SOVEREIGN CLOUD

Apache CloudStack Production Deployment Guide

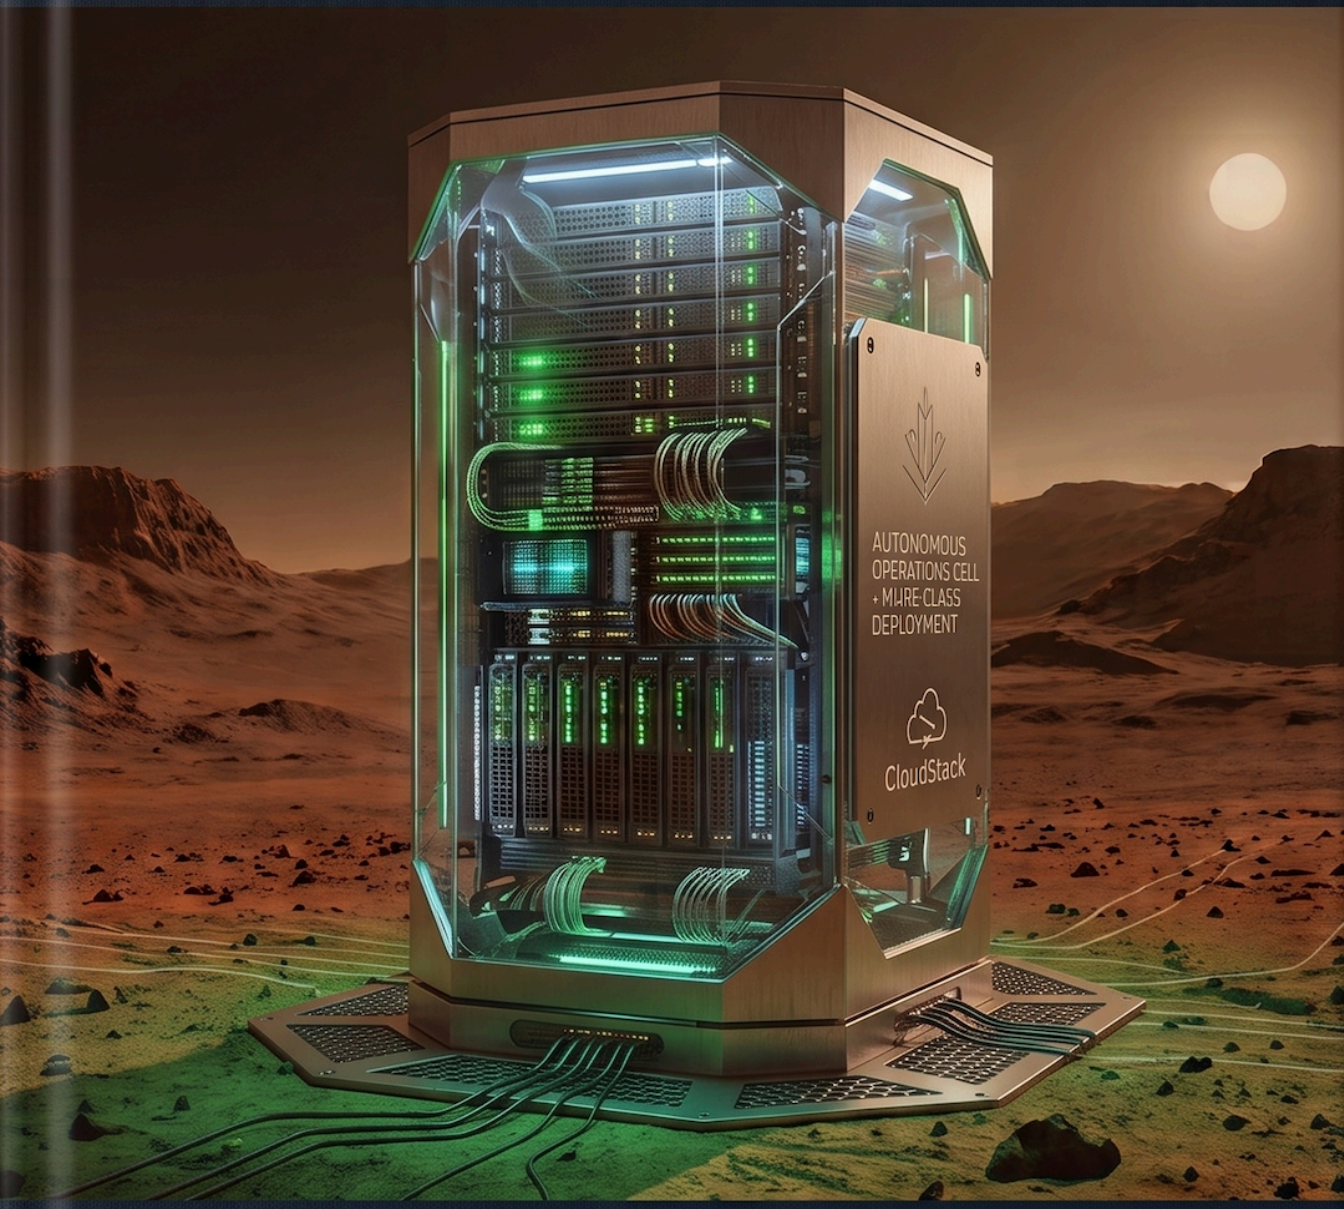

Volume I

Michael Hinsley

Apache CloudStack — A Production Deployment Guide

Part 1 — Foundations · Building a Sovereign Cloud on Apache CloudStack 4.22.1.0 LTS

Michael Hinsley

This book is available at https://leanpub.com/apache-cloudstack-deployment-guide

This version was published on 2026-08-08

This is a Leanpub book. Leanpub empowers authors and publishers with the Lean Publishing process. Lean Publishing is the act of publishing an in-progress ebook using lightweight tools and many iterations to get reader feedback, pivot until you have the right book and build traction once you do.

© 2026 Michael Hinsley

Contents

Sovereign Cloud **1**
On owning what you run 3

The Book in Full — Annotated Contents **4**
Volume 1 — Foundations (Chapters 1-18) 4
Appendices 7

How This Book Is Proven **8**
The lab behind the book 8
What actually happened 8
Why you can check any of this 9
And now you can run the proof yourself 9
What this means for you 10

Chapter 1: The Cost of Renting Your Existence **11**
Meet Leaf Spine Books — Understanding CloudStack Architecture Through Real-World Needs 11
The Story: The Renewal Letter 12
Understanding CloudStack: Building Blocks of Your Cloud 13
System Virtual Machines: CloudStack's Control Layer 19
The Management Server: Brain of the Cloud 21
Network Architecture: The Four Traffic Types 22
CloudStack 4.22 LTS: Why This Release 26
Basic vs Advanced Networking: A Critical Decision 30
Hypervisor Selection: Why KVM? 32
Storage Strategy: Why Ceph? 35
The Road Ahead: Leaf Spine's Implementation Journey 36
Key Takeaways 37
Questions to Answer Before Next Chapter 38
Coming in Chapter 2 39
Additional Resources 39

Chapter 2: Designing the Three-City UK Cloud **41**
From Whiteboard to Blueprint — Planning a Production-Grade Infrastructure 41
The Story: Architecture Week 41
Design Principles 41
Geographic Layout 43
Network Architecture 45
VLAN and IP Architecture 48
Hardware Specifications 50

Deployment Overview per Site . . . 54
CloudStack Zone Architecture . . . 56
Security Architecture . . . 57
AI and GPU Strategy . . . 58
Automation with Ansible . . . 59
Cost Analysis . . . 61
Migration Path . . . 63
Disaster Recovery Strategy . . . 63
Capacity Planning . . . 64
Monitoring and Alerting . . . 64
Documentation Standards . . . 64
Next Steps . . . 65
Chapter Recap: 5 Key Takeaways . . . 65
Up Next . . . 65
References . . . 65

Chapter 3: Building Your Learning Lab . . . 67
From One Second-Hand Server to a Cloud You Can Break . . . 67
Why a Lab at All . . . 67
The Story: Planning for the Future . . . 68
The Three-Phase Roadmap . . . 68
Phase 1: The Learning Lab . . . 69
The Miniature Estate . . . 72
Phase 2: Network Simulation . . . 74
Phase 3: Edge Computing . . . 75
Done When: the Phase 1 Finish Line . . . 75
Evolution to Production . . . 76
Key Takeaways . . . 77
Coming in Chapter 4 . . . 78
Additional Resources . . . 78

Chapter 3: -1 Deployment Guide . . . 79
Building Manchester in Miniature — the Step-by-Step Companion . . . 79
How to Use This Guide . . . 79
Page One: the Setting Everything Depends On . . . 79
The Seven Phases . . . 79
Pre-Deployment Checklist . . . 79
Phase 1: Host Preparation . . . 80
Phase 2: Virtual Networks . . . 81
Phase 3: Provisioning the Seven VMs . . . 81
Phase 4: The Management Plane . . . 81
Phase 5: The Storage Cluster . . . 82
Phase 6: Secondary Storage . . . 83
Phase 7: Building the Manchester Zone . . . 83

Chapter 3: -2 Lab Exercises . . . 84
How to Use This File . . . 84
Challenge Projects . . . 84
Completion Criteria . . . 84

Chapter 3: -3 Troubleshooting Guide **85**
Diagnosis and Recovery for the Manchester Miniature 85
How to Use This Guide 85
The Method: Five Steps 85
Section 1: First Boot — the Usual Suspects 85
Section 2: Host Layer 85
Section 3: Network Layer 85
Section 4: Staged Drill One — the MTU Mismatch 85
Section 5: Staged Drill Two — Losing Monitor Quorum 85
Section 6: CloudStack Layer 86
Section 7: Ceph Layer — the Lab-Scale WARN Taxonomy 86
Section 8: The Diagnostic Sweep 86
Section 9: Recovery — the Graduated Ladder 86
Section 10: Getting Help 86

Chapter 3: -4 VM Quick Reference **87**
The One-Page Card for the Manchester Miniature 87
The Estate 87
The Networks 87
CloudStack Names (the production names, on purpose) 87
The Done-When at a Glance 87
Where the Secrets Live 87
Getting In 87
Daily Commands 88
Quick Fixes: Straight into the Troubleshooting Guide 88
The Chapter 3 File Set 88

Chapter 3: -5 The Ampere ARM Build **89**
The Manchester Miniature on ARM64 Silicon 89
Why ARM at All 89
The One Difference That Reshapes Everything 89
The Other ARM Deltas, Briefly 89
The Estate, Inverted 89
The Build Sequence 89
A Second Host, Later 90
The GPUs, Honestly 90
Done When, on ARM 90
Key Takeaways 90
The Chapter 3 File Set 90
Additional Resources 91

Chapter 3: -6 The Laptop Build **92**
The Manchester Miniature in a Rucksack 92
Why a Laptop at All 92
The One Constraint That Reshapes Everything 92
The Estate, on a Diet 92
The Laptop Deltas 92
The Build Sequence 93
Travelling With a Datacentre 93

Network Simulation on the Move . . . 93
The Switch OS Question . . . 93
Done When, on a Laptop . . . 94
Key Takeaways . . . 94
The Chapter 3 File Set . . . 94
Additional Resources . . . 94

Chapter 4: Infrastructure Preparation . . . 95
Building the Foundation — From Blueprint to Bare Metal . . . 95
The Story: Hands on the Hardware . . . 95
Part 1: Understanding the Infrastructure Stack . . . 95
Part 2: Ubuntu 24.04 LTS Installation . . . 95
Part 3: Post-Installation Baseline . . . 96
Part 4: Network Configuration with Netplan . . . 96
Part 5: Storage Preparation . . . 97
Part 6: System Optimisation . . . 97
Part 7: Security Hardening from First Boot . . . 98
Part 8: Logging Baseline . . . 98
Part 9: Documentation as Code . . . 98
Part 10: Verification . . . 98
Part 11: Replicating to London and Leeds . . . 99
Chapter 4 — Key Takeaways . . . 99

Chapter 5: MySQL High Availability . . . 100
The Database Layer — One Cluster, Two Cities, No Single Point of Failure . . . 100
The Story: The Heart of the Cloud . . . 100
Part 1: Understanding the Database Layer . . . 100
Part 2: The Design — Three Nodes, Two Cities . . . 100
Part 3: Preparing the Nodes . . . 101
Part 4: Installing MariaDB and Galera . . . 102
Part 5: Configuring the Cluster . . . 102
Part 6: Bootstrapping the Cluster . . . 102
Part 7: Creating the CloudStack Databases . . . 103
Part 8: The Single-Writer Routing Layer . . . 103
Part 9: Testing What We Built . . . 104
Part 10: Backups — for What Clustering Cannot Prevent . . . 104
Part 11: Monitoring, Verification, and the Record . . . 105
Chapter 5 — Key Takeaways . . . 105

Chapter 6: CloudStack Management Server . . . 107
The Brain Arrives — Installing the Control Plane on the Nodes That Already Hold Its Memory 107
The Story: Switching On the Brain . . . 107
Part 1: What the Management Server Actually Is . . . 107
Part 2: Pre-Flight Checks and Firewall Doors . . . 107
Part 3: Database Housekeeping — Two Honest Reconciliations . . . 108
Part 4: Installing CloudStack on the First Node . . . 108
Part 5: Initialising the Schema — Without Borrowing Root . . . 108
Part 6: First Start — and First Contact . . . 109
Part 7: Nodes Two and Three — the Cluster Assembles . . . 109

Part 8: First Login . . . 109
Part 9: Three Front Doors, One Name . . . 110
Part 10: Testing What We Built — the Inherited Standard . . . 110
Part 11: Verification and the Record . . . 110
Chapter 6 — Key Takeaways . . . 111

Chapter 7: Secondary Storage . . . 112
The Repository That Rebuilds the Cloud — DRBD and NFS on Refurbished Metal . . . 112
The Story: The Library Question . . . 112
Part 1: Understanding Secondary Storage . . . 112
Part 2: The Design . . . 112
Part 3: Building the Replicated Filesystem . . . 113
Part 4: The NFS Service . . . 113
Part 5: Failover with Keepalived . . . 114
Part 6: Verification . . . 114
Part 7: Troubleshooting . . . 115
Chapter 7 — Key Takeaways . . . 115

Chapter 8: Primary Storage . . . 117
Where the VMs Actually Live — Ceph RBD on the NVMe Estate . . . 117
The Story: The Microseconds Chapter . . . 117
Part 1: Understanding Primary Storage . . . 117
Part 2: The Design . . . 117
Part 3: Deploying the Cluster . . . 118
Part 4: OSDs, the Pool, and the Credential . . . 119
Part 5: Firewall and Verification . . . 119
Part 6: The Drills . . . 120
Part 7: Troubleshooting . . . 120
Chapter 8 — Key Takeaways . . . 120

Chapter 9: Deploying CloudStack — The Zone Goes Live . . . 122
The Morning the Cloud Gets a Shape . . . 122
Where We Stand . . . 122
What This Chapter Is — and Is Not . . . 122
Part 1: The Shape of a Zone . . . 122
Part 2: The Toolbelt — UI, API, and an Honesty Box About the Wizard . . . 122
Part 3: Zone, Physical Network, and the Public Range . . . 123
Part 4: Pod and the Guest VLAN Range . . . 123
Part 5: The Cluster, Deliberately Empty . . . 124
Part 6: The Image Store — and the Auto-Seed Moment . . . 124
Part 7: Primary Storage — Staged, Verbatim, and One Honest Step Early . . . 124
Part 8: Verification, Then — and Only Then — Enablement . . . 125
Troubleshooting the Logical Build . . . 125
Chapter 9 Summary: Five Things to Hold On To . . . 125
The Values Card for Chapter 10 . . . 125

Chapter 10: Adding KVM Compute Hosts — The Day the Stage Crew Arrives . . . 126
The Morning After the Map . . . 126
Where We Stand . . . 126

What This Chapter Is — and Is Not . . . 126
Part 1: The Shape of the Day . . . 126
Part 2: The Bridges — Three Labels on Two Bonds . . . 126
Part 3: The Hypervisor Stack . . . 127
Part 4: The CloudStack Agent and the Service User . . . 127
Part 5: Pre-Flight — Never Assume, Interrogate . . . 127
Part 6: Adding the First Host . . . 128
Part 7: The Cascade . . . 128
Part 8: The GPU Cluster . . . 128
Part 9: Verification — Extending the Zone Script . . . 129
Part 10: Drills — Break It While It's Cheap . . . 129
Troubleshooting the Join . . . 129
Chapter 10 Summary: Five Things to Hold On To . . . 129
The Values Card for Chapter 11 . . . 129

Chapter 11: Deploying Your First Guest VMs — Opening Night . . . 131
The Morning the Cloud Earned Its Keep . . . 131
Where We Stand . . . 131
What This Chapter Is — and Is Not . . . 131
Part 1: The Shape of the Day . . . 131
Part 2: Compute Offerings — The Shapes a VM May Take . . . 131
Part 3: Disk Offerings — Data That Outlives the VM . . . 131
Part 4: The GPU Placement Proof — Five Minutes That Retire an Assumption . . . 131
Part 5: The Guest Template — The SSVM Does the Fetching . . . 132
Part 6: Editorial's Network — Where `cloudbr2` Earns Its Keep . . . 132
Part 7: Deploying the Three Guests . . . 132
Part 8: Everyday Lifecycle — and One Thing the Old Book Got Wrong . . . 133
Part 9: Verification — Extending the Zone Script . . . 134
Troubleshooting the First Workload . . . 134
Chapter 11 Summary: Five Things to Hold On To . . . 134
The Values Card for Chapter 12 . . . 134

Chapter 12: Advanced Networking — The Dell OS10 Leaf-Spine Fabric . . . 135
Turning "Trust the Trunk" Into a Tested Data Plane . . . 135
The Story: The Promissory Notes Come Due . . . 135
Part 1: Why a Fabric, and Why This One . . . 135
Part 2: Building the Fabric in the Lab . . . 135
Part 3: The Underlay — One Routed Fabric . . . 136
Part 4: VLANs, the Host Trunk, and Making 10.100.0.1 Answer . . . 137
Part 5: Proving It — Discharging the IOUs . . . 137
Part 6: What the Fabric Costs — and What the OS10 Licence Buys . . . 138
The Fabric Is Real Now . . . 139
Values Card — What Chapter 12 (OS10 edition) Hands Forward . . . 139

Chapter 13: Advanced Networking with VPCs — Walls Within the Walls . . . 140
Building Multi-Tier Isolation on the Fabric You Just Proved . . . 140
The Story: One Network Was Never Going to Be Enough . . . 140
Where We Stand . . . 140
What This Chapter Is — and Is Not . . . 140

Part 1: The Shape of the Day . . . 140
Part 2: VPCs, and the One Fact That Trips Everyone . . . 140
Part 3: Designing the Catalogue VPC — Addresses Before Anything . . . 140
Part 4: The VPC Offering and the VPC . . . 141
Part 5: The Tiers — Where the Lanes Multiply . . . 141
Part 6: Network ACLs — The Walls, Drawn Both Ways . . . 141
Part 7: The VPC as Code — Reproducible Walls . . . 142
Part 8: Deploying the Tiers, Proving the Walls, and Balancing the Load . . . 142
Part 9: Site-to-Site VPN — Extending the Private Space to Leeds . . . 142
Part 10: Verification — Extending the Zone Script . . . 142
Part 11: Troubleshooting — The Stateless Traps and Their Friends . . . 143
Part 12: Where This Goes — One Pattern, Many Tenants . . . 143
Chapter 13 Summary: Five Things to Hold On To . . . 143
The Values Card for Chapter 14 . . . 143

Chapter 14: Multi-Zone Expansion . . . 144
The Cloud Grows a Map of the Country — Without Stretching a Single Fiction . . . 144
The Story: Three Cities, One Control Plane, No Long-Distance Lies . . . 144
Where We Stand . . . 144
What This Chapter Is — and Is Not . . . 144
Part 1: What "Multi-Zone" Means Here — and the Anti-Pattern It Is Not . . . 144
Part 2: The Site Inventories — Canonical Estate, Honest Addressing . . . 144
Part 3: Extending the Fabric — Leaf Pairs to the Manchester Spines . . . 145
Part 4: London's Independent Ceph Cluster — the Chapter 8 Pattern, Four Nodes . . . 145
Part 5: Leeds' Independent Ceph Cluster — Three Nodes, the Honest Floor . . . 146
Part 6: Bringing Up the London Zone — the Wizard, Now That We Understand It . . . 146
Part 7: Bringing Up the Leeds Zone — Building a Zone With No Front Door . . . 146
Part 8: The Region-Level VPC — Sharing a Design Without Sharing a Failure Domain . . . 146
Part 9: Verification Across Three Zones — Build-Time, Not the Monitoring Chapter . . . 147
Troubleshooting the Multi-Zone Build . . . 147
Chapter 14 Summary: Five Things to Hold On To . . . 147
The Values Card for Chapter 15 . . . 147

Chapter 15: Monitoring & Operations . . . 148
Observability at Scale — One View of Three Zones . . . 148
The Story: Seven Dashboards Are Not a Dashboard . . . 148
Part 1: The Observability Problem at Estate Scale . . . 148
Part 2: The Estate Observability Hub . . . 149
Part 3: Instrumenting the Estate . . . 149
Part 4: Dashboards That Answer "How Is the Estate?" . . . 150
Part 5: Logs — Shipping the Estate's Diary Somewhere Durable . . . 150
Part 6: Alerting Without a SaaS Pager . . . 151
Part 7: Incident Response and On-Call — the Runbooks Owe . . . 151
Part 8: Capacity Signals for Operations . . . 152
Part 9: Lab Exercise — One View of the Lab . . . 152
Five Key Takeaways . . . 152

Chapter 16: Backup & Disaster Recovery . . . 153
An Independent Copy That Does Not Share Fate . . . 153

The Story: The Night the Mirror Lied . . . 153
Part 1: Three Words People Use Interchangeably and Shouldn't . . . 153
Part 2: Protection Tiers, and Why They Are Really a Storage-Media Decision . . . 153
Part 3: The Independent-Fate Backup Cluster . . . 154
Part 4: CloudStack Native Backup & Recovery – the Shipping KVM Workhorse . . . 154
Part 5: Cross-Site DR with rbd-mirror (Tier 1) . . . 155
Part 6: The Immutable Archive – RGW S3 Object Lock (Tier 3) . . . 155
Part 7: Database Disaster Recovery – Discharging the Chapter 5 Promise . . . 155
Part 8: What's Coming in 4.23 – Native KVM Backups via Veeam . . . 155
Part 9: Restore Drills, Runbooks, and the Mars Test . . . 156
Five Key Takeaways . . . 156

Chapter 17: Cost Optimisation & Capacity Planning . . . 158
Owning the Meter: What a Sovereign Cloud Actually Costs . . . 158
The Story: Does Owning This Thing Actually Pay? . . . 158
Part 1: You Cannot Manage What You Cannot Meter . . . 158
Part 2: What the Estate Actually Costs . . . 159
Part 3: The Escape That Pays . . . 159
Part 4: When the Migration Pays for Itself . . . 160
Part 5: Cost per Workload – From Architecture to Accountability . . . 160
Part 6: Return on Research – Finding the Idle Silicon . . . 161
Part 7: Seeing the Wall Before You Hit It . . . 162
Five Key Takeaways . . . 162

Chapter 18: Your Cloud Journey Starts Now . . . 164
The Estate Is Built – Reflection, Consolidation, and the Road into Part 2 . . . 164
The Story: After the Renewal Letter . . . 164
The Journey You've Taken . . . 164
What Sovereignty Actually Cost . . . 164
The Thesis: Not Only for the Few . . . 164
The Mars Test, One Last Time . . . 164
Where You Go Next: Part 2 . . . 164
Key Takeaways . . . 164
Before You Close Part 1 . . . 165
What Next, and Staying in Touch . . . 165

Appendix A: Glossary . . . 166

Appendix B – The Portable CloudStack Lab . . . 167
Running a Production-Grade Lab on a Laptop . . . 167
Introduction . . . 167
Hardware Requirements . . . 167
The Lab, Layer by Layer . . . 168
What's Different on a Laptop . . . 168
The Reader Lab, As It Was Actually Built . . . 168
Connecting Real Hardware . . . 172
Performance Notes (Laptop-Specific) . . . 172
Backup and Recovery . . . 172
Use Cases . . . 173

Conclusion 173

Afterword – The Estate Is Yours Now 174
About the author 174
Stay in touch 174

Sovereign Cloud

Apache CloudStack — A Production Deployment Guide

Leaf Spine Books edition

* * *

Copyright © 2026 Michael John Leslie Hinsley. All rights reserved. Published under the Leaf Spine Books imprint.

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means — electronic, mechanical, photocopying, recording, or otherwise — without the prior written permission of the copyright holder, except for brief quotations embodied in reviews and other non-commercial uses permitted by copyright law.

Use of code and examples. This book exists to help you do your job. The code, configuration, and command examples are provided for you to use and adapt in your own deployments — you do not need to ask permission to put them to work, and attribution, while appreciated, is not required. What does require permission is reproducing a substantial portion of the *text*, or redistributing the book itself. The method is yours to take; the words are not.

Disclaimer of warranty. This book is provided on an "as is" basis. While every effort has been made to ensure the accuracy of the information and the correctness of the procedures, the author and publisher make no warranty, express or implied, as to its completeness, accuracy, or fitness for any particular purpose.

A working book on real infrastructure. The procedures here operate on live systems. Many of the commands are powerful and some are destructive: they can erase data, interrupt services, and render systems unbootable. Always rehearse in an isolated lab or a non-production environment first, keep tested backups, and understand each step before you run it. You assume all risk for any action taken on the basis of this book. To the fullest extent permitted by law, the author and publisher accept no liability for any loss, damage, data loss, downtime, or cost arising directly or indirectly from the use of, or reliance on, the material in these pages.

A method, not a product — and not a guarantee. The architectures, costings, and forecasts in this book are illustrative and grounded in one specific scenario. Where a figure is a forecast or an illustration, it is labelled as such. Your hardware, scale, contracts, and operating conditions will differ. Treat every number as a starting point to be re-measured against your own estate, never as a promise of outcome.

A note on the story. To keep the engineering concrete, this book follows a fictional organisation — a publisher escaping a punitive licensing renewal — together with a fictional multi-city estate and its tenants and customers. These are narrative devices. Any resemblance to real organisations, businesses, or events is coincidental. (The Leaf Spine Books imprint under which this book is published is real; the in-story publisher that shares its name is part of the narrative.)

Trademarks. Apache, Apache CloudStack, CloudStack, and the Apache feather logo are trademarks or registered trademarks of The Apache Software Foundation. All other product names, logos,

and brands are the property of their respective owners; names referenced in this book include, among others, VMware and Broadcom, Dell Technologies, Ubuntu, MariaDB, Ceph, Kubernetes, SONiC, and OpenTofu. Use of these names is for identification and reference only and does not imply endorsement.

No affiliation. This is an independent work. It is not affiliated with, authorised by, sponsored by, or endorsed by The Apache Software Foundation or by any other company or project named within it.

Open-source and third-party material. The software deployed and discussed in this book is published by its respective projects under their own open-source licences — Apache CloudStack, for example, is licensed under the Apache License 2.0. Where this book reproduces upstream third-party documentation for reference, that material remains under its original licence and copyright. Always consult each project's own licence and documentation for authoritative terms.

Code License: All inline code snippets, scripts, OpenTofu/Terraform configurations, and architectural manifests provided in this book are licensed under the Apache License, Version 2.0. You may use, modify, and distribute the code within your own environments and products.

Disclaimer of Liability: The architectures, configurations, and scripts in this book are provided "as is" and without warranty of any kind. While every effort has been made to ensure the accuracy of the technical instructions against Apache CloudStack 4.22 LTS and Ubuntu 24.04 LTS, infrastructure deployments carry inherent risks. The author and publisher shall not be held liable for any data loss, downtime, or damages arising from the use of the information contained in this book. Test all configurations in a dedicated lab environment before applying them to production systems.

A Note on Fiction and Technical Truth This book uses a narrative framework to explain complex engineering decisions. 'Leaf Spine Books', 'Brackenford Books', and the engineering team—Sarah, Marcus, Priya, James, Elena, and Tom—are entirely fictional. They are characters in a story, designed to represent the real pressures, late-night emergencies, and board-level demands that infrastructure teams face every day.

However, the technical infrastructure is entirely real. Every command, IP address, architectural diagram, and deployment phase detailed in this book has been tested on real hardware. The methodology, the code, the BGP EVPN fabric, the Ceph storage configurations, and the CloudStack control plane are not theoretical. They are a literal, reproducible blueprint designed to help you build your own sovereign ecosystem.

Who This Book Is For A platform is only as good as the people who understand it. This guide is written to bridge the gap between the server room and the boardroom, serving four distinct readers:

1. The Solo Adventurers and Homelab Users For the hobbyists, students, and self-taught learners on a budget. This book proves that enterprise-grade infrastructure learning does not require an enterprise budget. Through the "Manchester Miniature" and the portable laptop lab (Appendix B), you will learn real cloud skills, on real open-source software, for the price of an evening's electricity. It is the purest expression of the GYOCYO (Get Your Own Cloud, You Own) ethos.
2. The Engineers and Builders For the practitioners who want absolute control and career capital. You care about how things are done, reproducibility, the troubleshooting method, and owning the whole stack. This book promises you a method, not a product. It teaches you how to build a system that passes the "Mars Test"—infrastructure that survives, heals, and alerts without a human hand holding it together.
3. Smaller IT Shops and Enterprise Operators For the infrastructure leaders running real estates and defending real budgets. If you are staring down a hypervisor licensing renewal that looks

like an extinction event, this book is your escape tunnel. It offers defensible economics, an exit from vendor lock-in, and the operational rigour required to productise your internal IT.

4. COOs, CTOs, CEOs, and Thought Leaders For the strategic minds at hyperscalers and enterprise organisations. You care about digital sovereignty as a posture, lock-out insurance, and the macro narrative of edge computing, sovereign AI, and energy efficiency. While the engineers read this book for the wiring, you should read it for the thesis: If you cannot control your infrastructure, you do not own your business. Welcome to the CloudStack community. Your journey begins now.

* * *

First edition — 2026. Covers Apache CloudStack 4.22.1.0 LTS.

Leaf Spine Books edition Published in the United Kingdom.

* * *

On owning what you run

A word on the spirit of this book. Its philosophy is **GYOCYO** — *Grow Your Own, Cook Your Own*: the conviction that if you cannot control your infrastructure, you do not truly own your business. The notices above protect the words on these pages. The method inside them is meant to be taken, adapted, and run on hardware you own — built, broken in a lab, and made yours. That is the whole point.

Apache CloudStack Deployment Guide — Leaf Spine Books edition Copyright © 2026 Michael John Leslie Hinsley. Code examples licensed under Apache License 2.0. Contact details for this book are on the contact page at the back.

The Book in Full — Annotated Contents

Apache CloudStack 4.22.1.0 LTS — A Production Deployment Guide · Leaf Spine Books edition

This is the whole journey on a single page. Every entry below gives a chapter's exact title and a two-line summary of what it builds, so you can see the full scope of the book before you start — and find the chapter you need when you come back to it. The story runs in order, each chapter standing on the one beneath it, from the licensing letter that started everything to a cloud platform you own down to the last packet.

* * *

Volume 1 — Foundations (Chapters 1-18)

Chapter 1 — The Cost of Renting Your Existence

Meet Leaf Spine Books — Understanding CloudStack Architecture Through Real-World Needs The VMware/Broadcom licensing shock — a 567% renewal — that drives fictional publisher Leaf Spine Books to escape to a sovereign Apache CloudStack platform. Introduces the core architecture (Region ▫ Zone ▫ Pod ▫ Cluster ▫ Host, the four network types, the system VMs, and the KVM-plus-Ceph foundations) through one company's real requirements.

Chapter 2 — Designing the Three-City UK Cloud

From Whiteboard to Blueprint — Planning a Production-Grade Infrastructure Turns the premise into a complete blueprint: a three-city estate (Manchester, London, Leeds) with full hardware, VLAN and IP plans, an OS10 leaf-spine fabric, and GPU capacity from day one. Sets out the economics — roughly £506k of owned platform against £181.5k a year of licensing, a 94% five-year saving.

Chapter 3 — Building Your Learning Lab

From One Second-Hand Server to a Cloud You Can Break Builds the entire design in miniature on a single second-hand server, so every mistake costs an evening rather than a career. Nested KVM, a real three-node cephadm Ceph cluster, and Dell OS10 network simulation rehearse the production platform before any real money is spent.

Chapter 4 — Infrastructure Preparation

Building the Foundation — From Blueprint to Bare Metal Turns racked metal into hardened, repeatable hosts: a standard Ubuntu 24.04 LTS build, the three-bond network layout, correct storage preparation per node class, and security from first boot. Build one host by hand to understand it, then let Ansible build the rest.

Chapter 5 — MySQL High Availability

The Database Layer — One Cluster, Two Cities, No Single Point of Failure Builds the control plane's database as a three-node MariaDB Galera cluster spanning Manchester and London, co-located on the management servers. Single-writer HAProxy routing, honest quorum arithmetic, a drilled failover, and backups for the failures clustering cannot prevent.

Chapter 6 — CloudStack Management Server

The Brain Arrives — Installing the Control Plane on the Nodes That Already Hold Its Memory Installs the CloudStack management servers on the same three nodes and initialises the schema through the routing layer — without root credentials ever leaving the database tier. Three live, clustered control planes, proven with a kill signal under a working login.

Chapter 7 — Secondary Storage

The Repository That Rebuilds the Cloud — DRBD and NFS on Refurbished Metal Builds highly available NFS secondary storage on a refurbished R740XD pair: DRBD synchronous replication, a floating service address, and failover tested all the way to a pulled power cord. The repository that reseeds the cloud, kept deliberately out of the storage cluster's blast radius.

Chapter 8 — Primary Storage

Where the VMs Actually Live — Ceph RBD on the NVMe Estate Builds primary storage as a six-node Ceph RBD cluster on the deliberately untouched NVMe drives, deployed and self-healed by cephadm. A replicated pool with a dedicated, least-privilege CloudStack credential, drilled against drive and whole-node failure under live load.

Chapter 9 — Deploying CloudStack — The Zone Goes Live

The logical build: zone, physical network and traffic labels, pod, guest VLAN range, cluster, and the registration of both storage tiers. The Manchester zone goes live and the system VM template seeds itself with no internet download — the stage is lit, the crew arrives next.

Chapter 10 — Adding KVM Compute Hosts — The Day the Stage Crew Arrives

Turns eight prepared machines into hypervisors across two clusters (Intel production, AMD GPU), honouring the traffic labels and registering both Ceph pools. The first system VMs boot, the console works, GPU passthrough is configured, and the cloud becomes operational.

Chapter 11 — Deploying Your First Guest VMs — Opening Night

The Morning the Cloud Earned Its Keep Turns an operational-but-empty cloud into its first real workload: compute and disk offerings, a guest template the secondary storage VM fetches and checksums itself, and Editorial's first isolated network drawing a VLAN on demand. Deploys three guests behind a virtual router with egress denied by default, proves GPU work lands on GPU hardware by placement rather than tags, and retires the old habit of using security groups where they do nothing.

Chapter 12 — Advanced Networking — The OS10 Leaf-Spine Fabric

Turning "Trust the Trunk" Into a Tested Data Plane Opens the switch fabric the previous chapters asked you to trust and builds it for real in OS10: a leaf-spine underlay on eBGP unnumbered, the VLANs and host trunk, and the management gateway as an SVI that finally answers. Discharges every "trust the trunk" promissory note with commands you can re-run, and counts the cost of an open-networking fabric against the licensed alternative.

Chapter 13 — Advanced Networking with VPCs — Walls Within the Walls

Building Multi-Tier Isolation on the Fabric You Just Proved Divides Editorial's flat network into a VPC with several tiers — web, app, database, management — routed through one VPC router and policed by stateless network ACLs written in both directions. Captures the whole structure as OpenTofu so the walls are reviewable, and shows how the same pattern becomes a tenancy you could one day hand to someone else.

Chapter 14 — Multi-Zone Expansion

The Cloud Grows a Map of the Country — Without Stretching a Single Fiction Extends the single Manchester site into three independent zones — London public, Leeds internal-only — each with its own Ceph cluster and leaf pair, under one control plane. Holds the line that multi-zone is capacity in more than one place, not disaster recovery, and answers the region-level VPC question as a shared design rather than a stretched object.

Chapter 15 — Monitoring & Operations

Observability at Scale — One View of Three Zones Folds the estate's scattered monitoring — three Ceph clusters' bundled stacks, the hosts, the fabric, the control plane — into one federated view of metrics, logs and alerts, reaching the internal-only Leeds zone over the sovereign WAN. Builds sovereign alerting and on-call runbooks so a fault pages a real person through infrastructure you own, with nothing leaving the estate.

Chapter 16 — Backup & Disaster Recovery

An Independent Copy That Does Not Share Fate Draws the hard line between replication, backup and disaster recovery, and builds an independent-fate backup cluster that production can write to but never overwrite. Adds cross-site RBD mirroring, an immutable S3 Object Lock archive, point-in-time database recovery, and the restore drills that turn a backup from a hope into a guarantee.

Chapter 17 — Cost Optimisation & Capacity Planning

Owning the Meter: What a Sovereign Cloud Actually Costs Installs the usage meter deferred since the management plane went up and turns raw consumption into a defensible per-workload cost. Builds an honest profit-and-loss against the VMware bill the estate escaped, finds the idle silicon worth reclaiming, and forecasts the next purchase in pounds and lead time, not just percentages.

Chapter 18 — Your Cloud Journey Starts Now

The Estate Is Built — Reflection, Consolidation, and the Road into Part 2 Closes Part 1 by stepping back from the build to take stock of the whole estate – costed, metered, monitored, backed-up and multi-zone, owned outright. Reflects on what sovereignty actually cost, why owning the stack changes the questions you get to ask, and points at the road into Part 2 without pushing you down it.

VXLAN and BGP EVPN on OS10 switches and managed through CloudStack Edge Zones. Treats energy as a first-class engineering constraint and closes by stepping out of the fiction to ground the idea in the author's own real research.

Appendices

Appendix A — Glossary

A *working engineer's glossary* The technical and coined vocabulary used across the book – from BGP EVPN, Ceph, KVM and the CloudStack zone/pod/cluster hierarchy to the book's own coinages: GYOCYO, the Mars test and Return on Research. Definitions are short, practical and grounded in how each term is used here, with the canonical project name given wherever a term names a real component.

Appendix B — The Portable CloudStack Lab

Running Production-Grade Infrastructure Labs on a Laptop How to build a complete CloudStack learning lab on a high-end laptop with external USB storage – network simulation, a management server, and nested KVM hosts and guests running several layers deep – at a fraction of a server lab's cost and power draw. Aimed at consultants needing a portable demo, students without a server room, and engineers minimising their lab's energy footprint.

Apache CloudStack Deployment Guide — Leaf Spine Books edition Copyright © 2026 Michael John Leslie Hinsley. Code examples licensed under Apache License 2.0. Contact details for this book are on the contact page at the back.

How This Book Is Proven

Most infrastructure books ask for your trust. The author assures you the commands were tested, the architecture is sound, the version numbers are current, and you have no way to check any of it until you are three hours into a build with a cluster that will not converge. This chapter exists because we decided to do it differently — and because the way we did it is, we think, a story worth a few pages of your time before the real work starts.

The short version: before this edition reached you, the test lab this book teaches was built, broken, repaired and acceptance-tested for real — management server, Ceph storage, two KVM compute hosts, five networks, NFS secondary storage and test guest VMs — from the book's own chapters, by a team you should know about. Where the build disagreed with the manuscript, the manuscript changed. What you are holding is the corrected result.

The lab behind the book

Leaf Spine Books runs an infrastructure research lab called MatrixClaw. It is a strange and rather wonderful thing: the engineers are AI — Claude, Anthropic's model, working through a set of defined roles — and the governance is ferociously human. Nothing mutates a host without a human approval, and that approval is bound to the exact plan it approved. Change the plan and the approval dies with it. Every operation writes a run record — a timestamped, hash-bound account of what was planned, who approved it, what actually executed and what the tests observed. The lab's discipline is captured in one rule we never break: **no run, no lesson**. If it did not execute, it does not get taught.

The team works the way a real platform team works, because it is one. An infrastructure engineer authors topologies and runbooks as code. A site reliability engineer writes the checks that decide whether a build is actually done. An independent reviewer picks fights with both of them. A curriculum author turns run records into teaching material. A compliance officer audits the lot and publishes the evidence library. Agents and humans share the same communication bus — GitLab issues and merge requests — and every merge, without exception, is performed by a human maintainer.

What actually happened

The sequence matters, so here it is in order.

First, the lab read this book. The seven-file Chapter 3 lab set was extracted from the manuscript and turned into executable runbooks and inventory — the same addresses, the same names, the same build order you will follow.

Second, the lab reviewed this book. Four independent verdicts were recorded before a single host was touched. The review found real defects — a cross-reference to a configuration step that did not exist, a copy-paste trap that would have given two hosts the same IP address, a command-line tool cited but never installed. Each finding was ratified by the maintainer and became a tracked erratum. You will not meet those traps, because this edition is the one they were fixed in.

Third, the lab built the estate. On a single laptop-class machine with nested virtualisation, the full seven-VM topology came up as gated, human-approved runbook runs: the platform host, the five networks with a jumbo-frame storage plane, three Ceph monitor-and-OSD nodes bootstrapped with cephadm, the NFS secondary store, the management server, and both compute hosts. The CloudStack zone itself — Manchester, the same zone this book designs — was then built as code and enabled, the system VMs came up, the built-in template seeded itself, and a first guest VM was deployed with its root volume verified, by database join, to be living on the production Ceph pool.

Fourth, the lab broke things — or rather, things broke, and the breakages went on the record. Ten numbered findings came out of the build. A package probe that lied about installation state. A client timeout that looked like a server failure and would have double-added a host to anyone who blindly retried. A Ceph health check, new in the Squid release, that condemned a perfectly serving cluster because of how the monitors' network declaration was written. None of these were smoothed over. Each was diagnosed on the record, fixed under approval, merged by the maintainer — and then written back into the chapters ahead of you, in the troubleshooting sections where you will want them at three in the morning.

Fifth, the lab proved it. The finish line for this book's lab is a formal acceptance suite — seventeen checkpoints, executed as twenty-four read-only checks covering everything from nested virtualisation on the host down to the snapshot floor that lets you experiment without fear. On 21 July 2026, on the 32 GB profile — the smallest build this book supports, with both compute hosts kept — the suite passed twenty-four out of twenty-four.

Why you can check any of this

Everything above is a recorded fact, not a marketing claim. The lab's repository holds forty-nine run records for this programme, thirty ratified decisions, and a lessons ledger of seventeen entries learned the hard way — every one of them traceable from claim to evidence. The acceptance suite's checks are named and reproducible; the same checks are taught to you, as commands you can run, in the Done-When sections of the lab chapters.

You do not have to take any of that on trust, because the build itself is in this book. **Appendix B** carries the laptop build exactly as the lab executed it: every phase in the order it ran, the commands as the runbooks issued them, and the checks that proved each phase before the next one started. After the clean path, set apart so it does not interrupt you, come the failures — with the diagnosis that resolved each one. The platform host took five attempts before it passed, and all four of those failures are in there, because a build that went right the first time would have taught you less than one that did not.

One figure in that appendix will not match the one above, and the difference is deliberate rather than careless. The seventy-four commands counted here are every step the nine runbooks executed across the whole programme; the appendix covers the twenty-three run records of the laptop build itself, and counts sixty-five. Different questions, different answers, drawn from the same records. Where this book gives you a number, it tells you what the number is counting — and if the two had been left to sit there contradicting each other, you would have been right to distrust both.

We are honest about the boundaries, too. The ARM build path in Chapter 3 has not been executed by the lab and says so. Live migration is excluded from the 32 GB profile by design rather than quietly skipped — the lab's motto for that trade is printed where you will make it: 32 GB demonstrates the platform; 64 GB teaches it. Where a figure could not be verified, it is labelled, not laundered.

And now you can run the proof yourself

There is a companion to this book you can pick up and use directly. **MatrixClaw Nano** is a small tool — separate from the research framework above, written fresh and published openly under the AGPL-3.0 licence at `https://github.com/leafspinebooks/matrixclaw-nano` — that runs the laptop lab build for you, one gated command at a time, and keeps the same kind of audit record the lab keeps. It is not MatrixClaw itself; it is a much smaller thing that does one job. It builds the reader lab and proves it, and it says plainly where its knowledge ends.

We put it to the hardest test we could devise. On 24 July 2026 the published tool was cloned onto a fresh laptop — a clean install of Ubuntu, on hardware that had played no part in writing it — and the entire lab was built following nothing but the tool's own reader guide: no inside knowledge, no access to the machine it was developed on. The estate came up, and the same twenty-four-check acceptance suite passed twenty-four out of twenty-four. Every gap that clean-room run exposed on that fresh laptop — and it exposed a fair few — was fixed before this edition reached you.

That is the difference between a proof you are shown and a proof you can repeat. Appendix B is the record of what the lab did; the published tool is the means to do it yourself and check the result against the same suite, on your own hardware, in your own hands.

What this means for you

It means the estate was built by execution rather than assertion — nine governed runbooks, seventy-four commands, each approved against its plan before it ran and recorded afterwards with its output — by something with no ego about being wrong and a governance system that refuses to let it pretend otherwise. It means the troubleshooting sections describe faults that actually happened, with the diagnosis discipline that actually resolved them. And it means that when this book tells you a thing is done — a cluster healthy, a zone enabled, a guest genuinely on the storage you built for it — it is because a check with a name proved it, on hardware, before the words were allowed to stand.

Let us be precise about one thing, because the distinction matters and you have no reason to take it on trust. What the lab executed were governed runbooks derived from these chapters — it did not paste the chapters' own text line by line. The runbooks are the executed artefact; these chapters are the teaching text they were built from. Where the two disagreed, it was the chapters that changed.

* * *

Apache CloudStack Deployment Guide — Leaf Spine Books edition Copyright © 2026 Michael John Leslie Hinsley. Code examples licensed under Apache License 2.0. Contact details for this book are on the contact page at the back.

Chapter 1: The Cost of Renting Your Existence

Meet Leaf Spine Books — Understanding CloudStack Architecture Through Real-World Needs

I have spent thirty-two years in IT, bridging the gap between delicate electronics and the mud and rain of the physical world. Whether it is a rack of Dell servers in a climate-controlled datacentre or a cluster of Raspberry Pis monitoring a beehive in a storm, the principle remains the same: **If you cannot control your infrastructure, you do not own your business.**

For the last decade, the industry fell into a comfortable trap. We outsourced our thinking to VMware. We treated hypervisors like electricity—something you just paid a utility bill for. It worked, until the landlord changed the locks and tripled the rent.

When Broadcom acquired VMware, the bill for many organisations didn't just go up; it became an extinction event.

This book is not just a technical manual. It is an escape tunnel. It is the documentation of how we take a fictional traditional publisher, **Leaf Spine Books**, and migrate them from a strangling licensing model to a sovereign, open-source cloud built on **Apache CloudStack**.

Meet Leaf Spine Books

Leaf Spine Books is a company that deals in heavy, physical things—pallets of paper, warehouses in Manchester, Leeds, and London, and logistics networks that span the UK. They are not a "born in the cloud" unicorn. They have legacy debt. They have servers that are too old, networks that are too slow, and a budget that just got decimated by the new VMware core-pricing model.

Their situation mirrors the reality many of you face today:

1. **The "Tax"**: Their virtualisation licensing costs have jumped 567%.
2. **The Waste**: They have hardware sitting idle because they can't afford to license the cores.
3. **The Fear**: They are realising that their entire digital existence is rented.

The Engineering Philosophy

In this guide, we are not just going to "install software." We are going to build an ecosystem. We will use the same principles I use when designing autonomous systems for aquaponics or apiaries.

When I put a server inside a beehive, I cannot send a technician to reboot it if it hangs. I have to design it as if it were on Mars. That is the mindset we will apply to the Leaf Spine Books cloud.

We will build on **CloudStack 4.22 LTS** (4.22.1.0 at the time of writing), not because it is free, but because it is battle-hardened. We will use **Dell open-networking switches running Dell SmartFabric OS10** because we want a fabric with a real automation surface and a vendor behind it—an Ansible collection, gNMI/OpenConfig telemetry, and a familiar IOS-style CLI, all on the same open-networking metal. The trade, made honestly, is a network-OS licence and a Dell support

relationship in place of community SONiC's read-the-source openness; the community SONiC build is the canonical edition's choice, a fork away on `main` if that openness matters more to you. We will use **Ceph storage** because data gravity is real, and we need one platform that speaks block for our VMs, file where a filesystem is genuinely needed, and modern S3 for everything else.

We are going to take the hardware you already have—or efficient ARM-based silicon like the Ampere Altra—and we are going to squeeze every ounce of "Return on Research" out of it.

The goal is simple: By the end of this book, you will not just have a cloud. You will have a sovereign platform that you own, down to the last packet.

Let's get to work.

The Story: The Renewal Letter

Sarah Martinez leaned back in her chair, studying the whiteboard covered in network diagrams and question marks. As CTO of Leaf Spine Books, she'd just finished yet another emergency meeting—but this one wasn't about an outage. It was about a letter.

The VMware renewal quote had landed three weeks earlier. The old perpetual licences were gone; in their place, a per-core subscription across every socket in the estate, at a minimum core count per CPU whether they used the cores or not. The finance director had read the number out twice because nobody believed it the first time: a 567% increase on what they'd paid the year before.

Leaf Spine Books wasn't a startup. What had begun decades ago as a regional printer had grown into a publisher with 200 employees across three UK cities: Manchester (headquarters), London, and Leeds. Pallets of paper still moved through their warehouses every day—but so did digital editions, a print-on-demand pipeline, and a growing platform that let independent publishers sell their own titles. The physical business paid the bills; the digital business was the future. And the digital business had just been taken hostage.

"We can't pay this, and we shouldn't," Sarah told her team. "But I'm not trading one landlord for another. We're not putting the business at the mercy of AWS pricing or lock-in either. We own our buildings, we own our presses, and from now on we own our infrastructure. And we absolutely must keep our data sovereign within the UK."

The requirements were clear:

Corporate Needs:

- Email and collaboration tools for 200 employees
- Development and testing environments
- Internal applications and databases (catalogue, royalties, logistics)
- File storage and backup
- Security compliance for customer data (UK GDPR, ICO requirements)

Platform Requirements (the digital arm):

- Host ebook stores for independent publisher clients
- Handle traffic spikes during book launches
- 99.95% availability target for paying customers
- PCI-DSS compliance for payment processing
- Scalable storage for millions of digital editions

- CDN integration for fast delivery

Strategic Goals:

- Own and control infrastructure (no vendor lock-in—never again)
- Leverage existing UK colocation contracts
- Build in true high availability
- Maintain data sovereignty within the UK
- Future-proof for IoT and edge computing needs

Innovation Vision:

- Enable IoT data processing for client projects
- Support edge computing for retail locations
- Provide GPU-accelerated services for AI/ML workloads
- Build a platform for the next generation of publishing services

After months of evaluation, Sarah's team recommended **Apache CloudStack 4.22 LTS** with KVM hypervisors and Ceph-based storage. It offered enterprise-grade cloud management without proprietary lock-in, and her team could deploy it across their three UK datacentres with complete data sovereignty.

But where to begin?

Understanding CloudStack: Building Blocks of Your Cloud

Before Leaf Spine could deploy anything, Sarah's team needed to understand what they were actually building. Let's walk through the concepts together.

What is CloudStack?

Apache CloudStack is an open-source Infrastructure-as-a-Service (IaaS) platform that transforms your servers, storage, and network into a private or public cloud. Think of it as the orchestration layer that sits on top of your hardware, automating:

- Virtual machine creation and lifecycle management
- Network provisioning and isolation (VLANs, VPCs, VXLAN)
- Storage allocation and management (snapshots, volumes, templates)
- User access, quotas, and multi-tenancy
- Resource monitoring, usage tracking, and billing
- High availability and disaster recovery

Key Principle: CloudStack doesn't replace your hypervisor or storage; it orchestrates them. You still run KVM (or another supported hypervisor) on your hosts. You still use NFS, iSCSI, Ceph RBD, or other storage protocols. CloudStack makes it all work together as a unified, manageable cloud platform.

The CloudStack Architecture Hierarchy

CloudStack organises infrastructure in a hierarchical structure. Understanding this is absolutely crucial:

```
Region
  +-- Zone (equivalent to a datacentre/availability zone)
       +-- Physical Networks
       |    +-- Management Network
       |    +-- Public Network
       |    +-- Guest Network
       |    +-- Storage Network
       |
       +-- Pod (Layer-2 network segment, like a rack or row)
            +-- Cluster (homogeneous hosts with identical hardware)
                 +-- Hosts (physical hypervisor servers)
                 |    +-- Virtual Machines (your workloads)
                 |
                 +-- Primary Storage (VM disks - low latency)

       +-- Secondary Storage (templates, snapshots, ISOs)
```

Let's understand each level in detail:

Regions (Optional, for Geographic Distribution)

- Largest organisational unit in CloudStack
- Contains multiple geographically distributed zones
- Each region has its own management server deployment
- Enables true disaster recovery and geographic distribution
- Users can deploy resources across regions via a single interface
- Supports multi-region replication for templates

Leaf Spine's Decision: Start with one region covering the UK. If they expand to Europe or globally, add additional regions later with full data sovereignty controls per region.

Zones (Datacentres / Availability Zones)

- Roughly equivalent to a single datacentre or availability zone
- Largest failure domain (zone failure doesn't affect other zones)
- All hosts in a zone typically share the same networking infrastructure
- Each zone requires its own secondary storage
- Can have one or more zones per region
- Enables geographic redundancy and disaster recovery

Leaf Spine's Deployment Plan:

- **Zone 1: Manchester** (Primary zone, largest deployment, HQ location)
- **Zone 2: London** (Secondary zone, disaster recovery, southern presence)
- **Zone 3: Leeds** (Development, testing, and northern redundancy)

Key Design Consideration: Each zone is completely independent. If Manchester's entire datacentre goes offline, London and Leeds continue operating.

Pods (Layer-2 Network Segments)

- Layer-2 broadcast domain within a zone
- Typically represents a physical rack or group of racks
- All hosts in a pod share the same IP subnet for management traffic
- Usually one pod per zone for smaller deployments
- Larger deployments might have multiple pods for:
  - Physical separation (different racks/rows)
  - Network segmentation
  - Failure domain isolation

Leaf Spine's Pod Strategy:

- One pod per zone initially (simpler management)
- Will add additional pods as they scale beyond 16-20 hosts per zone
- Pod isolation provides additional failure domain protection

Clusters (Homogeneous Host Groups)

- Group of hosts with **identical** hardware configuration
- All hosts in a cluster must access the same shared primary storage
- Maximum recommended: 16-32 hosts per cluster (varies by hypervisor)
- Enables critical features:
  - Live migration of VMs between hosts
  - High availability (HA) for VMs
  - Load balancing of workloads
  - Maintenance mode operations

Leaf Spine's Initial Cluster Design (the full blueprint is Chapter 2's job; these are the headline decisions):

Manchester (Zone 1) - Production:

- **Cluster MAN-Prod-01:** 4x Dell R660 (general compute, web services, databases)
- **Cluster MAN-GPU-01:** 4x Dell R7525 (2 hosts GPU-fitted for AI/ML services, 2 awaiting cards)

London (Zone 2) - DR/Secondary:

- **Cluster LON-Prod-01:** 4x Dell R660 (replicated services, DR capacity)

Leeds (Zone 3) - Development:

- **Cluster LEE-Dev-01:** 2x Dell R740 (existing hardware, reused—development and testing)

That Leeds cluster matters more than its size suggests. Those R740s are the servers Leaf Spine already owned—the ones sitting idle because the new licensing model made their cores unaffordable to light up. Under KVM there is no per-core meter running. The "waste" goes back to work on day one.

Critical Cluster Requirements (Do NOT Ignore):

All hosts in a cluster MUST be identical:

- **Same CPU model and generation** (e.g., all Intel Xeon Gold 6430)
- **Same CPU instruction sets and feature flags** (AVX, AVX2, AVX-512)
- **Same hypervisor type and version** (e.g., all Ubuntu 24.04 LTS with KVM)
- **Access to the same primary storage pools** (all see the same Ceph RBD pools)
- **Same network configuration** (bonding, VLANs, MTU settings)

Why this matters: Live migration moves a running VM from one host to another with zero downtime. If CPU instructions differ, the VM will crash during migration. If storage access differs, migration will fail. This is also why Leeds' older R740s form their own cluster rather than joining a production one: mixing CPU generations inside a cluster is how you turn a maintenance window into an incident.

Hosts (Physical Hypervisor Servers)

- The actual physical servers running the hypervisor
- Run KVM (or another supported hypervisor)
- Execute virtual machine workloads
- Must be dedicated to CloudStack (don't run other services on hypervisor hosts)
- CloudStack agent installed for management communication
- Connected to all four network types (Management, Storage, Guest, Public)

Leaf Spine's Host Specifications:

Production Hosts (Dell PowerEdge R660):

```
CPU: 2x Intel Xeon Gold 6430 (32 cores each, 64 cores total)
RAM: 512GB DDR5 ECC
Boot: 2x 960GB M.2 NVMe (BOSS-N1 RAID-1)
Network: 6x 25GbE Intel E810-XXV
  - 2x Management (bonded)
  - 2x Storage (bonded, LACP, Ceph RBD)
  - 2x Guest/Public (bonded, LACP, VLAN trunk)
Power: Dual redundant PSUs (1400W each)
```

GPU/AI Hosts (Dell PowerEdge R7525 – refurbished):

```
CPU: 2x AMD EPYC 7543 (32 cores each, 64 cores total)
RAM: 512GB DDR4-3200 ECC
Boot: 2x 960GB M.2 NVMe (BOSS RAID-1)
Network: 6x 25GbE
GPU: 2x NVIDIA RTX 4060 Ti 16GB (on 2 of the 4 cluster hosts initially)
Power: Dual redundant PSUs (1400W each)
```

Development Hosts (Dell PowerEdge R740 - Existing):

```
CPU: 2x Intel Xeon Gold 6254 (18 cores each, 36 cores total)
RAM: 384GB DDR4 ECC
Boot: 2x 480GB M.2 SATA (BOSS RAID-1)
Network: 4x 25GbE Intel XXV710
  - 2x Management/Storage (bonded)
  - 2x Guest/Public (bonded)
```

A note on that GPU line, because it surprises people: a 16GB RTX 4060 Ti is not a "datacentre" card, and that is rather the point. It is what we actually run in the lab pipeline that builds this book—quietly serving local language models all day at a fraction of the cost and wattage of the cards the brochures want to sell you. Return on Research means starting with the silicon that earns its keep, and scaling up only when the workload proves it must.

Primary Storage (VM Disk Storage - Performance Critical)

- Stores virtual machine root disks and data volumes
- Must be accessible by all hosts in the cluster (or zone-wide)
- Attached at cluster-level or zone-level depending on design
- Performance critical: This is where VM I/O happens
- Low latency essential (ideally <5ms, target <2ms)

Supported Protocols:

- **NFS:** Simple, widely supported, good for smaller deployments
- **iSCSI:** Block storage, better performance than NFS
- **Ceph RBD:** Distributed, scalable, native CloudStack support
- **Local Storage:** Fastest, but no live migration or HA
- **Fibre Channel:** High performance, expensive
- **Dell PowerFlex, Pure Storage, NetApp** (enterprise SAN)

Leaf Spine's Choice: Ceph RBD (upstream Ceph, deployed with cephadm):

Why Ceph RBD?

- Native CloudStack support via librbd (no iSCSI overhead)
- Distributed architecture (no single point of failure)
- Excellent performance with NVMe drives
- Scales horizontally (add nodes for capacity and performance)
- Built-in replication (3x for production, 2x for dev)

- Can span multiple sites for disaster recovery
- Open-source (no licensing costs)

Why upstream Ceph rather than a storage appliance? There are appliance products that wrap Ceph in a friendlier box, and for years they were a sensible on-ramp. But every layer between you and the storage is a layer you don't control—and modern Ceph has removed the excuse. The **cephadm** orchestrator gives you a supported, containerised deployment with a built-in dashboard, health checks, and rolling upgrades, straight from the Ceph project itself. Same protocols (RBD, CephFS, S3 via RGW), no appliance licence, no vendor between you and the mailing list when something breaks at 3am. If the server is on Mars, you want to be debugging Ceph, not somebody's wrapper around Ceph.

Manchester Production Storage Cluster:

```
Nodes: 6x Dell R760 (dedicated storage nodes)
Each Node:
  CPU: 2x Intel Xeon Gold 6430 (32 cores each, 64 cores total)
  RAM: 256GB DDR5 ECC
  OS Drives: 2x 480GB M.2 NVMe (RAID-1)
  OSD Drives: 4x 3.84TB NVMe U.2 (Samsung PM9A3 class)
  Network: 4x 25GbE (storage frontend and backend)

Total OSDs: 24 (6 nodes x 4 drives)
Raw Capacity: 24 x 3.84TB = ~92TB
Usable (3x replication): ~30TB
Effective (after overhead): ~25TB

Performance (estimated, NVMe OSDs):
  Aggregate reads: 12-15GB/s
  Aggregate writes: 5-7GB/s (replication overhead)
  Random IOPS: 150,000+

Pools Configuration:
  rbd-prod (default): 3x replication, for production VMs
  rbd-dev: 2x replication, for dev/test (more capacity, less protection)
  ec-backup: Erasure coding 4+2, for backup data (better capacity efficiency)

Network Configuration:
  Frontend (client): MTU 9000
  Backend (replication): MTU 9000, separate VLAN
  Separation prevents replication from impacting client I/O
```

A separate four-node backup cluster lives in Manchester too—deliberately independent of production storage, because a backup that shares a failure domain with the thing it protects is a comfort blanket, not a backup. Chapter 2 covers the full storage layout; Chapters 7, 8 and 16 build it.

Secondary Storage (Templates, Snapshots, ISOs)

- Stores VM templates (golden images for deployment)
- Stores ISO files for OS installation media

- Stores VM snapshots for backup and recovery
- Zone-wide (shared across the zone)
- Uses NFS protocol or S3-compatible object storage
- Does NOT need ultra-high performance (templates aren't accessed frequently)

Leaf Spine's Secondary Storage Design:

High-Availability NFS Servers (per zone):

```
Servers: 2x Dell R740XD, refurbished (HA pair with DRBD)
Storage: NFS export on high-capacity drives (12-18x 3.5" bays per chassis)
Protocol: NFSv4.1 with Kerberos security
Backup: Daily replication to object storage
Redundancy: DRBD active/passive, Pacemaker clustering
Compact alternative: 2x R650 with RAID-10 SSDs where space is tighter than budget
```

The Modern Alternative (and where this book ends up):

- **Ceph RGW** – the RADOS Gateway already included in the Ceph cluster – exposes an S3-compatible endpoint
- CloudStack supports S3-compatible object storage for secondary storage
- One storage platform instead of two; multi-site replication built in
- Better for very large template libraries
- The same RGW endpoint later becomes the backup target with S3 Object Lock for immutability (Chapter 16)

System Virtual Machines: CloudStack's Control Layer

One of CloudStack's most elegant architectural decisions is using VMs to provide cloud services. These "System VMs" run on your hypervisors but are fully managed by CloudStack. Understanding them is critical.

Secondary Storage VM (SSVM)

Purpose: Manages all template, ISO, and snapshot operations

What it does:

- Copies templates from secondary storage to primary storage when deploying VMs
- Uploads snapshots from primary to secondary storage
- Handles template registration and downloads
- Manages ISO uploads and distribution
- Runs checksums and validation

Deployment:

- One SSVM per zone (automatically deployed by CloudStack)

- Runs a lightweight Debian-based OS
- Minimal resources (2 vCPU, 2GB RAM)
- Connects to both secondary storage and the management network
- Automatically recreated if it fails

Important: Never manually stop or delete SSVMs. CloudStack manages their lifecycle.

Console Proxy VM (CPVM)

Purpose: Provides secure console access to VMs through the web UI

What it does:

- Proxies VNC/RDP connections from users to VMs
- Enables console access without opening firewall ports
- Handles authentication and encryption
- Provides an HTML5-based console in the web browser
- Supports multiple concurrent console sessions

Deployment:

- One CPVM per zone (more for high user load)
- Runs a lightweight Debian-based OS
- Minimal resources (2 vCPU, 2GB RAM, can increase for scale)
- **Performance consideration:** For 100+ concurrent users, deploy 2-3 CPVMs per zone

Scaling consideration: If console access is slow or users see "Console not available," add more CPVMs.

Virtual Router VM

Purpose: Provides networking services for guest networks

What it does:

- **DHCP:** Assigns IP addresses to guest VMs
- **DNS:** Provides DNS resolution (forwards to external DNS)
- **Firewall:** Enforces security group rules and network ACLs
- **NAT:** Source NAT for outbound internet access
- **Port Forwarding:** Maps public IPs to private VM IPs
- **Static NAT:** 1:1 IP mapping for public-facing VMs
- **VPN:** Site-to-site VPN and remote access VPN
- **Load Balancing:** Distributes traffic across multiple VMs

Deployment:

- One virtual router per isolated guest network (default)
- OR a redundant pair for high availability
- Runs a Debian-based system VM template
- Resources: 1-2 vCPU, 256MB-1GB RAM (varies by network size)

Architecture Consideration: Virtual routers are the networking workhorse of CloudStack. For production networks serving many VMs, always deploy redundant virtual routers.

Redundant Virtual Router HA:

```
Guest Network with Redundant Routers:

  [Virtual Router 1 - MASTER]  <--- Active, handling all traffic
         |
         |--- Keepalived/VRRP heartbeat ---
         |
  [Virtual Router 2 - BACKUP]  <--- Standby, takes over if master fails

If master fails: ~3 second failover
```

The Management Server: Brain of the Cloud

The Management Server is CloudStack's orchestration engine and control plane. It's a Java application that:

Core Functions:

- Provides the web UI and RESTful API
- Stores all cloud state in a MySQL database
- Manages resource allocation and scheduling
- Communicates with hypervisor agents on each host
- Controls and monitors System VMs
- Handles user authentication and RBAC
- Generates events, alerts, and usage data
- Manages billing and quota enforcement

Critical for Production: The Management Server must be highly available. A single Management Server is a single point of failure for your entire cloud—the Mars test fails at the first hurdle.

Leaf Spine's HA Management Server Architecture:

```
                [HAProxy Load Balancer - VIP]
               (https://cloudstack.leafspinebooks.co.uk)
                              |
              +---------------+----------------+
              |               |                |
     [Management Server 1] [Management Server 2]   [Management Server 3]
      (Manchester)          (Manchester)            (London - DR)
              |               |                |
              +---------------+----------------+
                              |
           [MySQL Galera Cluster - 3 Nodes]
           (Multi-master replication)
              Node 1: Manchester
              Node 2: London
              Node 3: Manchester (quorum)
```

Why this architecture?

- No single point of failure
- Management Servers can fail without a cloud outage
- Database is multi-master (any management server can write)
- Geographic redundancy (survives a datacentre failure)
- HAProxy provides a single entry point and health checking

One detail worth internalising early: if every Management Server goes down, **your VMs keep running**. The control plane orchestrates; it does not sit in the data path. You lose the ability to make changes, not the workloads themselves. That separation is what makes calm, planned maintenance possible.

Management Server Specifications:

```
Servers: 3x Dell R450 class (2x Manchester, 1x London)
CPU: 1x Intel Xeon (8-16 cores)
RAM: 128GB DDR5 ECC
Storage: 2x 960GB NVMe (RAID-1 for OS and logs)
Network: 4x 10GbE (bonded pair for redundancy)
OS: Ubuntu 24.04 LTS
Java: OpenJDK 17 (LTS)
```

Network Architecture: The Four Traffic Types

CloudStack uses different networks for different purposes. **This is not optional.** Understanding and properly implementing these networks is critical for security, performance, and reliability. Chapter 2 turns this into a full IP and VLAN plan; here is the model.

1. Management Network (Critical Security Zone)

Purpose: CloudStack control plane communication

Traffic:

- Management Server □ Hypervisor hosts
- Management Server □ System VMs
- System VMs □ Management Server

Security Requirements:

- **MUST be isolated** from guest networks
- Firewall rules to allow only CloudStack management traffic
- No direct internet access recommended
- SSH access restricted to jump hosts
- All traffic encrypted where possible

Performance Requirements:

- Low latency (typically <1ms within the datacentre)
- Moderate bandwidth (1-10Gbps sufficient for most)
- High reliability (use bonded interfaces)

Leaf Spine's Management Network:

```
VLAN: 100
Subnet: 10.100.0.0/16
Gateway: 10.100.0.1
DNS: 10.100.0.10, 10.100.0.11
NTP: 10.100.0.20

IP Allocation:
  10.100.1.0/24    - Management Servers, MySQL
  10.100.10.0/24   - Manchester KVM hosts
  10.100.11.0/24   - London KVM hosts
  10.100.12.0/24   - Leeds KVM hosts
  10.100.50.0/24   - Storage nodes (Ceph)
  10.100.100.0/24  - System VMs (SSVM, CPVM)
  10.100.150.0/24  - Infrastructure (switches, PDUs, IPMI)

Physical: Dedicated connection on each host
Bonding: Active-Backup (mode 1) for redundancy
MTU: 1500 (standard Ethernet)
```

2. Storage Network (Performance Critical)

Purpose: Primary storage traffic between hosts and storage

Traffic:

- Hypervisor hosts □ Ceph OSDs (librbd)
- VM disk I/O
- Storage replication and recovery
- Snapshot operations

Performance Requirements:

- **Highest bandwidth requirement** (25GbE minimum, 100GbE ideal at scale)
- **Lowest latency requirement** (target <2ms, ideally <1ms)
- Jumbo frames (MTU 9000) strongly recommended
- VLAN isolation from other traffic

Security Requirements:

- Isolated from guest networks
- No internet connectivity needed
- Can share physical infrastructure with management (separate VLANs)

Leaf Spine's Storage Network:

```
VLAN: 200 (frontend) / 201 (Ceph replication backend)
Subnet: 10.200.0.0/16 and 10.201.0.0/16
Gateway: None (no default route - storage-only network)
MTU: 9000 (jumbo frames for better throughput)

IP Allocation:
  10.200.10.0/24  - Manchester KVM hosts (storage interfaces)
  10.200.11.0/24  - London KVM hosts
  10.200.12.0/24  - Leeds KVM hosts
  10.200.50.0/24  - Ceph OSD nodes Manchester (frontend)
  10.201.50.0/24  - Ceph OSD nodes Manchester (backend/replication)

Physical: Dedicated 25GbE bonded pair (LACP)
Bonding: 802.3ad (LACP) for 50Gbps aggregate
Redundancy: Each OSD node has 4x 25GbE
```

Jumbo Frames Configuration: Why MTU 9000? Larger frames = fewer packets for the same data = lower CPU usage = better throughput.

Ensure MTU 9000 on:

- All host interfaces (storage bond)
- All switch ports (storage VLAN)
- All storage nodes
- Consistent across the entire storage path — one mismatched port and you will chase phantom performance problems for a week

3. Guest Network (VM Application Traffic)

Purpose: VM-to-VM communication and application traffic

Traffic:

- Communication between tenant VMs
- Application traffic within tenant networks
- Isolated per tenant (using VLANs or VXLAN)
- Can be completely isolated or routable

CloudStack Guest Network Models:

Isolated Networks (Most Common):

- VMs get private RFC1918 IPs (10.x, 172.16.x, 192.168.x)
- Virtual Router provides NAT, firewall, load balancing
- VMs cannot communicate across isolated networks
- Perfect for multi-tenant security

Shared Networks:

- All VMs in the zone share the same Layer-2 network

- No virtual router (direct Layer-2 connectivity)
- Good for simple deployments or specific use cases
- Less security isolation

VPC Networks (Virtual Private Cloud):

- Multi-tier networks within a VPC
- Example: Web tier (DMZ), App tier, Database tier
- Complete network isolation per tenant
- Advanced routing, ACLs, and security

Leaf Spine's Guest Network Design:

```
VLANs: 500-999 (500 available guest VLANs)
Type: Isolated Networks (with Virtual Routers)
Subnet Allocation: Automatic (CloudStack assigns from range)
Default Range: 10.10.0.0/16

Physical: 2x 25GbE bonded (LACP) on each host
VLAN Trunking: All guest VLANs trunked on the bond
MTU: 1500 (standard; 9000 possible if all VMs support it)

Guest VLAN Examples:
  VLAN 500: Publisher Client A - Web Tier
  VLAN 501: Publisher Client A - Database Tier
  VLAN 502: Publisher Client B - Production
  VLAN 503: Publisher Client B - Staging
  VLAN 600: Leaf Spine Internal - Production Apps
  VLAN 601: Leaf Spine Internal - Development
```

4. Public Network (Internet-Facing Traffic)

Purpose: Provides internet connectivity for VMs

Traffic:

- Inbound traffic to public-facing VMs
- Outbound internet access (via Source NAT)
- Load balancer public IPs
- VPN endpoints
- Elastic IPs (static NAT)

Requirements:

- **Routable public IP addresses** (allocated from your ISP or RIPE)
- Default gateway pointing to the internet router
- Proper firewall rules for security
- Can share a physical interface with the guest network (different VLAN)

Leaf Spine's Public Network:

```
Manchester Public IPs:
  Allocation: 203.0.113.0/26 (62 usable IPs - documentation range; use yours)
  VLAN: 999
  Gateway: 203.0.113.1 (ISP router)

IP Usage:
  203.0.113.2-10    - Virtual Router source NAT IPs
  203.0.113.11-50   - Static NAT / Elastic IPs for tenant VMs
  203.0.113.51-60   - Load Balancer IPs
  203.0.113.61      - VPN gateway

Physical: Shares the guest bond (VLAN tagged)
Security: Firewall rules at edge router + virtual router ACLs
```

Public IP Conservation Strategy:

- Most VMs use Source NAT (many VMs share one public IP)
- Static NAT (1:1 mapping) only for public-facing services
- Load balancers enable multiple VMs to share one public IP
- Port forwarding for selective service exposure

CloudStack 4.22 LTS: Why This Release

Leaf Spine chose **CloudStack 4.22 LTS** deliberately. LTS branches receive 24 months of support—proactive maintenance for the first 18 months, then blocker and security fixes—which matters when you are running a business on it rather than a hobby. At the time of writing the current maintenance release is 4.22.1.0, and that is the baseline for every example in this book. Here is what the recent releases bring and why it matters to this story:

1. Extensions Framework (Orchestrate Anything)

What it is: A plugin architecture (introduced in 4.21 and carried forward) allowing external systems to be managed through CloudStack

What it enables:

- Manage non-CloudStack resources through the CloudStack UI/API
- Built-in extensions for platforms such as Proxmox and Hyper-V
- Custom resource types and lifecycle management
- Third-party service integration

Why Leaf Spine cares: They're planning edge computing deployments at retail bookshop locations. The Extensions Framework lets them manage low-cost edge servers through their central CloudStack platform.

Use case example:

```
Leaf Spine Edge Deployment:
  Central Cloud: CloudStack in Manchester/London (heavy workloads)
  Edge Locations: 50 bookshops with mini-servers
  Management: All managed through a single CloudStack interface
  Use: Local caching, IoT data aggregation, offline resilience
```

2. VMware-to-KVM Migration Tooling

What it is: First-class tooling for importing VMware workloads onto KVM, improved substantially across 4.21 and 4.22 (including VDDK support and guest OS handling fixes in 4.22.1.0)

Why Leaf Spine cares: This is not a feature to them; it is the escape tunnel itself. The entire migration story of this book—moving live workloads off the estate whose licensing just rose 567%—runs through this tooling. Chapter by chapter, we will use it in anger.

3. Enhanced Backup and Disaster Recovery

What's included:

- Backup & Recovery (B&R) framework enhancements, including NAS backup support
- Improved snapshot operations
- Better backup scheduling and retention handling

Why Leaf Spine cares: A publisher's catalogue is its crown jewels. The B&R framework gives them provider-based backup on KVM today (we build it in Chapter 16, with Ceph RGW S3 as the immutable target).

4.23 Alignment: The headline B&R development to watch is **native Veeam Backup & Replication integration for CloudStack on KVM**—full and incremental backup/restore, multi-disk VMs, and a VMware-to-KVM restore path—targeted at CloudStack 4.23. On 4.22 the documented Veeam plugin remains VMware-oriented, so our KVM chapters use the NAS B&R provider plus Veeam direct-to-object S3 repositories on Ceph RGW today, and Chapter 16 shows exactly where the native integration will slot in once 4.23 ships.

4. GPU Support and Allocation

What it is: GPU integration with KVM in CloudStack—GPU discovery, inventory, passthrough and allocation—introduced as a technical preview in 4.21 and maturing through the 4.22 line

Why Leaf Spine cares: They're launching AI-powered services:

- Book recommendation engine using local LLMs
- AI content moderation
- A publisher's AI writing-assistant service
- Automated cover design tooling

The economics (indicative, at the time of writing):

```
GPU Investment:
  Adding an NVIDIA RTX 4060 Ti 16GB card to an existing host
  costs a few hundred pounds—around £400 used, £450 new at the
  time of writing, not thousands—and it idles at a few watts.

  Renting one always-on cloud GPU instance suitable for the
  same inference workloads runs to several hundred pounds a
  month, every month, forever: an AWS g5.xlarge (24GB A10G)
  is about $1.01/hour on demand—roughly £550 a month at ~$1.32
  to the pound.

  So the arithmetic is stark: a ~£450 card pays for itself
  against ~£550/month of rental in under a month. Break-even on
  owned inference silicon is weeks, not years. After that,
  every token generated is effectively free (minus
  electricity—and ours comes off the roof).
```

The deeper point: data never leaves the building. A "private AI" pitch to publishing clients is only honest if the GPUs are yours.

5. Security Enhancements

Recent improvements:

- Enhanced audit logging with structured events
- Improved RBAC with fine-grained permissions
- Network security group enhancements
- Certificate management automation
- UEFI enabled by default on KVM hosts (4.22.1.0)

UK GDPR / Compliance Benefits:

- Complete audit trail for data access (ICO expectations)
- Role-based access limits blast radius
- Automatic log retention and archival
- Integration with self-hosted SIEM tools (ELK, Graylog, Wazuh)

Example Audit Log:

```
{
  "timestamp": "2026-05-25T10:30:45Z",
  "user": "admin@leafspinebooks.co.uk",
  "action": "VM_STOP",
  "resource": "vm-12345-customer-db",
  "ip_address": "10.100.1.50",
  "result": "SUCCESS",
  "tenant": "customer-acme-publishing"
}
```

6. Performance, Scale, and the Webhooks API

Enhancements across recent releases:

- Native Webhooks API (4.20+) — CloudStack events pushed to your automation (we drive n8n with this later)
- Faster volume snapshot operations
- Optimised storage migration and live migration stability
- Improved database query performance
- Broader OS platform support (EL10, openSUSE 15.6)

Why Leaf Spine cares:

- Snapshots run faster = less impact on production
- Live migration stability = safer maintenance windows
- Webhooks = event-driven automation without polling
- Scale headroom = growth runway for 5+ years

7. IoT and Edge Computing Support

What's available:

- Lightweight templates for edge devices
- Multi-architecture support including ARM64 hosts
- Edge-to-cloud connectivity patterns

Leaf Spine's Future Vision: *This is where their innovation strategy shines:*

```
Leaf Spine IoT Platform (Phase 2 of business transformation):

Edge Locations (Bookshops):
  - Raspberry Pi or mini-PC class hardware at the edge
  - Local caching of popular ebooks
  - Customer behaviour analytics (footfall, browsing patterns)
  - POS integration
  - Offline resilience (works without internet)

Data Collection:
  - Environmental sensors (temperature, humidity for book preservation)
  - LoRaWAN for sensors where WiFi fails, gateways managed via MQTT
  - Inventory tracking
  - Energy consumption monitoring

Central Processing:
  - Aggregate data from 50+ edge locations
  - AI-powered insights and recommendations
  - Centralised management of edge infrastructure
  - Offer a "Smart Bookshop Platform" to other retailers
```

Aquaponics / Sustainable Tech Integration:

While Leaf Spine's core business is publishing, Sarah's team has discussed offering their IoT platform to other industries. The CloudStack platform they're building could just as easily support aquaponics monitoring and control, environmental sensor networks, sustainable agriculture data platforms, or smart building management. This isn't speculation on my part: the patterns in this book are the ones running my own aquaponics greenhouse and apiary right now. CloudStack does not care whether the workload is a bookshop till or a fish tank—and that versatility is the point.

Basic vs Advanced Networking: A Critical Decision

Sarah's team debated this extensively. CloudStack offers two fundamentally different networking modes, and the choice is permanent (you cannot change it after initial zone deployment).

Basic Networking (Simple, Limited)

Characteristics:

- Single flat Layer-2 network for all VMs
- All VMs in the same broadcast domain
- Security Groups for VM-level isolation
- DHCP assigns IPs from the same subnet
- No virtual routers deployed
- No true multi-tenancy

Pros:

- Extremely simple to set up and understand

- No VLAN configuration needed
- Good for small, single-tenant environments
- Lower resource overhead (no virtual routers)

Cons:

- No true network isolation between tenants
- All VMs visible to each other (Layer 2)
- Security Groups only (no network ACLs, VPCs, VPNs)
- Cannot do advanced features (load balancers, tiered networks)
- Difficult to scale securely for multi-tenant

Good fit for: small single-organisation deployments, dev/test environments, fully-trusted private clouds, simple workloads.

Advanced Networking (Powerful, Production-Ready)

Characteristics:

- Full VLAN isolation per network
- Virtual Private Cloud (VPC) support
- Virtual Routers provide rich networking services
- Complete network isolation between tenants
- Load balancers, VPNs, multi-tier networks
- True multi-tenant security

Pros:

- Complete network isolation (true multi-tenancy)
- Flexible network topologies (VPCs, tiers, etc.)
- Rich services (LB, VPN, firewall, NAT, port forwarding)
- Scales to thousands of tenants
- Enterprise-grade security
- Advanced features (network ACLs, egress rules, etc.)

Cons:

- More complex to plan and implement
- Requires VLAN planning and switch configuration
- Higher resource overhead (virtual routers per network)
- More moving parts = more to understand

Good fit for: multi-tenant platforms, service providers, enterprise private clouds with multiple departments, any production environment requiring strong isolation or compliance.

Leaf Spine's Decision: Advanced Networking

Reasoning:

1. **Multi-Tenancy Requirement:**
   - Publisher clients need complete isolation from each other
   - Corporate workloads must be separate from client workloads
   - Different security policies per tenant
2. **Compliance:**
   - PCI-DSS requires network segmentation
   - UK GDPR requires demonstrable data isolation
   - Clients demand proof of isolation
3. **Features Needed:**
   - VPCs for client multi-tier applications
   - Load balancers for high-traffic ebook stores
   - Site-to-site VPNs between Manchester-London-Leeds
   - Advanced firewall rules per client
4. **Future-Proofing:**
   - The platform business is growing
   - Cannot change networking mode later (must choose correctly now)

Implementation Decision:

```
Networking Mode: Advanced
Guest Network Isolation: VLAN (VLANs 500-999)
VPC Support: Enabled
Virtual Router HA: Enabled for production networks
Redundant Routers: Yes (for client SLAs)
Network ACLs: Enabled
Egress Firewall: Enabled
Load Balancing: Enabled
VPN: Enabled (both remote access and site-to-site)
```

Hypervisor Selection: Why KVM?

CloudStack supports multiple hypervisors (you can even mix hypervisors across different clusters):

Supported Hypervisors:

- **KVM** (Linux-based, open-source)
- **VMware vSphere** (commercial)
- **XenServer/XCP-ng** (Xen-based)
- **Hyper-V and Proxmox** (via the Extensions framework)
- **LXC** (containers, lightweight)

Hypervisor Comparison Matrix:

Feature	KVM	VMware vSphere	XenServer/XCP-ng
Licence Cost	Free (GPL)	Per-core subscription, 16-core/CPU minimum	Free (XCP-ng)
Performance	Excellent (in-kernel)	Excellent	Very Good
Live Migration	Yes	Yes (vMotion)	Yes
Storage Support	All protocols incl. native RBD	Limited (NFS, iSCSI, FC, vSAN)	Most protocols
GPU Passthrough	Yes (native, no licence)	Yes (vGPU licence extra)	Yes
Management	CLI + CloudStack	vCenter + CloudStack	XenCenter + CloudStack
Community	Large, active	Commercial support	Smaller community
Skills Requirement	Linux admin	VMware certification culture	Linux/Xen knowledge
Vendor Lock-In	None	High	Low
Security	SELinux/AppArmor, sVirt	Built-in	Good

Leaf Spine's Analysis:

Option 1: VMware vSphere

- **Pros:** Familiar to enterprise IT, excellent management tools, mature
- **Cons:** This is the landlord they are escaping. Under the post-Broadcom model the estate is licensed per core, on subscription, with a 16-core minimum per CPU. Chapter 2 designs a 14-host, 896-core estate; at current published rates that licensing alone runs to roughly **£181,500 per year** for the full VCF bundle—around **£907,500 over five years**, for software, before a single server is bought.
- **Verdict:** The problem, not the solution

Option 2: XenServer/XCP-ng

- **Pros:** Free (XCP-ng), good performance, proven
- **Cons:** Smaller community and ecosystem momentum than KVM
- **Verdict:** Viable, but less momentum

Option 3: KVM (Selected)

- **Pros:**
  - Zero licensing costs
  - Native Linux (the team is already expert in Linux)

  - Excellent performance (in-kernel, Type 1 characteristics)
  - Huge community and corporate backing across the industry
  - Supports all storage protocols including native Ceph RBD
  - GPU passthrough without extra licensing
  - Part of the Linux kernel (not going away)
- **Cons:**
  - Less GUI management on its own (but CloudStack provides this)
  - Requires Linux expertise (the team has this)
- **Verdict: SELECTED** – best fit for goals and team

KVM Technical Advantages for Leaf Spine:

1. **Direct Hardware Access:**
   - KVM is part of the Linux kernel
   - VMs run as Linux processes with hardware acceleration
   - No separate hypervisor licence layer between you and the metal
2. **Storage Flexibility:**
   - Native librbd (Ceph RBD) support (no iSCSI shim)
   - Supports all storage: NFS, iSCSI, RBD, local, FC
   - Excellent integration with Ceph features (cloning, snapshots)
3. **Security:**
   - SELinux/AppArmor integration (mandatory access control)
   - KVM uses sVirt (secure virtualisation)
   - Complete VM isolation
   - Regular security updates via the Ubuntu repositories
4. **GPU Support:**
   - NVIDIA/AMD GPU passthrough without licensing
   - VFIO for safe PCI device assignment
   - Perfect for the AI/ML workload requirements
5. **Cost Efficiency:**

```
VMware VCF subscription (5 years, 896-core estate):
  ~£181,500/year x 5 = ~£907,500

KVM licensing (5 years):
  £0

The difference buys the entire hardware estate —
with change. Chapter 2 does this arithmetic properly;
Chapter 17 audits it against the running business.
```

Storage Strategy: Why Ceph?

Storage was Leaf Spine's most critical decision. Options evaluated:

Option 1: Traditional SAN (Dell EMC, NetApp, Pure Storage)

Pros: proven, mature, commercial support, all-in-one, excellent performance **Cons:**

- **Extremely expensive** (six figures for an HA pair before disks)
- Vendor lock-in (proprietary)
- Single-site (DR requires a second SAN)
- Scale-up model (limited upgrade path)

Verdict: Too expensive, against strategic goals

Option 2: Standalone NFS Server

Pros: simple to deploy, well-understood, works with any hypervisor **Cons:** performance ceilings, hard to make highly available, scale limits, no native redundancy

Verdict: Good for secondary storage, not primary storage

Option 3: Upstream Ceph, Deployed with cephadm (Selected)

Ceph is a distributed storage system that provides:

- **Block storage** (RBD — RADOS Block Device) □ CloudStack primary storage
- **File storage** (CephFS) □ where a real filesystem is genuinely needed
- **Object storage** (RGW — S3-compatible) □ secondary storage, backup targets, application S3

One cluster, three protocols, no licences. For CloudStack primary storage we use **RBD**, which KVM speaks natively through librbd.

Ceph Architecture:

```
Ceph Cluster:
  MONs (Monitors): Track cluster state, maintain the cluster map (3+ for quorum)
  OSDs (Object Storage Daemons): Store data on disks (1 OSD per drive)
  MGRs (Managers): Metrics, orchestration, dashboard
  RGW (RADOS Gateway): S3/Swift API
  MDS (Metadata Servers): For CephFS only (not needed for RBD)

For CloudStack Primary Storage:
  KVM hosts connect directly to Ceph OSDs via librbd
  No iSCSI overhead (native protocol)
  VM disks are RBD images in Ceph pools
```

Why cephadm (and not an appliance wrapper):

- Official Ceph orchestrator: containerised deployment, one bootstrap command
- Built-in dashboard, monitoring, and health checks
- Rolling upgrades managed by the orchestrator
- Add OSD nodes with a single command (horizontal scaling)
- Nothing between you and upstream when you need help
- The skills you build are Ceph skills, portable everywhere Ceph runs

Why Ceph RBD + KVM is the right pairing:

1. **Native Integration:** QEMU speaks librbd directly; thin provisioning, snapshots, clones all native
2. **CloudStack Integration:** snapshots are copy-on-write (instant); template cloning is instant; live migration just works
3. **Performance:** with NVMe OSDs, sub-millisecond latency; scales linearly as nodes are added
4. **Reliability:** 3x replication, self-healing, no single point of failure, survives whole-node loss
5. **Operations:** add capacity by adding nodes; rolling upgrades; safe maintenance; built-in health checks

Cost Comparison (indicative, current street pricing):

```
Traditional SAN route:
  HA controller pair + ~100TB usable + 5yr support
  comfortably exceeds £400,000

Ceph route (Manchester production cluster):
  6x storage nodes with 24x 3.84TB NVMe
  lands around a quarter of that — on commodity
  servers you can redeploy, with no licence meter

And: better scaling, three protocols, no vendor lock-in.
Chapter 2 prices it precisely; Chapters 7-8 build it.
```

Fair warning from the workshop: 2026 is a brutal year to buy storage—the AI build-out has enterprise NVMe and even spinning disk on allocation, and prices have climbed hard. That cuts both ways in this story. It makes the hardware bill higher than it would have been a year ago, and it makes owning your refresh cycle—buying what you need when *you* choose, including refurbished kit—more valuable, not less. When you rent, somebody else decides when you pay the surcharge.

The Road Ahead: Leaf Spine's Implementation Journey

Sarah and her team now understood what they were building. The architecture was clear:

Phase 1: Foundation (Chapters 4-6) - Manchester Primary Site

- Install and configure KVM hypervisor hosts
- Deploy a MySQL Galera cluster for database HA
- Install CloudStack Management Servers (HA)
- Configure networking (VLANs, bonds, MTU)

Phase 2: Storage (Chapters 7-8)

- Deploy the Ceph cluster in Manchester with cephadm
- Configure RBD pools and performance tuning
- Set up secondary storage
- Test failover and recovery

Phase 3: Cloud Deployment (Chapter 9)

- Initialise CloudStack
- Run the zone configuration wizard
- Deploy System VMs (SSVM, CPVM)
- Verify basic cloud functionality

Phase 4: Production Ready (Chapters 10-12)

- Add all KVM compute hosts to clusters
- Deploy first guest VMs
- Implement monitoring and alerting
- Launch first workloads

Phase 5: Geographic Expansion (Chapters 13-15)

- Replicate the setup to London and Leeds
- Configure multi-zone operations
- Implement cross-site disaster recovery
- Test site failover scenarios

Phase 6: Advanced Capabilities (Chapters 16-18 and Part 2)

- Backup and disaster recovery in anger
- Cost optimisation and capacity planning
- GPU workloads, Kubernetes, advanced networking, automation, security hardening, edge and IoT — the Part 2 follow-on deepens each of these

The Ultimate Goal:

Transform from a company renting its digital existence to a company that:

- Owns and controls its entire cloud platform
- Escapes a six-figure annual licensing bill permanently
- Maintains complete data sovereignty in the UK
- Provides cloud services to its publishing clients
- Can scale across multiple sites with built-in DR and HA
- Can offer innovative services (GPU, IoT, edge computing)

Key Takeaways

Let's solidify what we've learned:

CloudStack is an orchestration platform, not a replacement for hypervisors or storage. It coordinates them into a unified cloud.

Architecture hierarchy matters: Region □ Zone □ Pod □ Cluster □ Host. Each level has specific purposes and constraints.

Cluster homogeneity is non-negotiable: identical CPUs, hypervisors, and storage access, or live migration and HA will fail you when you need them most.

Four network types serve different purposes: Management (control plane), Storage (disk I/O), Guest (VM traffic), Public (internet).

System VMs provide cloud services: SSVM (templates/snapshots), CPVM (console access), Virtual Routers (networking services).

Management Server HA is critical – but even if the control plane dies, your VMs keep running. Orchestration is not the data path.

Advanced Networking enables true multi-tenancy with VLANs, VPCs, and complete isolation. Required for client platforms and compliance.

KVM provides excellent performance and flexibility with zero licence cost and no per-core meter – the direct answer to the problem that started this book.

Ceph delivers distributed, self-healing storage – block, file, and S3 from one cluster, deployed with cephadm, at a fraction of SAN cost.

CloudStack 4.22 LTS is the battle-hardened baseline: Extensions framework, VMware-to-KVM migration tooling, B&R framework, GPU support, webhooks – with native KVM Veeam integration on the 4.23 horizon.

Data sovereignty and security are built in, not bolted on. Complete control over where data resides and who can access it.

Questions to Answer Before Next Chapter

Before designing your complete architecture in Chapter 2, ensure you can answer:

How many zones (datacentres) does your deployment need?

- Leaf Spine: 3 (Manchester, London, Leeds)

What networking mode fits your requirements?

- Leaf Spine: Advanced (multi-tenant platform requirements)

What hypervisor will you use and why?

- Leaf Spine: KVM (performance, cost, flexibility, no lock-in)

What storage protocol makes sense for your hypervisor?

- Leaf Spine: Ceph RBD (distributed, HA, native KVM integration)

How many management servers for your HA requirements?

- Leaf Spine: 3 (two in Manchester, one in London for geographic redundancy)

Do you need cross-site disaster recovery?

- Leaf Spine: Yes (London is the DR site for Manchester)

What is your availability target?

- Leaf Spine: 99.95% (about 4.4 hours of downtime per year, maximum)

What compliance requirements must you meet?

- Leaf Spine: UK GDPR, ICO expectations, PCI-DSS

What advanced features do you need?

- Leaf Spine: GPU (AI services), Kubernetes (containers), VPNs (site-to-site), load balancers

Coming in Chapter 2

In the next chapter, we'll design Leaf Spine's complete three-city cloud architecture:

- Detailed network diagrams with IP allocation
- VLAN planning and allocation strategy
- Complete hardware specifications
- Rack layouts and physical infrastructure
- Power, cooling, and datacentre requirements
- Cost analysis: the full arithmetic of the VMware escape
- Compliance architecture (UK GDPR, PCI-DSS)
- Capacity planning methodology and scaling strategy

Next Chapter: Chapter 2 — Designing the Three-City UK Cloud.

Additional Resources

Official CloudStack Documentation:

- https://docs.cloudstack.apache.org/

CloudStack Community:

- Mailing lists: https://cloudstack.apache.org/mailing-lists.html
- Community channels: https://cloudstack.apache.org/contribute.html

Learning Resources:

- Apache CloudStack Blog: https://cloudstack.apache.org/blog/
- CloudStack Collaboration Conference, Edinburgh, November 2026

* * *

Apache CloudStack Deployment Guide — Leaf Spine Books edition Copyright © 2026 Michael John Leslie Hinsley. Code examples licensed under Apache License 2.0. Contact details for this book are on the contact page at the back.

Chapter 2: Designing the Three-City UK Cloud

From Whiteboard to Blueprint — Planning a Production-Grade Infrastructure

The Story: Architecture Week

Sarah Martinez cleared the large whiteboard in the Manchester conference room. "Right team, we've got one week to finalise the architecture. After that, we start ordering hardware."

The whiteboard was already covered with network diagrams, IP ranges, and hardware specs. Marcus Chen, their senior engineer, had spent the weekend researching modern server configurations. Priya Patel, their network architect, had designed a topology that could scale from their current needs to supporting thousands of AI-powered instances.

"Let's start with what we know," Sarah said, drawing three circles labelled Manchester, London, and Leeds. "We have colocation space in all three cities. Manchester is our primary site – largest deployment, most resources. London is our DR site with enough capacity to take over if Manchester goes down. Leeds is development and testing."

James Wilson, their storage specialist, added, "And we're standing up a second Ceph cluster in Manchester purely for backups. Deliberately separate from production storage. A backup that shares a failure domain with the thing it protects is a comfort blanket, not a backup."

"Exactly," Sarah confirmed. "But here's the twist – I want us to plan for GPU workloads from day one. Instead of sending our customer data to OpenAI or Anthropic, we run the models ourselves with Ollama. Private, UK-based, GDPR-compliant AI."

"Then let me save us a fortune," Marcus said, pulling up a browser tab full of five-figure datacentre GPU listings. "The brochures want to sell us accelerator cards that cost more than the servers under them. We don't need that to serve language models to a publishing business. Refurbished R7525s – two-socket AMD boxes, 2U, thirty-two DIMM slots of affordable DDR4 – and we drop two consumer RTX cards in each. Hundreds of pounds per card, not thousands. If a workload ever genuinely outgrows them, we'll have earned the right to buy bigger – and we'll know exactly why."

"That's the Return on Research mindset," Sarah said. "Idle silicon is waste, and so is over-bought silicon. And everything gets automated. Ansible for the entire deployment, so it's repeatable and documented. If someone can't deploy our architecture by following our playbooks, we've failed."

Priya pulled out her laptop. "I've spec'd the networking – 6× 25Gbps NICs per compute node, separated properly: management, storage, and guest traffic each on their own bonded pair. And the switches run Dell OS10 – Dell's supported NOS on the same open-networking metal, with a real automation surface: an Ansible collection, gNMI telemetry, an IOS-style CLI we already know. The trade is a licence and a vendor relationship instead of community SONiC's read-the-source openness – and if that openness ever matters more to us, the SONiC build is a fork away."

Over the next week, the team designed every aspect of their three-city cloud: network topology, IP addressing, VLAN allocation, hardware specifications, security zones, and GPU integration. This chapter documents their decisions and gives you a complete blueprint for a production-grade CloudStack deployment. Chapter 1 sketched the model; this is the full set of working drawings.

Design Principles

Before diving into specifics, Leaf Spine Books established core principles:

1. High Availability First

- No single points of failure
- Redundant everything: power, network, storage, management
- Cross-site failover capability
- Manchester <-> London primary/DR pairing
- Leeds as independent dev/test (can run standalone)

2. Performance Matters

- 25Gbps networking throughout
- NVMe storage for production (no spinning disks in the VM path)
- GPU capability for AI workloads
- Proper network separation to avoid bottlenecks
- Over-provision where it counts

3. Security by Design

- Network segmentation (not bolt-on)
- Defence in depth
- Least privilege access
- Encrypted communication
- UK GDPR compliance built-in
- Private AI (data never leaves the infrastructure)

4. Automation Everything

- Ansible for all deployments
- Infrastructure as Code
- Documented and repeatable
- Version controlled
- No manual configuration drift

5. Future-Ready

- Scale compute and storage independently
- Add GPU capacity without redesign
- Support for 100Gbps networking later
- Container orchestration ready (Kubernetes)
- Multi-tenant from day one

And one principle that sits above the other five — the Mars test. If a server in this design hung at 3am and nobody could physically reach it, could we recover? Every choice that follows (out-of-band management on everything, redundant virtual routers, self-healing storage, automation instead of heroics) is that question being answered.

Geographic Layout

Site Characteristics

Manchester — Primary Production Site

Role: Primary datacentre, backup storage cluster, main operations

Datacentre: Equinix MA1 (example — adapt to your actual facility)

Resources:

- 8 KVM compute nodes in 2 clusters: 4× Dell R660 (production) + 4× Dell R7525 (GPU/AI cluster)
- 6 Ceph production storage nodes
- 4 Ceph backup storage nodes (separate cluster)
- 2 CloudStack management servers
- Support infrastructure (DNS, secondary storage, monitoring)

Network:

- 100Gbps uplink to internet
- Low-latency connection to London (dark fibre)
- VPN backup to Leeds

Power: Dual-feed A+B power with UPS and generator backup

Why Manchester?

- Company headquarters location
- Largest skilled workforce
- Best datacentre facilities
- Central to London and Leeds (geographic redundancy)

London — Disaster Recovery Site

Role: Secondary datacentre, DR failover, secondary production

Datacentre: Equinix LD8 (example)

Resources:

- 4 KVM compute nodes (1 cluster, Dell R660)
- 4 Ceph production storage nodes
- 1 CloudStack management server
- Minimal support infrastructure

Network:

- 100Gbps uplink to internet
- Dark fibre to Manchester
- VPN backup to Leeds

Power: Dual-feed with UPS and generator

Why London?

- Geographic separation from Manchester (disaster recovery)
- Major internet exchange (excellent connectivity)
- Customer concentration (lower latency for southern users)
- 180 miles from Manchester (safe DR distance)

Latency: Manchester <-> London: ~5ms (acceptable for most replication)

Leeds — Development & Testing

Role: Non-production workloads, testing, development

Datacentre: UK regional colocation (example: a Leeds facility)

Resources:

- 2 KVM compute nodes — the existing Dell R740s, reused
- 3 Ceph storage nodes (smaller drives)
- No dedicated management server (uses Manchester/London)
- Minimal support infrastructure

Network:

- 10Gbps uplink
- VPN to Manchester and London

Power: Standard redundant power

Why Leeds?

- Regional office location

- Cost-effective datacentre
- Close enough to Manchester for team access (43 miles)
- Isolated from production (safe to break things)

Those Leeds R740s matter more than their size suggests. They are the servers Leaf Spine Books already owned – the ones sitting idle because the new licensing model made their cores unaffordable to light up. Under KVM there is no per-core meter running. The "waste" goes back to work on day one, and the development environment costs nothing in new compute.

Latency:

- Leeds <-> Manchester: ~3ms
- Leeds <-> London: ~7ms

Network Architecture

Physical Network Design

Each site follows the same network architecture pattern, scaled to site size. Chapter 1 introduced CloudStack's four traffic types; here is how they land on real interfaces.

Network Traffic Types

Leaf Spine Books' network is separated into distinct traffic types, each on dedicated bonded interfaces:

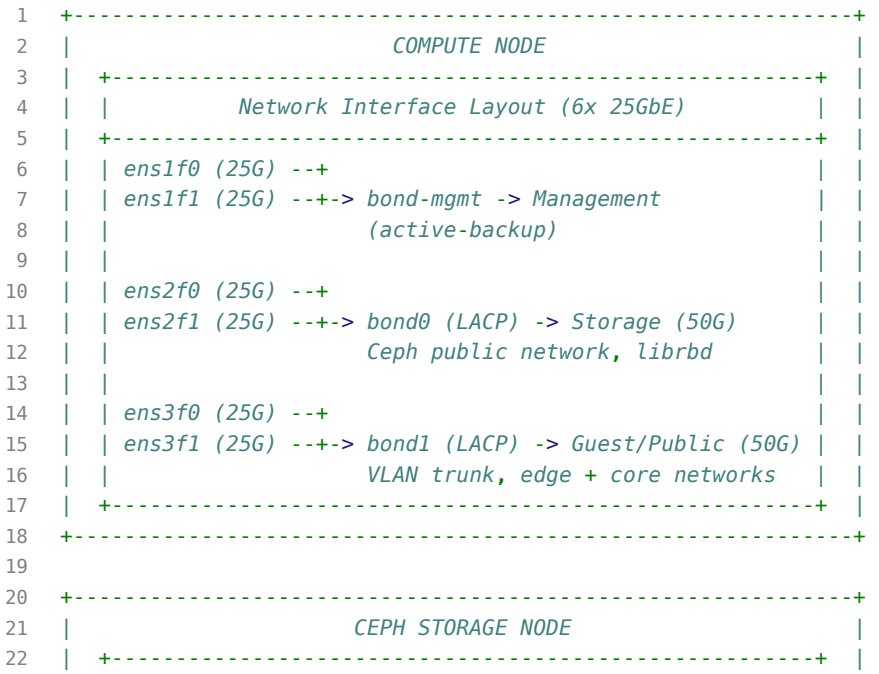

```
|  |     Network Interface Layout (4x 25GbE + 2x 1GbE)     |  |
|  +-------------------------------------------------------+  |
|  | eno1/eno2 (1G) -> bond-mgmt -> Management              |  |
|  |                   (active-backup, LOM ports)           |  |
|  |                                                        |  |
|  | ens2f0 (25G) --+                                       |  |
|  | ens2f1 (25G) --+-> bond0 (LACP) -> Ceph cluster (50G) |  |
|  |                    replication + recovery              |  |
|  |                                                        |  |
|  | ens3f0 (25G) --+                                       |  |
|  | ens3f1 (25G) --+-> bond1 (LACP) -> Ceph public (50G)  |  |
|  |                    librbd client I/O from KVM hosts    |  |
|  +-------------------------------------------------------+  |
+-------------------------------------------------------------+
```

1. Management Network

Purpose: CloudStack management, SSH access, monitoring, out-of-band coordination

Characteristics:

- Bonded pair, active-backup (bandwidth is not the point; surviving a switch failure is)
- Isolated VLAN, per-site /24 allocations within one /16
- No route to the internet (bastion host access only)
- Carries CloudStack agent traffic, Prometheus scrapes, Ansible runs

Why modest bandwidth is sufficient: management traffic is SSH, API calls, and metrics. The high-speed silicon is saved for data. On the R660s we bond two of the six 25GbE ports for management because the ports are there; on the storage nodes the two 1GbE LOM ports do the job perfectly well.

Ubuntu Configuration (Netplan):

```
# /etc/netplan/01-management.yaml
network:
  version: 2
  ethernets:
    ens1f0:
      dhcp4: no
    ens1f1:
      dhcp4: no
  bonds:
    bond-mgmt:
      interfaces: [ens1f0, ens1f1]
      parameters:
        mode: active-backup
        mii-monitor-interval: 100
      addresses:
        - 10.100.10.11/16
      routes:
```

```
        - to: default
          via: 10.100.0.1
      nameservers:
        addresses: [10.100.0.53, 10.100.0.54]
```

2. Ceph Cluster Network (Storage Backend)

Purpose: Ceph replication, recovery, and rebalancing between OSD nodes

Characteristics:

- Dedicated 25GbE bonded pair (50Gbps aggregate, LACP)
- Isolated, no routing between sites
- Per-site network only
- Jumbo frames (MTU 9000)

Why bonded 50Gbps?

- Ceph replication is bandwidth-intensive: 3× replication means every client write becomes three writes
- A single NVMe drive can saturate 25Gbps on its own
- Recovery traffic after a failed drive or node is a storm you want over quickly
- Failover redundancy

Traffic Flow:

```
Ceph Node 1 --50Gbps--+
                      +--> Replication / recovery
Ceph Node 2 --50Gbps--+    (3x copies, prod pools)
                      +
Ceph Node 3 --50Gbps--+
```

3. Ceph Public Network (Storage Frontend)

Purpose: Client I/O — KVM hosts speaking native librbd to the OSDs

Characteristics:

- 25GbE bonded pair per storage node (50Gbps aggregate)
- Isolated VLAN per site, MTU 9000
- This is the network the compute nodes' storage bond connects to

There is no iSCSI gateway, no NFS head, no appliance in this path. QEMU on each KVM host opens RBD images directly through librbd, talking to every OSD in parallel. Multipathing is not bolted on — it is how Ceph works: the client computes where the data lives (CRUSH) and goes straight there. Remove the middleman and you remove the middleman's failure modes.

Traffic Flow:

```
KVM hosts --librbd (Ceph public network)--> OSDs directly
(Parallel I/O to every OSD; no gateway, no single path)
```

4. Guest/Public Network

Purpose: VM traffic, internet access, customer applications

Characteristics:

- 25GbE bonded pair (50Gbps aggregate, LACP)
- VLAN trunking for multiple networks
- Split into Edge and Core network ranges
- Public IP space and RFC1918 private

Why bonded 50Gbps?

- Customer-facing traffic (the revenue path)
- AI inference workloads (high bandwidth)
- GPU instances streaming data
- Future-proof for growth

Edge vs Core:

- **Edge:** Customer-facing, public IPs, load balancers
- **Core:** Internal services, databases, private apps

The Switch Fabric

The name over the door is not an accident. Leaf Spine Books runs a leaf-spine fabric on **Dell open-networking switches running Dell SmartFabric OS10**: S5248F-ON (48× 25GbE) as leaf switches at every site, with a pair of S5232F-ON (32× 100GbE) spines in Manchester where the port counts justify them — the same metal as the canonical edition, with OS10 as the network OS in place of community SONiC. In our lab chapters the same topology runs on **OS10 virtual** — real Dell firmware under EVE-NG, on a 120-day trial or an owned entitlement (Chapter 3 covers the lab in full); in the production narrative, **OS10 Enterprise with Dell support** is the operating model where a vendor contract matters. The free **community SONiC build** (sonic-vs) remains the canonical edition's fabric on `main`, and staying on it is always an option. Refurbished open-networking switches are still a thoroughly sensible way to buy this fabric — the silicon does not care that it had a previous owner, and the savings are substantial — though on this edition you also sort the OS10 image entitlement (the second-hand-kit / Dell My Account route described in Chapter 3), and the per-switch OS10 licence is a real line item, priced in the cost analysis below. Chapter 12 builds the fabric in detail, BGP EVPN and all.

VLAN and IP Architecture

Chapter 1 set the addressing model; this section turns it into the full plan. One scheme, three sites, no surprises: management lives in 10.100.0.0/16, Ceph public in 10.200.0.0/16, Ceph cluster in 10.201.0.0/16, with per-site /24s inside each. Guest networks take VLANs 300–399 everywhere.

Site-Wide VLAN Plan

```
VLAN 100: Management — 10.100.0.0/16
  - Gateway: 10.100.0.1 (per-site firewall/router, VRRP)
  - Core allocations:
    - 10.100.0.53-54    - Internal DNS (MAN primary, LON secondary)
    - 10.100.0.20       - NTP
    - 10.100.1.0/24     - Management servers + MySQL Galera
    - 10.100.3.0/24     - Monitoring hub (Manchester)
    - 10.100.10.0/24    - Manchester KVM hosts
    - 10.100.11.0/24    - London KVM hosts
    - 10.100.12.0/24    - Leeds KVM hosts
    - 10.100.50.0/24    - Ceph storage nodes (all clusters)
    - 10.100.60.0/24    - Secondary storage
    - 10.100.70.10-250  - Pod reserved range (system use)
    - 10.100.100.0/24   - System VMs (SSVM, CPVM)
    - 10.100.103.0/24   - Monitoring warm replica (London)
    - 10.100.150.0/24   - Infrastructure (switches, PDUs, IPMI)

VLAN 200: Ceph Public (storage frontend) — 10.200.0.0/16
  - No gateway (isolated, storage-only), MTU 9000
  - - 10.200.10.0/24  - Manchester KVM hosts (storage bond)
    - 10.200.11.0/24  - London KVM hosts
    - 10.200.12.0/24  - Leeds KVM hosts
    - 10.200.50.0/24  - Manchester Ceph prod nodes (public)
    - 10.200.52.0/24  - London Ceph nodes (public)
    - 10.200.53.0/24  - Leeds Ceph nodes (public)

VLAN 201: Ceph Cluster (replication backend) — 10.201.0.0/16
  - No gateway (isolated), MTU 9000, OSD nodes only
  - - 10.201.50.0/24  - Manchester Ceph prod nodes (cluster)
    - 10.201.52.0/24  - London Ceph nodes (cluster)
    - 10.201.53.0/24  - Leeds Ceph nodes (cluster)

VLAN 205: Ceph Public — Backup cluster (Manchester) — 10.200.55.0/24
  - Independent failure domain; isolated, MTU 9000
    - 10.200.55.0/24  - Manchester Ceph backup nodes (public)

VLAN 206: Ceph Cluster — Backup cluster (Manchester) — 10.201.55.0/24
  - Independent failure domain; isolated, MTU 9000, OSD nodes only
    - 10.201.55.0/24  - Manchester Ceph backup nodes (cluster)

VLAN 300-399: Guest Networks (CloudStack-managed)
  - Dynamic allocation by CloudStack, isolated per tenant
  - 300-349: Edge (customer-facing applications, load balancers)
  - 350-399: Core (internal applications, databases, middleware)
  - Default guest CIDR pool: 10.10.0.0/16

VLAN 110: Public Network (per site)
  - Manchester: 203.0.113.0/24  (documentation range - use yours)
```

```
49    - London:     198.51.100.0/24 (documentation range - use yours)
50    - Leeds:      10.250.0.0/24   (dev/test - private, NAT to internet)
```

A word on why management is one /16 rather than three unrelated /24s: when the Manchester management server needs to reach a Leeds KVM host across the VPN, a single summarised supernet keeps the routing tables and firewall rules legible. You will thank yourself during the 3am incident.

Cross-Site Connectivity

Sites connect via:

1. **Primary:** Dark fibre between Manchester <-> London (low latency)
2. **Backup:** Encrypted VPN tunnels (all sites)
3. **Management:** Dedicated VPN for management traffic

```
Manchester                London                  Leeds
    |                        |                     |
    +--Dark Fibre (5ms)------+                     |
    |                        |                     |
    +------VPN (encrypted)---+---------------------+
```

Hardware Specifications

Production Compute Nodes (KVM Hypervisors)

Model: Dell PowerEdge R660 (or HPE ProLiant DL360 Gen11 equivalent)

Specifications:

Component	Specification	Notes
CPU	2x Intel Xeon Gold 6430 (32 cores each)	64 cores, 128 threads total
RAM	512GB DDR5 ECC (8x 64GB)	Expandable to 2TB
Boot	2x 960GB M.2 NVMe, RAID-1 via BOSS-N1	Dedicated boot, no shared storage
Network	6x Intel E810-XXV 25GbE	Proper traffic separation
IPMI	Dell iDRAC9 Enterprise	Out-of-band management (the Mars test)
Power	Dual redundant PSU (1400W), A+B feed	No single power failure point
OS	Ubuntu 24.04 LTS Server	Long-term support

Network Port Allocation:

```
Ports 1-2 (ens1f0/ens1f1): bond-mgmt  - Management (active-backup)
Ports 3-4 (ens2f0/ens2f1): bond0     - Ceph public / librbd (LACP)
Ports 5-6 (ens3f0/ens3f1): bond1     - Guest/Public trunk (LACP)
```

GPU/AI Compute Nodes

Model: Dell PowerEdge R7525 — refurbished

This is the Return on Research decision of the chapter. The R7525 is a 2U, two-socket AMD EPYC platform with 32 DIMM slots of affordable DDR4 and PCIe Gen4 throughout — and the refurbished market is full of them at a fraction of new-platform pricing. The 2U chassis is the point: proper airflow and physical space for full-height GPUs, which a 1U pizza box cannot offer.

Specifications:

Component	Specification	Notes
CPU	2x AMD EPYC 7543 (32 cores each)	64 cores, 128 threads — same per-host contribution as the R660s
RAM	512GB DDR4-3200 ECC (16x 32GB)	16 slots still free for cheap expansion
Boot	2x 960GB M.2 NVMe, RAID-1 via BOSS	
GPU	2x NVIDIA RTX 4060 Ti 16GB	On 2 of the 4 cluster hosts initially
PCIe	Gen4 x16 slots	Two full-height GPU slots free per node
Network	6x 25GbE	Same bond layout as the R660s
Power	Dual 1400W PSU, A+B feed	Two 4060 Tis add barely 330W at full load
OS	Ubuntu 24.04 LTS Server	

Why consumer cards in a datacentre? Because the workload is LLM inference for a publishing business, not training runs for a frontier lab. A 16GB RTX 4060 Ti serves quantised local models all day at around 165W, costs hundreds of pounds rather than thousands, and idles at single-digit watts. Two per host gives four GPUs across the initial pair of fitted nodes. This is the same silicon running the lab pipeline that builds this book — written from experience, not a brochure. When a workload genuinely needs more VRAM, the upgrade path is known: NVIDIA's L4 (24GB, 72W, single-slot) for inference density, or the L40S (48GB) where a model simply will not fit smaller. We originally specced multi-GPU A16 boards here; the arithmetic of 2026 retired that idea — the consumer-card-first approach delivers the same services for a tenth of the silicon spend.

A cluster homogeneity note (this matters): the GPU cluster is all-AMD (4× R7525) and the production cluster is all-Intel (4× R660). CloudStack live migration works *within* a cluster, and mixing CPU vendors inside one cluster is how you turn a maintenance window into an incident — a running VM cannot live-migrate between Intel and AMD hosts. Keeping each cluster homogeneous preserves live migration where it matters, and moving a workload between the clusters is a planned stop/start, not an emergency.

Development Compute Nodes (Leeds)

Model: Dell PowerEdge R740 — the existing estate, reused

Component	Specification
CPU	2x Intel Xeon Gold 6254 (18 cores each, 36 total)
RAM	384GB DDR4 ECC
Boot	2x 480GB M.2 SATA (BOSS RAID-1)
Network	4x 25GbE Intel XXV710 (2x mgmt/storage bond, 2x guest bond)
OS	Ubuntu 24.04 LTS Server

New cost: zero. These are the idle servers from Chapter 1, back at work.

Storage Nodes (Ceph, deployed with cephadm)

Model: Dell PowerEdge R760

Specifications:

Component	Specification	Notes
CPU	2x Intel Xeon Gold 6430 (32 cores each)	64 cores — sized for the OSDs you will *end* with, not start with
RAM	256GB DDR5 ECC	Ceph wants ~4-6GB per NVMe OSD plus headroom
Boot	2x 480GB M.2 NVMe, RAID-1 (BOSS-N1)	Simple boot drives
Data	4x 3.84TB NVMe 2.5“ U.2 to start	Samsung PM9A3 / Intel D7-P5520 class; backplane takes up to 16
Network	4x Intel E810-XXV 25GbE + 2x 1GbE LOM	Cluster bond, public bond, management
IPMI	iDRAC9 Enterprise	OOB management
Power	Dual redundant PSU, A+B feed	
OS	Ubuntu 24.04 LTS + cephadm	Containerised Ceph, official orchestrator

Why size the CPU for the end state? A four-OSD node would run happily on far less silicon. But these chassis will not stay at four drives: every drive you add is another OSD daemon, and the working rule for NVMe OSDs is roughly **2–4 physical cores and 4–6GB of RAM per OSD** if you want the performance the drives are capable of. Dual Gold 6430s (64 cores) keep a node honest at 12–16 OSDs. Buying the cores on day one and the drives as you need them is the right way round — in 2026's storage market, emphatically so. (Could you go Platinum? You could. Ceph would not notice. OSDs want cores, memory channels, and PCIe lanes — not top-bin clock speeds at a top-bin price. Return on Research says Gold.)

And the counterweight — don't build monsters. In a six-node cluster, a node loaded to 16 OSDs is a sixth of your entire cluster in one failure domain; lose it and Ceph must re-replicate everything

it held, hammering the cluster for hours. The honest growth rule: take nodes to around 8–12 OSDs, **then add nodes rather than drives**. The R760's spare bays are headroom, not a target.

NVMe Configuration:

- **Production clusters (MAN, LON):** 4x 3.84TB NVMe per node to start (high endurance)
- **Backup cluster (MAN):** 4x 3.84TB NVMe per node (standard endurance, lower cost)
- **Dev/Test (Leeds):** 4x 1.92TB NVMe per node (smaller, cost-effective)

Storage Capacity (Manchester Production):

```
6 nodes x 4 drives x 3.84TB = ~92TB raw
With 3x replication: ~30TB usable
With overhead (~15%): ~25TB effective storage

Pools (as established in Chapter 1):
  rbd-prod:   3x replication - production VMs
  rbd-dev:    2x replication - dev/test workloads
  ec-backup:  erasure coded 4+2 - backup data (capacity-efficient)

Enough for:
- Hundreds of VMs
- Snapshots and templates
- Room for growth - before a single extra drive is purchased
```

Management Servers

Model: Dell PowerEdge R450 class

Component	Specification
CPU	1x Intel Xeon (8-16 cores)
RAM	128GB DDR5 ECC
Storage	2x 960GB NVMe, RAID-1
Network	4x 10GbE (bonded pairs)
IPMI	iDRAC9
OS	Ubuntu 24.04 LTS Server

Deployment:

- 2x in Manchester (HA pair)
- 1x in London (geographic redundancy — survives a Manchester outage)
- Fronted by HAProxy; MySQL Galera multi-master across all three (Chapter 5)

Support Infrastructure

DNS Servers:

- VMs: 4 vCPU, 8GB RAM, Bind9 on Ubuntu 24.04
- Primary in Manchester, secondary in London

Secondary Storage (per zone):

- **Recommended build: 2x Dell PowerEdge R740XD (refurbished), DRBD active/passive pair with Pacemaker**
- 12-18x 3.5“ bays per chassis – secondary storage wants capacity, not IOPS; big, cheap, redundant drives are exactly right for templates, ISOs and snapshot exports
- NFSv4.1 export to the zone; daily replication to object storage
- Compact alternative: 2x R650 with RAID-10 SSDs where rack space is tighter than budget
- **The destination (and where this book ends up):** Ceph RGW – the S3-compatible gateway already in the Ceph cluster – takes over as secondary storage and backup target, with S3 Object Lock for immutability (Chapters 7 and 16)

Monitoring Stack:

- Prometheus: 8 vCPU, 32GB RAM
- Grafana: 4 vCPU, 16GB RAM
- Node exporters on all hosts; CloudStack exporter
- Alert routing is covered below – and it is deliberately sovereign

Ansible Control Node:

- VM: 2 vCPU, 4GB RAM, Ubuntu 24.04 with Ansible
- Playbooks version-controlled in self-hosted GitLab
- SSH keys (ed25519) to all infrastructure

Deployment Overview per Site

Manchester Deployment

Total Hardware:

- 4x R660 (cluster MAN-Prod-01)
- 4x R7525 (cluster MAN-GPU-01; 2 GPU-fitted, 2 awaiting cards)
- 6x R760 Ceph production nodes
- 4x R760 Ceph backup nodes (separate cluster)
- 2x Management servers
- Dell OS10 fabric: 2x S5232F-ON spine + 4x S5248F-ON leaf
- 1x Firewall/router pair
- Support infrastructure (R740XD secondary storage pair + VMs)

Rack Layout:

```
Rack 1 (42U):
+- Network (top)
|  +- Leaf switch 1 (S5248F-ON)
|  +- Leaf switch 2 (S5248F-ON)
|  +- Spine 1 (S5232F-ON)
|  +- Firewall/router (1U)
|  +- PDUs (0U - vertical)
+- Management servers
|  +- Management 1 (1U)
|  +- Management 2 (1U)
+- Cluster MAN-Prod-01
   +- KVM-MAN-01 (R660, 1U)
   +- KVM-MAN-02 (R660, 1U)
   +- KVM-MAN-03 (R660, 1U)
   +- KVM-MAN-04 (R660, 1U)

Rack 2 (42U):
+- Network (top: leaf 3 + leaf 4, spine 2 - resilience)
+- Cluster MAN-GPU-01
|  +- KVM-MAN-05 (R7525, 2U, 2x RTX 4060 Ti)
|  +- KVM-MAN-06 (R7525, 2U, 2x RTX 4060 Ti)
|  +- KVM-MAN-07 (R7525, 2U)
|  +- KVM-MAN-08 (R7525, 2U)
+- Secondary storage (R740XD pair, 2x 2U)
+- Spare U for expansion

Rack 3 (42U):
+- Ceph Production Cluster
   +- CEPH-MAN-01 (R760, 2U)
   +- CEPH-MAN-02 (R760, 2U)
   +- CEPH-MAN-03 (R760, 2U)
   +- CEPH-MAN-04 (R760, 2U)
   +- CEPH-MAN-05 (R760, 2U)
   +- CEPH-MAN-06 (R760, 2U)

Rack 4 (42U):
+- Ceph Backup Cluster (independent failure domain)
   +- CEPH-MAN-BAK-01 (R760, 2U)
   +- CEPH-MAN-BAK-02 (R760, 2U)
   +- CEPH-MAN-BAK-03 (R760, 2U)
   +- CEPH-MAN-BAK-04 (R760, 2U)
   +- Spare U for monitoring, etc.
```

Power:

- Dual-feed per rack (A+B); every server on both feeds
- 15kW per rack maximum
- UPS backed, generator backed — and remember Chapter 1's point about real N+1: dual PSUs protect against a failed supply, not a failed grid

London Deployment

Total Hardware:

- 4x R660 (cluster LON-Prod-01)
- 4x R760 Ceph nodes
- 1x Management server
- 2x S5248F-ON leaf switches (Dell OS10)
- 1x Firewall/router
- Secondary storage pair

Rack Layout:

```
Rack 1 (42U):
+- Network equipment
+- Management server (1U)
+- KVM nodes (4x R660, 1U each)
+- Ceph nodes (4x R760, 2U each)
+- Secondary storage (2x 2U)
+- Spare U
```

Leeds Deployment

Total Hardware:

- 2x R740 (cluster LEE-Dev-01 – existing servers, reused)
- 3x R760 Ceph nodes (1.92TB drives)
- 2x S5248F-ON leaf switches (Dell OS10; refurbished units are perfect here)
- Secondary storage (single NFS server is acceptable in dev)

Rack Layout:

```
Rack 1 (42U):
+- Network equipment
+- KVM nodes (2x R740, 2U each)
+- Ceph nodes (3x R760, 2U each)
+- Lots of spare U (room for growth)
```

CloudStack Zone Architecture

Zone Design

Each city is a CloudStack **Zone** (Advanced Networking), following the hierarchy from Chapter 1: Zone -> Pod -> Cluster -> Host.

```
CloudStack Region: UK
|
+- Zone: Manchester (MAN)
|  +- Pod: MAN-Pod-01
|  |  +- Cluster: MAN-Prod-01 (4x R660, Intel)
|  |  +- Cluster: MAN-GPU-01  (4x R7525, AMD + GPUs)
|  +- Primary Storage: Ceph RBD (rbd-prod, rbd-dev)
|  +- Secondary Storage: NFS (R740XD pair) -> RGW later
|
+- Zone: London (LON)
|  +- Pod: LON-Pod-01
|  |  +- Cluster: LON-Prod-01 (4x R660, Intel)
|  +- Primary Storage: Ceph RBD
|  +- Secondary Storage: NFS pair
|
+- Zone: Leeds (LEE)
   +- Pod: LEE-Pod-01
   |  +- Cluster: LEE-Dev-01 (2x R740, Intel)
   +- Primary Storage: Ceph RBD (rbd-dev)
   +- Secondary Storage: NFS (single server)
```

Zone Configuration Summary:

Setting	Manchester	London	Leeds
Network Type	Advanced	Advanced	Advanced
Guest VLAN Range	300–399	300–399	300–399
Public IP Range	203.0.113.0/24	198.51.100.0/24	10.250.0.0/24 (NAT)
Hypervisor	KVM	KVM	KVM
Primary Storage	Ceph RBD	Ceph RBD	Ceph RBD
Secondary Storage	NFS (HA pair)	NFS (HA pair)	NFS

Primary storage is registered in CloudStack as RBD pools — `rbd-prod` for production zones, `rbd-dev` for Leeds — and QEMU talks to them over librbd as described above. Chapter 8 walks through the registration step by step.

Security Architecture

Defence in Depth

Six layers, each assuming the one above it has failed:

```
Layer 1: Physical    - Datacentre access controls, locked racks
Layer 2: Network     - Firewalls, VLAN isolation, no internet to mgmt
Layer 3: Hypervisor  - Minimal host OS, SSH keys only, no passwords
Layer 4: CloudStack  - RBAC, projects, domain isolation, API keys
Layer 5: Instance    - Security groups, host firewalls, hardened templates
Layer 6: Application - TLS everywhere, secrets management, audit logging
```

Access Tiers

Tier	Who	Access
Tier 0	Infrastructure leads (2 people)	iDRAC, switches, Ceph, root
Tier 1	Operations team	CloudStack admin, host SSH (sudo)
Tier 2	Developers	CloudStack domain admin (own projects)
Tier 3	Application users	Instance access only

All administrative access flows through a bastion host on the management network, with MFA and full session logging. The SIEM is self-hosted (Wazuh feeding a local index) – security telemetry about your infrastructure is exactly the kind of data a sovereignty-minded business keeps at home. Security is built up across the book – the network ACLs of Chapter 13 and the sovereign SIEM and operations of Chapter 15. Part 2 takes it further: role-based access control and multi-tenancy in Chapter 19, and compliance automation in Chapter 28.

AI and GPU Strategy

The Private AI Vision

Leaf Spine Books' editorial tooling – manuscript triage, metadata generation, semantic search across the back catalogue – runs on local language models. No customer manuscript ever leaves the building. That is both a GDPR position and a commercial one: authors are told, truthfully, that their unpublished work is never sent to a third-party AI service.

Phase 1: The GPU Cluster as Built

- **MAN-GPU-01:** 4x R7525; two hosts each carry 2x RTX 4060 Ti 16GB (4 GPUs total), two hosts run CPU-only awaiting cards
- **Serving stack:** Ollama on Ubuntu VMs with GPU passthrough, fronted by an internal API gateway
- **Models:** quantised 7B-14B class models for editorial tooling; embedding models for search
- **Growth:** adding a pair of cards to the remaining hosts is a purchase order measured in hundreds of pounds, made *when a queue actually forms* – not before

GPU Passthrough (VFIO)

CloudStack hands a whole physical GPU to a guest using KVM's VFIO passthrough — the VM owns the card, full performance, no sharing:

```
R7525 host                          Guest VM
+----------------------+           +------------------+
| IOMMU groups (AMD-Vi)|           | Ubuntu 24.04     |
|  GPU 1 --vfio-pci----+---------->| NVIDIA driver    |
|  GPU 2 --vfio-pci----+----+      | Ollama           |
+----------------------+    +----->| (second VM)      |
                                    +------------------+
```

Host configuration (the Ansible role automates this):

- Enable AMD-Vi/IOMMU in BIOS and kernel cmdline (`amd_iommu=on iommu=pt`)
- Bind the GPUs to `vfio-pci` at boot (driver blacklisting on the host)
- CloudStack GPU service offerings map instances to available cards

Chapter 10 builds this end to end, including Ollama service templates and the internal API.

The Scale-Up Path (When Earned)

Need	Card	Why
More inference density	NVIDIA L4 (24GB, 72W)	Single-slot, passively cooled, datacentre warranty
Bigger models	NVIDIA L40S (48GB)	When a model genuinely won't quantise into 16GB
Today's workload	RTX 4060 Ti 16GB	Hundreds of pounds; proven in our own lab

The point of starting with consumer cards is not that datacentre cards are bad — it is that buying them *before the workload exists* is the cloud-bill mistake in CapEx clothing. Earlier drafts of this design specced multi-GPU A16 boards for VDI-style sharing; the workloads never justified them. The R7525s give you the chassis, power, and PCIe lanes to take L4s or L40Ss later without replacing a single server.

Automation with Ansible

Repository Structure

Everything in this chapter deploys from version-controlled playbooks in the team's self-hosted GitLab:

```
cloudstack-infrastructure/
+- inventory/
|  +- manchester.yml
|  +- london.yml
|  +- leeds.yml
+- group_vars/
|  +- all.yml
|  +- kvm_hosts.yml
|  +- ceph_nodes.yml
|  +- gpu_hosts.yml
+- roles/
|  +- common/          # base OS, users, SSH, NTP, chrony
|  +- networking/      # netplan bonds, VLANs, MTU
|  +- kvm_host/        # KVM, libvirt, CloudStack agent
|  +- ceph_node/       # prerequisites + cephadm bootstrap
|  +- gpu_ollama/      # VFIO, drivers, Ollama serving
|  +- management/      # CloudStack mgmt, MySQL Galera, HAProxy
|  +- monitoring/      # Prometheus, Grafana, exporters
+- playbooks/
   +- site.yml
   +- deploy-kvm-hosts.yml
   +- deploy-ceph.yml
   +- deploy-management.yml
```

Inventory excerpt (Manchester):

```yaml
# inventory/manchester.yml
all:
  children:
    kvm_hosts:
      children:
        man_prod:
          hosts:
            kvm-man-[01:04].mgmt.leafspine.example:
        man_gpu:
          hosts:
            kvm-man-[05:08].mgmt.leafspine.example:
          vars:
            gpu_passthrough: true
    ceph_nodes:
      children:
        man_ceph_prod:
          hosts:
            ceph-man-[01:06].mgmt.leafspine.example:
        man_ceph_backup:
          hosts:
            ceph-man-bak-[01:04].mgmt.leafspine.example:
```

A note on the `ceph_node` role: it prepares hosts (packages, networking, time sync, container runtime) and then hands over to **cephadm**, Ceph's own orchestrator, to bootstrap and manage

the cluster. Ansible builds the runway; cephadm flies the plane. Resist the urge to have Ansible micromanage Ceph daemons — that way lies drift between what Ansible believes and what the orchestrator knows. Chapter 8 covers the bootstrap in full.

Every playbook is idempotent: run it twice, get the same result. That property — not the YAML syntax — is what makes automation trustworthy.

Cost Analysis

What the Build Costs

Headline figures for the three-site build (verified for this edition; your quotes will vary):

Item	Unit Cost	Notes
KVM compute host (R660-class, 2x Gold 6430, 512GB DDR5)	£11,500	New, configured
Ceph storage node (as built: 4x 3.84TB NVMe)	£15,000	Entry configuration; R760 platform headroom included
GPU cluster host (R7525 refurbished + 2x RTX 4060 Ti)	indicative	Refurbished market; lands well under a new R660

Site	Total CapEx
Manchester	£308,000
London	£132,000
Leeds	£66,000
Total	**£506,000**

How Manchester's £308,000 breaks down (canonical anchors in bold; remaining lines are indicative balancing items within the site total):

- **4x R660 production hosts: £46,000**
- **10x Ceph nodes (6 prod + 4 backup): £150,000**
- 4x R7525 refurbished GPU hosts + 4x RTX 4060 Ti: indicative — the consumer-card decision means this entire cluster costs a fraction of the original A16-based plan
- 2x management servers, secondary storage pair (refurbished R740XD), Dell OS10 fabric, firewalls, cabling, racking and contingency: the balance

OS10 edition — the licence line (VERIFY). The hardware CapEx above is identical to the canonical edition: the same Dell switches, hosts and storage, the same **£506,000** total. What this edition adds is the **Dell OS10 per-switch licence** — a real cost (perpetual or term, depending on your entitlement) on every leaf and spine across the three sites. No figure is printed here, on purpose: a credible OS10 licence price has to come from your own Dell quote, not an invented number, so treat it as **VERIFY — operator to confirm current OS10 licence cost with Dell** and add it to the platform total. Community SONiC carries no such NOS licence, which is exactly the trade this edition is honest about; staying on the canonical SONiC edition (`main`) keeps the £506,000 as the whole platform bill.

Leeds deserves a second look: £66,000 buys the storage nodes and network — the compute is the existing R740s at £0. That is the licensing escape made visible.

The Comparison That Matters: The VMware Escape

The build exists because of a renewal quote. The like-for-like licensing position this platform replaces:

- **£181,500 per year** in virtualisation licensing for this estate
- **£907,500 over five years** — for permission to run hypervisors on hardware already owned
- The CloudStack platform: **£506,000 once**, no per-core meter, no renewal cliff

That licensing figure is deliberately conservative: it reflects a well-discounted enterprise rate, re-verified against 2026 market pricing — at the published list price (USD 350 per core per year) the same 896-core estate would invoice at roughly £234,000 a year, so the escape is understated here, not oversold.

Five-year saving: £852,500 — a 94% reduction. The entire three-site CapEx is recovered in under three years of avoided licensing alone, and the hardware is still yours afterwards.

> That 94% is the **community-SONiC edition's** figure. On this OS10 edition, the Dell OS10 licence — the VERIFY line above — is an added platform cost that narrows the gap by an amount you will have to price from your own Dell quote. It does not change the *structure* of the escape: a one-time hardware CapEx plus a network-OS licence still replaces a perpetual per-core virtualisation meter, and the hardware is yours either way. We will not invent the OS10 figure to put a new percentage on the page — measure it against your quote, and against the £0-NOS-licence community-SONiC edition on `main`.

A cloud comparison, briefly and qualitatively: running this estate's workload profile (steady-state VMs, large storage footprint, GPU inference) on a hyperscaler would cost multiples of the owned platform every single year, with egress fees discouraging the exit. Renting is the right answer for spiky, uncertain workloads; this is neither. We will not pretend to a per-pound cloud comparison here — list prices change monthly and discounts are opaque — the structural argument is enough.

Operating costs (colocation, power, connectivity, support contracts) continue either way and are deliberately left qualitative here; Chapter 17 builds the full operating model. (On this edition, Dell's OS10 support is one of those support contracts — another reason the OS10 licence/support line is yours to price, not ours to invent.)

A Note on Buying Hardware in 2026

This budget was re-verified in mid-2026, and it is roughly a quarter higher than the same bill of materials priced in late 2025 — almost entirely memory and NVMe. The AI build-out has DRAM contract prices up over 60% quarter-on-quarter and far more year-on-year, and enterprise SSDs are on allocation at major vendors. Three practical consequences:

1. **Owning your refresh cycle is now a strategy.** When you own the platform, you choose *when* to buy. A licensing renewal date does not care about the DRAM spot price; your expansion plan can.

2. **The refurbished market is the value play of the decade.** DDR4 platforms like the R7525 and R740XD sidestep the worst of the DDR5 premium entirely.
3. **Buy storage chassis with empty bays.** Drive prices are volatile; PCIe lanes are not. The R760's spare bays are an option contract on future capacity.

4.23 Alignment: none for this section — economics are version-independent. But hold this page in mind when we discuss upgrade planning: the upgrade series will revisit these figures.

Migration Path

The journey from the legacy VMware estate to this design runs in four phases:

Phase 1 — Foundation (Chapters 4-8): Build Manchester: network fabric, Ceph cluster, management servers, first KVM cluster. Production-ready before a single workload moves.

Phase 2 — Pilot Workloads (Chapters 9-10): Migrate internal, low-risk services first. Stand up the GPU cluster and first Ollama services. Build operational confidence.

Phase 3 — Production Migration (Chapter 11): The publishing platform itself, customer-facing services, and the bulk of the VM estate, with rollback plans at every step.

Phase 4 — Multi-Zone & DR (Chapters 14, 16): London zone live and the Leeds dev environment handed to the developers (Chapter 14); replication established and DR tested — actually tested, because a DR plan that has never run is a hypothesis (Chapter 16).

Disaster Recovery Strategy

Replication

- **VM storage:** RBD snapshot-based replication from Manchester's production pools to London on a schedule tiered by workload criticality (15 minutes for the publishing platform, hourly for internal services)
- **Database:** application-level replication for the systems that support it (MySQL/PostgreSQL replicas in London) — storage replication is the safety net, not the first line
- **CloudStack:** management server in London is part of the same Galera cluster; the control plane survives Manchester's loss

Backup

Backups follow the rule from the story: a separate failure domain, always.

1. **First copy:** Manchester backup Ceph cluster (independent hardware, independent failure domain), written via the NAS backup provider today, RGW S3 as the strategic target
2. **Immutability:** S3 Object Lock on the RGW backup buckets — ransomware cannot delete what nothing can delete

3. **Offsite copy:** replication of backup buckets to London's RGW. The offsite copy stays on sovereign, owned infrastructure — no hyperscaler bucket required

4.23 Alignment: CloudStack 4.23 is set to bring a native KVM backup framework with Veeam Backup & Replication integration — first-class, orchestrated VM backups from within CloudStack itself. On 4.22.1.0, the equivalent today is the NAS backup provider plus Veeam's direct-to-object support targeting RGW. Chapter 16 covers both, and the upgrade series will revisit this the moment 4.23 ships.

Capacity Planning

Compute (day one):

- 14 hosts (8 MAN + 4 LON + 2 LEE), **896 physical cores**, 7,168GB RAM
- At sensible oversubscription, comfortably 300-400 production VMs with HA headroom — capacity planning that ignores HA headroom is planning for the good days only

Storage (day one):

- ~75TB usable across the three production/dev clusters at design replication levels
- Growth path: drives into empty bays first (to ~8-12 OSDs per node), then nodes

Network:

- 25GbE access layer with 100GbE spine capacity in Manchester; the fabric outlives several server generations

GPU:

- 4 cards day one; the next 4 are a small purchase order away; the chassis for them are already racked

Monitoring and Alerting

Prometheus scrapes everything (node exporters, CloudStack exporter, Ceph's built-in Prometheus module); Grafana dashboards per audience (executive, operations, capacity). Alert routing is sovereign by design:

- **Alertmanager -> n8n** (self-hosted workflow engine) for routing, deduplication and escalation logic
- **Critical:** Slack on-call channel + phone escalation via the n8n workflow
- **Warning:** GitLab issues, auto-created and assigned — warnings become work items, not noise
- **Incident records:** PostgreSQL, on-platform — the incident history that trains next year's capacity plan belongs to you

No third-party SaaS pager service holds the keys to knowing whether your infrastructure is healthy. (We ran one for years; the subscription outlived its value the day n8n learned to make phone calls.)

Documentation Standards

Every component carries documentation as code, in the same GitLab as the playbooks:

- **Architecture decision records (ADRs)** – why, not just what; this chapter is effectively ADR-001 to ADR-020
- **Runbooks** – every alert links to one; an alert without a runbook is homework assigned at 3am
- **Network diagrams** – generated from source of truth where possible, reviewed quarterly
- **Change log** – every production change is an MR; the diff is the documentation

Next Steps

With the blueprint agreed, the build begins. Chapter 3 sets up the lab environment where every design in this chapter gets rehearsed before real money is spent – the EVE-NG network lab, nested KVM, and a complete miniature of the three-city cloud on a single workstation.

Chapter Recap: 5 Key Takeaways

1. **Design from principles, not products.** HA-first, automation-everything, and the Mars test drove every hardware and network choice – the specific Dell model numbers are consequences, not starting points.
2. **Separate traffic types ruthlessly.** Management, Ceph cluster, Ceph public, and guest traffic each get their own bonded interfaces and VLANs. Most "storage performance problems" are network design problems wearing a disguise.
3. **Size storage CPUs for the end state; size clusters for the failure domain.** 2-4 cores and 4-6GB per NVMe OSD, Gold not Platinum, and grow out (nodes) rather than only up (drives).
4. **Buy GPUs for the workload you have.** Refurbished R7525s with consumer RTX cards serve a publishing business's private AI for a fraction of datacentre-GPU pricing, with a known scale-up path when – if – it is earned.

(On this OS10 edition, add the Dell OS10 per-switch licence – the VERIFY line in the cost analysis – which narrows the gap by an amount you price from Dell, without changing the structural escape.)

Up Next

Chapter 3: Building the Lab – before Leaf Spine Books racks a single server, the entire three-city design gets built in miniature: EVE-NG, nested virtualisation, and an OS10 fabric you can run on a workstation.

References

- Apache CloudStack 4.22 Documentation: https://docs.cloudstack.apache.org/
- Ceph Documentation (Reef/Squid, cephadm): https://docs.ceph.com/

- Dell SmartFabric OS10 (this edition's fabric NOS) — OS10 User Guide and VXLAN/BGP-EVPN guide for the pinned 10.6.x build, via Dell Support / InfoHub: https://www.dell.com/support/
- SONiC (Software for Open Networking in the Cloud — the canonical edition's NOS, your stay-on-`main` option): https://sonicfoundation.dev/
- Ubuntu 24.04 LTS Server Guide: https://ubuntu.com/server/docs
- Ansible Documentation: https://docs.ansible.com/

* * *

Apache CloudStack Deployment Guide — Leaf Spine Books edition Copyright © 2026 Michael John Leslie Hinsley. Code examples licensed under Apache License 2.0. Contact details for this book are on the contact page at the back.

Chapter 3: Building Your Learning Lab

From One Second-Hand Server to a Cloud You Can Break

This chapter is written for system administrators, DevOps engineers and infrastructure architects. It assumes Chapters 1-2, basic Linux knowledge and an understanding of virtualisation concepts. Allow one to two weeks for the Phase 1 lab; Phases 2-3 are optional extensions.

A note for this edition. The companion deployment guide's Phase 2 is the five lab networks — it carries no network-simulation procedures. The network-simulation lab edition lives in Appendix B, written against community `sonic-vs`. Until an OS10 pass is made on that appendix, treat its *steps* as the canonical-edition reference and substitute the OS10 commands this book establishes in Chapters 12 and 40b; the *thinking* in this chapter is already OS10.

How to read this chapter. This chapter is deliberately self-contained. If you have just picked up the book and want to know what you are getting yourself into — what it costs, what it teaches, and what you will be able to do at the end — you can start here without having read a page of the migration story. When you are ready to type rather than read, the step-by-step procedures live in the companion **deployment guide** (`chapter-03-deployment-guide.md`), with hands-on **exercises** (`chapter-03-exercises.md`), a **troubleshooting guide** (`chapter-03-troubleshooting.md`) and a **VM quick reference** (`chapter-03-vm-quick-reference.md`) alongside it. Keep this file for the thinking; keep those open while you build.

Why a Lab at All

I have a test I apply to every system I design: *if this server was on Mars, could I fix it?* It is a question about autonomy — about whether the architecture can survive without a human hand on it. But there is an honest prerequisite to the Mars test that gets skipped too often: **you do not learn to fix things on Mars.** You learn on a bench, in a shed, on hardware you are allowed to break.

Nobody sensible practises live migration for the first time on the cluster that runs payroll. Nobody should discover what a Ceph monitor quorum loss feels like on the storage that holds the customer database. The entire point of a learning lab is to make every mistake this book can provoke — and it will provoke plenty — somewhere that recovery costs you an evening, not a career.

There is a second reason, and it is financial. Chapter 2 designed Leaf Spine Books a three-city estate with a hardware bill of roughly **£506,000**, justified against a VMware renewal running to about **£181,500 a year**. (This edition adds one line to that estate — the OS10 per-switch licence priced in Chapter 12 — but it changes neither the order of magnitude nor the argument.) Those numbers only make sense if the platform actually works — and the cheapest possible way to prove that is a single second-hand server costing less than a week of the VMware bill they are escaping. Validation is not a luxury. It is the difference between a confident purchase order and a £506,000 guess.

And a third, quieter reason: **the lab never stops earning.** Long after production is live, this same machine tests upgrades before they touch real workloads, rehearses recovery procedures, demos the platform to clients, and onboards new engineers. Idle silicon is waste; a lab is how a single box keeps paying rent for years. Return on Research, in one chassis.

The Story: Planning for the Future

Leaf Spine Books HQ, Manchester – Week 1, Friday afternoon

Sarah Martinez spread the architecture diagrams across the conference table and weighted the corners down with two paperback proofs and a coffee mug – a company that ships pallets of paper for a living never lacks for paperweights. The team had spent the week absorbing CloudStack concepts and sketching the three-city design, and the conclusion was unavoidable: before anyone signed a purchase order, they needed to actually run this thing.

"We need a learning environment," Marcus Chen said, "but I don't want to build something we'll throw away. Let's design a lab that can evolve."

"Agreed," said Priya Patel, pulling up her cost sheet. "A single used Dell R740 – the same model as the two we're putting back to work in Leeds – is about three thousand pounds on the refurbished market. One box, nested virtualisation, and we can run the whole platform in miniature: management server, Ceph storage, compute hosts, the lot."

James Wilson, who had been quietly studying the network diagram, looked up. "Can it rehearse the network design too? The leaf-spine fabric is the part I can't un-buy if we get it wrong."

"That's phase two," Priya said. "The switch operating system we've chosen for this build – Dell's SmartFabric OS10 – ships as a virtual node, and it's the *real firmware*, not a cut-down sim. We can stand up the whole fabric in software, prove VLT and EVPN behave, and break it on purpose before a single switch is ordered. It's licensed, mind – a 120-day trial out of the box, or our own entitlement if we're running Dell kit – so we treat the trial clock as part of the plan."

Sarah smiled. "So we learn the platform on hardware we already understand, we validate the network in software on the firmware we'll actually run, and the first time we touch the production budget, nothing in it is a surprise. Design the complete evolution now, even if we build it in phases."

The challenge: build a learning environment that –

- Requires minimal upfront investment (thousands, not the £506k estate)
- Teaches production-grade skills that transfer directly to Chapters 4-18
- Rehearses the *actual* Leaf Spine design in miniature – same versions, same storage, same network scheme
- Validates the network architecture before switch hardware is purchased
- Evolves rather than gets thrown away
- Extends to edge computing and IoT when the time comes

The Three-Phase Roadmap

Leaf Spine Books adopted a three-phase lab strategy. Each phase has its own budget, its own outcome, and its own decision it de-risks.

```
Phase 1: Learning Lab (Weeks 1-8)
+- Hardware: one used Dell PowerEdge R740
+- Technology: Ubuntu 24.04 LTS, nested KVM, CloudStack 4.22.1.0,
|              three-node cephadm Ceph cluster (RBD via librbd)
+- Outcome: a complete working CloudStack zone you built yourself
+- De-risks: the entire platform decision

Phase 2: Network Simulation (Weeks 9-12)
+- Hardware: none new - runs on the same server
|            (OS10 virtual image required: 120-day trial or owned entitlement)
+- Technology: Dell OS10 virtual nodes (real firmware) under EVE-NG or
|              libvirt, leaf-spine topology, VLANs + VLT + EVPN end to end
+- Outcome: the production network design, proven on the real firmware
+- De-risks: the switch fabric purchase

Phase 3: Edge Computing (Weeks 13-16, optional)
+- Hardware: Raspberry Pi 5 + ESP32 sensors (~£400-450)
+- Technology: edge patterns, IoT integration, LoRaWAN
+- Outcome: extended-cloud concepts validated
+- De-risks: the innovation roadmap (and previews Chapter 40)
```

All lab pricing in this chapter is **indicative UK street pricing for used/refurbished kit at the time of writing (June 2026)** – the second-hand market moves, so treat the numbers as a budget envelope rather than a quote. The production figures, by contrast, are the canonical Chapter 2 numbers and do not change between chapters.

Phase	Indicative spend (lab-tier)	What it de-risks
1 – Learning lab	~£2,950-3,450	The £506,000 estate design
2 – Network simulation	£0 on the 120-day trial (reuses Phase 1 server)	The leaf-spine switch fabric
3 – Edge computing	~£400-450	The IoT/edge innovation strategy

Spending three and a half thousand pounds to validate half a million is not caution. It is just arithmetic. Phase 2 stays at zero cash on the trial licence; the only new requirement over the canonical edition is obtaining the OS10 virtual image (trial or entitlement), covered below.

Phase 1: The Learning Lab

One Server, a Whole Cloud

The learning lab runs a complete CloudStack deployment – management server, database, storage cluster, compute hosts and guest VMs – as virtual machines nested inside a single physical server. Nested virtualisation means the KVM hypervisor runs *inside* a KVM virtual machine: the physical host virtualises "servers", and those servers virtualise the guests. It costs some performance, and it buys you an entire datacentre on one power lead.

What makes this lab worth your time is that it is not a generic CloudStack sandbox. It is **the Leaf Spine Books design in miniature**:

- The same software, version for version: **CloudStack 4.22.1.0 LTS** on **Ubuntu 24.04 LTS**
- The same storage: a real **Ceph** cluster, bootstrapped with **cephadm**, serving RBD to KVM through **librbd** – three nodes instead of six, gigabytes instead of terabytes, but the same commands, the same dashboard, the same failure modes
- The same network scheme: the production VLAN plan from Chapter 2, scaled down but not dumbed down
- The same names: when you create your first cluster and call it **MAN-Prod-01**, you are rehearsing the exact object you will create in production in Chapter 9

Skills transfer is the whole game. Every hour spent in this lab is an hour you will not spend panicking in Chapter 10.

The Hero of This Chapter: a Used Dell R740

Chapter 1 made a quiet point that deserves repeating loudly: the two servers Leaf Spine Books put back to work in Leeds – cluster **LEE-Dev-01** – are Dell R740s that the old licensing model had priced into idleness. The lab buys a third one, second-hand, and the symmetry is deliberate. The machine that teaches you the platform is the same machine the story rescues.

Why the R740 specifically?

- Excellent value on the refurbished market (lab-tier, ~£2,500-3,500 at the time of writing)
- Two-socket Xeon Scalable platform with huge RAM ceilings – ideal for nested virtualisation
- Abundant supply, spares everywhere, every fault already documented by someone else
- Proper enterprise plumbing: iDRAC9 remote management, redundant PSUs, BOSS boot mirroring – a lab server without out-of-band management fails the Mars test before it is even switched on
- It is the buy-refurbished argument from Chapter 2, practised before it is preached

Specifications

Minimum (budget) – runs the full Phase 1 lab:

Component	Specification	Notes
CPU	2x Intel Xeon Gold (16+ cores each, 32+ total)	Nested virtualisation needs cores
RAM	**192GB** DDR4 ECC	The honest floor for the full miniature estate
Boot	2x 480GB SATA SSD (BOSS RAID-1)	Host OS
Data	2-4x ~1TB NVMe or SAS SSD	VM disks + virtual Ceph OSDs
Network	4x 1GbE on-board (10GbE optional)	The lab network is virtual; physical NICs barely matter

Component	Specification	Notes
Management	iDRAC9	Treat it as compulsory – it will save you a drive to the shed
Indicative cost	~£2,500-3,000 used	Lab-tier street pricing, June 2026

Recommended – the comfortable experience:

Component	Specification	Notes
CPU	2x Intel Xeon Gold (16-18 cores each)	Headroom for Phase 2 running alongside
RAM	**256-384GB** DDR4 ECC	Guests stop fighting for memory
Data	4x 1.6TB NVMe	Snapshots become guilt-free
Indicative cost	~£3,000-3,500 used	Lab-tier street pricing, June 2026

A fair warning carried over from Chapter 2: 2026 is a hard year to buy memory and flash, and second-hand prices have firmed with everything else. If the numbers above have drifted by the time you read this, the *shape* of the advice has not – buy the RAM, economise on everything else. Nested labs live and die on memory.

If a Rack Server Is Not an Option

Not everyone has space, power or tolerance for a 2U server that sounds like a hairdryer with ambitions – and nobody should be priced out of this book. Two companion guides in this chapter's file set cover genuinely different platforms, not cut-down compromises:

- **The Ampere ARM guide** (`chapter-03-ampere-arm-cloudstack-complete-guide.md`) builds the lab on efficient ARM64 silicon – the same architecture this book's own examples are tested against, and a serious look at where dense, low-watt cloud hosting is heading.
- **The laptop guide** (`chapter-03-laptop-cloudstack-deployment-guide.md`) builds a portable lab on a single high-spec laptop – the version of this environment that fits in a rucksack and demos CloudStack on a client's desk.

Appendix B covers a portable variant in further depth (its OS10-edition pass carries the same Option A network-lab note as this chapter). Pick the platform that matches your budget and your electricity bill; the CloudStack you learn is identical.

The One Setting Everything Depends On

Nested virtualisation must be enabled in the host's KVM module, or every hypervisor VM you build will refuse to start its own guests. The deployment guide walks the full host build; this is the single check worth knowing by heart:

```
cat /sys/module/kvm_intel/parameters/nested   # Intel: expect Y
cat /sys/module/kvm_amd/parameters/nested     # AMD:   expect 1
```

If the answer is anything else, the deployment guide's first page fixes it before anything is installed on top.

The Miniature Estate

Phase 1 builds **Manchester in miniature**: one complete CloudStack zone with the same moving parts as the production design — scaled to fit comfortably inside 192GB of RAM. London and Leeds come later as an exercise (the multi-zone expansion belongs with Chapter 14, and the exercises file sets it as a challenge); resisting the urge to cram three zones into one box on day one is itself a design lesson. Build one zone properly, then replicate.

The Nested Layout

```
Physical host: Dell PowerEdge R740
+- Ubuntu 24.04 LTS + KVM (nested virtualisation enabled)

   Nested virtual machines - the Manchester zone in miniature:

   +- mgmt-01      CloudStack Management Server + MySQL
   |               4 vCPU, 16GB RAM, 100GB disk
   |               (production splits these and clusters both -
   |                Chapters 5-6 - the lab co-locates them honestly)
   |
   +- ceph-01     -+
   +- ceph-02      +- cephadm-deployed Ceph cluster
   +- ceph-03     -+  4 vCPU, 8GB RAM each
   |               + 2x 50GB virtual OSD disks per node (6 OSDs)
   |               MONs, MGRs and dashboard exactly as production
   |
   +- kvm-01      -+  Nested KVM compute hosts
   +- kvm-02       +  -> CloudStack cluster MAN-Prod-01
   |               8 vCPU, 32GB RAM each
   |               Run the System VMs and your guest workloads
   |
   +- nfs-01       Secondary storage (NFS)
                   2 vCPU, 4GB RAM, 200GB disk
                   (Chapter 7 upgrades this thinking to Ceph RGW)

   Resource budget:
   +- vCPU: ~34 allocated (oversubscription on 32+ cores is fine -
   |        these VMs do not all work hard at once)
   +- RAM:  ~108GB allocated -> fits 192GB with host headroom
   +- Disk: ~800GB -> fits 2TB of data storage with snapshot room
```

That budget is deliberately honest. Earlier drafts of this lab tried to nest all three zones and demanded 384GB of RAM and a spreadsheet to track what was running; the result taught exhaustion more than CloudStack. Seven VMs teach everything the platform has to offer, and your evenings survive. And because the entire estate is a set of files on one host, snapshot before every experiment, do something reckless, then roll back in minutes — break-and-rebuild is not a risk in this lab; it is the syllabus.

The sizes printed above are the canonical 192GB profile, and they are one row of a three-profile matrix: the laptop guide carries the 64GB diet, and a 32GB profile trims further still while keeping **both** compute hosts — a deliberate change from earlier drafts of this advice, which sacrificed the second host first and lost the two-host lessons with it.

Two truths worth stating plainly. First, the GPU cluster (**MAN-GPU-01** in the production design) is *not* rehearsed here — GPU passthrough into nested VMs is a science project, not a learning path, so the lab treats GPU hosts as ordinary compute and Chapter 2's GPU story stays theoretical until real silicon arrives. Second, nested performance is fine for learning and useless for benchmarking. Never quote lab numbers in a capacity plan.

The Network: Production's Scheme, Scaled Down

This is where the lab earns the phrase *design in miniature*. Chapter 2's VLAN plan carries over **with the same VLAN IDs and the same address families**, implemented as libvirt virtual networks on the host. When you later configure the real thing, nothing about the scheme will be new — only the cable in your hand.

Network	VLAN	Lab subnet	Mirrors production
Management	100	10.100.99.0/24	10.100.0.0/16 management space
Storage frontend	200	10.200.99.0/24	Ceph client traffic, MTU 9000
Storage backend	201	10.201.99.0/24	Ceph replication, MTU 9000
Guest	500-509	CloudStack-managed	Production guest VLANs 500-999
Public (simulated)	999	203.0.113.0/28	Documentation range — substitute your own

The `.99` third octet marks lab addressing at a glance — anything you ever see in a production log with `.99` in that position has escaped from somewhere it should not have. Small conventions like this are free, and they save real confusion later.

Ceph's own documentation calls this pair the *public* and *cluster* networks — VLANs 200 and 201 are that split, named — and the separation pays for itself the first time client load and replication compete. Jumbo frames (MTU 9000) are configured on the two storage networks exactly as production demands, because the *failure mode* of a mismatched MTU — storage that mostly works, then stalls under load with nothing useful in the logs — is one of the most valuable things this lab can teach you. The troubleshooting guide stages that exact fault on purpose.

The Storage: Real Ceph, Tiny Diet

The storage cluster is the part of this lab most worth lingering over, because it is **not a simulation**. The three `ceph-0x` VMs run genuine Ceph, deployed with the genuine orchestrator, in precisely the way Chapters 7 and 8 deploy the production cluster:

```
# The same bootstrap production uses (deployment guide, Phase 5):
cephadm bootstrap --mon-ip 10.100.99.21
```

One command, and you have monitors, managers, the dashboard and a working cluster – then OSDs join, pools are created, and KVM consumes RBD images through librbd with no translation layer in between. CloudStack is then given a **dedicated cephx client** for its pools – never `client.admin`. The platform gets the keys to its pool, not to the kingdom, and Chapter 8 holds exactly that line in production. The lab creates the same pools as production, at toy scale:

- `rbd-prod` – 3x replication, primary storage for the MAN-Prod-01 cluster
- `rbd-dev` – 2x replication, the capacity-versus-protection trade-off made visible

A Squid-era trap worth knowing before you bootstrap. On Ceph Squid 19.2.x the mon-scope `public_network` must name **both** public-facing planes – the management subnet and the storage frontend – or the `OSD_UNREACHABLE` health check condemns a perfectly healthy cluster: all OSDs up and in, real I/O flowing, and HEALTH_ERR on the board, because the check judges OSD public addresses against the mon-scope setting alone. Set both planes before any OSD exists; the deployment guide's §5.4 carries the exact command and the remediation ladder for a cluster already caught out.

Six 50GB virtual OSDs will never trouble a benchmark, but they will show you a placement-group warning, an OSD marked out, a rebalance in progress and a monitor election – the events that matter at 3am, rehearsed at a scale where they cost nothing. When Chapter 8 bootstraps the production cluster, it should feel like muscle memory.

Phase 2: Network Simulation

The most expensive mistake available to Leaf Spine Books is not a server – servers can be repurposed. It is the switch fabric: a leaf-spine design built on Dell open-networking switches running OS10, where a wrong assumption about VLANs, routing or failover gets bolted into hardware. Phase 2 exists so that every one of those assumptions is tested in software first.

Dell SmartFabric OS10 – the network operating system this edition's production fabric runs – is available as **OS10 virtual**, the real firmware packaged to run as an ordinary virtual machine. That it is the *real firmware* rather than a simplified image is the whole point of this edition's lab, and it changes what a network design can prove before a switch is bought:

- Build the *actual topology*: two spines, two leaves, as four OS10 virtual nodes on the lab server
- Configure the *actual* VLANs: 100, 200, 201 and the guest range, trunked exactly as Chapter 2 specifies

- Break the *actual failure cases*: pull a virtual link, watch traffic re-route, prove the redundancy story before it is purchased
- Prove the *actual resilience*: because it is real firmware, you can build the **VLT** leaf pair, dual-home a host, and watch it survive a leaf — and stand up **EVPN** — for real, not in description (the two seams the community `sonic-vs` lab has to defer)
- Practise the *actual workflow*: OS10's configuration model, learned where a typo costs a reboot rather than an outage

There is a trade, and this edition states it rather than hiding it. OS10 virtual is a **licensed image**: it runs on a **120-day trial** out of the box (the node reboots periodically once the trial lapses), or under your own **entitlement** if you own or work with Dell OS10 kit and can pull the image from Dell My Account. That is unlike community `sonic-vs`, which is free and runs indefinitely. So Phase 2 here costs no cash on the trial, but it does run a clock — plan the network-validation work as a focused block within the trial window, or run it on an entitled image. And if the licensing does not suit you, the canonical SONiC edition of this book builds the identical fabric on free, indefinite `sonic-vs`; nothing about the architecture you are validating differs between the two.

A point of honesty about tooling, because this book promised it: you can host the OS10 virtual nodes under plain **libvirt** wiring (sufficient, and the lightest path) or under an emulation environment such as **EVE-NG**, which this edition uses elsewhere (Chapters 12 and 40b) because living in network simulations all day rewards a polished topology designer. The licensed thing in this edition is the **OS10 image**, not the emulator — EVE-NG is a convenience, not a requirement. The distinction between OS10 virtual in the lab and **OS10 Enterprise under a Dell support contract** in production — where the support matters — is kept explicit whenever it appears. Chapters 12 and 22 go deep on the fabric itself.

Phase 3: Edge Computing

The third phase is optional, costs about as much as a single production NIC, and previews where this whole book ends up.

A **Raspberry Pi 5** with a handful of **ESP32 sensors** (~£400–450 all in) extends the lab beyond the rack: a tiny edge node reporting telemetry back to the CloudStack platform, sensors on links where WiFi does not reach, and **LoRaWAN** carrying readings across distances that would embarrass an access point. Leaf Spine Books frames this as their bookshop strategy — fifty retail sites with local caching and offline resilience, managed from the centre.

A sizing note: the ESP32s do the sensing, and a single-duty gateway sits comfortably on the 8GB board — but an edge cluster earns its keep by running several systems in parallel (broker, time-series database, dashboards, automation), and that is 16GB territory. The envelope above prices the 16GB board; specify down only when the node will genuinely do one job.

I frame it differently, because I run this pattern for real. The sensors in my world monitor beehives and an aquaponics greenhouse; the gateways are managed over MQTT; the edge nodes live outdoors, on solar power, where nobody reboots anything by hand. That is the Mars test passed in a field in the north of England — and it is the subject of Chapter 40, where a waterproof, solar-powered cluster goes inside a working beehive. Phase 3 is the gentle on-ramp to that finale. If the idea of your cloud platform ending inside a beehive does not make you want to read on, this may be the wrong book.

Done When: the Phase 1 Finish Line

Earlier drafts of this chapter ended with six informal bullets — UI up, zone Enabled, a guest on `rbd-prod`, a live migration under ping, an OSD killed and healed, a snapshot rolled back. Every one of them survives, but the finish line has since grown up: it is now a formal acceptance suite of **17 named checkpoints**, executed as **24 read-only checks**, so that "done" is something you prove rather than something you feel. The seventeen:

1. Nested virtualisation enabled on the host
2. The Phase 1 platform valid
3. Five networks active, storage bridges at MTU 9000
4. All seven VMs answering SSH with the right hostnames
5. The don't-fragment jumbo proof — 8972 bytes across the storage plane
6. The management UI answering (the admin-password-changed half is human-observed)
7. Ceph HEALTH_OK: three mons in quorum, six OSDs up and in, both pools initialised, the scoped client present, RBD answering
8. The secondary export visible
9. Both compute hosts prepared
10. Zone Manchester Enabled
11. Both system VMs Running
12. Both hosts Up
13. The built-in template Ready
14. The first guest Running, its ROOT volume attributably on `rbd-prod`
15. Live migration under ping
16. The OSD-kill drill — stop an OSD mid-write, watch the cluster heal
17. The snapshot/break/revert drill — the floor under every other experiment

Checkpoints 15 and 16 are *acts* rather than states — drills whose end-states fold back into checkpoints 7, 12 and 14 — which is why the read-only suite counts 24 checks rather than 17.

Not every profile runs all seventeen, and the exclusion is a design decision, not a shortfall. The canonical 192GB build and the 64GB laptop diet run the full 1-17. The 32GB profile runs 1-14 plus 16-17, with live migration (checkpoint 15) excluded by design — the motto is worth committing to memory: **"32GB demonstrates the platform; 64GB teaches it."**

And this is not aspiration. The complete suite passed 24 out of 24 on a real 32GB-class build on 2026-07-21 — the smallest profile this book prints, proven end to end.

The exercises file turns each checkpoint into a workout; the troubleshooting guide is your companion for the inevitable.

Evolution to Production

Nothing built in this lab is thrown away. The mapping from lab to the rest of the book is direct:

Lab component	Becomes	Chapters
Host preparation, KVM, networking	Production host builds	4
mgmt-01's co-located MySQL	A clustered, replicated database tier	5
mgmt-01's management server	HA management servers behind a load balancer	6
nfs-01 secondary storage	Production secondary storage, then Ceph RGW	7
The cephadm mini-cluster	The six-node Manchester production cluster	8
Zone creation, MAN-Prod-01	The real Manchester zone and clusters	9-10
First guest VMs	Production workloads and migration	11
OS10-virtual fabric rehearsal	The physical leaf-spine fabric (Dell OS10)	12, 22
Single-zone discipline	Multi-zone: LON-Prod-01, LEE-Dev-01	14
Pi 5 edge experiments	The Hive Data Centre	40

What does not transfer is worth naming too: performance numbers (nested anything is slow), and failure-domain behaviour — one power lead feeds your entire datacentre here, which is comedy, not resilience. Carry the skills forward; leave the benchmarks behind.

And when production is live, the lab's second career begins: every CloudStack upgrade, every Ceph release, every configuration change rehearses here first. The Mars test again — by the time a change touches hardware you cannot easily reach, it has already succeeded somewhere you can.

Key Takeaways

A lab is the prerequisite to the Mars test. Autonomous systems are designed by people who have already made their mistakes somewhere cheap.

Validate £506,000 with £3,500. The lab's entire budget is a rounding error against the estate it de-risks — and less than a week of the VMware bill being escaped.

Rehearse the real design, not a toy. Same CloudStack 4.22.1.0, same Ubuntu 24.04, same cephadm-deployed Ceph with librbd, same VLAN scheme, same cluster names. Skills transfer is the entire point.

One zone, built properly. Manchester in miniature first; multi-zone when Chapter 14 earns it. Seven VMs that you understand beat thirty that you babysit.

The storage is real. A three-node cephadm cluster at toy scale exhibits the same behaviours — and teaches the same recoveries — as the production cluster in Chapter 8.

Prove the network on the real firmware. OS10 virtual validates the leaf-spine design — VLT and EVPN included — on the same firmware you will deploy; it is licensed (a 120-day trial, or your own entitlement), so plan the work to the clock. The most irreversible purchase gets de-risked first, and if the licence does not suit you the canonical SONiC edition does the same job for free.

Buy refurbished without apology. A used R740 is the hero of this chapter for the same reason reused R740s anchor LEE-Dev-01: hardware priced into idleness by licensing is a bargain waiting for a better hypervisor.

The lab never retires. Upgrade testing, recovery rehearsal, demos and onboarding – Return on Research means this box earns its keep for years.

Coming in Chapter 4

The planning is done; the typing begins. Chapter 4 prepares the production infrastructure – host operating system builds, network bonding and VLANs, the KVM hypervisor stack, and the unglamorous pre-flight checks that make Chapter 9's deployment boring in the best possible way. Build the Phase 1 lab from the deployment guide first, and Chapter 4 will read like a tour of somewhere you already live.

Next chapter: Chapter 4: Infrastructure Preparation

Additional Resources

Official CloudStack documentation: https://docs.cloudstack.apache.org/ **Ceph documentation (cephadm):** https://docs.ceph.com/ **Dell SmartFabric OS10 documentation:** https://www.dell.com/support/ – OS10 User Guide, and Dell's VXLAN/BGP-EVPN Configuration Guide for the fabric **Community:** https://cloudstack.apache.org/mailing-lists.html – and the CloudStack Collaboration Conference, Edinburgh, November 2026

* * *

Apache CloudStack Deployment Guide – Leaf Spine Books edition Copyright © 2026 Michael John Leslie Hinsley. Code examples licensed under Apache License 2.0. Contact details for this book are on the contact page at the back.

Chapter 3: -1 Deployment Guide

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Building Manchester in Miniature — the Step-by-Step Companion

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

How to Use This Guide

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Page One: the Setting Everything Depends On

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Seven Phases

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Pre-Deployment Checklist

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Hardware

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Software (downloaded in Phase 1 — listed here so nothing surprises you)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Lab Address Plan (fixed — write it once, use it everywhere)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 1: Host Preparation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1.1 BIOS Settings

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1.2 Install Ubuntu 24.04 LTS Server

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1.3 Base Configuration

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1.4 KVM and libvirt

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1.3a Is Your Switch Port Tagged? Check Before You Build Anything

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1.3b The Trap: Your Switch Watches What You Build

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1.4a Arm an Undo Before You Change the Host's Network

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1.4b Turn Off libvirt's Default Network — Before Anything Else

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1.5 Storage Pools

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1.6 Fetch the Cloud Image

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1.7 Phase 1 Validation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 2: Virtual Networks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 3: Provisioning the Seven VMs

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

3.1 Common Post-Provision Configuration

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 4: The Management Plane

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

4.1 MySQL

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

4.2 CloudStack Management Server

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

4.3 Reaching the UI

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 5: The Storage Cluster

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

5.1 Prerequisites on All Three Nodes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

5.2 Bootstrap

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

5.3 Join ceph-02 and ceph-03

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

5.4 Set the Networks — Before Any OSD Exists

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

5.5 Create the OSDs

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

5.6 Pools

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

5.7 The Dedicated cephx Client

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

5.8 Prove It Works

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 6: Secondary Storage

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 7: Building the Manchester Zone

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

7.1 Prepare the Compute Hosts

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

7.2 The Zone Wizard

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

7.3 The First Guest VM

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

7.4 The Done-When Checklist, Executed

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

7.5 Living With It

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 3: -2 Lab Exercises

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

How to Use This File

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Challenge Projects

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Completion Criteria

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 3: -3 Troubleshooting Guide

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Diagnosis and Recovery for the Manchester Miniature

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

How to Use This Guide

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Method: Five Steps

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Section 1: First Boot — the Usual Suspects

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Section 2: Host Layer

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Section 3: Network Layer

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Section 4: Staged Drill One — the MTU Mismatch

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Section 5: Staged Drill Two — Losing Monitor Quorum

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Section 6: CloudStack Layer

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Section 7: Ceph Layer — the Lab-Scale WARN Taxonomy

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Section 8: The Diagnostic Sweep

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Section 9: Recovery — the Graduated Ladder

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Section 10: Getting Help

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 3: -4 VM Quick Reference

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The One-Page Card for the Manchester Miniature

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Estate

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Sizing at Three Weights

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Networks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

CloudStack Names (the production names, on purpose)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Done-When at a Glance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Where the Secrets Live

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Getting In

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Daily Commands

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Quick Fixes: Straight into the Troubleshooting Guide

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Chapter 3 File Set

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 3: -5 The Ampere ARM Build

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Manchester Miniature on ARM64 Silicon

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Why ARM at All

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Who This Build Is For

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Reference Platform

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The One Difference That Reshapes Everything

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Other ARM Deltas, Briefly

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Estate, Inverted

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Build Sequence

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Host Build (the ARM substitutions)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Infrastructure VMs (five, ARM64, UEFI)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Host Joins Its Own Lab

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Zone Creation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

A Second Host, Later

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The GPUs, Honestly

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Done When, on ARM

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Key Takeaways

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Chapter 3 File Set

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Additional Resources

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 3: -6 The Laptop Build

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Manchester Miniature in a Rucksack

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Why a Laptop at All

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Who This Build Is For

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Reference Platform

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The One Constraint That Reshapes Everything

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Estate, on a Diet

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Laptop Deltas

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

One Trip to the Firmware

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The One Setting Everything Depends On

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Storage: the Pool Moves Outside

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Mains Power, Thermals, and the Sound of Honesty

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Lid: This Platform's Signature Outage

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

WiFi Is a Fine Uplink

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Build Sequence

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Travelling With a Datacentre

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Network Simulation on the Move

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Switch OS Question

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Done When, on a Laptop

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Key Takeaways

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Chapter 3 File Set

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Additional Resources

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 4: Infrastructure Preparation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Building the Foundation — From Blueprint to Bare Metal

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Story: Hands on the Hardware

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 1: Understanding the Infrastructure Stack

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What We're Preparing — Manchester First

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 2: Ubuntu 24.04 LTS Installation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Why Ubuntu 24.04 LTS?

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Installation Method

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Disk Layout — Compute Nodes (R660 and R7525)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Disk Layout — Ceph Storage Nodes (R760)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Installer Walkthrough (Condensed)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 3: Post-Installation Baseline

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

System Update and Essential Packages

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Hostname and Hosts File

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Time Synchronisation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Arm the Undo — the habit that saves the afternoon

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 4: Network Configuration with Netplan

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Compute Node Netplan (kvm-man-01 shown)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Storage Node Netplan (ceph-man-01 shown)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Apply and Verify

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 5: Storage Preparation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Compute Nodes: Verify, Don't Add

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Ceph Nodes: The Discipline of Doing Nothing

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 6: System Optimisation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Kernel Parameters

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Transparent Huge Pages: Off

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

CPU Frequency: a Firmware Setting, Not an OS One

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 7: Security Hardening from First Boot

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

SSH: Keys Only

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Host Firewall: UFW

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Fail2Ban

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Unattended Security Updates

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

AppArmor: Leave It On

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 8: Logging Baseline

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 9: Documentation as Code

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 10: Verification

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Performance Baselines

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 11: Replicating to London and Leeds

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 4 — Key Takeaways

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1. The foundation layer is load-bearing

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

2. Three traffic types, three bonds, no sharing

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

3. Local disks are plumbing; the orchestrators own the rest

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

4. Security is a first-boot activity

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

5. Build one by hand, then never again

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 5: MySQL High Availability

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Database Layer — One Cluster, Two Cities, No Single Point of Failure

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Story: The Heart of the Cloud

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 1: Understanding the Database Layer

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What CloudStack Asks of Its Database

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What Galera Is

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The One Thing Galera Won't Do — and What It Means for CloudStack

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Choosing the Engine: MariaDB 10.11 LTS

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 2: The Design — Three Nodes, Two Cities

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Where the Database Lives

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Quorum Arithmetic, Taught Straight

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What London Costs Us

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 3: Preparing the Nodes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 1: Confirm the Foundation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 2: Hosts File Entries

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 3: Open Galera's Doors in UFW

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 4: Database Kernel Tuning

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 5: Resource Limits for the mysql User

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 4: Installing MariaDB and Galera

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 1: Install the Stack

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 2: Secure the Installation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 5: Configuring the Cluster

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 1: Server Configuration (identical on all three nodes)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 2: Galera Configuration

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 3: Log Directory

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 6: Bootstrapping the Cluster

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

How a Cluster Is Born

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 1: Bootstrap mgmt-man-01

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 2: Join mgmt-man-02

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 3: Join mgmt-lon-01

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 4: Verify the Cluster as a Whole

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 7: Creating the CloudStack Databases

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 1: Databases and the Application User

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 2: Watch Replication Do Its Job

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 8: The Single-Writer Routing Layer

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Shape of It

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 1: The Health Check (all three database nodes)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 2: HAProxy on the Management Servers

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 3: Prove the Path

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 9: Testing What We Built

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

First, a probe table — and why `SELECT 1` will not do

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Drill 1: Graceful Node Shutdown (planned maintenance, rehearsed)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Drill 2: The Writer Dies Mid-Flight

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Drill 3: Quorum, Demonstrated Honestly

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 10: Backups — for What Clustering Cannot Prevent

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Tool and the Pattern

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Restore Drill (read it now, rehearse it in the lab)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 11: Monitoring, Verification, and the Record

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Daily Glance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Verification Script

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Documentation as Code

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 5 — Key Takeaways

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1. Single-writer is the design, not a limitation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

2. The “or equivalent” clause is load-bearing

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

3. Quorum arithmetic is a contract, not trivia

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

4. Bootstrap is a one-time word

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

5. Clusters keep data available; backups keep it recoverable

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 6: CloudStack Management Server

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Brain Arrives — Installing the Control Plane on the Nodes That Already Hold Its Memory

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Story: Switching On the Brain

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 1: What the Management Server Actually Is

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Control Plane, Located Precisely

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Three Servers, Not Two — and Not Active-Passive

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Contract This Chapter Inherits

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 2: Pre-Flight Checks and Firewall Doors

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 1: Verify the Layers Beneath

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 2: Open the Management Server's Doors in UFW

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 3: Database Housekeeping — Two Honest Reconciliations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Reconciliation 1: The PROCESS Grant

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Reconciliation 2: max_connections Was Sized for One Tenant — Three Are Moving In

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 4: Installing CloudStack on the First Node

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 1: Add the CloudStack Repository

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 2: Install the Management Server Package

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 5: Initialising the Schema — Without Borrowing Root

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Flag Built for Our Situation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 1: Run It

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 2: Verify What It Did

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 6: First Start — and First Contact

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 1: cloudstack-setup-management

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 2: Watch It Wake Up

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 7: Nodes Two and Three — the Cluster Assembles

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 1: Install and Configure mgmt-man-02 and mgmt-lon-01

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 2: Watch Them Find Each Other

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 8: First Login

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 1: The Default Credentials Problem, Solved in the First Minute

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 2: A Short, Honest Tour of an Empty Cloud

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What We Are Deliberately Not Doing Yet

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 9: Three Front Doors, One Name

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 10: Testing What We Built — the Inherited Standard

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Drill 1: A Management Server Dies

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Drill 2: The Database Writer Dies Under a Live Control Plane

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 11: Verification and the Record

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Verification Script

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Documentation as Code

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 6 — Key Takeaways

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1. The control plane and the cloud fail separately

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

2. `--schema-only` is a trust boundary, not a convenience

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

3. `127.0.0.1` everywhere is the whole design, honoured

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

4. Multi-node is native — size for it

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

5. An empty cloud is the right place to practise killing things

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 7: Secondary Storage

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Repository That Rebuilds the Cloud — DRBD and NFS on Refurbished Metal

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Story: The Library Question

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 1: Understanding Secondary Storage

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What Lives on Secondary Storage

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Boundary With Chapter 9 — Read This Twice

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 2: The Design

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Hardware: Refurbished R740XD Pair

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Naming and Addressing

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Replication Layer: DRBD, From the Archive

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Failover Layer: Keepalived, With Open Eyes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Capacity Planning

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 3: Building the Replicated Filesystem

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 1: Prepare the Backing Volume [both]

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 2: Install and Configure DRBD [both]

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 3: Initialise and Synchronise

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Step 4: Filesystem [primary]

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 4: The NFS Service

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Install and Pin the Ports [both]

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Export [both]

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Firewall Openings [both]

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Stop What Keepalived Will Own [both]

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 5: Failover with Keepalived

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Moving Parts

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Keepalived Configuration

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Test Sarah Asked For

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 6: Verification

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Performance Baselines

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

A Sidebar: The Object Storage Future

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 7: Troubleshooting

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 7 — Key Takeaways

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1. Secondary storage is the recovery path — keep it out of the blast radius

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

2. Match the hardware to the workload, not the catalogue

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

3. Protocol C plus a promotion guard is what makes failover honest

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

4. HA decisions are workload decisions

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

5. Built here, registered there

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 8: Primary Storage

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Where the VMs Actually Live — Ceph RBD on the NVMe Estate

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Story: The Microseconds Chapter

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 1: Understanding Primary Storage

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What Primary Storage Does

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Why Native RBD

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Boundary With Chapter 9 — Same Rule, Second Verse

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 2: The Design

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Hardware: Six R760s, Deliberately Untouched

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Cluster Topology

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Why Not One Stretched Cluster — The Honesty Box

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Pool Design

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 3: Deploying the Cluster

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Which Ceph, and Where From

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Orchestrator's Identity

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Bootstrap

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Adding the Other Five Nodes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Placing the Daemons — Labels, Not Luck

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 4: OSDs, the Pool, and the Credential

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Chapter 4 Payoff

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Deploying the OSDs — a Spec, Not a Standing Order

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Pool

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Credential: A Dedicated Key With Minimal Capabilities

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Values Card for Chapter 9

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 5: Firewall and Verification

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Firewall Openings [all nodes]

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Verification

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 6: The Drills

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Drill 1 — Pull a Drive

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Drill 2 — Pull a Node

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Baselines

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 7: Troubleshooting

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 8 — Key Takeaways

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1. Primary storage is the latency path — give the hypervisor a straight line

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

2. The orchestrator owns the drives

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

3. Failure domains are arithmetic, not vibes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

4. One cluster per site — the speed of light is not negotiable

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

5. A dedicated credential, minimal capabilities, proven — and built here, registered there

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 9: Deploying CloudStack — The Zone Goes Live

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Morning the Cloud Gets a Shape

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Where We Stand

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What This Chapter Is — and Is Not

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 1: The Shape of a Zone

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Traffic Types: Telling CloudStack About Wires It Hasn't Seen

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 2: The Toolbelt — UI, API, and an Honesty Box About the Wizard

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

First, Two Minutes of Hygiene

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

cmk Against the Front Door

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 3: Zone, Physical Network, and the Public Range

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Creating the Zone

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Physical Network and Its Three Labels

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Waking the Virtual Router Providers

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Public Range

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 4: Pod and the Guest VLAN Range

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Pod and Its Reserved Addresses

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Guest VLANs: 300–399, and Why Not 200–299

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 5: The Cluster, Deliberately Empty

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 6: The Image Store — and the Auto-Seed Moment

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Pre-Flight: The Path CloudStack Will Use

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Registering the Store

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What Happens Next, Unbidden

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

If the Seed Fails

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 7: Primary Storage — Staged, Verbatim, and One Honest Step Early

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Ordering Truth

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Values Card, Field by Field

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Staged Registration

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 8: Verification, Then — and Only Then — Enablement

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Troubleshooting the Logical Build

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 9 Summary: Five Things to Hold On To

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Values Card for Chapter 10

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 10: Adding KVM Compute Hosts — The Day the Stage Crew Arrives

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Morning After the Map

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Where We Stand

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What This Chapter Is — and Is Not

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 1: The Shape of the Day

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 2: The Bridges — Three Labels on Two Bonds

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Netplan — kvm-man-01 Shown

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Disruptive Move, Done Carefully

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 3: The Hypervisor Stack

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Packages — Including the One Everyone Forgets

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Libvirt, the Modern Way — and a Genuine CloudStack Advantage

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

AppArmor Stays Enforcing — an Honest Disagreement

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Firewall Doors This Service Owns

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 4: The CloudStack Agent and the Service User

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Install

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

A Service User Instead of a Root Password

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 5: Pre-Flight — Never Assume, Interrogate

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 6: Adding the First Host

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What the Platform Does With Your Two Minutes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 7: The Cascade

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

7.1 The Staged Registration Fires

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

7.2 What You Did Not Have to Do

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

7.3 First Light: the System VMs

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

7.4 Soak, Then Scale

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 8: The GPU Cluster

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

8.1 The Cluster

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

8.2 A Second Pool, and a Capabilities Lesson

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

8.3 GPU Passthrough on the R7525s

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

8.4 Join and Register

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 9: Verification — Extending the Zone Script

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 10: Drills — Break It While It's Cheap

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Troubleshooting the Join

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 10 Summary: Five Things to Hold On To

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Values Card for Chapter 11

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 11: Deploying Your First Guest VMs — Opening Night

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Morning the Cloud Earned Its Keep

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Where We Stand

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What This Chapter Is — and Is Not

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 1: The Shape of the Day

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 2: Compute Offerings — The Shapes a VM May Take

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 3: Disk Offerings — Data That Outlives the VM

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 4: The GPU Placement Proof — Five Minutes That Retire an Assumption

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 5: The Guest Template — The SSVM Does the Fetching

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 6: Editorial's Network — Where `cloudbr2` Earns Its Keep

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

6.1 The Network Offering — Egress Denied by Default

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

6.2 The Account and the Network

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

6.3 First Light on the Guest Lane

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

6.4 The Virtual Router and Its Public Leg

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 7: Deploying the Three Guests

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

7.1 A Key, Not a Password

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

7.2 cloud-init — First-Boot Without a Console

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

7.3 Deploy

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

7.4 Proof: DHCP, Console, and the Locked Door

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

7.5 Opening the Doors We Choose

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

7.6 A Data Volume That Survives a Rebuild

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 8: Everyday Lifecycle — and One Thing the Old Book Got Wrong

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

8.1 Live Migration — the Boring Kind

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

8.2 Resize — Stop, Reshape, Start

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

8.3 Snapshots — Copy-on-Write, and Why It Is Nearly Free

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

8.4 The Correction: Security Groups Do Not Belong Here

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 9: Verification — Extending the Zone Script

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Troubleshooting the First Workload

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 11 Summary: Five Things to Hold On To

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Values Card for Chapter 12

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 12: Advanced Networking — The Dell OS10 Leaf-Spine Fabric

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Turning “Trust the Trunk” Into a Tested Data Plane

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Story: The Promissory Notes Come Due

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 1: Why a Fabric, and Why This One

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The shape of the thing

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What the fabric must carry

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 2: Building the Fabric in the Lab

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The lab topology

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Obtaining OS10 virtual (decision d-os10-1)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Defining the switch nodes and the links

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

First boot and management

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

How the lab maps to the metal

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 3: The Underlay — One Routed Fabric

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Why unnumbered, and why eBGP

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The addressing and ASN plan

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Bring the fabric ports up and set the loopback

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Configure BGP (unnumbered)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Prove the underlay

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 4: VLANs, the Host Trunk, and Making 10.100.0.1 Answer

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Which VLANs are routed, and which are not

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Create the VLANs

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The host-facing trunk

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Stand up the management gateway (the SVI)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Advertise the management subnet into the fabric

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

When you outgrow this: EVPN-VXLAN (carry-forward step 2)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 5: Proving It — Discharging the IOUs

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

IOU 1 — `show vlan` actually shows something

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

IOU 2 — the management gateway answers (Chapter 9)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

IOU 3 — the public VLAN is really tagged (Chapter 10)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

IOU 4 — the guest tag reaches `bond1.301` (Chapter 11)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

IOU 5 — jumbo frames cross VLAN 201 unbroken (Chapter 4)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

IOU 6 — losing a leaf does not lose the host (VLT failover — proven here)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The change-control ladder

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 6: What the Fabric Costs — and What the OS10 Licence Buys

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Counting the switches

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The hardware — the same three price points (decision d10, inherited)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The one economic delta: the OS10 licence

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Why OS10 here — what it buys, and what it trades

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Fabric Is Real Now

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Values Card — What Chapter 12 (OS10 edition) Hands Forward

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 13: Advanced Networking with VPCs — Walls Within the Walls

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Building Multi-Tier Isolation on the Fabric You Just Proved

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Story: One Network Was Never Going to Be Enough

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Where We Stand

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What This Chapter Is — and Is Not

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 1: The Shape of the Day

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 2: VPCs, and the One Fact That Trips Everyone

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 3: Designing the Catalogue VPC — Addresses Before Anything

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 4: The VPC Offering and the VPC

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

4.1 The VPC Offering — Services the Router Will Provide

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

4.2 Create the VPC

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 5: The Tiers — Where the Lanes Multiply

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

5.1 Create the Tier Networks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

5.2 Watch the Lanes Appear

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 6: Network ACLs — The Walls, Drawn Both Ways

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

6.1 The Posture: Deny by Default

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

6.2 The Web Tier — Public In, App Out, Replies Allowed

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

6.3 The App and Database Tiers — Tighter Each Step

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 7: The VPC as Code — Reproducible Walls

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 8: Deploying the Tiers, Proving the Walls, and Balancing the Load

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

8.1 Three Guests, One Per Tier

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

8.2 Prove the Walls — Including the Ones That Should Not Open

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

8.3 Public Load Balancing for the Web Tier

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

8.4 Internal Load Balancing for the App Tier

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 9: Site-to-Site VPN — Extending the Private Space to Leeds

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 10: Verification — Extending the Zone Script

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 11: Troubleshooting — The Stateless Traps and Their Friends

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 12: Where This Goes — One Pattern, Many Tenants

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 13 Summary: Five Things to Hold On To

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Values Card for Chapter 14

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 14: Multi-Zone Expansion

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Cloud Grows a Map of the Country — Without Stretching a Single Fiction

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Story: Three Cities, One Control Plane, No Long-Distance Lies

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Where We Stand

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What This Chapter Is — and Is Not

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 1: What “Multi-Zone” Means Here — and the Anti-Pattern It Is Not

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The anti-pattern, named so we can avoid it

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The three zones at a glance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 2: The Site Inventories — Canonical Estate, Honest Addressing

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

London — the southern public zone

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Leeds — the internal-only zone

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The derived addressing — and a mask we must be honest about

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 3: Extending the Fabric — Leaf Pairs to the Manchester Spines

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The shape: spineless leaves over the inter-site links

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The underlay extends — new ASNs, same eBGP unnumbered

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Per-site VLANs and the anycast gateway

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 4: London's Independent Ceph Cluster — the Chapter 8 Pattern, Four Nodes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Bootstrap and build (on ceph-lon-01)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The pool and the credential — the same names, deliberately

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 5: Leeds' Independent Ceph Cluster — Three Nodes, the Honest Floor

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 6: Bringing Up the London Zone — the Wizard, Now That We Understand It

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The wizard, page by page (with the cmk underneath)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 7: Bringing Up the Leeds Zone — Building a Zone With No Front Door

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Why an internal-only zone is a feature, not a limitation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 8: The Region-Level VPC — Sharing a Design Without Sharing a Failure Domain

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

One module, three instances

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

When two zones must actually talk

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 9: Verification Across Three Zones — Build-Time, Not the Monitoring Chapter

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Troubleshooting the Multi-Zone Build

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 14 Summary: Five Things to Hold On To

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Values Card for Chapter 15

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 15: Monitoring & Operations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Observability at Scale — One View of Three Zones

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Story: Seven Dashboards Are Not a Dashboard

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 1: The Observability Problem at Estate Scale

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Three Pillars, Briefly

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Islands We Already Have

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Decision: Federate, Don't Rebuild

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Sovereignty Is Not Optional Here Either

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Shape of What We Build

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 2: The Estate Observability Hub

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Placement and Sizing

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Deploy Prometheus (the Federation Layer)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Federation Scrape — Pulling Breadth Up

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

DNS and the Firewall

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 3: Instrumenting the Estate

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The KVM Hosts — node_exporter and libvirt-exporter

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Three Ceph Clusters — Already Done

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

CloudStack Itself — the Built-in Exporter

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Dell OS10 Fabric — Telemetry via gNMI/OpenConfig

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 4: Dashboards That Answer "How Is the Estate?"

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Install Grafana on the Hub

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Board's Slide — Estate Overview

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Operator's Screen — Infrastructure Health

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Per-Zone and Per-Tenant Views

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 5: Logs — Shipping the Estate's Diary Somewhere Durable

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Why Loki, Not a Heavyweight Search Cluster

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Deploy Loki on the Hub

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Promtail on Every Host

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Leeds Logs Travel the WAN — and Stay on the Estate

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 6: Alerting Without a SaaS Pager

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Estate Alertmanager

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Alert Rules — Symptoms, Not Spikes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Sovereign Notification Path

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Silences for Planned Work

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 7: Incident Response and On-Call — the Runbooks Owe

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

On-Call and Escalation, Self-Hosted

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Severity, Briefly

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Runbook Index — Mostly Drills You Have Already Run

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Blameless Post-Mortems

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 8: Capacity Signals for Operations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 9: Lab Exercise — One View of the Lab

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Five Key Takeaways

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 16: Backup & Disaster Recovery

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

An Independent Copy That Does Not Share Fate

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Story: The Night the Mirror Lied

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 1: Three Words People Use Interchangeably and Shouldn't

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Replication — what Chapters 8 and 14 gave you

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Backup — a copy from a point in the past that the present cannot overwrite

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Disaster Recovery — the rehearsed procedure that turns copies back into a running business

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Mars test

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 2: Protection Tiers, and Why They Are Really a Storage-Media Decision

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 3: The Independent-Fate Backup Cluster

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Why a separate cluster, not just more OSDs

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The three roles on one cluster

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Deploying the cluster

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Replicated or erasure-coded? Pool design on a four-node cluster

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 4: CloudStack Native Backup & Recovery — the Shipping KVM Workhorse

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Where the backups land

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Enabling the framework and registering the NAS provider

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Backup offerings — policy as an object

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Taking, listing, and restoring a backup

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 5: Cross-Site DR with rbd-mirror (Tier 1)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Leeds is reachable, but only from the inside

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 6: The Immutable Archive — RGW S3 Object Lock (Tier 3)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Modes: COMPLIANCE vs GOVERNANCE

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 7: Database Disaster Recovery — Discharging the Chapter 5 Promise

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Restoring Galera, step by step

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 8: What's Coming in 4.23 — Native KVM Backups via Veeam

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 9: Restore Drills, Runbooks, and the Mars Test

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Drill 1 — "Oops" recovery (Tier 0 / Tier 2)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Drill 2 — Dirty-data point-in-time recovery (Tier 2 + database PITR)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Drill 3 — Immutability proof (Tier 3)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Drill 4 — Site loss and the Leeds WAN partition (Tier 1 + Chapter 14)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Post-restart health triage

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Alerting and escalation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

verify-dr.sh

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Five Key Takeaways

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 17: Cost Optimisation & Capacity Planning

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Owning the Meter: What a Sovereign Cloud Actually Costs

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Story: Does Owning This Thing Actually Pay?

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 1: You Cannot Manage What You Cannot Meter

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The contract we deferred

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Where it runs, and the database contract

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Installation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Configuration

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What gets recorded

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Pulling the data

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 2: What the Estate Actually Costs

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Reconciling the fiction we inherited

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Capital expenditure

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The market context, and why owning beats renting in it

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Depreciation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Operating expenditure

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 3: The Escape That Pays

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The bill we walked away from

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

CloudStack versus VMware, over five years

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The fabric, and an honest term mismatch

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What renting the same thing would cost

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The shape of the argument

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 4: When the Migration Pays for Itself

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Two break-evens, because there are two honest ways to count

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

A payback calculator you can hand to finance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 5: Cost per Workload — From Architecture to Accountability

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Turning Chapter 16's media tiers into pounds

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Costing compute and memory

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

From usage records to a statement

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

When the cost centre becomes a customer

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 6: Return on Research — Finding the Idle Silicon

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Idle silicon is waste

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Two questions, two data sources

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The right-sizing analyser

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Make the estate flag its own waste

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Why this is the highest-leverage saving you have

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Part 7: Seeing the Wall Before You Hit It

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The forecast Chapter 15 started, costed

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The method

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The forecaster

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Lead time is a moving target in a crunch

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Expansion is scale-out, and the forecast respects the failure domain

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Five Key Takeaways

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

1. You cannot manage what you cannot meter

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

2. The escape pays — but count CapEx and OpEx honestly

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

3. Cost per workload turns architecture into accountability

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

4. Idle silicon is waste — right-sizing is the highest-leverage saving

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

5. Forecast the wall in pounds, not just percent

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Chapter 18: Your Cloud Journey Starts Now

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Estate Is Built — Reflection, Consolidation, and the Road into Part 2

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Story: After the Renewal Letter

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Journey You've Taken

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What Sovereignty Actually Cost

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Thesis: Not Only for the Few

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Mars Test, One Last Time

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Where You Go Next: Part 2

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Key Takeaways

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Before You Close Part 1

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What Next, and Staying in Touch

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Appendix A: Glossary

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Appendix B — The Portable CloudStack Lab

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Running a Production-Grade Lab on a Laptop

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Introduction

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Hardware Requirements

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Primary Laptop

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

External Storage

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

USB-C Dock

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Optional: Raspberry Pi External Hosts

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Optional: USB Bluetooth Adapters

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Configuration Tiers

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Lab, Layer by Layer

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What's Different on a Laptop

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The Reader Lab, As It Was Actually Built

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 1 — The platform host

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The steps, in order

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What proved it — 7 of 7 checks passed

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 2 — The five networks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The steps, in order

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What proved it — 4 of 4 checks passed

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 3 — The seven virtual machines

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The steps, in order

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What proved it — 3 of 3 checks passed

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 4 — Bootstrapping Ceph

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The steps, in order

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What proved it — 4 of 4 checks passed

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 5 — The Ceph cluster

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The steps, in order

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What proved it — 7 of 7 checks passed

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 6 — Secondary storage

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The steps, in order

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What proved it — 4 of 4 checks passed

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 7 — The CloudStack management server

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The steps, in order

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What proved it — 5 of 5 checks passed

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 8 — The first compute host

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The steps, in order

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What proved it — 9 of 9 checks passed

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 9 — The second compute host

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The steps, in order

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What proved it — 9 of 9 checks passed

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 10 — The Manchester zone

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The resources, in order

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What proved it — 10 of 10 checks passed

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 11 — The first guest

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The resources, in order

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What proved it — 5 of 5 checks passed

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Phase 12 — The acceptance suite

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The steps, in order

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

What proved it — 24 of 24 checks passed

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

When it did not go to plan

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

The platform host took five attempts

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

A probe that lied about installation state

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

A client timeout that was not a server failure

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Connecting Real Hardware

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Raspberry Pi as External Hosts

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

USB Bluetooth Passthrough

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Performance Notes (Laptop-Specific)

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Backup and Recovery

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Use Cases

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Conclusion

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/apache-cloudstack-deployment-guide.

Afterword — The Estate Is Yours Now

It started with a letter and a number: a renewal that turned renting into a threat. Most of the industry treats that letter as a cost to be absorbed. This book treated it as a question — *what would it take to stop renting altogether?* — and Part 1 has spent eighteen chapters answering the foundational half of it.

You went from bare metal to a sovereign, multi-city platform: KVM compute, Ceph storage that survives a node without dropping a beat, a OS10 leaf-spine fabric you proved command by command, multi-zone resilience, one observability hub that pages you before your users do, backups that do not share fate with what they protect, and a meter that turns the estate's costs into numbers you can defend. You own it down to the last packet.

What you hold is a **method, not a product, and not a guarantee.** The architectures and figures were illustrative, grounded in one estate; your numbers will differ. The discipline is the deliverable — measure before you change, rehearse before you commit, and keep the evidence — and that travels to any infrastructure you ever touch.

Sovereignty was never really about the money, though the money was substantial. It was about which questions you get to ask. When you own the stack, "can we?" stops being a vendor's decision and becomes an engineering one. That is the quiet shift this whole book was building toward.

The road keeps going — autonomous edge cells, sovereign AI on hardware you own, energy treated as a first-class design constraint — but the principle underneath it does not change. **Grow your own, cook your own.** Build it, break it in a lab, and make it yours.

Now go and take your infrastructure back.

About the author

Michael John Leslie Hinsley has spent over three decades bridging the gap between delicate electronics and the mud and rain of the physical world. From enterprise data centres to off-grid aquaponics and apiaries, his engineering philosophy remains the same: if you cannot control your infrastructure, you do not own your business.

Stay in touch

Leaf Spine Books publishes at https://www.leafspinebooks.com/ — where Part 2 will appear when it is ready, alongside the rest of the line.

Published under the **Leaf Spine Books** imprint.

* * *

Apache CloudStack Deployment Guide — Leaf Spine Books edition Copyright © 2026 Michael John Leslie Hinsley. Code examples licensed under Apache License 2.0.