

Brought to you by:

NUTANIX[™]

Private Cloud

for
dummies[®]
A Wiley Brand



Build your own
private cloud

Make hybrid easy with
the correct platform

Deploy automation
tools and services

Nutanix
Special Edition

Ed Tittel

About Nutanix

Nutanix is a global leader in cloud software and hyperconverged infrastructure solutions, making infrastructure invisible so that IT can focus on the applications and services that power the business. Companies around the world use Nutanix Enterprise Cloud software to bring one-click application management and mobility across public, private, and distributed edge clouds so they can run any application at any scale with a dramatically lower total cost of ownership. The result is organizations that can rapidly deliver a high-performance IT environment on demand, giving application owners a true cloud-like experience.

Learn more at www.nutanix.com or follow [@nutanix](https://twitter.com/nutanix) on Twitter.



Private Cloud

Nutanix Special Edition

by Ed Tittel

**for
dummies**[®]
A Wiley Brand

Private Cloud For Dummies®, Nutanix Special Edition

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2021 by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, Dummies.com, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. Nutanix and the Nutanix logo are trademarks or registered trademarks of Nutanix. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN 978-1-119-71955-7 (pbk); ISBN 978-1-119-71959-5 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

We're proud of this book and of the people who worked on it. For details on how to create a custom *For Dummies* book for your business or organization, contact info@dummies.biz or visit www.wiley.com/go/custompub. For details on licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

Some of the people who helped bring this book to market include the following:

Project Manager: Martin V. Minner

Production Editor: Siddique Shaik

Associate Publisher: Katie Mohr

Nutanix Review Team: Allan Waters,

Senior Managing Editor: Rev Mengle

Mayank Gupta, Sahil Bansal,

Business Development Representative:
Karen Hattan

John Williamson, Ines Marjanovic,

Steve Carter, Greg White,

Krishnan Badrinarayanan

Table of Contents

INTRODUCTION	1
A Cloud Computing Bestiary	1
Cloud Management: On- and Off-Premises	2
How This Book Is Organized	3
Icons Used in This Book	3
Beyond the Book	4
CHAPTER 1: The State of Cloud Adoption	5
Oh IT, Where Art Thou?	5
Business drivers for change	6
Top IT challenges	6
The path to IT modernization	7
Understanding the Cloud	8
Public and private clouds	8
Cloud basics and benefits	8
Who's using cloud tech?	9
Cloud Adoption: Facts and Figures	9
Don't Forget Private Clouds	10
CHAPTER 2: Why Private Cloud?	13
Digging Deeper into Cloud Architectures	13
Who's Got the Keys?	14
Understanding Cloud Metrics	14
If you can't measure it	15
Understanding cloud costs	15
Understanding cloud security	15
Sources of cloud flexibility	16
Capabilities Abounding	16
HCI Works for Private Cloud	18
CHAPTER 3: Building a Private Cloud	19
HCI Basics	19
Why use HCI?	19
Compute	21
Storage	21

	Networking	21
	Virtualization.....	22
	HCI lifecycle management	22
	The Right Infrastructure Is Key	22
	Private Cloud Pieces and Parts	23
	Nutanix HCI = Firm Foundation	26
CHAPTER 4:	Automating the Cloud	27
	Cloud Automation Basics	27
	Benefits of automation	29
	Workflows, policies, and processes	29
	Selecting suitable targets	30
	Services and Tools for Private Cloud	30
	Automation examples	31
	The self-service advantage	32
	Tools and automation together	33
CHAPTER 5:	Exploring Hybrid and Multi-Cloud Vigor	35
	What IS Hybrid and Multi-Cloud?	35
	Why Hybrid and Multi-Cloud?	37
	Benefits of hybrid and multi-cloud	37
	Hybrid and multi-cloud strategy	39
	Extending a Private Cloud	40
	Types of Hybrid and Multi-Clouds.....	40
	Public service providers	40
	Managed service providers	41
	Nutanix Clusters	41
CHAPTER 6:	Nine Nuggets of Private Cloud Wisdom.....	43

Introduction

Look up in the sky. You may see many different kinds of clouds up there: cirrus, cumulus, nimbus, stratus, and more. This book doesn't talk about those kinds of clouds. Instead, think about the clouds that provide computing services, draw electricity, and are attached to the Internet.

A Cloud Computing Bestiary

The most basic beast in this menagerie is the cloud, as it refers to cloud computing. This describes a computing environment wherein computing resources — especially data storage, networking, and compute power — are available on demand without requiring active, direct management from its users. Generally, cloud users are “here” (from their perspective, anyway), and cloud computing lives “somewhere else” (often in a datacenter somewhere).

As it turns out, numerous kinds of computing clouds exist, too (often in interesting combinations, just like real clouds). Here are some terms you'll need to ponder and understand to grasp this book's story.

Clouds come in two basic flavors, depending on who owns, controls, and operates them:

- » **Public cloud:** A public cloud is available to users on demand and by subscription. It's the kind of cloud you have heard a lot about, and comes from many providers, including Amazon (Amazon Web Services, or AWS), Microsoft (Azure), Google (Google Cloud Platform, or GCP), and so on.
- » **Private cloud:** A private cloud serves only one company or organization, which owns and controls at least all of its content. Because a third party can host a private cloud in its datacenters, or the owner can host it in its own datacenter (or centers), a private cloud need not reside on premises. But it can, and often does, do just that.

Clouds also come in various numbers, where more than one cloud involves some kind of combination. Here's what's up with that:

- » **Single cloud:** This describes the situation where a business or organization uses only a single provider or solution for cloud applications or infrastructure. This arrangement is pretty unusual in today's marketplace, if not unheard of.
- » **Multi-cloud:** A business or organization uses two or more cloud providers or solutions for its cloud applications or infrastructure. Some multi-clouds offer little or no interoperability between their constituent clouds. Other multi-clouds offer the ability to migrate workloads and data between constituent clouds. The former is called a *low-interoperability multi-cloud*, and the latter is — you guessed it! — a *high-interoperability multi-cloud*.
- » **Hybrid cloud:** A name-worthy variation on multi-cloud, a *hybrid cloud* enables one or more applications to run seamlessly across its constituent clouds with easy orchestration and control.

Cloud Management: On- and Off-Premises

Most organizations of any size use many clouds, public and private, across their various departments and functions. But the best of all possible worlds is one in which all clouds work together in harmony and present a single, consistent, and coherent face to their users and the IT professionals in charge of setting them up and keeping them running.

A global, overarching view of cloud computing in any organization is best approached and managed as an “enterprise cloud environment.” That means using tools and solutions that unify operations across all clouds, and deliver interoperability for workloads for the whole enterprise.

This puts on-premises and off-premises clouds and their applications and data assets under a single “pane of glass,” with a single point of control for all applications and data. Whether clouds are public, private, or a combination of the two, this provides the

power needed to manage costs, deliver consistent security, and meet service-level agreement (SLA) requirements that apply to application and service delivery.

The goal is to place and manage workloads where they'll provide the best results (and best economics) and to move them when and wherever they might be needed. This book talks a lot about this.

How This Book Is Organized

The chapters in this book explain the ins and outs of private clouds and how they fit best into the Enterprise Cloud vision:

- » **Chapter 1: “The State of Cloud Adoption”** reviews the state of the marketplace, including uptake and adoption rates, IT challenges, and trends in private cloud use.
- » **Chapter 2: “Why Private Cloud?”** explains why enterprises find private clouds worth building and using.
- » **Chapter 3: “Building a Private Cloud”** explores private clouds' technology underpinnings in hyperconverged infrastructures, as well as their typical pieces and parts.
- » **Chapter 4: “Automating the Cloud”** digs into the key role that automation plays in cloud computing, and especially what it can do in a private cloud.
- » **Chapter 5: “Exploiting Hybrid and Multi-Cloud Vigor”** shows you how hybrid and multi-cloud lets organizations manage workloads anywhere and everywhere to help them maximize cloud return on investment (ROI).
- » **Chapter 6: “Nine Nuggets of Private Cloud Wisdom”** lays out value propositions and benefits that make private cloud worthwhile.

Icons Used in This Book

Each *For Dummies* book includes visual widgets called *icons* in its margins. Their job is to flag text that's especially noteworthy, interesting, or perhaps alarming. Or not, as you'll see (and decide for yourself).



REMEMBER

This flags something that might be worth revisiting, or thinking about further, as you work with a private cloud.



TECHNICAL
STUFF

I can't help it. Sometimes I have to peel back the covers and dive into technical details. When you see this icon, it means "tech-talk ahead." Feel free to skip it if you start nodding off.



TIP

As the icon indicates, this icon flags information that's on-target and timely, intended to help you get the most out of your private cloud learning, experiments, and adventures.



WARNING

This icon flags text that provides information to help keep you out of trouble and avoid unnecessary cost and effort. Don't say I didn't warn you.

Beyond the Book

Between these covers, I've provided as much information about private clouds as I can cram into a short book. If you find yourself hungry for more, please visit <https://nutanix.com> and search for "private cloud." You'll find a wealth of additional information, including whitepapers, e-books, videos, blog posts, and more. Eat up!

IN THIS CHAPTER

- » Understanding IT's impetus to put the cloud to work
- » Digging into cloud spending and adoption
- » Making the best of public, private, and hybrid and multi-clouds

Chapter **1**

The State of Cloud Adoption

Many sources, including experts, pundits, IT professionals, and technology vendors, all rightfully and accurately claim that “the cloud changes everything.” Indeed, the use of cloud technology has remade and reshaped almost everything about how information technology is built and consumed.

And indeed, cloud adoption and consumption are mind-boggling in all of their facets and assets. This chapter looks at how IT works nowadays and shows why cloud technologies play such an important role in how business does IT — and also in how IT does business.

Oh IT, Where Art Thou?

In today's modern economy, technology plays a vital role. The IT acronym expanded reads “information technology.” Increasingly, the cloud (in all of its forms) plays a starring role in making IT cost-effective, efficient, and enabling business or mission success. As the economy becomes increasingly dependent on technology (and IT), it's no coincidence that pundits call it “the digital economy.”

Business drivers for change

In fact, business needs to adapt and change if it is to survive, or perhaps even thrive, in this new digital economy. IT can provide the means for such changes and help businesses and organizations reinvent themselves for the modern age.

Properly chosen and deployed, IT can enable a range of positive and desirable business outcomes. First, the flexibility and reach of cloud-based technologies can foster rapid innovation to grow revenue and create new business opportunities. Second, the capability and scope of cloud-based technologies can delight clients or customers. This helps businesses keep existing customers and attract new ones, while advocating on behalf of customers to deliver positive experiences. And last, cloud-based technologies' support for automation helps simplify IT operations while making them faster and more efficient.

Top IT challenges

Without a doubt, IT must have its act together. IT requires a consistent and coherent view of technology needs across the whole organization. IT must also act quickly and decisively. In fact, IT faces certain challenges before it can claim “mission accomplished” in today’s digital economy. Here’s what IT must do:

- » Break departments, business units, and other divisions in the company out of their individual technology silos.
- » Simplify and streamline overly complex, manually managed, and poorly integrated infrastructures.
- » Reduce the time it takes to provision and deploy applications.
- » Prevail in the ongoing struggle to keep applications and data secure and compliant with rules, regulations, policies, and best practices.
- » Overcome slow infrastructure rollouts and learn to scale better for growth (both scale-up and scale-out).
- » Make operations predictable and reliable to avoid excess costs and unscheduled downtime.
- » Overcome the lack of time, budget, and resources so it can itself innovate while also supporting and fostering innovation throughout the business.

The path to IT modernization

Figure 1-1 depicts a workable path to achieve IT modernization.

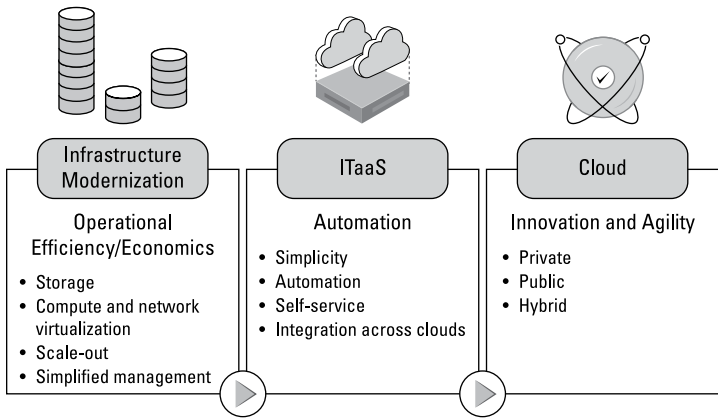


FIGURE 1-1: IT modernization starts with infrastructure and proceeds from there.

To begin with, organizations must modernize their IT infrastructure to achieve operational efficiency and make best use of the savings and efficiencies that cloud-based options can deliver. This applies to all aspects of computing technology, including storage, compute, and network resources. It also takes advantage of virtualization across the board, provides scale-out capability, and works under a simplified, coherent set of management and monitoring tools.

Next comes “IT as a service” (ITaaS). This relies on simplicity by design and automation to deliver self-service access to IT on demand, along with integration across all clouds in use in the organization (private and public, multi, and hybrid).

Through deployment of a truly Enterprise Cloud environment that seamlessly combines private, public, and hybrid and multi-cloud elements, the organization achieves agility and the ability to innovate and grow as needs and circumstances dictate. Here comes nerd-vana!

Understanding the Cloud



REMEMBER

A cloud is a computing environment that offers computing resources on demand — including data storage, networking, and compute power — without requiring active and direct setup and management from its user (or consumer).

Cloud computing is the on-demand availability of computer system resources, especially data storage and computing power, without direct active management by the user. The term is generally used to describe datacenters available to many users over the Internet. In those datacenters, cloud computing relies on hyper-converged infrastructure (HCI). HCI is purpose-built to provide resources — namely, storage, networking, and compute — on flexible, scalable, highly virtualized hardware. Though a cloud offers resources to consumers as and when they need them, behind that offer stands one or more datacenters on which those resources run. Also, behind the scenes are the tools, controls, and software to define, configure, run, and manage those resources. That's what makes clouds work.

Public and private clouds

A *public cloud* comes from an independent, third-party vendor. The biggest of these are Amazon Web Services (AWS) and Microsoft Azure. They maintain compute resources (storage, networking, and compute) that customers can use as they need them. Organizations that use a public cloud share its compute resources in common in a “multi-tenant” environment.

A *private cloud* is a computing environment owned by a particular organization. It features public cloud benefits and capabilities but is hosted in the organization's own datacenter or via a third-party provider (in a “single-tenant” environment). Private clouds are known for reliability, scalability, and security. That's why they are so often used to run sensitive or important enterprise workloads.

Cloud basics and benefits

Cloud technology, both public and private, offers virtually unlimited access to compute resources that include storage, networking, and compute power. Public cloud lets buyers trade CapEx for OpEx because they don't need to buy, house, power, cool, or maintain underlying hardware.

Public cloud users pay only for what they use for as long as they use it, so demand goes up and down, and charges track usage. Businesses can buy resources they need during peak periods or end-of-cycle upticks, then fall back to lower levels during periods of reduced demand. Both private and public clouds benefit equally from coherent, simplified management, automation, and scalability (public clouds offer a nearly infinite top end where growing private clouds costs money).

Private clouds deployed on-premises are subject to CapEx but provide the ultimate in security, confidentiality, compliance, and control. However, private clouds set up as “single-tenant” environments through a hosted private cloud provider can offer the best of both. Today, they’re increasingly likely to represent an attractive deployment scenario for private clouds.

Who’s using cloud tech?

The quick answer is “everybody.” Most surveys show that correlation between organization size and cloud use is fading fast (the bigger, the more likely to use cloud technology). As far back as 2015, a Deloitte study, “Small Business, Big Technology,” showed 85 percent of SMBs believed cloud was key to rapid growth.

Today, cloud is built into Windows 10 in OneDrive (iCloud plays a similar role for MacOS). The cloud is ubiquitous in IT for concerns of all shapes and sizes.

Cloud Adoption: Facts and Figures

The research firm ITCandor reports that global spending on cloud services by quarter has increased dramatically since 2005. Here are some of ITCandor’s salient findings:

- » \$246 billion was spent globally for the 12-month period from July 1, 2018 to June 30, 2019.
- » The Americas (U.S., Canada, and Mexico) spent \$122 billion of that total during that time.
- » EMEA (Europe, Middle East, and Africa) spent \$71 billion.
- » Asia Pacific accounts for the remaining \$53 billion.

The trend for all plots (and their aggregate) has a steeply rising slope, more so as it passes 2015. Gartner forecasts a global total of more than \$266B in 2020, \$355B in 2022 (<https://www.gartner.com/en/newsroom/press-releases/2019-11-13-gartner-forecasts-worldwide-public-cloud-revenue-to-grow-17-percent-in-2020>).

In the context of overall global IT spending, Gartner (Jan. 2020) estimates for 2019–2021 are (<https://www.gartner.com/en/newsroom/press-releases/2020-01-15-gartner-says-global-it-spending-to-reach-3point9-trillion-in-2020>):

- » 2019: 3.74 trillion (0.5 percent growth)
- » 2020: 3.87 trillion (3.4 percent growth)
- » 2021: 4.01 trillion (3.7 percent growth)

Alas, those estimates may not hold owing to the worldwide coronavirus pandemic. But the trend is clear: Cloud spending consumes a growing portion of overall IT outlays and is rising much faster than overall IT spending (more than 12 percent versus approximately 4 percent).

Don't Forget Private Clouds

Now that public clouds have been around almost two decades, businesses have a much keener appreciation for when and where public clouds do — and don't — make sense.

Today, some organizations are repositioning workloads from the public cloud onto private clouds. They may do so because of security concerns (or stringent compliance or contractual requirements). They may do so because of exacting performance requirements or service-level agreements (SLAs), where even a few milliseconds can make a difference. They may choose to run workloads on private networks because they involve sensitive, proprietary, or even classified data. Most of all, they may choose to run predictable, frequent workloads “on-prem” because it's cheaper that way. Cost still reigns supreme!

As organizations decide where workloads should run, the kinds of resources and data they use will often dictate (or at least strongly

suggest) what should go public and what private. Beyond cost, application or service attributes will determine which choice makes most sense. In general, organizations should make best use of resources on hand, understanding that on-premises capabilities are already bought and paid for. Using them as much as possible maximizes return on investment (ROI).

But because both private and public clouds have benefits and good uses, organizations that adopt a hybrid and multi-cloud model can use both. If a private cloud connects with a public one, businesses can run workloads in both worlds. Starting with the private cloud, the business can “burst” into the public cloud when added demand calls for more resources. That entails a high level of compatibility between the software that runs in those clouds and the services they both use. Thus, it’s vital to make sure that your private cloud is built to make public cloud connections simple and straightforward (that is, it doesn’t take countless consulting hours to pull this off).

IN THIS CHAPTER

- » Understanding public and private clouds and their interoperation
- » Making the most of cloud metrics and monitoring
- » Seeing why hyperconverged infrastructure makes for peachy private clouds

Chapter 2

Why Private Cloud?

The chapter title voices the main question this book seeks to illuminate and answer. A private cloud, on its own, can provide some of the primary features associated with public clouds. But a private cloud also has limitations that you can't ignore or overlook. Otherwise, they can come back and bite you. This chapter shows you what's involved in considering and possibly choosing a private cloud for your organization.

Digging Deeper into Cloud Architectures

In the sky, it's said that no two clouds are exactly alike. On the ground (and in the datacenter), there's some truth to that assertion, too. Multi-cloud environments mean multiple clouds, invariably from different providers, each with its own interfaces, requirements, and idiosyncrasies.



REMEMBER

Nearly all hybrid and multi-clouds offer limited or no interoperability between their component clouds. It takes special platforms, forethought, and careful engineering before hybrid and multi-clouds can migrate workloads and data between their clouds easily and quickly. The first kind is called a *low-interoperability hybrid multi-cloud*, and the second is a *high-interoperability hybrid multi-cloud*. Given the choice, organizations choose to go high every time.

Because hybrid and multi-cloud are increasingly status quo for most businesses and organizations, taking steps to ensure interoperability is a huge deal. As you consider the clouds you have and the ones you're thinking about getting, it's vital to ask about — and ensure — that those clouds can interact and exchange data (plus software and workloads) as needed. Otherwise, you could face difficulties in dealing with “bursting” out of a private cloud and into a public one. Ditto for moving data, workloads, and software between one public cloud and another.

Who's Got the Keys?

The question is certainly simple. Here's hoping the answer is, too. But for public clouds, some responsibilities and controls are shared between the provider and the consumer. Buyers must make arrangements, including payment, for the services they want before they can start consuming resources (compute, networking, and storage) that the public cloud delivers in astonishing quantity and variety.

For private clouds, if resources are on premises, there's no doubt who's in control: You are! But you're also completely and solely responsible for CapEx, maintenance, management, licenses, power and cooling, and all the costs involved in managing your own assets.

Understanding Cloud Metrics

Once you're in the clouds and using resources, consumption correlates with cost. That definitely needs to be tracked. But other metrics are of great interest to buyers and users as well. These include a great many things, including latency, response time, availability and reliability, security, throughput (aka bandwidth), capacity, scalability, and more. For a great discussion on this topic, see *Guiding Metrics* in “The Cloud Service Industry's 10 Most Critical Metrics” (<https://guidingmetrics.com/content/cloud-services-industrys-10-most-critical-metrics/>).



TIP

Nearly all these metrics — except cost, though it may still be tracked as a “chargeback” of some kind — are important for private as well as public clouds. Acquiring, metering, monitoring, and leveraging what's observed is essential to good cloud management.

If you can't measure it . . .

As the old saying goes “If you can't measure it, you can't manage it.” Thus, when arranging for consumption of cloud resources, it's essential to understand what's being measured and how those measurements can best be monitored. The users want a positive experience, and you'll want to make sure service-level agreements (SLAs) are met and minimum acceptable performance levels are honored.



TIP

It's extremely important to understand how monitoring data is made available and logged. You must stay on top of that data to get your money's worth for cloud services of any kind.

Understanding cloud costs

Because cloud services come and go with demand, staying on top of associated costs is key to avoiding unpleasant surprises. Think: “What happens to my costs if the website explodes with activity?” “What if orders go through the roof?”



WARNING

Some providers send consolidated bills once a month or once per quarter. If you wait for the bill to show up, it may include unexpected costs that fit neither your budget nor your chief financial officer's willingness to pay. Cost monitoring in real-time at regular intervals is essential to avoid payment problems. Also, reserving resources means paying for always-on compute, storage, or networking capabilities in the cloud, but at much lower costs. If you manage your costs carefully and plan in advance for predictable demands, you can use more expensive on-demand resources to cover demand spikes as they happen.



TIP

Dealing with cloud service provider bills is more art than science. Make sure you understand how to read and interpret what's on your bills and work with your salesperson to get the best possible deals. This is key to avoiding both “sticker shock” and unpleasant, unexpected payment demands.

Understanding cloud security

In the public cloud, businesses and organizations share responsibility for security with the cloud provider. In the private cloud, the onus is on the owner. This means protecting data, applications, endpoints, and infrastructure from threats and vulnerabilities. It also means protecting data privacy, in motion and at

rest, particularly when such data falls under one or more regulatory frameworks — General Data Protection Regulation (GDPR), Sarbanes-Oxley Act (SOX), Payment Card Industry (PCI), Health Insurance and Portability Act (HIPAA), and so on.

Organizations with private clouds should look to platform hardening, security auditing and reporting, and protection from network threats as part of a “zero-trust” outlook that prevents cyberattacks and avoids data loss or unwanted data disclosures.

Sources of cloud flexibility

What gives the cloud environment its elasticity is its ability to provide more resources on demand for as long as they’re needed, and then to spin them down and turn them off when they’re not. This helps manage costs on public clouds but is equally important on private clouds where it helps to keep active resources in sync with actual consumption, too.

Another powerful source of flexibility comes from hybrid cloud, a highly interoperable form of multi-cloud. Because this cloud architecture permits quick and easy migration of workloads, software, and data from the private cloud to the public cloud, it allows organizations to stick to what they own when it suffices and “burst” out onto the public cloud when more resources are needed. This confers incredible flexibility to scale up (and out) when responding to peak demands, handling unusually large data sets for analytics or machine learning, performing experiments or usability tests, and other “outside the box” computing action.

Capabilities Abounding

Private cloud computing offers many of the primary benefits and capabilities that make the public cloud so attractive and valuable. In a private cloud, businesses and organizations can leverage these capabilities, also found in public clouds:

- » **Self-service:** Modern private clouds include self-service portals so that users can request resources and applications from a catalog of available items and register requests for new or missing items. This cuts down on the need to interact with IT staff. Items often can be granted using automated workflows.

- » **Scalability:** Just as on the public side, private clouds can scale resources up and down when needed. For entirely private clouds, this capability is limited to the total resources available from the organization's datacenters and IT infrastructure. But with a hosted ("single-tenant") private cloud, organizations can burst through that ceiling or use their private cloud just like a public cloud.
- » **Provisioning and configuration of virtual machines (VMs):** Private clouds can install, set up, and configure VMs (or platforms, or even infrastructures) just like the "big boys" do on public clouds.
- » **Automation:** The tools that provide virtualized computing resources, networking, and application/service access in public clouds work on private clouds, too.
- » **Chargebacks:** Organizations allocate costs for private cloud resources using chargebacks levied on their users and their parent departments or divisions. This feature lets internal costs flow to those who incur them and helps keep IT solvent and self-funding.

Given such benefits for private clouds, it's important to understand their limitations and drawbacks. You can overcome or work around such issues, but you can't ignore them:

- » **Legacy infrastructure:** The wrong infrastructure hampers cloud computing. This occurs in low-interoperability situations where a private cloud is not able to interact with public clouds easily (or at all). Likewise, software that handles storage, servers, virtualization, and networking can introduce unwanted complexity and added integration efforts. This can be especially vexing when using storage area network (SAN) or network attached storage (NAS) arrays: upgrading them requires time, effort and (too often, lots of) money.
- » **More software needed:** Connecting to the public cloud from a private cloud requires additional software that may not already be integrated into your IT infrastructure. Organizations with multiple IT teams may be able to build their own private clouds, but smaller teams may struggle to learn and use such new and complex tools.
- » **Vendor lock-in:** Legacy on-demand infrastructures work in a vendor-specific virtualization framework. No matter the vendor or its hypervisor, organizations violently resist being forced to use a single vendor's solutions.

» **Over-packaging/purchasing:** Certain private cloud vendors may bundle multiple solutions that result in unwanted deployment issues and complexity. Businesses may also balk at extra costs associated with such bundles.

The best solution is a purpose-built private cloud environment designed to interoperate with all public clouds but that imposes no specific hypervisor or container architectures.

HCI Works for Private Cloud

Hyperconverged infrastructure (HCI) is well-suited at hosting private clouds. It allows organizations to use on-premises assets and infrastructures but leaves the door open to grow into a highly interoperable, hybrid and/or multi-cloud in the future.

Organizations must make sure their current infrastructure can support future hybrid and multi-cloud initiatives. Ultimately, this means replacing standalone storage arrays (NAS or SAN), physical servers, proprietary virtualization, and legacy networking elements used in typical three-tier architectures. In particular, the separation of process logic (“business rules”), computer data storage, and data access poses the kinds of interoperability issues that hybrid cloud seeks to avoid. That is, there’s no worry about moving to HCI if you have been using a SAN for 20 years or so. HCI has been around for more than a decade and is routine to learn and use. HCI has been approved for big mission-critical apps such as SAP HANA, so making the transformation today is no big thing.

With HCI, separated subsystems merge. Thus, they eliminate silos that divide separated systems and functions. This architecture supports proper private clouds and includes:

- » Reliable security configurations and audits
- » Data-at-rest encryption
- » Micro-segmentation
- » Built-in data protection, backup, and disaster recovery
- » Rapid, automated deployment
- » IT as a service (IaaS)
- » Reduced operating expenses and improved return on investment (ROI)

- » Understanding hyperconverged infrastructure basics
- » Getting infrastructure right
- » Digging into private cloud's key pieces and parts

Chapter 3

Building a Private Cloud

All clouds, public and private, must deliver computing resources to their users. This is the eternal triad: compute, storage, and networking. Delivering such resources depends on an underlying infrastructure. That infrastructure uses underlying physical systems and a variety of specialized software to stand up and deliver virtual machines, virtualized networks, and numerous kinds of “X-as-a-service” offerings. Solving for X turns out to be an interesting exercise in understanding what cloud computing does.

HCI Basics

Hyperconverged infrastructure (HCI) combines common datacenter hardware elements. That is, HCI incorporates locally attached storage resources, networking devices, and compute resources with purpose-built, intelligent software to create flexible building blocks for the same resources that the cloud offers — namely, compute, storage, and networking.

Why use HCI?

HCI modernizes legacy IT infrastructures. That is, HCI replaces discrete servers, storage networks, and storage arrays. This results in a simplified, streamlined IT infrastructure with lower

total costs of ownership (TCO), increased performance, and greater IT productivity. That's a win-win-win! Figure 3-1 shows a stack of legacy gear at the left replaced by one or more powerful HCI boxes at the right.

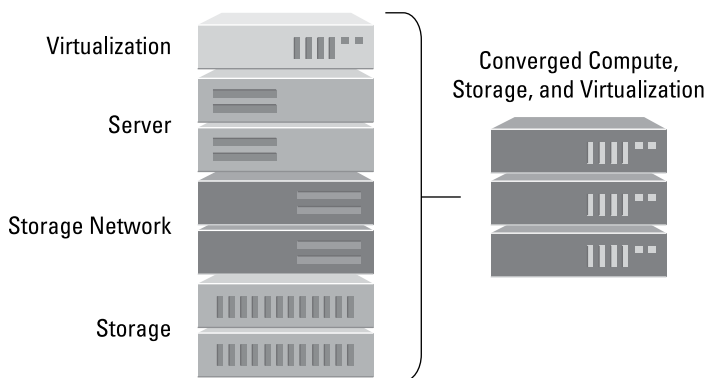


FIGURE 3-1: One or more HCI devices replace a whole stack of legacy gear.

HCI assembles a complete datacenter stack, including compute, storage, storage networking, and virtualization. An HCI platform replaces complex, legacy infrastructure, and runs on turnkey, industry-standard nodes (the right-hand boxes shown in Figure 3-1) that can start small, then scale one node at a time. Purpose-built software runs on each node to distribute all operating functions across all of them — collectively called “a cluster” — to deliver outstanding performance and superior resilience.

HCI features two primary components:

- » **The distributed plane** runs on a cluster to deliver storage, virtualization, and networking services to guest applications. These might run in virtual machines (VMs) or in similar, isolated virtualized runtime environments called *containers*.
- » **The management plane** provides a single, consistent administration interface for all HCI resources. It obviates separate management solutions for servers, storage networks, storage, and virtualization (legacy networks).



A *container* is a standard software component that includes code and all the resources it depends on within a single run-time package. This lets applications run quickly and reliably from one computing environment to another. A container can launch and run quickly, on its own, in any compatible OS.

Compute

Delivering compute resources in HCI means providing virtual compute resources, where a hypervisor places workloads in a cluster. To provide those compute resources (and to support applications and services), the cluster uses constituent HCI CPU elements and memory (RAM) to stand up and position as many VMs or containers as local workloads require on the nodes where they're placed. The hypervisor decides where (and for how long) to run VMs, to best meet current demand.

Storage

HCI intelligent software provides a distributed storage fabric across a cluster. Flexible, distributed storage is the real foundation for HCI technology. It provides local storage services for applications running in the cluster (on a VM or container), and permits pooling of faster flash and slower hard disk drives across the cluster. In fact, the storage layer protects data against hardware or device failure. Also, you don't need discrete networked storage solutions from third-party vendors (such as HPE, Dell-EMC, or NetApp). This HCI storage includes high-end capabilities to accelerate performance, deliver data reduction and protection, and more.



Some HCI implementations do a better job with Software-Defined Storage (SDS, the basis for HCI storage services) than others. A truly distributed filesystem enables resiliency and high performance at any scale. When choosing an HCI platform this is a vital area of functionality to research and experiment with, to make sure it meets your needs.

Networking

A key benefit of HCI is that it obviates the need for separate, dedicated storage networks. Virtualized networking in HCI includes software-defined networking (SDN) and network function virtualization (NFV) capabilities. In particular, an enterprise-grade

HCI solution offers service chains to enable network feature virtualization, along with network visualization and optimization capabilities.

Virtualization

HCI supports a flexible and powerful platform on which to run application services. Consequently, the HCI environment revolves around virtualization as a primary capability. An HCI environment can support virtual machines or containers in which applications run. By supporting any major hypervisor, HCI fits right into the application management framework you're already using.



TECHNICAL
STUFF

Another benefit of supporting any major hypervisor — most enterprises use more than one — is that there's no need to maintain specific vendor attachments because you *must* run their software. You can switch from one hypervisor to another, as you see fit.

HCI lifecycle management

The consistent and coherent management plane built into HCI helps to streamline and simplify IT operations. It coordinates all the different devices, firmware and management tools in your environment under a “single pane of glass.” This tracks firmware and software across the cluster. It also validates and automates updates, firmware, and device or software dependencies. HCI lifecycle management (LCM) works over the entire IT lifecycle from initial setup and deployment, through maintenance, expansion and multi-cloud-integration, to decommissioning and replacement of old, outdated nodes with faster, newer, more capable ones.

The Right Infrastructure Is Key

When it comes to putting HCI in place, indeed, choosing the right infrastructure is the key to success. Above all, it's essential to use an HCI that can not only support a private cloud, but that can also easily and seamlessly integrate and interoperate with multiple public clouds. This is the foundation for a workable, sustainable hybrid cloud that does not lock you into a specific cloud vendor. For example, VMware only works with AWS. Nutanix offers true multicloud with AWS, and will have Azure, Google, and its own Xi cloud shortly.

Private Cloud Pieces and Parts

Many elements go into building a private cloud to support an overarching enterprise cloud vision that also encompasses one or more public clouds and perhaps one or more hosted private clouds as well. The following bullets lay them out:

- » **Data services:** A cloud-capable HCI includes file services to support various filesystem protocols, such as SMB 2.1 for Windows and NFS v4 for Linux/Unix. It also provides object storage, using an industry-standard S3-compatible REST application programming interface (API) to handle huge volumes of unstructured data. And it implements a block service to provide applications with iSCSI block storage access when needed, including physical devices or VMs with specific requirements.
- » **Networking:** A top-notch HCI offers complex virtualized networking capabilities at the cluster level, sufficient to support full-blown PaaS and IaaS configurations. Multiple clusters can share a common VLAN, with individual subnets for each cluster. This makes complex networking within and between datacenters simple and straightforward. For maximum bandwidth, top-of-the-line HCI supports Fiber Channel (up to dual 32Gb links) and Dual 10Gb Ethernet connections. It also integrates with one or more SD-WAN solutions to replace MPLS with broadband WAN connections.
- » **Automation:** IT admins can define reference images, templates, and blueprints for virtualized applications, services, and infrastructures. These can be instantiated quickly and automatically using ready-to-run scripts through the HCI management facility. Such a facility also handles dynamic scheduling to make workload, VM, and volume placement and migration decisions to maximize efficiency and performance. Likewise, automated provisioning and orchestration capabilities make services easy to stand up and tear down, based on activity, demand, and consumption policies.
- » **Self-service options:** In many ways, self-service is central to smooth, rapid delivery of applications and services for any kind of cloud. A top-shelf HCI delivers a full range of

self-service options that include a self-service portal where users go to request access to pre-defined resources, to provision and customize resources available to them from a resource catalog, and to request entire workflows (which involve complex combinations of resources, services, and applications). This capability allows users to drive more of their IT consumption and takes some of that load off IT staff.

- » **Security:** A state-of-the-art HCI environment must be hardened by default and makes use of the well-known security of *least privilege* (never grant more access, rights, or permissions than needed). It implements a “defense-in-depth” security model (the architecture is broken into multiple security layers, each with its own set of defenses and protections).

To provide maximum safety and security for users a top-notch HCI uses two-factor authentication (2FA) and encrypts all data at rest. Its developers should follow a security development lifecycle (SecDL) — security is built in at all stages of development, from design and development to testing and hardening. Threat modeling helps to assess and mitigate customer risk from code changes. SecDL testing should be automated during development. All security-related code modifications must be timed during minor releases to minimize customer risk.

- » **Metering and cost governance:** A top-notch HCI environment offers thorough monitoring of private cloud resource consumption and can generate and deliver chargebacks to users. Working in a multi- or hybrid cloud environment requires ongoing analysis and optimization of third-party costs (either single-tenant hosted private cloud, or per-provider public cloud). Some HCI makers offers special tools to help monitor and manage overall cloud costs, and to provide out-of-the-box visibility into private cloud costs. See also this blog post for a thorough overview of cost analysis, monitoring and management: “28 Actions You Can Take to Manage Cloud Costs” (<https://www.nutanix.com/blog/28-actions-can-take-manage-cloud-costs>).

- » **Compliance:** The best HCI systems security baselines exceed U.S. Department of Defense (DoD) requirements. These systems also comply with a broad suite of security evaluation programs, including TAA, Suite B, 508, FIPS 140-2, the Common Criteria, and more. When it comes to regulatory

compliance, such HCI systems support various data protection mechanisms, including 2FA, access and audit controls, encryption for data in motion and at rest, and more. But companies must understand which regulatory frameworks require their compliance and perform a thorough analysis of those requirements. Only then can they determine how best to meet those requirements. Please see the blog post “5 Keys to Cloud Compliance in 2019” to get you started on this journey (<https://www.nutanix.com/blog/5-keys-cloud-compliance-2019>).

» **Disaster recovery and more:** Top-shelf HCI offers capabilities across the stack to address high availability, data protection, resiliency, and disaster recovery/business continuity coverage. To help organizations avoid the complexity of buying and integrating separate services or solutions for all these things, HCI should engineer high availability and data protection into its platforms.

VMs and applications should be resistant to hardware failures (at the node level), thanks to redundant power supplies, memory, and CPU. The ability to access storage in the event of individual disk failures comes from a distributed data architecture that automatically rebuilds affected VMs and their associated data. The more nodes in a cluster, the faster this goes.

Controller and hypervisor upgrades are also touch-free and non-disruptive. A feature called “tunable redundancy” causes synchronous replication across multiple nodes in a cluster (2-node minimum, obviously). Tunable redundancy also lets organizations set levels of protection based on application SLAs, and enable policies at the VM level. This ties further into availability domains that determine optimal placement for replicas based on awareness of available hardware resources. (To learn more about tunable redundancy, see http://stevenpoitras.com/the-nutanix-bible/#data_protection. For more on availability domains, see http://stevenpoitras.com/the-nutanix-bible/#avail_domains.)

Backup occurs at the VM level, either to a remote (private) site or to the public cloud (a special cloud connector may be required). Backups should be WAN optimized by default, with data de-duplicated and compressed before transmission. Differential backup means only actual changes need to

traverse any networks. For many customers, this is enough capability to kill the need for other backup solutions. VSS (Microsoft's Volume Shadow Service) is integrated into native HCI snapshots, so backups for Microsoft workloads including Microsoft SQL, Exchange, and other platforms remain application consistent. A top-notch HCI can, however, integrate with third-party backup solutions when that capability is needed.

The HCI environment must include enterprise-grade disaster recovery capabilities. Then, VMs can be replicated locally or to another site on a schedule tailored for each workload. Support for one-to-many, many-to-one, and many-to-many replications means protection for all workloads across all sites. The DR process can be automated completely, using REST APIs, command line interfaces, and HCI-based management, to ensure minimal downtime when a disaster occurs. The best HCI vendors are working on continuous availability across different metro areas, to boost uptime and availability.

Nutanix HCI = Firm Foundation

Nutanix HCI solutions are engineered to enable enterprises to start out with on-premises components and scale up their datacenters as they deem appropriate. But by deliberate design, Nutanix HCI also extends easily and readily into multi-tenant public clouds as needed. It can integrate with hosted single-tenant private cloud technologies, too. That's what makes Nutanix HCI a firm and viable foundation for a vision of enterprise cloud architecture that accommodates private, public, multi-, and hybrid cloud environments.

IN THIS CHAPTER

- » Working through automation basics
- » Understanding automation in the private cloud
- » Selecting suitable targets for automation
- » Looking at Nutanix automation tools and services
- » Making the most of Nutanix self-service

Chapter 4

Automating the Cloud

Automation means turning over tasks to software. On the one hand, that frees IT staff to work on things that can't be automated — such as creating automation scripts. On the other hand, it means that automation software can be called upon whenever it's needed. This might occur when certain things happen (event-driven triggers), or at regular, scheduled intervals (scheduled triggers), or when an IT person decides to manually fire off an automation sequence (manual triggers) to handle some particular task. It may not sound like much, but it provides great capability and power, as this chapter shows.

Cloud Automation Basics

Cloud automation uses tools and services to eliminate tedious, time-consuming manual processes involved in working with cloud technologies. Manual processes require IT staff to sit down at their PCs, open up various tools or consoles, and then drive user interfaces and enter inputs one command or instruction at a time. Such commands, instructions, actions, keystrokes, and so on can instead be captured and stored in a file that does what a human used to do (or a good imitation).

What kinds of things can be automated? In general, just about anything a person can do on a computer — given the right tools and sufficient time, ingenuity, and testing — an automation script can do also. Working in a datacenter or on a cloud computing environment (which may mean the same thing for an on-premises private cloud) generally involves working with virtual machines or containers, applications or services, file systems and storage, workloads, and so on.



REMEMBER

In computing, a *workload* represents some processing activity that consumes resources: compute, storage, and networking. A workload encompasses an application or service that must be run, along with some number of users who will connect to and interact with that application or service. A workload might run only for a specific time, or until it has been idle for a while (timeout interval).

When it comes to working in a datacenter or with a cloud computing environment, though, typical tasks suitable for automation include:

- » **Managing resources:** Before they can be consumed, compute, networking, and storage must be named and defined, configured, provisioned or allocated, have access granted, and more. Automation has this covered.
- » **Setting up virtual machines (VMs) or VM clusters:** Individual virtual servers or collections of servers must be named and defined, the operating system installed, access granted, applications and services installed, and so forth. The whole process might be automated, or just bits and pieces, depending on how IT wants to work it.
- » **Creating virtual networks:** Networks need names, addresses, routing, security, protocols, and more. Automation does this stuff, too.
- » **Deploying workloads:** This is where users get hooked up with applications or services they can use. Here again, automation has it covered.

These are basic example of infrastructure automation. Nutanix can auto-right-size resources (add or remove CPU or RAM as needed), and auto-send emails, Slack messages, or even IT service management (ITSM) tickets. Nutanix can integrate shell

scripts and API calls with private cloud automation and even trigger POST commands with webhooks.

Benefits of automation

Sure, automation relieves IT staff of tedious tasks. This is important because it means people can spend their time doing other things (including automating more tasks). But there's more to automation than skipping over the boring, repetitive parts, however valuable that may be. This includes:

- » Automation runs at computing speeds, which are much faster than human speeds — 10X, 100X, or better!
- » Automation handles complex as well as simple tasks. The more complex the automated task, the bigger the benefits from avoiding human time delays (and errors).
- » Automation must be tested, tested, tested until it works correctly, as expected and desired. Once deployed into production, it pays off *big* in increased reliability.
- » Automation works on a variety of triggers. Thus, it can respond to security events, device errors or failures, traffic or consumption thresholds, and spending data.

Overall, automation makes for faster, more responsive, more reliable, and more efficient systems. It improves both IT effectiveness and return on investment (ROI).

Workflows, policies, and processes

Automation handles all kinds of tasks. IT's most visible responsibility may be to define and deliver workloads, so users can access applications and services. Indeed, these can — and probably should — be automated. But, other important IT work serves to define policies of all kinds.

Such policies include acceptable use, access controls, role-based security profiles, and much, much more — see the SANS library of information security policy templates (<https://www.sans.org/security-resources/policies/>) for one key area that policy covers. The same is true for IT processes, which usually include onboarding new employees, assigning a job role (with its policies, applications, and access rights), delivering new office or mobile computing equipment, creating or expanding Nutanix clusters,

defining workflows and self-service portals and projects, and more.

All these things happen regularly, perhaps even frequently. This makes them ripe targets for automation.

Selecting suitable targets

In general, frequency of occurrence is a reliable metric for selecting tasks, activities, or responses suitable for automation. The more often something occurs, the more beneficial automating it can be. Other criteria for selection must also include:

- » **Time is money (or risk):** The speed at which automation works makes it ideal to initiate incident response, handle errors or failures, and dampen spending spikes.
- » **Sending up red flags:** A variation on the preceding item, to make sure IT workers know what's going on, too.
- » **Enforcing policy:** Given the importance of automation in defining policy, it makes sense to automate policy activity, including potential problems or violations.
- » **Defining and delivering workloads:** This is a great way to make applications and services available to users. With access to self-service capabilities (covered later in this chapter) users themselves help make IT more efficient.

Services and Tools for Private Cloud

In the private cloud environment, more is better when it comes to automating activities that occur frequently, perhaps even many times a day. If such a common set of activities can be identified, then they can also be automated.

Often, this depends on the ability of an HCI or private cloud environment to take a series of events or alerts and permit the system to intercept and record them, then replay them as needed, with a set of specific inputs to customize them.

Such sequences of alerts and activities may be called *playbooks*, because they associate an alert, event or activity with the start of a sequence (the “trigger”). They continue by recording the series of

actions that the playbook should take in response to that trigger. For each action, its details can be supplied. The actions a playbook can call on are many and varied. This might involve sending email or a Slack message, making API calls, or creating new automation scripts. Thus, playbooks can interact with other systems such as a configuration management database (CMDB) or help/support desk ticketing facilities. As an added bonus, such automation makes sure all systems in the private cloud or hyperconverged infrastructure (HCI) environment use the same definitions and configurations.

Because APIs are crucial for automation, it's essential to understand which ones are available. The primary APIs available to Nutanix automation include:

- » **REST API:** A set of web APIs for “representational state transfer” that permit resource access via specific URLs. REST exposes every capability of the Prism UI (the primary Nutanix management interface) to drive Nutanix actions. In addition, this means users can employ tools such as Saltstack, Puppet, vRealize Operations, System Center Orchestrator, Ansible, and so on to create custom workflows for Nutanix. There's even an HTML5 GUI so IT staff can create Nutanix automations as web pages.
- » **Various command line interfaces (CLIs):** These include the aCLI (Acropolis CLI) and nCLI (Nutanix CLI). Together they provide command-line access to most Nutanix features and functions. (To learn more about aCLI, see <https://portal.nutanix.com/page/documents/details/?targetId=Command-Ref-AOS-v510:man-acli-c.html>. For more on nCLI, see <https://portal.nutanix.com/page/documents/details/?targetId=Command-Ref-AOS-v51:Nutanix-Command-Line-Interface-Reference>.)
- » **Scripting interfaces for scripting languages such as PowerShell, go, Python and more:** The Nutanix GitHub repository (<https://github.com/nutanix>) is a good place to start looking for more information.

Automation examples

The Nutanix Bible (<https://nutanixbible.com>) is a comprehensive and nicely illustrated online document that explains the

company's systems and capabilities. It is particularly informative about automation. Interested readers will find a good discussion of Playbooks, the automated “scripts” within Prism Pro X-Play that show how to create automations, step-by-step.

In particular, the section of the Bible entitled “Book of Acropolis” is key to automation, because it covers the primary Nutanix tools for workload and resource management (usually key foci for automation).

The Automation repository on the Nutanix GitHub page (<https://github.com/nutanix/Automation>) includes a series of automations from the community well worth investigating. These include scripts for configuring a basic cluster, for working with VMs, and for working with blueprints (import, move, save).



TIP

In the Nutanix Calm subsystem, a *blueprint* defines a model for the details of how an application runs on the cloud. A blueprint is a kind of application recipe that encompasses application architecture and infrastructure elements, along with provisioning and deployment steps, application binaries, command steps, monitoring endpoints, remediation steps, licensing and monetization data, and policy definitions. In short, it's a complete collection of everything that's needed to fire off an application in the Nutanix environment. Thus, whenever a blueprint is executed, it spawns an application. Read more on this fascinating topic in *Nutanix Calm: The What and How of Calm Blueprints* (<https://next.nutanix.com/blog-40/nutanix-calm-the-what-and-how-of-calm-blueprints-31480>).

The self-service advantage

One core capability that developers and business users love about the public cloud comes from its ability to provision applications and virtual machines without making IT get involved. Nutanix Self-Service has this same goal — namely, to bring this simplicity and ease of access to private clouds. In a nutshell, self-service lets users deploy applications on demand based on policies set by IT administrators. It's a powerful case for automation all by itself.

The Nutanix Self-Service Portal is a key ingredient of the company's private cloud offering. Admins create a catalog of projects, then use Active Directory or LDAP integration to assign users and resources (including compute, storage, and networking) to those

projects. When users log in to the environment with their credentials, what they see depends on their access rights (usually based on job roles). They can choose projects they wish to join through the self-service portal. This simplifies access for developers and users alike and lowers the burden on IT.

At present self-service works within Prism Central in the Nutanix private cloud environment. It exposes resources from a single Nutanix advanced hypervisor cluster. In Prism, Nutanix supports full role-based access controls (RBAC), with distinct roles for Prism Central administration, self-service administration, and project users.



TIP

You'll find Nutanix's Prism Self-Service Overview (<https://portal.nutanix.com/page/documents/details/?targetId=SSP-Admin-Guide-v55:ssp-ssp-overview-for-ssp-admin-c.html>) helpful in digging into more self-service details. Prism and its interface are also described in the Nutanix Bible (<https://nutanixbible.com>) — search for “Prism.”

Tools and automation together

Behind the scenes, what makes self-service and resource management work in the Nutanix private cloud environment depends a great deal on built-in automation. Working with Calm and Prism Pro — especially PrismPro's X-Play tool — provides an excellent demonstration of what Nutanix does with automation. This also exposes the automation capabilities Nutanix makes available to administrators and IT staff.

IN THIS CHAPTER

- » Exploring and exploiting hybrid and multi-cloud benefits
- » Crafting the correct hybrid and multi-cloud strategy
- » Bringing hybrid and multi-cloud into your organization

Chapter 5

Exploring Hybrid and Multi-Cloud Vigor

In a nebulous sense, a hybrid cloud is an architecture that combines on-premises IT — which could be traditional legacy infrastructure or a private cloud — with off-premises services or resources from a public cloud (such as Azure, Amazon, or Alibaba). That said, the “outside connection” element in hybrid and multi-cloud may also be a cloud service provider (CSP) or even a software-as-a-service (SaaS) provider, such as Oracle, SAP, Salesforce, and so on.

Many possible connections can go into hybrid and multi-cloud. Indeed, there are many ways to construe or define a hybrid and multi-cloud, so this chapter digs into these foggy constructs.

What IS Hybrid and Multi-Cloud?

A more exact definition of hybrid and multi-cloud starts with the idea of a service created from a combination of different clouds (shown in Figure 5-1). These could include public and private clouds, plus CSPs. Using a three-tiered application stack, the

following division of residence might make sense for a hybrid and multi-cloud implementation of such a service:

- »» The presentation service might be on a public cloud.
- »» The application services might reside on a managed private cloud.
- »» The database service might reside on-premises.

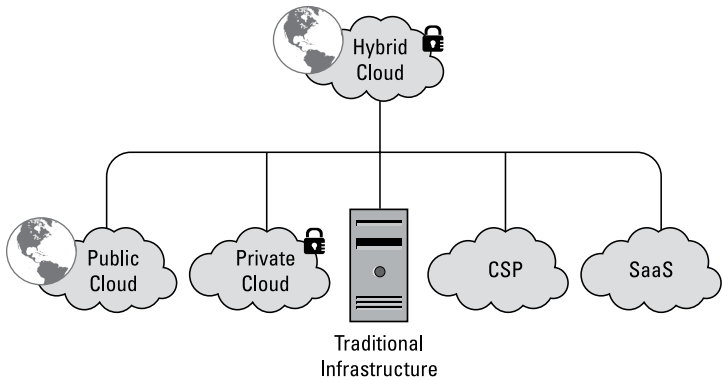


FIGURE 5-1: A hybrid cloud combines elements from on- and off-premises.

A THREE-TIERED APPLICATION STACK

This is tried-and-true technique for dividing an application into layers because each has separate concerns, capabilities, and responsibilities. The *presentation tier* is also known as a user interface or client application and handles user interaction. The *business logic tier* is also known as the application server (or application services) and handles data collection, processing, analysis, visualization, and all the other “stuff” that applications do. The *data storage tier* is also known as the database server (or database services). It handles only reading and writing of database records, as dictated by the application services tier. And there you have all three tiers, no icing needed.

In a three-tiered architecture it makes sense to run presentation services on a public cloud. Public clouds are easily accessible anywhere and everywhere, so that service will be readily available to users.

Given that an application can contain proprietary logic, algorithms, or data, it's reasonable to run it on a private cloud. This gives its owner the flexibility to grow and shrink resources with demand, on a pay-as-you-go basis, and leaves the owner firmly and fully in charge. Last, because proprietary data is the lifeblood of the organization, it makes sense to run its containing database on-premises. Capisce?

Why Hybrid and Multi-Cloud?

For many good reasons, it makes sense to evolve an organization's IT infrastructure in the hybrid and multi-cloud direction. But this takes planning and forethought and may require changes across everything IT has and does right now.

That's because elements that combine to create a hybrid and multi-cloud must be highly interoperable. This allows data and workloads to migrate around all of its pieces and parts with minimal friction and resistance. This is not something organizations can take for granted, or that happens spontaneously. Fact is: The odds are better than even that some of your on-premises IT infrastructure must change to support hybrid and multi-cloud capabilities. Ditto for existing public cloud (and private cloud, if you've already got it).



WARNING

Differences in APIs, data, image formats, and architectures pose interesting (and serious) challenges to workload migration. Sure, information wants to be free, but freedom of migration is tough!

Benefits of hybrid and multi-cloud

Given that thought, time, effort, and money must all be expended to take the hybrid and multi-cloud route, why bother? This move offers numerous benefits, which means considerable gains to offset any pain. These benefits include:

- » **Flexibility and agility:** Properly architected, hybrid and multi-cloud pays considerable dividends in business agility. Organizations gain ready access to resources that can

support new applications, and handle development, test, and what-if experiment projects. If something unexpected happens, you can flex your infrastructure to squeeze it in. If there's sufficient interoperability among constituent parts, you can move workloads from on-premises to in-cloud locations, too.

» **Elasticity:** Business and industries move in cycles. Thus, they may move through periods of peak and weak demand. Take retail, for example: Activity spikes near year's end during the holiday shopping season. Sales may fall off in following quarters. A hybrid and multi-cloud model lets organizations adapt to changing resource demands. In the same vein, some applications may show large swings in demand (such as AI or analytics workloads when processing large data sets, versus performing regressions or statistical analyses). What's nice about the hybrid and multi-cloud is that it can supply huge swaths of resources on demand, and release them when they're no longer needed. This lowers overall costs.

» **Easy workload migration:** The ability to easily migrate workloads to the right cloud (or another cloud) without having to re-architect apps means organizations can use the right cloud resources for their workloads. In particular, this supports "cloud bursts" for peak demands.

Nutanix offers complete "lift and shift" application migration across your entire IT infrastructure. That includes your datacenters and all of your private and public clouds. This removes the sting from potential migration problems and hurdles.

» **Speed new product and service delivery:** Because it lifts barriers that can impede business and development teams, hybrid and multi-cloud boosts the pace of innovation. It isn't just easier to create and deploy new digital services. The creative minds behind building and selling them can do more, quicker, and better. This can create new opportunities, and even new markets.

» **Control costs:** The hybrid and multi-cloud lets organization optimize application use and performance. Its pay-as-you-go model reduces CapEx for infrastructure and datacenter. This avoids the "capital sink" that occurs when on-premises assets sit idle (it's much cheaper to not use cloud resources than real ones, and vice versa).



REMEMBER

- » **Avoid vendor lock-in:** For those who adopt a cloud-first or cloud-only IT approach, it can be difficult to avoid locking in with one or more cloud vendors. It is costly to extract your data from the cloud, so it's wise to exercise restraint before pouring all your data into the cloud. A hybrid and multi-cloud, properly architected, gives your data more freedom (and much greater ease) of movement.
- » **Avoid obsolescence:** The faster an organization can access new technologies, the bigger the boost to its competitive advantage (think AI, for example). Large public clouds innovate quickly and offer competing services and deals. A hybrid and multi-cloud approach lets you use the best available technology so you can *carpe diem*.

Hybrid and multi-cloud strategy

To deliver the best returns on IT investment, organizations need a kind of “cloud operating system” that can manage, monitor, and orchestrate across all environments — on-premises, and all flavors of private and public clouds — using one set of tools.



TIP

A hybrid and multi-cloud integrates both control and data planes. Control plane integration provides a single point of control for all cloud assets, public and private. The data plane handles migration of applications and data between clouds. A truly workable solution enables both planes so organizations can build and operate a true hybrid and multi-cloud.

Indeed, it's best to choose a hybrid and multi-cloud based on a careful strategy of election and selection:

- » Choose one framework — this becomes the “cloud operating system” — that enables your organization to manage workloads both on-premises and in the cloud.
- » Update the on-premises systems and infrastructures to support that framework (to handle the “on-prem” side of hybrid and multi-cloud). This means private cloud, too.
- » Choose *only* those public clouds and CSPs that also support that framework.

Extending a Private Cloud

Ambitious cloud architects want to exploit the hybrid and multi-cloud for a variety of interesting reasons. They include capacity bursting, workload optimization, test/development, data analytics, and failover support for disaster recovery or business continuity. The best of hybrid and multi-cloud is possible when it's easy for workloads to migrate from an on-premises infrastructure — best case, already in a private cloud — into the public cloud or additional hosted private cloud resources from a single-tenant CSP offering.

The ultimate goal is to obtain seamless workload and data mobility between private and public cloud infrastructures. With the right solution in place, organizations can leverage their “cloud operating system” choice to take advantage of a public cloud provider's storage, compute, networking, and perhaps even data analytics services. The private cloud is where things start, but the public cloud serves as a kind of overflow catchment to permit organizations to handle peak demands, test/development scenarios, or one-off and periodic compute- and storage-intensive processing tasks (big data analyses, research projects, end-of-cycle data collection and reporting, and so on).

Types of Hybrid and Multi-Clouds

Given the prolific numbers of cloud combinations, it's inevitable that hybrid and multi-clouds should come in many forms. These generally depend on the business model and offerings available from the company that provides an organization with cloud services, and how that consuming organization takes advantage of what a provider has to offer.

Public service providers

In the hybrid and multi-cloud strategy stated earlier in this chapter, the rubric for selecting a cloud provider reads “Choose *only* those public clouds and CSPs that also support that framework” that matches the organization's “cloud operating system” choice. This means that data and workloads migrate easily between the organization's own private cloud and its public cloud connections. This is highly desirable, and should not only be inquired about,

but should also be tested and verified with simple, basic proof-of-concept testing. Ideally, all API calls and automation tools will work, and data and workloads move easily and quickly from private to public cloud (and vice versa). You can't know until you check!

Managed service providers

The same thing that's true for public cloud providers is also true for MSPs and CSPs. Realizing the promise of hybrid cloud means that organizations must test and verify the ability to migrate workloads, move data, and run the same automation and management tools on the provider's cloud that work so well in the private cloud. Once again: This poses challenges, so organizations should expect to expend some time and effort in making sure that their multi-cloud environment really works as a hybrid cloud is supposed to.

Nutanix Clusters

Nutanix Clusters provides tools and technologies to interlink private and public clouds for seamless interoperation. Thus, it delivers powerful workload and data mobility between public and private cloud infrastructures. Nutanix Clusters is a service that provides the underlying hybrid and multi-cloud platform needed to extend on-premises environments to one or more (usually more) public clouds.

Nutanix Clusters consolidates management across all constituent clouds in hybrid and multi-cloud environments. Organizations can easily access all cloud resources, set up and deploy new instances, and more. They can also take advantage of centralized governance and control that continuously optimizes resource consumption across all cloud environments — private, public, and even on the network edge (an increasingly important portion of the computing landscape).

The key benefits of Nutanix Clusters include:

- » Lift-and-shift of applications from on-premises to hybrid cloud means users benefit from using public cloud services without having to spend time and effort in re-architecting applications.

- » Easy “burstability” from on-premises into a hybrid cloud environment by leveraging public cloud flexibility.
- » Unifying skillsets and management tooling between on-premises and public cloud.
- » Complete license portability between on-premises and hybrid cloud lets organizations use existing on-premises investments in the hybrid cloud and benefit from cost efficiency.



TIP

Nutanix also offers cost efficient consumption through an on-demand hibernate or resume feature that sends your data into cheap long-term storage and shuts down expensive compute instances when you don’t need them. You do not have to pay for hybrid and multi-cloud compute instances 24/7/365 with Nutanix Clusters — pay only as you use them.

In addition, Nutanix offers a collection of SaaS services that include the following, all of which help to extend cloud capabilities, private and public:

- » **Beam:** Monitor, manage, and optimize cloud expenditures while ensuring security compliance across multi-cloud environments.
- » **Leap:** Obtain single-click access to natively integrated, cloud-based disaster recovery capabilities. Xi Leap leverages and extends Nutanix’s built-in snapshot and backup mechanisms.
- » **Frame:** Enable users to access secure, software-defined virtual desktop workspaces in any cloud environment. Xi Frame brings VDI onto any Internet-connected device.

Within the organization, the hybrid and multi-cloud vision rests on a capable and flexible private cloud as its foundation. Nutanix HCI and its clusters, together with Nutanix software, define a capable cloud OS that can — and does — integrate with public cloud providers, CSPs, and other cloud players to offer a true hybrid and multi-cloud experience. For more information and guidance, please consult the Nutanix-sponsored e-book *Designing and Building a Hybrid Cloud* (<https://www.nutanix.com/go/designing-and-building-a-hybrid-cloud>). It picks up where this book leaves off.

IN THIS CHAPTER

- » Understanding that “the cloud changes everything”
- » Seeing what’s in the cloud lets you manage, optimize, and automate it, too
- » Building on the right private cloud leads naturally to a true hybrid cloud

Chapter 6

Nine Nuggets of Private Cloud Wisdom

This chapter offers nine nuggets of wisdom to guide you on your private cloud journey:

- » **The cloud revolution:** Use of cloud technologies has forever changed the way people design, build, and consume computing services. The Internet blurs the distinction between “here” and “there” (or rather, between “here” and “in the cloud”).
- » **Trends in cloud adoption:** Most organizations are taking a hybrid or multi-cloud approach to computing. In this emerging marketplace, private clouds play an increasingly important role. See the “Flexera 2020 State of the Cloud Report” (<https://resources.flexera.com/web/pdf/report-state-of-the-cloud-2020.pdf>) for all kinds of fascinating details.
- » **Benefits of private cloud:** Private clouds are like public ones, except owners must absorb CapEx and OpEx costs to run their own infrastructures rather than a pay-as-you-go model. Benefits include IT-as-a-service (ITaaS); improved latency, privacy, and confidentiality; self-service portals for easy application access; reliable security configurations and

audits; data-at-rest encryption; built-in data protection, backup, and disaster recovery; rapid, non-disrupted deployment of applications and updates; reduced operating expenses; and improved return on investment (ROI).

- » **Visibility:** Private clouds depend on the ease of management, monitoring, and reporting that a “single pane of glass,” management overview brings. Nutanix HCI and its private cloud software offers a powerful dashboard that lets IT administrators view, manage, and control the environment from one place.
- » **Cost controls:** Whether on premises or off, IT needs to know what users are consuming and what it’s spending on internal and external resources and services. From the “Nutanix 2019 Cloud Usage Report,” (<https://www.nutanix.com/go/cloud-usage-report-2019>) resource elimination and rightsizing deliver the biggest savings to cloud users. Nutanix helps companies get more from private clouds via charge-backs and usage reports.
- » **Managing workloads:** Automation takes the drudgery out of repetitive tasks. It creates a platform on which organizations can build complex, reusable infrastructures, services, and applications. Recognizing, defining, and managing workloads is the real key to cloud success, whether private, public, multi-, or hybrid.
- » **Extending the private cloud:** Picking the right private cloud essentially defines a cloud operating system that controls access to resources, provides visibility and management, and supports automation. Organizations that choose the right framework for their cloud OS can manage workloads both on-premises and in the cloud. They should update their on-premises systems and infrastructures to support that framework, too — in a private cloud — and choose public clouds and cloud service providers (CSPs) that support that framework.
- » **Surviving slings and arrows:** A modern IT environment must be able to handle errors, failures, disconnects, and crashes. Look for a solution that offers built-in snapshot and backup mechanisms to protect data and virtual machines. Makes sure your hardware is engineered for fault tolerance to withstand component failures with redundant power supplies.
- » **Exploit the Nutanix advantage:** Nutanix wants to help organizations modernize datacenters so they can run applications at any scale, on premises or in the cloud. For more information, please visit [Nutanix.com](https://www.nutanix.com).

NUTANIX™



Your clouds. All together now.

Effortlessly move apps and data between all your clouds
and build a true hybrid cloud.

nutanix.com/together

These materials are © 2021 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

Embrace the advantages of a private cloud

The goal of any cloud is to place and manage workloads where they'll provide the best results — and best economics — and to move them whenever and wherever they might be needed. This book is your guide to the ins and outs of private clouds and how they fit best into the hybrid and multi-cloud vision.

Inside...

- Explore the cloud revolution
- Manage and optimize workloads
- Make the most of cloud metrics
- Manage costs and resources
- Automate the private cloud
- Craft the correct hybrid and multi-cloud strategy

NUTANIX™

Ed Tittel is an author, trainer, and consultant with more than 100 technology books to his credit.

Go to **Dummies.com™**
for videos, step-by-step photos,
how-to articles, or to shop!

ISBN: 978-1-119-71955-7

Not For Resale



for
dummies®
A Wiley Brand

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.