

LEARNING MADE EASY



2nd HPE
Special Edition

Multicloud Storage

for
dummies[®]
A Wiley Brand



Accelerate digital
transformation

—
Make data center
infrastructure cloud-ready

—
Optimize storage across
clouds and data centers

Brought to
you by:



**Hewlett Packard
Enterprise**

**Douglas Ko
John Boutelle**

About Hewlett Packard Enterprise

Hewlett Packard Enterprise is an industry-leading technology company that enables customers to go further, faster. With the industry's most comprehensive portfolio, spanning the cloud to the data center to workplace applications, our technology and services help customers around the world make IT more efficient, more productive, and more secure.

Hewlett Packard Enterprise and Intel Corporation

Hewlett Packard Enterprise and Intel have long been strong collaborators. Intel's industry-leading server, storage, and networking components are used in HPE servers and storage – including Intel® Xeon® processors and Intel SSDs – offering customers the performance, orchestration, and security features they need to better manage shared resources and modernize their data centers. Together, HPE and Intel provide organizations with a range of innovations that are designed for optimized performance in a hybrid infrastructure environment.



Multicloud Storage

2nd HPE Special Edition

by Douglas Ko and John Boutelle

**for
dummies**[®]
A Wiley Brand

Multicloud Storage For Dummies® , 2nd HPE Special Edition

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

Copyright © 2018 by John Wiley & Sons, Inc., Hoboken, New Jersey

Some cited material © Forrester Research, Inc.

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, The Dummies Way, Dummies.com, Making Everything Easier, 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. 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.

ISBN 978-1-119-51634-7 (pbk); ISBN 978-1-119-51689-7 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

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.

Publisher's Acknowledgments

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

Development Editor: Elizabeth Kuball

Copy Editor: Elizabeth Kuball

Acquisitions Editor: Amy Fandrei

Editorial Manager: Rev Mengle

Business Development Representative:
Karen Hattan

Production Editor: G. Vasanth Koilraj

Special Help: Gavin Cohen;

Valerie Glover; Sandeep Karmarkar;

Shanta Kohli; Manda Mafy;

Lawrence Miller, CISSP; David Wang;

and Simon Watkins

Table of Contents

INTRODUCTION	1
About This Book	1
Foolish Assumptions	2
Icons Used in This Book	2
Beyond the Book	3
Where to Go from Here	3
CHAPTER 1: Moving to Multicloud: Simply a Smarter Strategy	5
Identifying Multicloud's Role in Digital Transformation	6
Modernizing Your Infrastructure with a Multicloud Strategy	8
Moving to Multicloud Storage: It Just Makes Sense	9
CHAPTER 2: Overcoming Barriers to Multicloud Adoption	13
Using Multicloud to Debunk Common Cloud Concerns	14
Addressing Frequently Overlooked (But Important) Considerations	19
Organizational and cultural challenges: Is IT ready for multicloud?	20
Educational challenges: How much will people have to learn?	20
Financial considerations: How solid is the business case?	20
CHAPTER 3: Architecting the Multicloud Data Center	21
Your Data Center Lives On: Addressing New Challenges in the Multicloud Era	22
Keeping the lights on and meeting service-level agreements	22
Integrating or future-proofing for cloud	22
Finding and fixing performance and availability issues	23
Managing costs	24
Overcoming Challenges with Predictive Analytics and Flash Storage Innovations	25
Eliminating the app-data gap	25
Pinpointing problems anywhere in the stack	25

	Anticipating resource needs.....	26
	Increasing uptime via installed-base learning.....	26
	Visualizing your environment.....	27
	Reining in costs	27
	Choosing Cloud-Ready Infrastructure	28
	Native support for cloud.....	28
	Support for any workload, anywhere.....	28
	Global visibility, management, and automation	29
	Pay-as-you-go	29
	Flash for performance.....	29
	Predictive analytics for simplicity and reliability	30
	Easy migration of data between clouds.....	30
CHAPTER 4:	Evaluating Multicloud Storage Services.....	31
	The Four Waves of Cloud Adoption: Where Are You?.....	31
	The Evolution of Cloud Storage Services	32
	Object storage	32
	Cloud block storage	33
	Cloud-ready data protection	33
	True multicloud storage volumes	34
	Market Impact of Multicloud Storage Services	34
	Checklist: Core Requirements for Multicloud Storage.....	36
	Enterprise-grade features.....	36
	Easy mobility: Move data where and when you need it	37
	Visibility: Insights from monitoring and predictive analytics	38
CHAPTER 5:	Ten Multicloud Storage Use Cases.....	39
	Expanding Cloud Backup Options to Cut Costs.....	39
	On-Ramping/Off-Ramping Data	40
	Clustering Mission-Critical SQL Databases in the Cloud	41
	Accelerating Development, Automating DevOps	41
	Moving an Application Using Docker Containers	42
	Using the Cloud for Disaster Recovery	42
	Cloud Bursting for Free (Almost).....	43
	Switching Between Public Cloud Providers in a Flash	43
	Monitoring and Predicting Usage Costs	44
	Using Cloud Storage like You'd Use Your Storage Area Network.....	44

Introduction

Most companies would love to move enterprise applications to the public cloud and cut the costs, complexities, and limitations of current infrastructure.

Major public cloud providers, such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud have done well in enabling web, mobile, and content applications in the cloud. These are cloud-native applications designed to use cloud “object” storage, and object storage is very popular for use cases such as backup and disaster recovery. However, enterprise applications typically require “block” storage. Although native cloud block storage exists, it has several shortcomings, particularly around reliability, durability, and lack of data mobility. So, enterprise applications have largely stayed on-premises or in private clouds.

Now there’s an innovative new way to move enterprise applications to the public cloud while actually reducing risks and trade-offs. It’s called multicloud storage, and it’s an insanely simple, reliable, secure way to deploy your enterprise apps in the cloud and also move them between clouds and on-premises infrastructure, with no vendor lock-in. Multicloud storage allows you to simplify your infrastructure, meet your service-level agreements, and save a bundle.

About This Book

Multicloud Storage For Dummies, 2nd HPE Special Edition, consists of five short chapters that explore the following:

- » How the multicloud storage model aligns with modern business and IT initiatives (Chapter 1)
- » Common barriers to cloud adoption and how a multicloud storage model addresses them (Chapter 2)
- » How to build a multicloud data center (Chapter 3)
- » What to look for in multicloud storage services (Chapter 4)
- » Real-world multicloud use cases (Chapter 5)

Foolish Assumptions

It's been said that most assumptions have outlived their usefulness, but I assume a few things nonetheless!

Mainly, I assume that you are a cloud or storage architect with some knowledge of cloud and storage technologies. As such, this book is written primarily for technical readers. With that said, this book takes nothing for granted, so I make sure all the terms are defined and all the concepts are clearly explained so that you can fully grasp what an inspired concept multicloud storage is and what it can do for your organization.

If any of these assumptions describes you, then this book is for you! If none of these assumptions describes you, keep reading anyway. It's a great book, and when you finish reading it, you'll know quite a bit about multicloud storage!

Icons Used in This Book

Throughout this book, I occasionally use special icons to call attention to important information. Here's what to expect:



REMEMBER

This icon points out information you should commit to your non-volatile memory, your gray matter, or your noggin — along with anniversaries and birthdays!



TECHNICAL
STUFF

You won't find a map of the human genome here, but if you seek to attain the seventh level of NERD-vana, perk up! This icon explains the jargon beneath the jargon!



TIP

Tips are appreciated, never expected — and I sure hope you'll appreciate these tips. This icon points out useful nuggets of information.



WARNING

These alerts point out the stuff your mother warned you about (well, probably not), but they do offer practical advice to help you avoid potentially costly or frustrating mistakes.

Beyond the Book

There's only so much I can cover in 48 short pages, so if you find yourself at the end of this book, thinking "Where can I learn more?" just go to www.hpe.com/us/en/storage/multi-cloud.html.

Where to Go from Here

With our apologies to Lewis Carroll, Alice, and the Cheshire cat:

"Would you tell me, please, which way I ought to go from here?"

"That depends a good deal on where you want to get to," said the Cat — er, the Dummies Man.

"I don't much care where . . .," said Alice.

"Then it doesn't matter which way you go!"

That's certainly true of *Multicloud Storage For Dummies*, which, like *Alice in Wonderland*, is also destined to become a timeless classic!

If you don't know where you're going, any chapter will get you there — but Chapter 1 might be a good place to start! However, if you see a particular topic that piques your interest, feel free to jump ahead to that chapter. Each chapter is written to stand on its own, so you can read this book in any order that suits you (though I don't recommend upside down or backward).

I promise you won't get lost falling down the rabbit hole!

4 Multicloud Storage For Dummies, 2nd HPE Special Edition

IN THIS CHAPTER

- » Connecting the dots: multicloud, modernization, and digital transformation
- » Harnessing multicloud to modernize storage infrastructure
- » Understanding why multicloud storage just makes more sense

Chapter 1

Moving to Multicloud: Simply a Smarter Strategy

We know that enterprises are embracing the public cloud in record numbers. According to the Rightscale 2018 *State of the Cloud Report*, 92 percent of enterprises have at least some applications in the public cloud.

And we know that interest in the multicloud model is also soaring. According to the same report, 81 percent of enterprises now have a multicloud strategy, and on average organizations are now leveraging five clouds.

In this chapter, I give you a closer look at what's driving multicloud adoption from several perspectives: the role of multicloud in digital transformation, the urgency of modernizing the storage infrastructure with multicloud, and the capabilities that make multicloud a smarter option for enterprise application storage.

Identifying Multicloud's Role in Digital Transformation

Been to any industry conferences lately? Chances are good that there were multiple sessions and keynotes about the urgent need to modernize your data center infrastructure. What's behind that? Two words: digital transformation.

Business leaders are eager to take advantage of the opportunities of the digital age. The Internet of Things, cognitive computing, artificial intelligence, augmented reality, machine learning, consumerization of business applications, and automation of business processes all represent multi-billion-dollar market opportunities. That's why 92 percent of enterprises are engaged in some form of digital transformation today, according to a Capgemini report.

All of these digital opportunities depend on high-performance, scalable, reliable IT infrastructure. But all too often, infrastructure is an obstacle to digital transformation rather than a catalyst. You see the evidence in reports of stalled or delayed initiatives. According to a recent TechValidate survey, 64 percent of business leaders say their transformation initiative is behind schedule, and a MuleSoft survey found that only 18 percent of executives are confident they'll succeed in meeting digital transformation goals in the next 12 months.

That puts enormous pressure on IT managers to modernize their infrastructure so that they can run a wide range of workloads that power digital initiatives.

And that's where the cloud comes in. The cloud offers a unique opportunity to modernize infrastructure while also cutting costs. More specifically, the private cloud can provide operational efficiencies that translate to reduced operational expenditures (OpEx), while the public cloud can help cut capital expenditures (CapEx) by reducing the amount of infrastructure you need to purchase. And the move to a multicloud model can increase the simplicity, agility, and cost-effectiveness of infrastructure provisioning — and thereby serve as an accelerant to the digital transformation agenda.

So, why isn't cloud adoption translating to greater speed and success in executing digital transformation? Because many

companies, in their haste to exploit the advantages of the multi-cloud model, overlook some of the nuances that directly impact cloud costs, capabilities, and implementation timetables. And storage is a prime example.

SORTING OUT THE CLOUD LINGO: WHAT'S WHAT?

As the cloud model has caught on, we've all witnessed a big bang-style explosion in the buzzwords related to cloud computing. Unfortunately, there are significant differences in how these terms are used. It seems that every stakeholder or vendor attempts to spin the definition in a way that showcases their products or expertise, and the terms become meaningless, like "open computing" back in the 1990s. In particular, many people have started using the terms *multicloud* and *hybrid cloud* interchangeably, and that's going to lead to trouble. So, let's look to the U.S. National Institute of Standards and Technology (NIST) for some objective definitions of what's what.

- **Private cloud** is cloud infrastructure operated solely for a single organization, whether managed internally or by a third party, and hosted either internally or externally.
- **Public cloud** renders cloud services over a network that is open for public use.
- **Hybrid cloud** is a composition of two or more clouds (private, community, or public) that remain distinct entities but are bound together, offering the benefits of multiple deployment models.
- **Infrastructure as a Service (IaaS)** offers computing infrastructure as a service to subscribers. IaaS clouds offer additional resources, such as raw block storage as a service, to subscribers.

Multicloud is the use of multiple cloud computing services in a single heterogeneous architecture to reduce reliance on single vendors, increase flexibility through choice, mitigate against disasters, and so on. Multicloud differs from hybrid cloud in that it refers to multiple cloud services rather than multiple deployment modes (public, private, and legacy). Multicloud uses multiple cloud providers (Amazon Web Services, Azure, internal IT, and so on) for multiple workloads.

Yes, you can cut the CapEx of storage systems by moving to multicloud. But how will you minimize the cost of transferring data back and forth among clouds? How will you use the cloud to protect your data while keeping up with end-user performance requirements? How will you maximize the cost benefits of multicloud for backup or long-term data retention for regulatory or government compliance requirements? Until you have smart answers to those and other questions about data storage, your digital transformation journey may be in jeopardy of hitting some unexpected roadblocks.

Modernizing Your Infrastructure with a Multicloud Strategy

Storage infrastructure is particularly critical to digital transformation initiatives because it houses application data — which is the lifeblood of the enterprise. Data is the raw material for understanding customers, identifying new market opportunities, and creating the innovative new software, products, and services that deliver competitive advantages.

Storage infrastructure is also vital because the flow of data between the infrastructure and the application directly impacts the performance experienced by end users — and today's end users have zero patience for sluggish applications, services, and business processes.

So, if you're looking to modernize your storage infrastructure, why not look to the cloud (or, more correctly, the clouds)? You could use private cloud storage to minimize OpEx, leverage public cloud storage to minimize CapEx, and use multicloud services to optimize performance. Right?

But there's a problem: When it comes to storage, not all clouds are ready for all applications.

Public clouds typically offer object storage because it is massively scalable, which is great for content and web apps that store documents, videos, music, or social media content. But when it comes to mission-critical business applications, such as customer relationship management (CRM) and enterprise resource planning (ERP), object storage is not up to the task.



Object storage manages data as objects, unlike *file storage* (which manages data in a hierarchical file structure) and *block storage* (which manages data as blocks in sectors and tracks). The public cloud provides low-cost, high-capacity object storage, which makes it appealing for long-term storage of backup and archival data, as well as easy access to that data. A recent ESG analyst survey reports that when it comes to leveraging the cloud in specific ways to support cost-effective IT modernization, backup and disaster recovery are two of the most popular uses for cloud storage today. This is in large part due to the operational agility, efficiency, and appealing cost structure that a secure offsite data repository provides. To learn more about object storage and block storage, see Chapter 4.

Business applications require the enterprise-grade features, flexibility, and performance provided by block storage. Block storage is usable by almost any application, file, database, or file system, and it delivers the low latency needed for business applications.

So, companies that are looking to the cloud for storage need to look for block storage alternatives for their enterprise applications — but they also need the option of mixing public and private clouds. They need to be able to move data at will between private clouds and public clouds to optimize costs, performance, reliability, security, and so on — which can now be delivered in multicloud storage.

Moving to Multicloud Storage: It Just Makes Sense

With all the myriad cloud types and cloud offerings out there, why go with a multicloud strategy for application data storage? Because it combines the best attributes of multiple cloud options and takes away many of the trade-offs to which IT and application administrators have become accustomed.

You can move enterprise application data to public clouds (such as Amazon Web Services [AWS], Microsoft Azure, and Google Cloud), transfer it between them, or migrate it back to your data center, your on-premises private cloud, or a private cloud from a trusted third-party service provider — with no vendor lock-in.

When you decide to move data, there are no data migration or costly egress charges — you just flip the switch from the

multicloud portal, and the connection is switched to the new cloud provider instantly, without moving a single byte of data! Plus, if you decide to move data off the public cloud and back to your own data centers, you can do so easily without any egress charges. And you can manage your storage volumes through a simple web portal just as you do with AWS or Azure, but with data durability that's orders of magnitude higher.

As enterprises look to maximize IT efficiency through “hybrid IT” — leveraging both the public cloud and on-premises infrastructure — modern multicloud storage offerings can deliver the following advantages:

- » **Enterprise-grade features:** Multicloud storage offerings are built with an understanding of the requirements of enterprise applications and can deliver enterprise-grade reliability, security, and data protection in a service that is simple and easy to use. Compared to traditional storage, they can provide far better data durability, greater copy data management capabilities, stronger security mechanisms, and more.
- » **Easy mobility:** Data mobility is normally a major challenge in the cloud. You already have your own data centers and infrastructure. But how do you incorporate the cloud and move data back and forth without diminishing the value of those investments? Multicloud storage can give you a fast, simple on-ramp to the cloud and all its associated benefits, while allowing your data to move quickly and seamlessly among different private and public clouds — without physically migrating the data or incurring data egress charges.
- » **Global visibility:** Through a single, simple-to-use portal, you can track your current usage and estimate future costs easily. Eventually, you'll be able to apply advanced analytics to track and monitor all your multicloud services so you'll have global visibility into your entire infrastructure, whether it sits in the cloud or your own data centers.

With today's native cloud storage, the only viable option is to install additional third-party monitoring tools and incur licensing costs. Most of these tools are designed for the cloud with little to no visibility into your own data centers. With global visibility, everything from predicting and recommending hardware upgrades in your data center to recommending and optimizing data placement based on end-user locations and usage patterns becomes possible.

- » **Painless backup and data protection:** Backing up and securing data is simplified because you can harness disk-to-disk-to-cloud (D2D2C) capabilities, you can leverage more flexible and less expensive options for long-term retention (LTR), and you can take instant snapshots in seconds not hours, allowing you to protect your data quickly and more often. With snapshots, you don't need to make full copies of your data and pay for storing them. Instead, you only pay for incrementally changed data — which is typically just a small percentage of your full data set.
- » **Disaster recovery (DR):** The ability to store data remotely on cloud-based storage for disaster recovery — for example, as an alternative for traditional tape (especially where investments in tape are fully depreciated and due for refresh) — is a top driver for cloud object storage adoption. By using the cloud for DR, you can avoid the cost and complexity of building and maintaining an off-site infrastructure, freeing up resources to work on other projects, becoming better protected if something happens to your facility, and meeting your business and regulatory requirements for off-site backup retention and disaster recovery.
- » **LTR of backups for compliance:** LTR is critical for compliance in many industries, including healthcare, pharmaceutical, energy, government, finance, and aerospace. Multicloud storage overcomes the challenges of using tape storage for LTR because data retrieval no longer needs to be planned/scheduled meticulously ahead of time, there is transparent on-demand access, operating costs are far lower, and reliability is higher because failure rates due to human error are all but eliminated.
- » **Broad range of use cases:** Multicloud storage is extremely versatile, allowing you to on-ramp data to the cloud, off-ramp it back to your on-premises data center, deploy infrastructure for software development and testing, implement backup and disaster recovery via the cloud, move an application using Docker containers, lift and shift with less pain and porting, monitor and predict usage costs, and much more. I get into specific use cases in Chapter 5.



TIP

According to Forrester Research, “Cloud maturity is not a one-lane road; it’s a multilane highway. Cloud services have matured to the point that they can replace, augment, and host an increasingly wide range of enterprise workloads.”

- » Addressing cloud concerns with multicloud capabilities
- » Considering often overlooked considerations

Chapter 2

Overcoming Barriers to Multicloud Adoption

There's no question that the multicloud model holds considerable promise for achieving high-priority business goals such as improving the performance and agility of enterprise applications, cutting total infrastructure costs, and hitting the accelerator on digital transformation initiatives.

The new question is, can the multicloud model actually help debunk any remaining concerns and misconceptions about cloud computing? When you examine these apprehensions head on, you find that multicloud can be instrumental in overcoming objections and broadening your view of the value of the cloud. Forrester Research reports that enterprise cloud computing adoption accelerated in 2016 and predicts that it will do so again in upcoming years.

In this chapter, you explore the cloud concerns that multicloud can address and the frequently overlooked considerations in implementing a cloud strategy in terms of organizational, educational, and financial impacts.

Using Multicloud to Debunk Common Cloud Concerns

The mere thought of moving enterprise applications to the cloud makes many IT managers and administrators break into a cold sweat. But fear and anxiety aren't good reasons to avoid action. There was a time when people also believed it wasn't safe to ride in an elevator without a human operator. Some cloud concerns are justified, but many are simply myths and misconceptions. Let's take a closer look:

» **Clouds may increase the risk of data loss.** This misconception stems from the fact that cloud-based block storage services historically have had significant differences in their data durability (the likelihood of data loss) when data is not backed up. In some cases, the annual failure rate can be as high as 1 in 500, which is clearly unacceptable for enterprise applications.

The assumption is that multicloud storage services would be even less likely to provide reliable, consistently low data loss rates for enterprise applications. However, the reality is that multicloud services can bring enterprise-grade reliability and data protection to the public cloud and deliver measured, proven data durability that is millions of times higher than cloud-native block storage.

» **Vendors won't offer enterprise-grade support.** Enterprise-grade support requires extensive expertise built on many years of real-world experience. But traditional cloud support models have largely been based on a "do-it-yourself" (DIY) approach via online support forums and knowledge base articles.

Today, multicloud service providers are prioritizing support as an integral part of their service. Many offer enterprise-grade capabilities such as support delivered by experienced technical support engineers; deep visibility into overall health, automated actionable reporting, and proactive troubleshooting; 24/7 service options for mission-critical deployments; and more.

» **Data gravity increases in the cloud.** As workloads scale, data is typically the hardest component of a workload to move. The metric that describes this inertia is *data gravity*. It's easy to presume that data gravity would increase with multicloud models because you're talking about moving data not only to the cloud but also between clouds, as well as between clouds and on-premises data centers.

However, the reality is that multicloud services can actually "grease the skids" in data mobility as customers move workloads between their data centers and different cloud providers. For example, multicloud storage can remove data gravity by acting as a single repository for all your cloud providers while providing easy mobility — making the transitions not only faster, but easier as well — because there's no need to migrate when moving between clouds.

» **Cloud services create vendor lock-in.** You don't have to be a cynic to suspect that cloud service providers might try to include "fishhooks" in the contract such as data egress charges that effectively lock you in to their service by making it so difficult and expensive to move your data back from their cloud to your data center that you really have no choice but to stick with their offering.

However, you'll discover that some multicloud service providers recognize that it's in their interest to increase your agility rather than lock you in. Multicloud service providers have an economic incentive to keep you nimble. For cloud service providers, the future favors the flexible.

» **Cloud storage will never be secure enough.** It's true that when you push enterprise applications to the public cloud, you may be pushing sensitive data outside of your direct control — and that could open you up to increased security risks. But let's not neglect the other side of the coin — because the multicloud model can provide new ways to solve security and data protection problems.

For example, multicloud services can test security mechanisms across huge populations of users and possible attacks. This makes it possible for users to benefit from the experience and remedies of millions, as opposed to the efforts of a single organization trying to address all possible security threats.

In addition, multicloud services can allow you to take advantage of predictive analytics — which can make it possible for cloud security systems to identify critical security information, evaluate the significance of that information, analyze user behavior, and pinpoint risky activities to help keep your enterprise applications secure.

Multicloud also can incorporate data encryption technology, such as 256-bit Advanced Encryption Standard (AES) encryption, adding a layer of protection to privacy and confidentiality for sensitive data. Moreover, multicloud makes it possible to utilize and extend the existing firewall and network architecture to cloud-based services, providing additional security with relatively little cost and effort.

» **Backup and recovery options will be limited.** Many companies are especially concerned about the prospect of offsite cloud backup and recovery — both from a data protection perspective and from a cost perspective. Data must not be compromised, and ongoing costs can add up quickly in the public cloud when you're charged for every gigabyte sent, stored, and retrieved. In addition, companies want to move data *seamlessly* to the cloud with no separate cloud gateway or virtual appliance required, ideally using their existing backup and enterprise apps. They don't want to worry about additional physical footprint or management overhead.

Today's multicloud storage offerings are capable of addressing all of these concerns. In fact, multicloud storage can actually increase your backup and recovery options. For example, multicloud makes it possible to adhere to the "3-2-1" best practice for hybrid IT data protection (three copies of data, two copies on two different types of media, one copy offsite) allowing you to seamlessly and securely move backup data to public, private, and hybrid clouds or on-premises object storage — which in turn enables long-term retention, archive, and reliable disaster recovery that is simple and efficient. Cloud-based data protection is also cost-efficient because it can provide automated, native tiering of backup data and intelligent deduplication, reducing storage and access costs by up to a factor of 20. I talk more about cloud data protection in Chapter 4.

In addition, some multicloud services come with a rich set of data management capabilities, such as instant backups, that don't impact production windows and performance, and instant thin clones so you can quickly create zero-copy clones for test/dev, analytics, and bursting. And by providing fast, easy restoration, they can help you meet accelerated recovery point objective and recovery time objective requirements without impacting production workloads.

- » **Surprise fees will blow up the business case.** One of the most common complaints you hear from users about the public cloud is that sooner or later they receive monstrous "surprise bills" that are several times larger than they're expecting. Stories about monthly bills spiking as much as five-fold are all too common. That kind of volatility can quickly kill the business case you carefully crafted for adopting a cloud strategy.

The root cause of these awful surprises is poor monitoring and tracking tools and lack of best practices on cloud use. Multicloud services can track your current usage, estimate future usage, charge only for the resources you use, and provide features that help minimize the bill (such as charging only for newly changed data rather than full copies) compared to cloud-native storage.

- » **The "black box penalty" will increase costs.** Simply moving applications to the cloud doesn't make all your problems go away. Cloud services are often like a black box — you can't see inside. The result is a black box penalty of spiraling costs where troubleshooting of issues is nearly impossible, leaving you no choice but to purchase and install additional third-party monitoring tools and licenses. Most of these tools are designed for the cloud and provide little or no visibility into your own data centers.

This concern naturally applies to multicloud storage offerings as well. However, multicloud services are now available that provide visibility whether your data is in the cloud or on-premises. It also allows you to see up the stack

into the virtualization layer, and this level of visibility is also being extended to the entire stack, including the network, servers, storage, and even the application itself.

» **You'll need separate tools and processes to take advantage of analytics and automation.**

Two of the primary drivers for moving to the cloud for many organizations are lowering costs and increasing agility. Analytics and automation can play a huge role in accomplishing those goals — if they're easily accessible. Some cloud offerings require separate tools or new ways of working, but others are specifically designed to take advantage of analytics and intelligent automation using your existing tools and processes. Current offerings can also take analytics to new levels — for example, by using predictive analytics to anticipate and prevent issues across the stack, so you can optimize data placement and resource usage, and uncover opportunities for savings.

» **“Cloud-first” and “all-in-cloud” strategies are best.** Lots of companies are looking to move every application and data set into the cloud as quickly as they can. There are some problems with this approach:

- *There is the very real possibility that you'll get it wrong — placing all your eggs in one basket and realizing too late that there was a better option.* You may inadvertently sacrifice the ability to mix on-premises and cloud-based resources for better economics and business agility.
- *A cloud-first or all-in strategy can create wholesale changes to personnel or skill set requirements.* And meeting the new requirements of the cloud doesn't happen overnight.

» The fact is, there's still plenty of value in keeping some applications and data on your internal, on-premises infrastructure — and buying on-premises infrastructure to satisfy immediate needs — even if you have a cloud-first strategy.

Table 2-1 compares some of the features and capabilities of multicloud and native cloud storage.

TABLE 2-1 Multicloud Storage versus Native Cloud Storage

	Multicloud Storage	Native Cloud Block Storage
Enterprise-grade		
Reliability	Assured enterprise storage reliability	99.95%
Data durability	Millions of times more durable	0.1% to 0.2% annual failure rate
Snapshots and clones	Instant and efficient	Full and slow
Data mobility		
Cloud on-ramp	Yes	Manual migration appliance
Multicloud support	Yes	No, costly egress charges
Hybrid/private cloud portable	Yes	No, manual migration only
Global visibility via predictive analytics		
See and manage public cloud and on-premises	Yes	No, requires third-party tools
Predictive analytics	Yes	No, requires third-party tools
Cross-stack problem isolation	Yes	No, requires third-party tools

Addressing Frequently Overlooked (But Important) Considerations

IT leaders look at the potential of multicloud services and see a host of new possibilities for transformation and increasing agility. But all too often, critical aspects of harnessing the cloud are overlooked, including organizational readiness, educational requirements, and financial implications. I take a closer look at a few of these in this section.

Organizational and cultural challenges: Is IT ready for multicloud?

Adopting multicloud is a journey, not an action. It's going to require skills your IT staff may not currently possess; it's going to change your processes and procedures; it may even require a new organizational structure. You've got to ask, "How ready is my organization?" You'll need to carefully assess both your organizational maturity for cloud adoption and the attitudes of current IT staff about moving to multicloud — because their viewpoints and mindsets can ultimately determine the success or failure of any cloud adoption initiatives.

Organizational change also creates cultural considerations. Change can be disruptive, and it's important to understand how IT staff will react to that disruption. Do you have a culture that thrives on exploring new concepts and new ways of working? If not, have you considered incentives that could improve acceptance of the change or facilitate better collaboration among team members?

Educational challenges: How much will people have to learn?

There are a couple of key questions to ask in assessing educational issues. How much training or retraining will be required to obtain the needed skillsets or best-practice knowledge for the transition to multicloud, and how will IT staff obtain those skills?

Equally important, it pays to give some thought to who among the existing staff is best qualified to take on a leadership role in implementing a multicloud model. Who has the right combination of skills, positive attitude, and ability to communicate with and educate others, including both technical staff and business leaders?

Financial considerations: How solid is the business case?

It's one thing to understand how multicloud will save money. It's quite another to do the math and determine precisely how much it will save in CapEx and OpEx, and over what period of time — and to incorporate your economic analysis into a comprehensive business case. Many individual financial factors come into play — from the cost of making your data center infrastructure cloud-ready to the cost of training and the incremental costs of supporting new services — and all of them need to be taken into account.

IN THIS CHAPTER

- » Addressing data center challenges in a multicloud environment
- » Leveraging predictive analytics and flash storage
- » Focusing on simplicity, reliability, performance, and mobility

Chapter 3

Architecting the Multicloud Data Center

With cloud infrastructure, compute resources are ephemeral. They come and go, turn on and off, burst up and shut down. However, storage needs to be much more persistent because people and organizations tend to want to keep their data for ever longer periods. As a result, the sheer volume of data and the storage required to keep it continues to grow and accumulates gravity.

In a multicloud environment, storage infrastructure becomes the lifeblood of your critical data. As such it must become more fluid, allowing data to move more freely among storage resources, whether they reside in your data center or in a public or private cloud. That means that when you're architecting a multicloud data center, you need to keep storage considerations front and center — and deploy storage infrastructure that is truly cloud-ready. Because the future belongs to the fast and flexible.

This chapter describes the key challenges and opportunities for architecting a data center that is both storage-centric and multicloud-ready.

Your Data Center Lives On: Addressing New Challenges in the Multicloud Era

No matter how excited you are about the possibilities of the multicloud model, and regardless of how far along you are in adopting a multicloud strategy, your data center isn't going anywhere any time soon. You're going to need to maintain and manage your on-premises data center infrastructure for many years to come.

Many workloads are not yet cloud-ready, many IT departments are not yet multicloud-skilled, and many business leaders are not yet ready to embrace multicloud as a safe, cost-effective alternative — particularly for extremely sensitive corporate or customer data. Let's take a closer look at the challenges that data center managers face in a multicloud environment.

Keeping the lights on and meeting service-level agreements

While the opportunities of the digital age beckon, the realities of meeting current obligations remain unchanged. It seems ironic: Virtually every enterprise today is embarking on a digital transformation initiative — in fact, a recent Progress global survey found that 96 percent of organizations see digital transformation as critical or important. Yet 80 percent to 99 percent of current IT budgets are still allocated to traditional IT activities, also known as “keeping the lights on” (or, simply, KTLO).

IT departments need to continue delivering on current service-level agreements that specify metrics for such things as performance, availability levels, service request fulfillment time frames, and so on. In particular, they need to continue meeting the requirements for applications, such as fast load times, data accessibility, problem resolution, and a whole host of user experience metrics.

Integrating or future-proofing for cloud

Many organizations continue to struggle with data center modernization or cloud-readiness initiatives. For example, integrating pools of existing resources — such as compute instances, storage volumes, and networking infrastructure — and making them more future-proof for the cloud era can be daunting. What's

needed is a more structured approach to understanding and dealing with data center challenges in the multicloud era, and broader adoption of new innovations that can expedite cloud readiness.

Finding and fixing performance and availability issues

Application performance has become increasingly critical to both end users and data center managers. Slow applications mean lower productivity, and in a time when acceleration of business processes is the key to competitiveness, no business can afford slow applications.

In a recent survey by Oxford Economics and Nimble Storage, almost half of employees said they lose more than 10 percent of their workdays — about 48 minutes — just waiting for software to load. IT decision makers feel that pain too: 43 percent said they lose 11 to 30 minutes each workday because of delays they encounter while trying to use applications (see Figure 3-1).

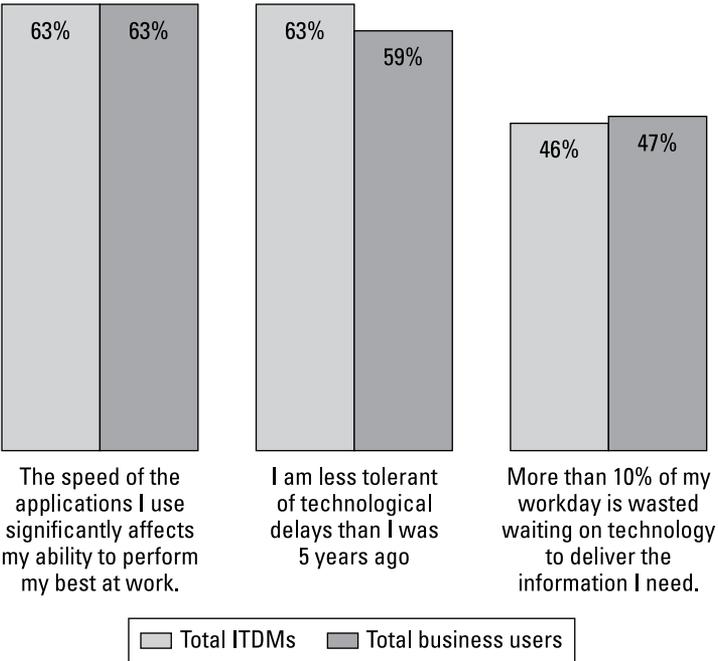


FIGURE 3-1: IT decision makers and business users are tired of sluggish applications.

Businesses are also under increasing pressure to meet the instant gratification requirements of today's digital natives and “Millennialize” their applications. In the same Oxford/Nimble survey, more than three-quarters (77 percent) of Millennials say that suboptimal application performance affects their ability to achieve their personal best, compared with just half of Baby Boomers and 72 percent of Generation Xers. In fact, half of Millennials say they've stopped using a cloud-based application because it runs too slowly — significantly more than other groups.



REMEMBER

Of course, if applications become unavailable because of infrastructure downtime, performance is a secondary concern. Various recent research shows that the average cost of an hour of downtime can be as much as half-a-million dollars, and this will only increase with the continued digitization of industries.

Managing costs

The delicate balancing act between CapEx and OpEx gets even more complicated for data center managers in a cloud environment. The challenge is to minimize total infrastructure costs, even as total infrastructure capacity scales to unprecedented levels.

Cloud-based object storage is an increasingly attractive option for reducing costs because it provides OpEx efficiency. You don't need to invest in hardware to have a scalable and secure offsite data repository for backup, disaster recovery, and other vital use cases.

In short, the physical systems you purchase must be both cloud-enabled and future-ready, and the infrastructure you subscribe to as a service must have predictable costs. That means there must be

- »» No “shock fees”
- »» No black box penalties
- »» No outrageous data egress charges
- »» No unexpected downtime (with its associated hard dollar and intangible costs)
- »» No vendor lock-in

Overcoming Challenges with Predictive Analytics and Flash Storage Innovations

New innovations in predictive analytics and flash storage infrastructure have arrived to address data center challenges — old and new — and the multicloud model is putting these innovations to work right now.

Eliminating the app-data gap

When users experience application performance issues, they have hit what's known as the *app-data gap*, any slowdown or disruption of the delivery of data to applications. We've all experienced it — that frustrating delay that forces you to wait . . . and wait . . . for . . . your . . . application . . . to . . . do . . . something.

But closing the app-data gap is not as simple as adding more high-performance storage systems such as flash arrays. In fact, 54 percent of application performance issues are not related to storage, according to a recent study by Nimble Labs Research. The app-data gap can be caused by issues anywhere across the entire stack (storage, networks, servers, and software). And that means not only are blazing-fast systems required, but businesses also need predictive insights into their entire infrastructure stack if they want to get ahead of problems.

Predictive analytics anticipate and prevent virtually any barrier that slows down data velocity or leads to costly downtime before the issue occurs. By combining these capabilities with blazing-fast all-flash or hybrid-flash arrays and converged, integrated infrastructure systems, companies can not only close but obliterate the app-data gap.

Pinpointing problems anywhere in the stack

Today's predictive analytics solutions collect more sensor data points than there are stars in our galaxy. They use data science and machine learning to analyze and correlate trillions of sensor data points to find problems and resolve complex infrastructure issues. Analytics solutions can diagnose and prevent problems even outside of storage. It's like having a team of data scientists

watching over your infrastructure so that it runs perfectly without your having to lift a finger.

With predictive analytics, nine out of ten problems can be automatically detected by collecting and analyzing billions of sensor data points from each storage array. Non-storage problems, misconfigurations, and other user errors can be quickly diagnosed and resolved, resulting in higher availability levels.

Predictive analytics can also accelerate root-cause analysis and cut hours of tedious manual troubleshooting. Administrators can see across storage, networks, servers, and virtual machines; view correlated analysis on-demand to quickly resolve issues (even when they're not related to storage); and avoid all the vendor finger-pointing that slows down problem resolution. With predictive analytics, there shouldn't be a need to talk to a support engineer asking for your name, support contract number, and whether you turned on your storage — all the information is already there. For the relatively few problems that do require manual assistance, you can go straight to a Level 3 expert who uses “pre-collected” data to rapidly resolve the issue.

Anticipating resource needs

You can also take advantage of analytics to predict future infrastructure needs. Analytics solutions can accurately forecast capacity, performance, and bandwidth needs based on historical data, and correlate and match similar consumption patterns across the entire install base. They can even identify potential resource ceilings in your environment as your usage increases — and tell you how to avoid them.

Additionally, analytics solutions can provide prescriptive guidance to ensure optimal long-term performance of your entire infrastructure stack. Planning guesswork is eliminated by leveraging installed-base learning and statistical modeling to precisely predict future requirements.

Increasing uptime via installed-base learning

Using predictive analytics has resulted in measured availability levels over 99.9999 percent across a storage vendor's entire

installed base of customers. How is this extremely high system reliability achieved? It starts with features that are built into the storage platform: no single point of failure, dual controllers that allow for non-disruptive upgrades, a fault-tolerant software architecture, and extremely robust data integrity including triple+ parity redundant array of independent disks (RAID) and end-to-end integrity validation. But the breakthrough innovation is the addition of predictive analytics.

For any new problem experienced across the company's installed base, the analytics solution uses pattern matching algorithms and continuously searches health signatures across all systems. If a signature is detected, the analytics solution will either prevent the problem from occurring or proactively resolve it, even if the problem is outside of storage. There are no false alerts as machine learning normalizes performance behavior across the installed base. Each system continually gets smarter, learning from the installed base, and downtime events are increasingly prevented.

Visualizing your environment

With a combination of predictive analytics and intuitive dashboards, you can have complete visibility through the cloud to all information you need to maintain a resilient environment and ensure smooth operations. Executive dashboards give you peace of mind that everything is running perfectly and alert you to things you need to know, such as performance, capacity, and efficiency metrics. Correlated visualization can give you a view of what's happening across the stack, from applications to storage, so you can quickly see and resolve issues before they impact end users.

Reining in costs

Previous economic models required you to pay for resources you weren't actually using. The multicloud storage model makes it possible to tie costs directly to actual resources used — so there's no waste. You also get visibility into usage before, during, and after. And with the use of predictive analytics, you can accurately predict your total costs even before you deploy your storage volumes. You can monitor usage on a midmonth cycle and reconcile against your end-of-the-month bill.

Choosing Cloud-Ready Infrastructure

Because your data center isn't going anywhere, it makes sense to optimize your on-premises storage infrastructure for the cloud model. In other words, it's a good idea to select and deploy "cloud-ready" infrastructure — storage systems that work seamlessly with a wide range of clouds, that allow you to easily move data between your data centers and the cloud providers of your choice, and that evolve with your cloud strategy.

So, what are the key attributes to consider in selecting cloud-ready infrastructure? Focus on flexibility, simplicity, reliability, performance, and mobility — and make sure these operate the same way in and out of the cloud. Here's how that translates into specific capabilities.

Native support for cloud

Cloud-ready storage infrastructure should simplify working with a wide range of cloud services. That means it should have native support for the cloud model, with storage APIs that integrate with cloud-native APIs from AWS, Azure, and others. It shouldn't require additional hardware and software in the cloud, or any additional on-premises equipment to act as a bridge or gateway into the cloud. It should move data securely and efficiently with encryption, change block tracking, deduplication, and compression. Moreover, it should simplify the scale-up process both in the cloud and on-premises so you can balance resources and capacity in the places they're most needed.

Support for any workload, anywhere

You should be able to run the same production workloads on your storage arrays or in the cloud without substantial cost or performance penalties. You should be able to spin up and speed up secondary uses such as test/dev, analytics, and cloud bursting with efficient data transfers and cloning. And you should have the option to back up your data to cloud and restore it back on-premises or in the cloud on demand to meet your disaster recovery objectives.

Global visibility, management, and automation

Multicloud storage should become the new normal by allowing you to do more and manage less. For example, you should be able to see and manage your data across clouds and your data centers to predict and prevent problems on your infrastructure. You should be able to build, deploy, and manage your apps on any cloud or your infrastructure. And your cloud-ready infrastructure should integrate with automation tools such as containers and configuration managers used by DevOps and CloudOps people.

Pay-as-you-go

You should have access to flexible, consumption-based pricing models that allow you to pay for on-premises storage the same way you would in the cloud. That means you should pay for storage on a per-gigabyte, per-month basis, with the ability to grow or shrink your footprint as needed, on demand, without financial penalties. Equally important, you should be able to leverage the OpEx advantages of the cloud model. Why incur enormous CapEx for a secondary disaster recovery site when you could purchase the capability as a service and free up capital?

Flash for performance

Speed is of the essence for workloads in the cloud, so storage infrastructure should enable a level of data velocity that eradicates the app-data gap. In most cases, this means it should be based on flash technology.

For primary workloads, such as enterprise applications, an all-flash approach may make the most sense. For other primary or even secondary workloads, hybrid flash arrays that allow you to change the service level of any volume at any time may be a better choice. Whether you choose all-flash, hybrid flash, or secondary flash, flash technology gives you the combination of speed and scalability you need to build a multicloud environment.

Predictive analytics for simplicity and reliability

In the era of big data, there's no reason to accept infrastructure issues as a given. Cloud-based predictive analytics solutions are available that can anticipate and resolve problems before they impact workloads, application users, or your business. These solutions have proven to automatically predict and resolve up to 86 percent of all issues.

Easy migration of data between clouds

Your data center infrastructure should make it simple to move data from anywhere to anywhere. From your data center to a public cloud. From a public cloud to a private cloud. Between clouds. Between multiple clouds and your data center. And this level of data mobility should not be expensive. It's your data. You should be able to put it wherever you want — painlessly.

IN THIS CHAPTER

- » Catching the cloud adoption wave
- » Tracing the evolution of cloud storage services
- » Recognizing the market impact of multicloud storage
- » Identifying key multicloud storage requirements

Chapter 4

Evaluating Multicloud Storage Services

What's the big picture of cloud adoption today, and where do multicloud storage services fit in that picture? What types of multicloud storage services are available today? What's coming next? How will new multicloud storage offerings impact the Infrastructure as a Service (IaaS) market? And, most important, what types of features and capabilities should you be looking for as you evaluate the various multicloud storage services available today? Read on. I answer all those questions in this chapter.

The Four Waves of Cloud Adoption: Where Are You?

As you contemplate your move to multicloud, it pays to start by understanding the 30,000-foot view of cloud adoption. This will give you insights into your company's *cloud maturity*, or readiness to take full advantage of the benefits of multicloud services. Here are the four primary waves of cloud adoption, according to

an October 2016 Forrester Research paper titled *Take the Wheel: Build Your Cloud Computing Strategic Plan Now*:

- » **Wave 1: Systems of Engagement:** The first phase is moving apps and systems used directly by employees — such as email, web, and mobile apps — to the cloud. Most companies now design these content-driven applications to be cloud-native, and they've been running on private or public clouds for years.
- » **Wave 2: Systems of Insight:** The next phase of cloud adoption involves moving analytics and business intelligence capabilities to the cloud. Most companies have at least dipped a toe in the water and explored the possibilities, but widespread adoption is still in its early stages.
- » **Wave 3: Systems of Record:** The third phase, which is just starting to emerge today, is moving critical business resources such as transactional databases to the cloud. This transition adds significant complexity to cloud considerations and requires enterprise-grade capabilities that traditionally haven't been available in public clouds. See Chapter 2 to learn more.
- » **Wave 4: Systems of Innovation:** A fourth phase, still hovering elusively just beyond the reach of today's technology, is the use of clouds as the raw material for reinventing business processes. This phase will include the migration of new enterprise workloads to the cloud.

The Evolution of Cloud Storage Services

Cloud storage was already a thing even before the term *cloud computing* was coined. It started back in 1983 when CompuServe offered its consumer users a small amount of disk space that could be used to store any files they chose to upload. Since then, several new iterations of cloud storage have emerged, leading up to the multicloud storage services we're beginning to see today, including those covered in the following sections.

Object storage

These services are based on a proven, durable architecture that manages data as objects, as opposed to file hierarchies. Object

storage is commonly used in the first wave of cloud adoption (systems of engagement): social, mobile, and other cloud-native applications. This provides a layer of abstraction that makes management easier for administrators, but requires additional work programming to specific APIs.

Today, the majority of cloud storage services leverage an object storage architecture, including Amazon S3, Microsoft Azure Blobs (object storage), and many others. And their use of object storage is expanding quickly to new use cases. For example, it's now both possible and practical to use object storage for offsite backup, long-term retention, disaster recovery, and more.

Cloud block storage

This more recent innovation in cloud presents cloud servers with logical storage that can access data in blocks as if it were attached to physical storage. The first major cloud block storage offering was Amazon's Elastic Block Store (EBS), which provides persistent block storage volumes for use with Amazon EC2 instances in the Amazon Web Services (AWS) cloud. In addition to Amazon EBS, virtual block storage offerings in the market today include Azure Disks from Microsoft Azure, Google Cloud Platform Persistent Disks, and the Cloud Block Storage from DigitalOcean.

However, several challenges continue to plague cloud block storage offerings, including durability issues, lack of enterprise-grade features and data services, single host connectivity, limited scalability, lack of mobility, and the notorious black box penalty.

Cloud-ready data protection

This novel concept provides multi-tiered, end-to-end protection as you evolve to hybrid IT, natively linking your primary, secondary, and private/public cloud storage. This lets you easily leverage the economics and agility of the cloud for long-term backup or disaster recovery (DR), while continuing to use your on-premises infrastructure for fast, reliable operational recovery.

Cloud-ready data protection solutions such as HPE Cloud Bank Storage enable your cloud backup and DR for your enterprise apps and send, store, and retrieve only unique data. This cuts costs while still providing scalability and reliability for cloud storage and DR. The cloud-ready data protection model supports a variety of object cloud storage options, including AWS S3 and Microsoft

Azure Blob storage (Hot and Cool Tiers), and can also provide private cloud options. Users can also move their data to a lower-cost tier for archiving (for example, from AWS S3 to AWS Glacier) and retain the data for an indefinite period of time. You can also recover the data from the cloud to supported systems either on-premises or in other clouds.

True multicloud storage volumes

In addition to the problems mentioned earlier, cloud block storage services that are currently on the market have only been available for use within the service provider's public cloud. Today's enterprises want more flexibility, less risk of vendor lock-in, and more enterprise-grade capabilities so that they can move enterprise applications to the public cloud without having to worry.

To that end, service offerings such as HPE Cloud Volumes (formerly Nimble Cloud Volumes) provide multicloud block storage, combining simplicity with enterprise-grade durability and feature sets. With the Cloud Volumes service, a customer can attach and use compute virtual machines in AWS or Microsoft Azure, the same way they would with the native storage volume offerings available within those public clouds. You just provision the storage volume from the HPE Cloud Volumes console, select your desired storage volume size, I/O operations per second (IOPS) performance, product tier, and the AWS or Azure instances you want to attach.

Market Impact of Multicloud Storage Services

A recent IDC Market Note offered an analysis of how multicloud storage services such as HPE Cloud Volumes would affect the overall IaaS public cloud market. To summarize the findings, IDC wrote that HPE Cloud Volumes “brings clear benefits to public cloud IaaS customers and existing cloud service providers . . . [and] also represents an interesting new dimension of growth for infrastructure vendors.”

More specifically, the IDC report highlighted three broad areas of impact:

- » **End customers will have a broader set of cloud storage volume options.** HPE Cloud Volumes introduces an alternative for virtual block storage for AWS and Microsoft Azure customers . . . and introduces competition in a previously uncontested feature space — virtual block storage. The fact that customers can now choose between the native option and a third-party alternative will increase the rate of improvements and enhancements in cloud storage volume products — driven by competition.
- » **Public cloud service providers get a boost to their ecosystems and see increased enterprise adoption:** HPE Cloud Volumes introduces new capabilities to the public cloud environments that it supports. The service-level agreements (SLAs), visibility, and 24x7 support, backed by a proven technology provider in the cloud storage industry, will reduce the barriers for migration of business-critical and performance-sensitive applications to these public cloud environments. These will add to the rate of public cloud adoption by enterprises.
- » **The broader infrastructure vendor community has a potential direction for growth in the cloud.** For the broader infrastructure community, HPE Cloud Volumes demonstrates a new direction of investments and growth in a cloud-centric world. If successfully executed, infrastructure vendors can leverage their specific areas of expertise to build “as a service” offerings that can be delivered in existing public cloud ecosystems. This may evolve into a scenario in which multiple infrastructure vendors build and deliver specific IaaS components to customers, within the major public cloud environments. The net impact of these will be increased choice for end customers, increase in direct vendor participation in the public cloud market, and increasing rate of innovation and progress in each specific public cloud IaaS component.

Checklist: Core Requirements for Multicloud Storage

With all the cloud storage options out there, and with the pace of innovation around cloud storage accelerating, what are the key features and capabilities you really should insist upon today? Read on.

Enterprise-grade features

When you're looking to move enterprise applications to the public cloud, take a hard look at the prospective service provider's enterprise-class capabilities:

- » **Data durability:** Ask what the annual failure rate is. If it's anywhere in the ballpark of 0.2 percent, as it is for most cloud storage services, *run!* Multicloud storage services are available with data durability rates millions of times better.
- » **High availability — guaranteed:** Again, insist on seeing quantified, proven metrics for availability — and don't settle for less than 99.95 percent on your SLA. It also pays to understand what types of events may be excluded from these calculations so there won't be any surprises.
- » **Enterprise-class performance:** Look for a service that is capable of delivering low latencies and tens of thousands of IOPS, even in the face of widely varying workloads — without causing a “noisy neighbor” problem. And make sure it can actually deliver the performance that it is configured for.
- » **Copy management (fast, easy clone creation):** Look for flash-optimized snapshot capabilities that can provide instant, high-performance copies or clones in large numbers. The ability to create both read-only and read/write clones, and retain them without performance impacts can help save space and money while speeding up many use cases, such as DevOps, analytics, and bursting.
- » **Strong data protection and security:** On the data protection side, look for offerings that harness object storage for backup. You should be able to back up to the cloud for cost-effective long-term retention and DR, and integrate with your primary array GUI for end-to-end protection from the data center to the cloud. And by deploying highly efficient de-duplicated data transfer, you can reduce time, costs, and

network bandwidth for cloud storage by more than 20 times compared to standard public cloud services.

In addition, thoroughly evaluate the offering's malware protection, advanced persistent threat (APT) defenses, policy and permission controls, remote wipe, and 256-bit Advanced Encryption Standard (AES) encryption capabilities. Also look at the extent to which the service leverages predictive analytics. Analytics can enable cloud security systems to evaluate the significance of critical security information, analyze user behavior and find suspicious activity, and pinpoint other risky behavior before a breach occurs.

Easy mobility: Move data where and when you need it

Assess exactly how many hoops you're going to have to jump through to move your data. For example, do you have to re-architect your applications before you can move them to the cloud? Is there an intuitive dashboard that doesn't require you to be a storage architect just to move data? Can you migrate data among the clouds or back to your data center at will? Can you containerize the application itself and move it, too? Do you need to purchase third-party devices and learn entirely new skillsets?

- » **Interoperability (native support for cloud, ability to move data between clouds):** If you want to move data back and forth across your data center, public clouds, and private clouds, you need infrastructure that has native support for the cloud application programming interfaces (APIs), as well as current data center workflows. You should also be able to manage your storage through a simple web portal — just as you do with AWS or Azure. Compression and change block tracking technologies make the on/off-ramp transfer very efficient. Your data is mobile between clouds without the need for data migration.
- » **Minimal data gravity (greater flexibility):** You'll need to determine whether the cloud storage service will increase or reduce data gravity. If the service provides mobile attachment/detachment of storage volumes from public cloud compute instances and lets you attach it to another cloud provider, chances are, it will diminish data gravity, which is what you want. Low data gravity makes moving applications much easier and less risky.

» **No vendor lock-in:** Data is the greatest source of vendor lock-in in the cloud. When data is stored with one public cloud provider, it's sometimes very difficult and expensive to move it anywhere else. Data becomes trapped because you can only use or access it within the confines of that cloud provider. Make sure there's no huge effort or expense required to move data between cloud providers or back in-house. Also, check the contract carefully for data egress charges and make sure you understand them clearly — they should be minimal.

Visibility: Insights from monitoring and predictive analytics

You need global visibility and insights across the entire infrastructure stack, no matter where your data lives. Make sure the prospective service provider can deliver advanced monitoring capabilities and harness predictive analytics to improve visibility:

» **Cloud and data center monitoring capabilities:** Unless you're pursuing an all-cloud, all-the-time strategy, you'll need the ability to monitor and track both your cloud and data center data storage and data movement activities.

» **Predictive analytics to predict, track, and optimize storage:** According to IDC, "cloud-based predictive analytics are not yet considered a requirement in enterprise storage, but there are undeniable benefits to its use: higher performance, better reliability and availability, more efficient management, a vastly improved support experience, and a much better understanding of how your storage is performing on a variety of different metrics that inform daily administration as well as future planning." So, ask prospective service providers whether — and to what extent — they leverage predictive analytics. Citing IDC one more time: "Within just a few short years, it will be a requirement, and customers can expect more mature and comprehensive implementations . . . to provide better overall value."



TIP

A multicloud storage service should be its own separate cloud, not running on top of another cloud service. This allows you to separate cloud storage from cloud compute capabilities, enables mobility, and helps you avoid vendor lock-in.

IN THIS CHAPTER

- » Creating cost-effective new cloud backup options
- » Lifting and shifting apps and data
- » Enabling DevOps and Docker containers
- » Bursting capacity to meet fluctuating demand levels
- » Switching between clouds and controlling costs
- » Leveraging enterprise-grade storage in the cloud

Chapter 5

Ten Multicloud Storage Use Cases

In this chapter, I explore some of the noteworthy ways in which companies are using multicloud storage today. It turns out that the use cases are broader and more creative than you may have guessed. Here's a partial list.

Expanding Cloud Backup Options to Cut Costs

A recent report from ESG says that 42 percent of enterprises use cloud for backup and archive. That number is poised to go much higher with the arrival of new multicloud offerings such as HPE Cloud Bank Storage. These services provide highly efficient data transfer and storage to create low-cost, flexible, safe alternatives for cloud backup. They cut the cost and increase the performance and scalability of object storage, making it possible to harness the cloud for long-term retention and archiving of critical data.

With HPE Cloud Bank Storage, you can move backup data to the cloud of your choice with no separate cloud gateway or virtual appliance required. You can choose the public and/or private cloud service providers that best fit your requirements. And you can empower database administrators (DBAs) and storage administrators to control end-to-end protection of their applications from the data center to the cloud via familiar native tools.

Equally important, you can orchestrate the optimal balance between performance and cost with automated, intelligent storage tiering — from on-premises arrays to the public cloud. You can mitigate the risk of downtime with cloud disaster recovery (DR), while avoiding the cost and complexity of building a secondary site. And you can ensure secure transmission and storage in the cloud with in-flight/at-rest encryption and data immutability. As data ages and access becomes infrequent, it can be moved to even more cost-efficient cloud storage for archiving or extended retention.

On-Ramping/Off-Ramping Data

The multicloud storage model makes it easy to *lift and shift*, or move data and applications from the data center to the cloud. Several innovations have had a profound impact on simplifying the process.

First, storage clouds that are built for enterprise applications and data have emerged. These storage clouds

- »» Include enterprise-class reliability and security features
- »» Use compression and change block tracking technologies to accelerate migration
- »» Provide intuitive interfaces that enable data transfer with the push of a button
- »» Offer direct compatibility across clouds, which wasn't possible before

Second, new multicloud storage offerings also provide replication with speed and efficiency that wasn't possible before. You can create instant clones by the dozens in no time and push them to the cloud. There is now an economical consumption model that

makes multicloud storage eminently practical: You pay as you go, with no prohibitive data egress charges and no vendor lock-in.

Finally, multicloud storage services provide consistent data services. When you move data into the cloud and between clouds, you get consistent performance and reliability without having to change anything.

Clustering Mission-Critical SQL Databases in the Cloud

SQL databases often contain the enterprise's most valuable, critical, and sensitive data, and any SQL failure could spell disaster. Unfortunately, clustering to protect data is typically not an option with cloud block storage because data can't be shared — cloud block storage can't be accessed by two compute instances, which is a fundamental requirement for clustering.

However, the clustering capabilities provided by a multicloud storage service make it possible to host mission-critical SQL databases in the cloud — with peace of mind. Clustering can provide redundancy so that even if the server-side fails, data remains available. Clustering can also give multiple compute instances access to the same storage while maintaining data protection, shared access to storage, and all the other benefits of the multicloud model.

Accelerating Development, Automating DevOps

With multicloud storage, developers can now easily create data sets and clones for building, testing, and deploying apps into production. Developers can make multiple copies — even a hundred or more — in an instant. So, if the company doesn't want to invest in more CapEx for development resources, it can adopt the multicloud storage model, copy data to the cloud, and give developers and testers fast access to the resources they need, and pay only for what's actually used. And now that the data is portable, databases, file servers, file shares, and other applications that rely on them can be portable, too.

Another way to speed up software development is to automate infrastructure provisioning. This requires support for public APIs that enable DevOps teams to simplify existing workflows, and that capability is built into multicloud storage offerings today. According to a May 2016 report by 451 Research titled *Automate or Die*, “In order to sustain growth, everything in the business must be process-driven and automated in software to the furthest extent possible.”

Moving an Application Using Docker Containers

Docker containers are suddenly everywhere. Containers pack multiple user-space instances into one receptacle, and that makes it easier to move entire applications — not just application data.

Enterprise IT and DevOps teams want to extend this portability to enterprise-class applications and workloads, and Docker containers make that possible. Now you can build, ship, and run with persistent data anywhere, without sacrificing production performance or storage efficiency, and without having to retrain staff.

Using the Cloud for Disaster Recovery

A recent ESG analyst survey reports that when it comes to leveraging the cloud to support cost-effective IT modernization, DR is one of the most frequent uses for cloud storage today. This is in large part due to the efficiency and appealing cost structure that a secure offsite data repository provides. And with the advent of cloud-ready data protection, cloud-based DR has become even more practical and cost-efficient.

With HPE Cloud Bank Storage, for example, backup data is self-describing to make it independent of the host storage system. In the unlikely event of the system being unavailable or failing, the data can be restored using another system for local or cloud-based disaster recovery. Backup data can be put to work by reading the backup data through another system. This ensures that an independent copy of the backup data is available for recovery testing, checking compliance, data mining, and other purposes.

By allowing enterprises to put data in the cloud — or in multiple clouds — the cloud-ready data protection model provides new options for DR while bringing down the cost. Instead of maintaining a separate DR site with redundant compute infrastructure that doesn't get used unless there's a disaster (or you're testing your DR capabilities), now you can simply move data copies to the cloud and avoid all that CapEx.

Cloud Bursting for Free (Almost)

Cloud bursting is simply a model in which any unexpected spike in demand for computing capacity (burst) is delivered by public cloud resources. Typically, this model is used in scenarios where an unanticipated event creates a sudden surge in demand, but cloud bursting is also used for foreseeable spikes such as the end of a quarter, a seasonal sale (like “Black Friday”), and so on.

In the past, cloud bursting could be an expensive and somewhat risky proposition, particularly for enterprise applications and data. However, with new multicloud storage innovations this is no longer the case. It is now possible to create many copies of data using cloning features without the need to pay for multiple copies of data. You can make many copies and host them where the demand is, making it possible to spin up needed infrastructure resources with very little incremental cost.

Another interesting version of this use case is also emerging: the ability to leverage more compute resources for a limited but predictable period of time — for example, to test a new ecommerce concept or to perform analytics on a data set — and then scale back to normal levels. The multicloud storage model also accommodates this use case with low and predictable incremental costs.

Switching Between Public Cloud Providers in a Flash

In the past, moving to the cloud sometimes meant the wholesale migration of data and applications; it required the purchase of additional hardware to handle migrations; and it exposed companies to unanticipated costs. Multicloud storage can give you the

freedom to compare and contrast cloud service providers, pick the ones that make the most economic sense for your workloads, move whatever portion of your data and apps you want — and not worry about getting locked in.

You can even move data between multiple public and private cloud providers to meet specific service-level agreements or respond to downtime issues. For example, if Provider A is experiencing a server outage, you can quickly shift data to Provider B for any amount of time, with minimal incremental cost. You're just changing the connection, so there is no data migration and no egress charge.

Monitoring and Predicting Usage Costs

By integrating with predictive analytics capabilities, multicloud storage services can now perform automated monitoring and tracking of the resources you're actually using, so you know what your costs are going to be. You can see precisely how much capacity you're using in any given period of time, so there are no surprises when you receive an invoice. Multicloud storage services can also help you estimate future usage based on any number of variables that you provide, and resize capacity on-the-fly to meet future requirements.

Using Cloud Storage like You'd Use Your Storage Area Network

Does your company produce its own electricity? Does it maintain its own water treatment facility? No? Then why should it have to acquire and maintain all the physical systems needed to store data? The multicloud model makes it easy to on-ramp data to the cloud and still get enterprise-grade data capabilities and consistent data services — so you can actually start using the cloud the way you used to use storage area networks, but for a lot less money.

Accelerating next



Hewlett Packard
Enterprise

CLOUD READY WHEN YOU'RE READY

Storage that Evolves with Your Cloud Strategy



HPE Storage is cloud ready. Combined with HPE multicloud storage services you now have the flexibility to run any workload you want, anywhere — on-premises or in the cloud. Built with native cloud integration to deliver efficiency and simplicity. Turn your cloud strategy and vision into reality.

Learn more at hpe.com/storage/multicloud



© Copyright 2018 Hewlett Packard Enterprise Development LP.

These materials are © 2018 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

Leverage multicloud storage across all your data centers

Multicloud storage services empower you to move your apps and data across public cloud and your on-premises data centers so you can optimize application performance, control costs, maximize reliability, and ensure data security. Make your data center cloud-ready and remove costs and complexities of extra hardware and software. Pick any workload, from production to backup, to run anywhere you want. Simplify and automate operations and rein in cloud spend with AI, predictive analytics, and DevOps tools.

Inside...

- Reduce “data gravity” to improve data mobility
- See why cloud block storage in the public cloud isn’t “enterprise-grade”
- Run primary flash to backup in the cloud
- See what matters most for a cloud-ready infrastructure
- Find out which multicloud use cases address your needs



Hewlett Packard Enterprise

Douglas Ko is a Senior Strategy and Market Development Manager at HPE. He is an expert in data storage and cloud computing and is interested in the role of cloud computing in enabling digital transformation at enterprises and service providers. **John Boutelle** is a technical writer and has worked with hundreds of enterprises worldwide, including Apple, HPE, Nokia, Oracle, Sun Microsystems, and VeriSign.

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

for
dummies[®]
A Wiley Brand



Also available
as an e-book

ISBN: 978-1-119-51634-7
Not for resale

WILEY END USER LICENSE AGREEMENT

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