



PRENTICE
HALL

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

Contents

Foreword.....	xxi
Preface.....	xxiii
Chapter 1: Xen—Background and Virtualization Basics	1
Common Uses and Benefits of Virtualization.....	2
Types of Virtualization	5
Emulation.....	6
Full Virtualization.....	7
Paravirtualization	8
Operating System Level Virtualization	9
Other Types of Virtualization.....	11
Overview of Virtualization Types	12
Virtualization Heritage.....	13
The IBM Mainframe	14
Virtualization on Commodity Hardware	15
Virtualization Extensions for x86	15
Xen Origins and Time Line	15
Other Virtualization Systems for Commodity Hardware.....	18
Emulation.....	18
Full Virtualization.....	19

Paravirtualization	21
Operating System Virtualization	23
Popular Virtualization Products	24
Summary.....	25
References and Further Reading	26
Chapter 2: A Quick Tour with the Xen LiveCD	27
Running the LiveCD.....	28
Step 1: Downloading the LiveCD Image and Creating the CD	29
Step 2: Choosing a Domain0 Image from the GRUB Menu	30
Step 3: Logging In and the Desktop	31
Step 4: Creating Guests	33
Step 5: Deleting a Guest.....	38
Step 6: Interacting with Your Guests	38
Step 7: Testing Your Networking	41
Too Many Guests	44
Summary.....	44
References and Further Reading	45
Chapter 3: The Xen Hypervisor.....	47
Xen Hypervisor	48
A Privileged Position	50
Protection Rings	50
Domain0.....	51
Xen Boot Options	54
Choosing an OS for Domain0	59

xend	60
Controlling xend	60
xend Logs	62
xend Configuration	63
XenStore	67
Summary.....	73
References and Further Reading	73
Chapter 4: Hardware Requirements and Installation of Xen Domain0	75
Xen Domain0 Processor Requirements.....	76
Intel VT.....	77
AMD-V.....	77
HVM.....	78
Hardware Device Support and Recommendations.....	78
Disks and Controllers	78
Networking Devices.....	80
Graphics Devices	80
Power Management	81
Help for Unsupported Hardware	81
Memory Requirements.....	81
Choosing and Obtaining a Version of Xen	83
Open Source Distributions	83
Commercially Supported Options	84
Methods of Installing Domain0 Hosts	86
Common Prerequisite: The Grand Unified Boot Loader (GRUB).....	87
Linux Distributions.....	87
OpenSUSE.....	88
CentOS	91

Ubuntu.....	98
Xen from Binary Packages.....	101
Gentoo	105
XenExpress.....	112
Non-Linux Domain0 Installations	114
Building from Source	116
Summary.....	118
References and Further Reading.....	118
Chapter 5: Using Prebuilt Guest Images.....	121
Introduction to DomU Guests.....	122
Guest Images	122
Operating System Kernels.....	123
Configuration Files	123
Working with Prebuilt Guest Images.....	128
Types of Guest Images	128
Downloading Prebuilt Guest Images.....	130
Mounting and Booting Prebuilt Images	131
Downloading Compressed File System Images	146
Converting Images from Other Virtualization Platforms	161
Summary.....	162
References and Further Reading.....	163
Chapter 6: Managing Unprivileged Domains.....	165
Introduction to the xm Utility	166
Prerequisites for Running the xm Utility	166
Generic Format of an xm Command.....	167

The <code>xm list</code> Subcommand	169
Basic List Information.....	169
Listing Information about a Specific Guest	171
long Option.....	172
Label Option.....	173
The <code>xm create</code> Subcommand.....	174
Prerequisites for <code>xm create</code>	174
Simple Examples of <code>xm create</code>	175
Guest Configuration Files	178
Python Format.....	178
Common Configuration Options	179
S-Expression (SXP) Format.....	180
Path to Configuration Files.....	181
Diagnosing Problems with Guest Creation.....	182
Dry Run	182
Console Output.....	183
Sample Problems.....	184
Automatically Starting DomUs	191
Shutting Down Guest Domains	193
<code>xm shutdown</code>	193
<code>xm reboot</code>	196
<code>xm destroy</code>	198
Pausing Domains	199
<code>xm pause</code>	200
<code>xm unpause</code>	200
Interacting with a Guest Nongraphically	201
<code>xm console</code>	202
SSH.....	204

Interacting with a Guest Graphically	204
X Forwarding with SSH.....	205
Configuration of SSH Server and Client.....	205
VNC.....	207
Virtual Frame Buffer and Integrated VNC/SDL Libraries.....	210
Freenx.....	212
Remote Desktop.....	213
Summary.....	215
References and Further Reading.....	216
Chapter 7: Populating Guest Images	217
Hardware Virtual Machine (HVM) Guest Population.....	218
Populating a Guest Image from a Disc or Disc Image (Windows XP Example)....	218
Automated Population with <code>virt-install</code>	225
Paravirtualized (PV) Guest Population	228
OpenSUSE: YaST Virtual Machine Management.....	229
CentOS/Fedora: <code>virt-manager</code>	233
Debian/Ubuntu: <code>debootstrap</code>	242
Gentoo: <code>quickpkg</code> and <code>domi</code> Scripts	246
Xen Express	256
Guest Image Customization	266
Customizing Hostnames.....	266
Customizing Users.....	267
Customizing Packages and Services.....	268
Customizing the File System Table (<code>/etc/fstab</code>)	268
Converting Existing Installations.....	270
Summary.....	274
References and Further Reading.....	274

Chapter 8: Storing Guest Images	277
Logical Volumes	278
Basic LVM Usage	279
Resizing Images.....	282
Image Snapshots Using Copy on Write	286
Network Image Storage Options	287
iSCSI	288
ATA over Ethernet (AoE).....	293
NFS.....	297
Comparing Network Storage Options.....	300
Guest Image Files	301
Preparing Compressed tar Image Files.....	301
Preparing Disk Image Files.....	302
Preparing Guest Partition Image Files	312
Mounting Disks and Partition Images.....	314
Summary.....	316
References and Further Reading	316
Chapter 9: Device Virtualization and Management	319
Device Virtualization.....	320
Paravirtualization of Devices	320
Full Virtualization of Devices.....	321
No Virtualization.....	321
Backends and Frontends.....	322
Backend Information in XenStore.....	323
Frontend Information in XenStore.....	325

Granting Control of a PCI Device	326
Identifying a PCI Device.....	326
Hiding a PCI Device from Domain0 at Boot.....	327
Manually Unbinding/Binding a PCI Device at Runtime.....	328
Granting a PCI Device to Another Domain.....	329
Exclusive Device Access Versus Trusted Driver Domains	331
Exclusive Device Access.....	331
Trusted Driver Domains	332
<i>Problems Using Trusted Driver Domains</i>	333
Device Emulation with QEMU-DM	334
Future Directions	335
More Devices.....	336
Smart Devices	336
Summary.....	336
References and Further Reading	337
Chapter 10: Network Configuration	339
Network Virtualization Overview.....	340
Designing a Virtual Network Topology	341
Bridging, Routing, and Network Address Translation.....	343
Frontend and Backend Network Drivers and Naming.....	347
Overview of Network Configuration in Xen	349
High-Level Steps.....	349
Xend Configuration File	350
Guest Domain's Configuration File	352
Details of Bridging Mode.....	354
Bridging Configuration Example	355
Testing Results.....	361

Details of Routing Mode.....	364
Routing Configuration Example	365
Testing Results	371
Details of NAT Mode	373
NAT Configuration Example.....	373
Testing Results	379
Configuring Purely Virtual Network Segments.....	382
Configuring dummy0	383
Testing dummy0	385
Configuring Dummy Bridge.....	385
Testing Dummy Bridge.....	388
Assigning MAC Addresses to Virtual Network Interfaces	389
MAC Addresses	389
Specifying or Generating a MAC Address for a Guest Domain	390
Assigning IP Addresses	391
Using an External DHCP Server to Obtain an IP for a Guest Domain	392
Manually Assigning an IP to a Guest Domain.....	392
Using an Internal DHCP Server to Obtain an IP for a Guest Domain.....	393
Handling Multiple Network Interfaces in a Domain	394
Handling Multiple Network Interfaces in a driver domain	394
Handling Multiple Network Interfaces in a Guest Domain.....	396
vnet—Domain Virtual Network	399
Installing vnet	400
Running vnet.....	401
Summary.....	403
References and Further Reading	403

Chapter 11: Securing a Xen System	405
Structuring Your System for Security	406
Special Purpose Virtual Machines	406
Creating Virtual Network Segments.....	407
Securing the Privileged Domain	407
Removing Software and Services	407
Limiting Remote Access.....	408
Limiting the Local Users.....	412
Move Device Drivers into DriverDomains.....	412
Firewall and Network Monitors.....	413
Running a Firewall with iptables	413
Snort.....	419
Obtaining Snort.....	419
Snort and Network Intrusion Detection Mode	420
Mandatory Access Control with sHype and Xen Security Modules	422
sHype	423
Xen Security Modules (XSM)	432
DomU Security	433
Running VMs Only When Needed.....	434
Backing Up Virtual Machine Images.....	434
Summary.....	435
References and Further Reading	436
 Chapter 12: Managing Guest Resources	 437
Accessing Information about Guests and the Hypervisor.....	438
xm info.....	438
xm dmesg.....	443
xm log	444

xm top.....	446
xm uptime.....	449
Allocating Guest Memory	449
Shadow Page Tables.....	451
Balloon Driver	451
Improving Stability with Swap Space	454
Managing the Allocation of Guest Memory	454
Managing Guest Virtual CPUs.....	458
Comparing Virtual, Logical, and Physical Processors	458
HVM VCPU Management.....	459
VCPU Subcommands.....	460
When to Manually Administer VCPUs.....	462
Tuning the Hypervisor Scheduler	463
Weight and Cap.....	463
Protection from Misbehaving Guests	464
Using the Credit Scheduler Command	465
Choosing a Guest IO Scheduler	466
Noop Scheduler	466
Deadline Scheduler	466
Anticipatory Scheduler (as)	467
Complete Fair Queuing Scheduler (cfq).....	467
Using IO Schedulers	467
Summary.....	469
References and Further Reading.....	469
Chapter 13: Guest Save, Restore, and Live Migration.....	471
Representing the State of a Virtual Machine.....	472

Basic Guest Domain Save and Restore.....	473
xm save.....	474
xm restore.....	476
Possible Save and Restore Errors.....	478
Types of Guest Relocation.....	479
Cold Static Relocation.....	480
Warm Static (Regular) Migration.....	481
Live Migration.....	482
Preparing for xm migrate.....	484
Configuring xend.....	485
Proximity of Sources and Destinations on the Network.....	488
Network-Accessible Storage.....	489
Guest Domain Configuration.....	489
Version and Physical Resource Requirements.....	491
Experience with xm migrate.....	491
xm migrate.....	491
Using xm migrate for Warm Static Migration.....	492
Using xm migrate for Live Migration.....	494
Possible Migration Errors.....	497
Summary.....	498
References and Further Reading.....	498
Chapter 14: An Overview of Xen Enterprise Management Tools.....	499
Programmatic Interfaces to the Xen Hypervisor.....	500
Libvirt.....	500
Xen-CIM.....	501
Xen API.....	501
Legacy Interfaces to Xend.....	502

Citrix XenServer Enterprise, Standard and XenExpress Editions.....	502
Virtual Iron.....	504
IBM Virtualization Manager.....	506
Enomalism.....	507
virt-manager.....	509
XenMan.....	513
Managing Multiple Systems.....	518
Summary.....	518
References and Further Reading.....	519
Appendix A: Resources.....	521
Xen Community.....	522
XenWiki.....	523
Xen Mailing Lists and Bug Reporting.....	524
Xen Summits.....	525
Xen Source Code.....	526
Academic Papers and Conferences.....	528
Distribution-Specific Resources.....	530
Appendix B: The xm Command.....	531
Appendix C: Xend Configuration Parameters.....	537
Appendix D: Guest Configuration Parameters.....	541
Appendix E: Xen Performance Evaluation.....	545
Xen Performance Measurements.....	546
Repeatability of the Xen Team's Results.....	546
Xen and Virtual Web Hosting.....	548

Comparing XenLinux to Native Linux on Older PC Hardware	550
Xen on x86 Versus IBM zServer	551
Performance Isolation in Xen	553
Performance of Xen Virtual Network and Real Network	556
Summary.....	558
Index.....	559