQRM181
 - Overview181
 - General Concepts for the Xen/openQRM Mix182
 - Plug-ins and Licensing184
 - Installing openQRM187
 - System Requirements188
 - Installing openQRM 3.1.x Server190
 - Installing the openQRM Xen Plug-in193
 - Managing Xen with openQRM196
 - Provisioning with openQRM-Pro201
- Enomalism203
 - Overview of Enomalism204
 - Installing Enomalism205
 - System Requirements205
 - Installation Walkthrough206
 - Using Enomalism to Manage Xen206
- Project ConVirt and XenMan210
 - Overview of ConVirt211
 - Installing ConVirt212
 - System Requirements212
 - Installation212
 - Using ConVirt to Manage Xen214
- Summary220
- Solutions Fast Track220
- Frequently Asked Questions222

Chapter 6 Deploying a Virtual Machine in Xen	225
Introduction	226
Workload Planning and Virtual Machine Placement	226
Memory	226
CPU	227
Network	228
Installing Modified Guests	229
Installing Red Hat Enterprise Linux 4	229
Installing Unmodified Guests	235
Installing Red Hat Linux Enterprise 5	236
Installing Windows Guests	240
Windows Guest Installation	240
Physical-to-Virtual Migrations of Existing Systems	245
P2V Migration	245
Importing and Exporting Existing Virtual Machines	246
Exporting XenVMs	247
Importing XenVMs	249
Summary	252
Solutions Fast Track	252
Frequently Asked Questions	255
Chapter 7 Advanced Xen Concepts	257
Introduction	258
The Virtual Split Devices Model	258
Advanced Storage Concepts	259
High-Performance Solutions for Xen	259
iSCSI Integration with Xen	260
Copy-on-Write	263
DmUserspace	264
UnionFS	265
Advanced Networking Concepts	266
Bridging VLANs	266
Creating Interface Bonds for	
High Availability and Link Aggregation	267
Routing, Forwarding, and Other Network Tricks	269
Building a Xen Cluster	270
XenVM Migration	279
XenVM Backup and Recovery Solutions	282

- Options for Backing Up Your XenVM282
- Making Xen Part of Your Disaster Recovery Plan283
- Full Virtualization in Xen283
 - The New Processors with
 - Virtual Extensions (VT-x and AMD-V)284
- Summary288
- Solutions Fast Track288
- Frequently Asked Questions290

Chapter 8 The Future of Virtualization. 293

- Introduction294
- The Unofficial Xen Road Map294
 - Performance and Scalability295
 - NUMA-Aware Architecture296
 - Multicore Processors298
 - Smart I/O299
 - Operating System Support301
 - Support in Linux Distributions301
 - Xen and Microsoft302
 - Other HVM Guests303
 - Beyond the x86 CPU Architecture303
 - IA-64 Feature Sync with x86304
 - Porting to PowerPC304
 - Porting to the UltraSPARC Architecture306
 - Architecture Enhancements307
 - Control Tools307
 - Virtual Hard Disk Images and XenFS308
 - Virtual Device Enhancements309
- Virtual Infrastructure in Tomorrow’s Data Center311
 - Technology Trends Driving
 - Improvements in Virtualization311
 - Hardware Economies of Scale311
 - Multicore and Multithreaded Computing312
 - Solutions for Small and Medium-Sized Businesses316
 - Integrated Computing316
 - Data Center in a Box317
 - Large Enterprises318
 - Reliability and Availability318
 - Security320

Compliance	323
The Magic Recipe: Other Hardware and Software Virtualization	
Trends	324
Increasing Density Further with Blade Servers	325
Storage Virtualization	325
Network Virtualization	326
Summary	327
Solutions Fast Track	327
Frequently Asked Questions	329
Appendix A Glossary of Terms	331
Appendix B Other Virtualization	
Technologies and How They Compare to Xen	343
Index	357

Foreword

The open source Xen project, led by Ian Pratt of the University of Cambridge and XenSource Inc., will arguably have a greater impact on the enterprise software industry than Linux has had. The Xen hypervisor is a lightweight, high-performance, secure virtualization platform that is now collaboratively developed by more than 20 of the industry's largest enterprise IT vendors as an open industry standard for virtualization. Its architecture has tremendous advantages over existing virtualization technologies: It has broad hardware support through its reuse of existing operating systems such as Linux to safely virtualize I/O for other guests; it offers superb resource partitioning for performance isolation and security; and it can be implemented either as a virtualization platform or as an integrated component within an operating system. Xen has been ported to a wide range of hardware architectures, including x86, x86_64, the Intel Itanium, IBM's PowerPC, the SGI Altix, and the ARM 9. It also is ideally suited to hardware-accelerated virtualization. For these reasons, and because Xen is freely available in source code form, every major OS vendor has either adopted Xen or the Xen architecture as a core component of the next major OS release for the x86 platform.

As a catalyst of change in the IT industry, the Xen project needs to scale its knowledge and skill base so that a competent IT pro can easily get a handle on the technology and how to deploy and use it. Cool technology can be adopted only as fast as human users can acquire the understanding and skills to use it. Perhaps more important, however, is the fact that the Xen project relies on the innovation of the community to continue the development of its feature set. It is therefore a great

pleasure and also a great relief to introduce this timely, thorough, and highly accessible guide to the art of Xen virtualization. This book demystifies Xen by placing it in a practical context that any IT pro who wants to get something working will immediately understand, while also providing a thorough grounding in the architecture and implementation of the hypervisor. It also takes an important step beyond the basics of Xen by offering a detailed tutorial on how to use the definitive platform implementation of Xen, XenEnterprise from XenSource. Although enthusiasts may want to dig into Xen by building it from source, most readers will be delighted to find included with this book a CD containing the powerful free XenExpress bare-metal hypervisor from XenSource ready to install for production virtualization of Linux and Windows guests.

It is my hope that this book will achieve two aims: encourage a new generation of contributors to the Xen project and foster broad adoption of the Xen hypervisor as a ubiquitous open standard for virtualization.

—*Simon Crosby*
CTO and Founder
XenSource Inc.

About the CDs

This book includes two CDs to help you consolidate Windows, Linux, or mixed deployment with XenExpress 3.2. This free, production-ready virtualization platform enables you to quickly get started with Xen virtualization. With XenExpress you can host up to four virtual servers running Linux or Windows (Intel VT or AMD-V hardware-assist virtualization technology is required for Windows) on a broad range of standard server hardware. XenExpress supports dual-socket servers with up to 4 GB of RAM, offering you all of the base performance, tools, and easy-to-use features of XenEnterprise. Easily installed and seamlessly upgradable to XenEnterprise or XenServer, XenExpress is the ideal on-ramp to Xen.