

Foreword by
Martin Glassborow, aka Storagebod,
storage industry expert



Rethinking Enterprise Storage

A Hybrid
Cloud Model

Marc Farley

Visit us today at

microsoftpressstore.com

- **Hundreds of titles available** – Books, eBooks, and online resources from industry experts
- **Free U.S. shipping**
- **eBooks in multiple formats** – Read on your computer, tablet, mobile device, or e-reader
- **Print & eBook Best Value Packs**
- **eBook Deal of the Week** – Save up to 60% on featured titles
- **Newsletter and special offers** – Be the first to hear about new releases, specials, and more
- **Register your book** – Get additional benefits



Hear about it first.

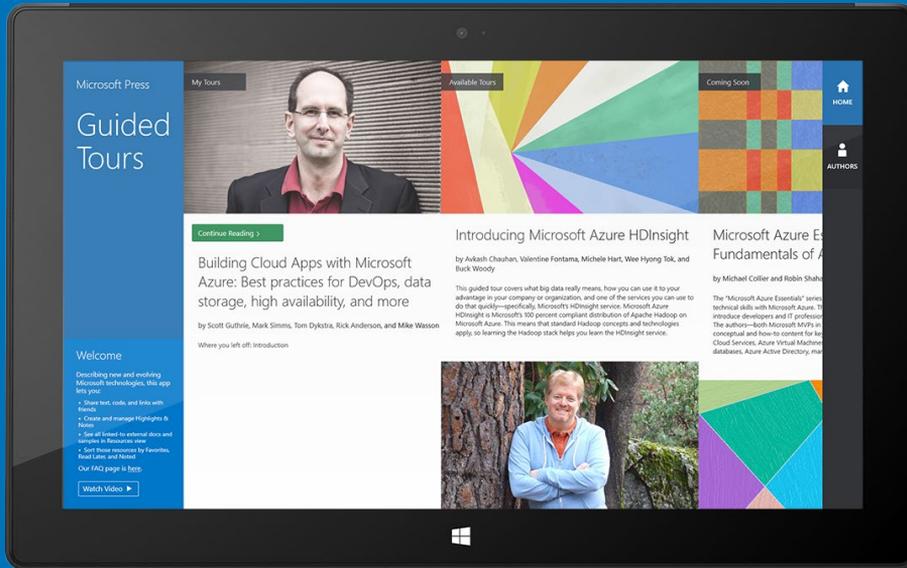


Get the latest news from Microsoft Press sent to your inbox.

- New and upcoming books
- Special offers
- Free eBooks
- How-to articles

Sign up today at MicrosoftPressStore.com/Newsletters

Wait, there's more...



Find more great content and resources in the Microsoft Press Guided Tours app.



The [Microsoft Press Guided Tours](#) app provides insightful tours by Microsoft Press authors of new and evolving Microsoft technologies.

- Share text, code, illustrations, videos, and links with peers and friends
- Create and manage highlights and notes
- View resources and download code samples
- Tag resources as favorites or to read later
- Watch explanatory videos
- Copy complete code listings and scripts





From technical overviews to drilldowns on special topics, get *free* ebooks from Microsoft Press at:

www.microsoftvirtualacademy.com/ebooks

Download your free ebooks in PDF, EPUB, and/or Mobi for Kindle formats.

Look for other great resources at Microsoft Virtual Academy, where you can learn new skills and help advance your career with free Microsoft training delivered by experts.

Microsoft Press

PUBLISHED BY
Microsoft Press
A Division of Microsoft Corporation
One Microsoft Way
Redmond, Washington 98052-6399

Copyright © 2013 Microsoft Corporation

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

Library of Congress Control Number: 2013939540
ISBN: 978-0-7356-7960-3

Microsoft Press books are available through booksellers and distributors worldwide. If you need support related to this book, email Microsoft Press Book Support at mspinput@microsoft.com. Please tell us what you think of this book at <http://www.microsoft.com/learning/booksurvey>.

"Microsoft and the trademarks listed at <http://www.microsoft.com/about/legal/en/us/IntellectualProperty/Trademarks/EN-US.aspx> are trademarks of the Microsoft group of companies. All other marks are property of their respective owners."

The example companies, organizations, products, domain names, email addresses, logos, people, places, and events depicted herein are fictitious. No association with any real company, organization, product, domain name, email address, logo, person, place, or event is intended or should be inferred.

This book expresses the author's views and opinions. The information contained in this book is provided without any express, statutory, or implied warranties. Neither the authors, Microsoft Corporation, nor its resellers, or distributors will be held liable for any damages caused or alleged to be caused either directly or indirectly by this book.

Acquisitions Editor: Anne Hamilton

Developmental Editor: Carol Dillingham

Project Editor: Carol Dillingham

Editorial Production: Christian Holdener, S4Carlisle Publishing Services

Technical Reviewers: Sharath Suryanarayan, Maurilio Cometto, and Guru Pangal

Copyeditor: Andrew Jones

Indexer: Jean Skipp

Cover: Twist Creative • Seattle

Contents at a glance

	<i>Foreword</i>	<i>ix</i>
	<i>Introduction</i>	<i>xi</i>
	<i>Next steps</i>	<i>xv</i>
CHAPTER 1	Rethinking enterprise storage	1
CHAPTER 2	Leapfrogging backup with cloud snapshots	11
CHAPTER 3	Accelerating and broadening disaster recovery protection	25
CHAPTER 4	Taming the capacity monster	43
CHAPTER 5	Archiving data with the hybrid cloud	57
CHAPTER 6	Putting all the pieces together	67
CHAPTER 7	Imagining the possibilities with hybrid cloud storage	81
	<i>Index</i>	<i>97</i>

Contents

<i>Foreword</i>	<i>ix</i>
<i>Introduction</i>	<i>xi</i>
<i>Next steps</i>	<i>xv</i>

Chapter 1 Rethinking enterprise storage 1

The hybrid cloud management model	1
The transformation of enterprise storage with cloud storage services	3
The constant nemesis: data growth	3
Increasing the automation of storage management	4
Virtual systems and hybrid cloud storage	4
Reducing the amount of data stored	5
Best practices or obsolete practices?	7
Doing things the same old way doesn't solve new problems	7
Introducing the hybrid cloud storage architecture.	8
Change the architecture and change the function	8
Summary.	9

Chapter 2 Leapfrogging backup with cloud snapshots 11

The inefficiencies and risks of backup processes	11
The many complications and risks of tape	12
Backing up to disk.	15
Virtual tape: A step in the right direction	15
Incremental-only backup	16

What do you think of this book? We want to hear from you!
Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

microsoft.com/learning/booksurvey

	Dedupe makes a big difference	17
	For the love of snapshots	17
	A big breakthrough: Cloud snapshots	18
	Fingerprints in the cloud	19
	Comparing cloud snapshots	20
	Looking beyond disaster protection	22
	Summary	23
Chapter 3	Accelerating and broadening disaster recovery protection	25
	Minimizing business interruptions	25
	Planning for the unexpected	26
	Disaster recovery with the Microsoft HCS solution	30
	Introducing the metadata map	31
	Recovery times with the hybrid cloud storage solution	33
	Windows Azure Storage as a recovery service	38
	Redundancy as a service: local and geo-replication	39
	Location-independent recovery	39
	Summary	40
Chapter 4	Taming the capacity monster	43
	The need for flexible storage	43
	Migrating data with server virtualization technology	43
	Thin provisioning brings relief	45
	Storage architectures: Scale-up, scale-out, and scale-across with cloud storage as a tier	47
	Scale-up and scale-out storage	47
	Scale-across storage	48
	Separating dormant data from active data with cloud-as-a-tier	49
	The life cycles of fingerprints	50

CiS designs for efficient working set storage.	53
Data reduction and tiering within the CiS system	53
Summary.	54
Chapter 5 Archiving data with the hybrid cloud	57
Digital archiving and electronic discovery	57
Protecting privacy and ensuring integrity and availability	59
Policies for managing data archives	59
Storage options for data archives	59
Archiving with the Microsoft HCS solution	61
Data archiving with Windows Azure Storage	61
Compliance advantages of Windows Azure Storage	62
Integrated archiving with the Microsoft HCS solution	62
A closer look at data retention policies with the Microsoft HCS solution	62
Meeting regulatory requirements for privacy, data integrity, and availability	65
Archiving data from ROBO locations	66
Summary.	66
Chapter 6 Putting all the pieces together	67
The complete picture of hybrid cloud storage	67
The system of fingerprints and pointers	68
Understanding hybrid cloud storage performance	71
Deployment scenarios for the Microsoft HCS solution	74
Summary.	78
Chapter 7 Imagining the possibilities with hybrid cloud storage	81
Thanks to VMs, everything done in data centers today can be done in the cloud tomorrow	81
Infrastructure virtualization	82

Data portability in the hybrid cloud	84
Migrating applications and copying data	84
Can you get there from here?	85
Recovery in the cloud	86
Big Data and discovery in the cloud	88
Summary.....	89
<i>Appendix</i>	91
<i>Glossary</i>	93
<i>Index</i>	97

What do you think of this book? We want to hear from you!
Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

microsoft.com/learning/booksurvey

Foreword

When I started my career in IT, storage was incredibly boring and something that most people tried to avoid. Enterprise data storage was the domain of strange people interested in tracks, cylinders, and data placements; they did not write code—they were the forgotten people.

Twenty-five years or so later, storage is neither boring nor straightforward. Data growth flows at exponential rates; structured data has been joined by unstructured data, the Facebook generation creates extensive social content in unprecedented quantities, and the enterprise is looking not only at how they store but also how they derive value from this content in the form of Big Data analytics. And somewhere along the line, I became a storage person—a StorageBod if you will.

We are at the centre of the storm brought on by cloud computing, and the promise of infinite scale and elasticity are changing the questions asked of enterprise storage. The certainty of managing data storage with enterprise arrays from the big five storage vendors is gone. There are now many possible answers to a problem that has moved away from simply being a case of how much capacity we require to store our application's data. Instead, we are thinking about how to balance user and business requirements in the context of flat-lining IT budgets. Should all our data be stored off-premises in the cloud or should we look at everything being stored in-house? Should all our data be stored in an object store? If so, whose?

This ambiguity brings increasing levels of complexity to the storage world. Data will live in many places on many different platforms and how we manage it, access it, and secure it for the enterprise is the next big question to be answered in storage.

Martin Glassborow

Blogger, *Storagebod.com*

June 2013

Introduction

Just as the Internet has fundamentally changed many industries, cloud computing is fundamentally changing the information technology industry, including infrastructures such as enterprise data storage. This book is about one of the new infrastructure game changers—a storage architecture called hybrid cloud storage that was developed by a company called StorSimple, now a part of Microsoft, as a way to integrate cloud storage services with traditional enterprise storage. Hybrid cloud storage is a completely different approach to storing data with a single comprehensive management system covering data through its entire life cycle, including active and inactive states as well as backup and archive versions. IT teams with cloud-integrated storage arrays running in their data centers use cloud storage as a data management tool and not simply as additional storage capacity that needs to be managed. That concept takes a little time to fully understand and it's why this book was written.

The audience for this book includes all levels of IT professionals, from executives responsible for determining IT strategies to systems administrators who manage systems and storage. The book explains how hybrid cloud storage changes the ways data protection is accomplished without tape backup systems; how disaster recovery works with data that is stored in the cloud; how cloud services are used to facilitate capacity management; and how the performance of data stored in the cloud is managed. Several applications for hybrid cloud storage are discussed to help IT professionals determine how they can use the Microsoft hybrid cloud storage (HCS) solution to solve their own storage problems. The last chapter is a hypothetical look into the future that speculates how this technology might evolve.

Conventions

The following naming conventions are used in this book:

- **The Microsoft HCS solution** The hybrid cloud storage solution discussed in this book combines a StorSimple-designed Cloud-integrated Storage system with the Windows Azure Storage service. This combination is referred to throughout the book as “the Microsoft HCS solution.”
- **Hybrid cloud boundary** The term is used in this book to identify the aspects of hybrid cloud that create a separation between computing on-premises and computing in the cloud. Physical location, bandwidth

availability, and latency are examples of things that can form a hybrid cloud boundary.

- **The IT team** The term refers to all the employees and contractors that work together to manage the technology infrastructure of an organization.

Sidebars are used throughout the book to convey information, ideas, and concepts in a less formal fashion or to draw attention to tangential topics that I thought might be interesting to readers. Sidebars are easy to identify by being offset from the rest of the text with a shaded background. An example of a sidebar is in Chapter 1, “Rethinking enterprise storage,” in the section “The hybrid cloud management model.”

Acknowledgments

Even a short book like this one has many contributors. I’d like to thank a number of people who helped make this book happen. Maurilio Cometto for his kind patience, Sharath Suryanarayan for his experience and perspective, Guru Pangal for his encouragement, Gautam Gopinadhan for his depth of knowledge, Mark Weiner for his unwavering support, Ursheet Parikh for his vision and faith, and Carol Dillingham for her insights and guidance throughout.

Errata & book support

We’ve made every effort to ensure the accuracy of this book. Any errors that have been reported since this book was published are listed on our Microsoft Press site at oreilly.com:

<http://aka.ms/HybridCloud/errata>

If you find an error that is not already listed, you can report it to us through the same page.

If you need additional support, email Microsoft Press Book Support at mspinput@microsoft.com.

Please note that product support for Microsoft software is not offered through the addresses above.

We want to hear from you

At Microsoft Press, your satisfaction is our top priority, and your feedback our most valuable asset. Please tell us what you think of this book at:

<http://aka.ms/tellpress>

The survey is short, and we read every one of your comments and ideas. Thanks in advance for your input!

Stay in touch

Let's keep the conversation going! We're on Twitter: <http://twitter.com/MicrosoftPress>.

Next steps

We hope this book piques your interest in the Microsoft hybrid cloud storage (HCS) solution. If you want to learn more about implementing the Microsoft HCS solution in your own enterprise, please visit the following site, where you can read case studies and request a demo:

<http://www.microsoft.com/StorSimple>

To connect with the author or other readers of this book, check out:

- Marc Farley's blog, "Hybrid Cloud Storage": *<http://blogs.technet.com/b/cis/>*
- The book's website: *<http://blogs.technet.com/b/cis/p/rethinkingenterprise-storage.aspx>*
- StorSimple on Twitter: *<https://twitter.com/StorSimple>*
- Marc Farley on Twitter: *<https://twitter.com/MicroFarley>*

Rethinking enterprise storage

The information technology (IT) world has always experienced rapid changes, but the environment we are in today is bringing the broadest set of changes that the industry has ever seen. Every day more people are accessing more data from more sources and using more processing power than ever before. A profound consequence of this growing digital consumption is that the corporate data center is no longer the undisputed center of the computing universe. Cloud computing services are the incubators for new applications that are driving up the demand for data.

IT managers are trying to understand what this means and how they are going to help their organizations keep up. It is abundantly clear that they need the ability to respond quickly, which means slow-moving infrastructures and management processes that were developed for data center-centric computing need to become more agile. Virtualization technologies that provide portability for operating systems and applications across hardware boundaries are enormously successful, but they are exposing the limitations of other data center designs, particularly constraints that hinder storage and data management at scale.

It is inevitable that enterprise storage technologies will change to become more scalable, agile, and portable to reflect the changes to corporate computing. This book examines how storage and data management are being transformed by a *hybrid cloud storage architecture* that spans on-premises enterprise storage and cloud storage services to improve the management capabilities and efficiency of the organization. The Microsoft hybrid cloud storage (HCS) solution is an implementation of this architecture.

The hybrid cloud management model

As a subset of hybrid cloud computing, hybrid cloud storage has received far less attention from the industry than the larger dialogue about how to enable hybrid applications. Nevertheless, pragmatic IT leaders are anticipating new hybrid cloud management tools to help them improve their IT operations. Hybrid cloud storage is an excellent example of this type of hybrid management approach that uploads data and

metadata from on-premises storage to the cloud, fulfilling the roles for a number of storage and data management practices.

Don't just take it from me

Another example of the power of hybrid management is the Hyper-V Recovery Manager which is described in an article written by John Joyner and published on the TechRepublic website titled "Hyper-V Recovery Manager on Windows Azure: Game changer in DR architecture." The article can be found by following this link: <http://www.techrepublic.com/blog/datacenter/hyper-v-recovery-manager-on-windows-azure-game-changer-in-dr-architecture/6186>. In the article Joyner explains how a cloud-based service controls the operations of on-premises systems and storage.

As a management abstraction, hybrid cloud management can provide centralized monitoring and control for on-premises and in-cloud systems and applications. If there are going to be applications and data that span on-premises and in-cloud resources, it only makes sense that there will be a need for management tools that facilitate those applications. Figure 1-1 depicts a hybrid cloud management model where three separate on-premises data centers are exchanging management information with resources and management services running in the cloud.

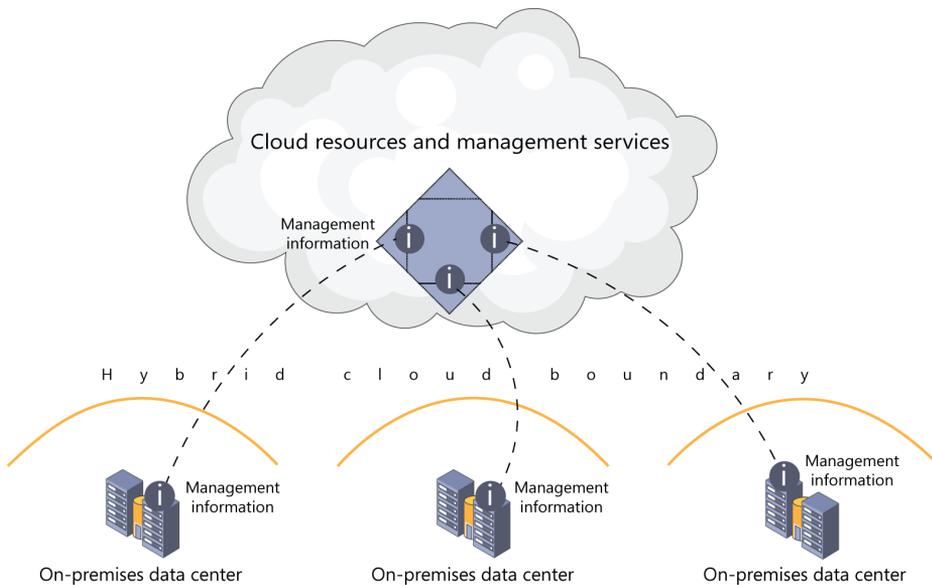


FIGURE 1-1 Three on-premises data centers exchange management information with cloud resources and management services across the hybrid cloud boundary.

The transformation of enterprise storage with cloud storage services

Storage has been an integral part of information technology from its inception and will continue to be throughout the cloud computing transformation that is underway. That's because all the data we create, use, and share has to be stored somewhere if it is to have more than fleeting value. A lot of this data is stored in corporate data centers, but a rapidly growing percentage is being stored in the cloud.

Enterprise storage architectures will need to adapt to this reality and integrate with cloud storage. Just as cloud services have changed the ways we consume data, they will also change how we store, manage, and protect it. It is short-sighted to think of cloud storage merely as big disk drives in the sky when there is so much compute power in the cloud to do interesting things with it. If it is possible to find information needles in data haystacks using data analytics, it is certainly possible to discover new ways to manage all that data more effectively. For example, the implementation of *erasure coding* in Windows Azure Storage demonstrates how advanced error-correction technology can also be used to effectively manage cloud storage capacity.

But the advancements in enterprise storage won't all be cloud-resident. In fact, many of the most important changes will occur in on-premises storage management functions that take advantage of hybrid cloud designs. The section "Change the architecture and change the function," later in this chapter, examines how extending traditional storage architectures with the addition of cloud storage services makes familiar storage management functions much more powerful.

The constant nemesis: data growth

IDC's Digital Universe study estimates that the amount of data stored worldwide is more than doubling every two years, so it is no surprise that managing data growth is often listed as one of the top priorities by IT leaders. IT professionals have ample experience with this problem and are well aware of the difficulties managing data growth in their corporate data centers. Balancing performance and data protection requirements with power and space constraints is a constant challenge.

IT leaders cannot surrender to the problems of data growth, so they need a strategy that will diminish the impact of it on their organizations. The hybrid cloud storage approach discussed in this book leverages cloud storage to offload data growth pressures to the cloud. Storage, which has always had an integral role in computing, will continue to have a fundamental role in the transformation to hybrid cloud computing—for its primary functionality (storing data) as well as its impact on those responsible for managing it.

Increasing the automation of storage management

Historically, storage management has involved a lot of manual planning and work, but as the amount of data continues to grow, it is clear that the IT team needs more automated tools in order to work more efficiently. This book describes how hybrid cloud storage enables higher levels of automation for many different tasks. Chapter 2, “Leapfrogging backup with cloud snapshots,” for instance, examines how hybrid cloud storage technology virtually eliminates the manual administration of one of the most time-consuming IT practices—backup.

People expect that their data will always be available when they want it and are unhappy when it isn't. Traditional data center solutions that provide high-availability with remote data replication are resilient, but have high equipment, facilities, and management costs—which means there's a lot of data that companies can't afford to replicate. Automated off-site data protection is an excellent example of a storage management function that is much more affordable with hybrid cloud storage. Chapter 3, “Accelerating and broadening disaster recovery protection,” explores this important topic.

Virtual systems and hybrid cloud storage

IT teams use virtualization technology to consolidate, relocate, and scale applications to keep pace with the organization's business demands and to reduce their operating costs. Hypervisors, such as ESX and ESXi from VMware and Hyper-V from Microsoft, create logical system images called virtual machines (VMs) that are independent of system hardware thereby enabling IT teams to work much more efficiently and quickly.

But virtualization creates problems for storage administrators who need more time to plan and implement changes. The storage resources for ESX and ESXi hypervisors are Virtual Machine Disk Format (VMDK) files, and for Hyper-V hypervisors, they are Virtual Hard Disk (VHD) files. While VMs are rapidly moved from one server to another, moving the associated VMDKs and VHDs from one storage system to another is a much slower process. VMs can be relocated from one server to another without relocating the VMDKs and VHDs, but the process of load balancing for performance usually involves shifting both VMs and VMDKs/VHDs. Data growth complicates the situation by consuming storage capacity, which degrades performance for certain VMs, and forces the IT team to move VMDKs/VHDs from one storage system to another, which can set off a chain reaction of VMDK/VHD relocations along the way. Hybrid cloud storage gracefully expands the capacity of storage, including VMDKs and VHDs, eliminating the need to move them for capacity reasons. By alleviating the pressures of data growth, hybrid cloud storage creates a more stable environment for VMs.

Data portability for hybrid cloud computing

VM technology is also an essential ingredient of cloud computing. Customers can instantly provision cloud computing resources as virtual machines running in the cloud without spending capital on equipment purchases. This gives the development organization a great deal of flexibility and allows them to test their work in a way they couldn't afford with their own equipment in their own data centers. The result is rapid application development that brings innovations to market faster.

Organizations want to develop software in the cloud and deploy it there or in their data centers, or in both places, using the hybrid cloud model. For example, Microsoft Windows Azure provides an environment that allows customers to deploy applications running on Windows Server 2012 with Hyper-V on Azure virtual machines.

If VMs can run either on-premises or in the cloud, companies will want a way to copy data across the hybrid cloud boundary so it can be accessed locally ("local" in the cloud context means both the VM and data are located in the same cloud data center). However, if copying data takes too long, the hybrid cloud application might not work as anticipated. This is an area where hybrid cloud storage could play a valuable role by synchronizing data between on-premises data centers and the cloud. Chapter 7, "Imagining the possibilities with hybrid cloud storage," discusses future directions for this technology, including its possible use as a data portability tool.

Reducing the amount of data stored

Considering that data growth is such a pervasive problem, it makes sense for storage systems to run processes that reduce the amount of storage capacity consumed. Many new storage arrays incorporate data reduction technologies, and the hybrid cloud storage design discussed in this book is an example of a solution that runs multiple data reduction processes—on-premises and in the cloud.

Know your storage math

Many of the advancements in storage and data management today are based on advanced mathematical algorithms for hashing, encoding, and encrypting data. These algorithms tend to assume that there is enough processing power available to not impact system performance and that the data being operated on is stored on devices with sufficient performance so bottlenecks can be avoided. Much of the design work that goes into storage systems today involves balancing the resources used for serving data with the resources used for managing it.

So, if data growth has been a problem for some time, why hasn't data reduction been used more broadly in enterprise storage arrays? The answer is the performance impact it can have. One of the most effective data reduction technologies is *deduplication*, also known as *dedupe*. Unfortunately, dedupe is an I/O intensive process that can interfere with primary storage performance, especially when device latencies are relatively high as they are with disk drives. However, enterprise storage arrays are now incorporating low-latency solid state disks (SSDs) that can generate many more I/O operations per second (IOPS) than disk drives. This significantly reduces the performance impact that dedupe has on primary storage. The Microsoft HCS solution discussed in this book uses SSDs to provide the IOPS for primary storage dedupe.

Chapter 4, "Taming the capacity monster," looks at all the various ways the Microsoft HCS solution reduces storage capacity problems.

Solid State Disks under the covers

SSDs are one of the hottest technologies in storage. Made with nonvolatile flash memory, they are unencumbered by seek time and rotational latencies. From a storage administrator's perspective, they are simply a lot faster than disk drives.

However, they are far from being a "bunch of memory chips" that act like a disk drive. The challenge with flash is that individual memory cells can wear out over time, particularly if they are used for low-latency transaction processing applications. To alleviate this challenge, SSD engineers design a number of safeguards, including metadata tracking for all cells and data, compressing data to use fewer cells, parity striping to protect against cell failures, wear-leveling to use cells uniformly, "garbage collecting" to remove obsolete data, trimming to remove deleted data, and metering to indicate when the device will stop being usable.

SSDs manage everything that needs to be managed internally. Users are advised not to use defrag or other utilities that reorganize data on SSDs. They won't perform faster, but they will wear out faster.

Best practices or obsolete practices?

The IT team does a great deal of work to ensure data is protected from threats such as natural disasters, power outages, bugs, hardware glitches, and security intrusions. Many of the best practices for protecting data that we use today were developed for mainframe environments half a century ago. They are respected by IT professionals who have used them for many years to manage data and storage, but some of these practices have become far less effective in light of data growth realities.

Some best practices for protecting data are under pressure for their costs, the time they take to perform, and their inability to adapt to change. One best practice area that many IT teams find impractical is disaster recovery (DR). DR experts all stress the importance of simulating and practicing recovery, but simulating a recovery takes a lot of time to prepare for and tends to be disruptive to production operations. As a result, many IT teams never get around to practicing their DR plans.

Another best practice area under scrutiny is backup, due to chronic problems with data growth, media errors, equipment problems, and operator miscues. Dedupe backup systems significantly reduce the amount of backup data stored and help many IT teams successfully complete daily backups. But dedupe systems tend to be costly, and the benefits are limited to backup operations and don't include the recovery side of the equation. Dedupe does not change the necessity to store data off-site on tapes, which is a technology that many IT teams would prefer to do away with.

Many IT teams are questioning the effectiveness of their storage best practices and are looking for ways to change or replace those that aren't working well for them anymore.

Doing things the same old way doesn't solve new problems

The root cause of most storage problems is the large amount of data being stored. Enterprise storage arrays lack capacity "safety valves" to deal with capacity-full scenarios and slow to a crawl or crash when they run out of space. As a result, capacity planning can take a lot of time that could be used for other things. What many IT leaders dislike most about capacity management is the loss of reputation that comes with having to spend money unexpectedly on storage that was targeted for other projects. In addition, copying large amounts of data during backup takes a long time even when they are using dedupe backup systems. Technologies like InfiniBand and Server Message Block (SMB) 3.0 can significantly reduce the amount of time it takes to transfer data, but they can only do so much.

More intelligence and different ways of managing data and storage are needed to change the dynamics of data center management. IT teams that are already under pressure to work more efficiently are looking for new technologies to reduce the amount of time they spend on it. The Microsoft HCS solution discussed in this book is a solution for existing management technologies and methods that can't keep up.

Introducing the hybrid cloud storage architecture

Hybrid cloud storage overcomes the problems of managing data and storage by integrating on-premises storage with cloud storage services. In this architecture, on-premises storage uses the capacity on internal SSDs and HDDs, as well as on the expanded storage resources that are provided by cloud storage. A key element of the architecture is that the distance over which data is stored is extended far beyond the on-premises data center, thereby providing disaster protection. The transparent access to cloud storage from a storage system on-premises is technology that was developed by StorSimple and it is called Cloud-integrated Storage, or CiS. CiS is made up of both hardware and software. The hardware is an industry-standard iSCSI SAN array that is optimized to perform automated data and storage management tasks that are implemented in software.

The combination of CiS and Windows Azure Storage creates a new hybrid cloud storage architecture with expanded online storage capacity that is located an extended distance from the data center, as illustrated in Figure 1-2.

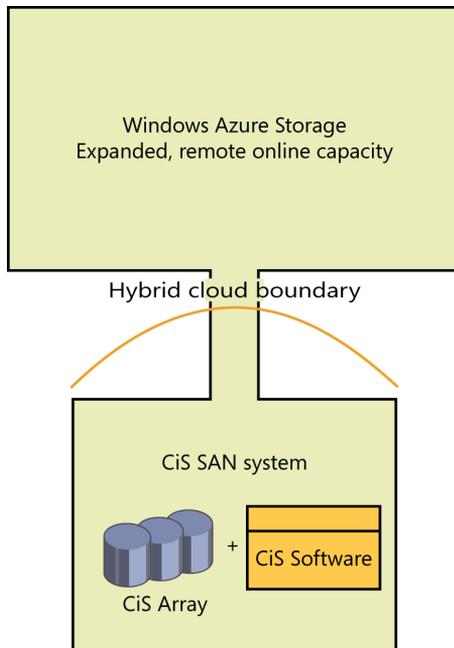


FIGURE 1-2 In the hybrid cloud storage architecture, the CiS SAN system accesses the expanded capacity available to it in Windows Azure Storage over an extended distance.

Change the architecture and change the function

CiS performs a number of familiar data and storage management functions that are significantly transformed when implemented within the hybrid cloud storage architecture.

Snapshots

CiS takes periodic snapshots to automatically capture changes to data at regular intervals. Snapshots give storage administrators the ability to restore historical versions of files for end users who need to work with an older version of a file. Storage administrators highly value snapshots for their efficiency and ease of use—especially compared to restoring data from tape. The main limitation with snapshots is that they are restricted to on-premises storage and susceptible to the same threats that can destroy data on primary storage.

Implementing snapshots in a hybrid cloud storage architecture adds the element of extended distance, which makes them useful for backup and disaster recovery purposes. Cloud snapshots are the primary subject of Chapter 2, “Leapfrogging backup with cloud snapshots.”

Data tiering

CiS transparently performs *data tiering*, a process which moves data between the SSDs and HDDs in the CiS system according to the data’s activity level with the goal of placing data on the optimal cost/performance devices. Expanding data tiering with a hybrid cloud storage architecture transparently moves dormant data off site to the cloud so it no longer occupies on-premises storage. This transparent, online “cold data” tier is a whole new storage level that is not available with traditional storage architectures, and it provides a way to have archived data available online.

Thin provisioning

SAN storage is a multitenant environment where storage resources are shared among multiple servers. *Thin provisioning* allocates storage capacity to servers in small increments on a first-come, first-served basis, as opposed to reserving it in advance for each server. The caveat almost always mentioned with thin provisioning is the concern about over-committing resources, running out of capacity, and experiencing the nightmare of system crashes, data corruptions, and prolonged downtime.

However, thin provisioning in the context of hybrid cloud storage operates in an environment where data tiering to the cloud is automated and can respond to capacity-full scenarios on demand. In other words, data tiering from CiS to Windows Azure Storage provides a capacity safety valve for thin provisioning that significantly eases the task of managing storage capacity on-premises.

Summary

The availability of cloud technologies and solutions is pressuring IT teams to move faster and operate more efficiently. Storage and data management problems are front and center in the desire to change the way data centers are operated and managed. Existing storage

technologies and best practices are being questioned for their ability to support data-driven business goals.

A new architecture called hybrid cloud storage improves the situation by integrating on-premises storage with cloud storage services providing both the incremental allocation of cloud storage as well as remote data protection. Extending the traditional on-premises storage architecture to include cloud storage services enables much higher levels of management automation and expands the roles of traditional storage management functions, such as snapshots and data tiering, by allowing them to be used for backup and off-site archiving.

The rest of the book explores the implementation of the Microsoft HCS solution and how it fundamentally changes how data and storage management is done.

Leapfrogging backup with cloud snapshots

When catastrophes strike IT systems, the IT team relies on backup technology to put data and systems back in place. Systems administrators spend many hours managing backup processes and media. Despite all the work that they do to prepare for the worst, most IT team members worry about how things would work out in an actual disaster.

IT professionals are hoping that cloud storage can make things easier and more reliable. The Microsoft hybrid cloud storage (HCS) solution promises to alleviate many of the chronic problems that have plagued backup with a new hybrid cloud technology called *cloud snapshots*. This chapter discusses existing backup technologies and practices that have been used for years and explains how cloud snapshots can significantly improve and simplify data protection.

The inefficiencies and risks of backup processes

If cloud storage had existed decades ago, it's unlikely that the industry would have developed the backup processes that are commonly used today. However, the cloud didn't exist, and IT teams had to come up with ways to protect data from a diverse number of threats, including large storms, power outages, computer viruses, and operator errors. That's why vendors and IT professionals developed backup technologies and best practices, to make copies of data and store them off site in remote facilities where they could be retrieved after a disaster. A single "backup system" is constructed from many different components that must be implemented and managed correctly for backup to achieve its ultimate goal: the ability to restore the organization's data after a disaster has destroyed it.

Many companies have multiple, sometimes incompatible, backup systems and technologies protecting different types of computing equipment. Many standards were developed over the years, prescribing various technologies, such as tape formats and communication interfaces, to achieve basic interoperability. Despite these efforts, IT teams have often had a difficult time recognizing the commonality between their backup systems. To many, it is a byzantine mess of arcane processes.

Technology obsolescence is another difficult aspect of data protection. As new backup storage technologies are introduced, IT teams have to manage the transition to those technologies as well as retain access to data across multiple technologies. This tends to be more problematic for long-term data archiving than backup, but it is a consideration that weighs on IT teams nonetheless.

Disaster recovery is the most stressful, complex undertaking in all of IT. Recreating replacement systems from tape backups involves many intricate details that are very difficult to foresee and plan for. Doing this without the usual set of online resources is the ultimate test of the IT team's skills—a test with a very high bar and no chance for a retry. Most IT teams do not know what their own recovery capabilities are; for example, how much data they could restore and how long it would take. When you consider how much time, money, and energy has been invested in backup, this is a sad state of affairs for the IT industry. Data growth is only making the situation worse.

The many complications and risks of tape

Magnetic tape technology was adopted for backup many years ago because it met most of the physical storage requirements, primarily by being portable so that it could be transported to an off-site facility. This gave rise to a sizeable ecosystem of related backup technologies and services, including tape media, tape drives, autoloaders, large scale libraries, device and subsystem firmware, peripheral interfaces, protocols, cables, backup software with numerous agents and options, off-site storage service providers, courier services, and a wide variety of consulting practices to help companies of all sizes understand how to implement and use it all effectively.

Tape media

Tape complexity starts with its physical construction. In one respect, it is almost miraculous that tape engineers have been able to design and manufacture media that meets so many challenging and conflicting requirements. Magnetic tape is a long ribbon of multiple laminated layers, including a microscopically jagged layer of extremely small metallic particles that record the data and a super-smooth base layer of polyester-like material that gives the media its strength and flexibility. It must be able to tolerate being wound and unwound and pulled and positioned through a high-tension alignment mechanism without losing the integrity of its dimensions. Manufacturing data grade magnetic tapes involves sophisticated chemistry, magnetics, materials, and processes.

Unfortunately, there are many environmental threats to tape, mostly because metals tend to oxidize and break apart. Tape manufacturers are moving to increase the environmental range that their products can withstand, but historically, they have recommended storing them in a fairly narrow humidity and temperature range. There is no question that the IT teams with the most success using tape take care to restrict its exposure to increased temperatures and humidity. Also, as the density of tape increases, vibration during transport has become a factor, resulting in new packaging and handling requirements. Given that tapes

are stored in warehouses prior to being purchased and that they are regularly transported by courier services and stored off-site, there are environmental variables beyond the IT team's control—and that makes people suspicious of its reliability.

Tape's metallic layer is abrasive to tape recording heads and constantly causes wear and tear to them. Over time the heads wear out, sometimes much faster than expected. It can be very difficult to determine if the problem is head wear, tape defects, or dirty tape heads. Sometimes the only remedy is to replace both the tape heads and all the tapes. The time, effort, and cost involved in managing wear-and-tear issues can be a sizeable burden on the IT group with no possible return on that investment to the organization. Tape aficionados are very careful about the tapes they buy and how they care for them, but many IT leaders no longer think it is worthwhile to maintain tapes and tape equipment.

Media management and rotation

Transporting tapes also exposes them to the risk of being lost, misplaced, or stolen. The exposure to the organization from lost tapes can be extremely negative, especially if they contain customer account information, financial data, or logon credentials. Businesses that have lost tapes in-transit have not only had to pay for extensive customer notification and education programs, but they have also suffered the loss of reputation.

Backup software determines the order that tapes are used, as well as the generation of tape names. Unfortunately, tapes are sometimes mislabeled which can lead to incomplete backup coverage, as well as making restores and recoveries more challenging. It sounds like a simple problem to solve, but when you consider that multiple tapes may have been used as part of a single backup job and that some tapes (or copies of tapes) are off site and cannot be physically checked, it turns out that there is not always a fast way to clear up any confusion.

Tape rotation is the schedule that is used by backup software to determine which tapes should be used for the next backup operation. If an administrator improperly loads the wrong tape in a tape drive, the backup software may not run, which means new data is not protected. Conversely, the backup software may choose to overwrite existing data on the tape, making it impossible to recover any of it. A similar problem occurs when a backup administrator erroneously deletes tape records from the backup system's database or erases the wrong tapes. Backup only works correctly when the database used to track data on tape accurately reflects the data that is recorded on tapes.

These sorts of problems are well-known to backup administrators and are more common than one might think. Backup administration and tape management tends to be repetitive, uninteresting work which sets the stage for operator oversights and errors. This is the reality of tape backup and it is why automated data protection with the Microsoft HCS solution from Microsoft is such an important breakthrough. It removes the responsibility for error-prone processes from people who would rather be doing something else.

When you look at all the problems with tape, it is highly questionable as an infrastructure technology. Infrastructures should be dependable above all else and yet, that is the consistent weakness of tape technology in nearly all its facets.

How many full copies do you really need?

Most tape rotation schemes make periodic full copies of data in order to avoid the potential nightmare of needing data from tapes that can't be read. The thinking is that tapes that were recently written to will be easier to recover from and that the backup data will be more complete. The simplest rotation scheme makes full copies once a week on the weekends and then once a day during workdays. Sometimes IT teams use other rotation schemes that include making full copies at monthly or longer intervals.

One of the problems with making full backup copies is that the operation can take longer to finish than the time available to get the job done. When that happens, system performance can suffer and impact productivity. Obviously, being able to skip making full copies would be a big advantage, which is how the Microsoft HCS solution does it.

Synthetic full backups

An alternative to making full backup copies is to make what are called *synthetic full copies*, which aggregate data from multiple tapes or disk-based backups onto a tape (or tapes) that contains all the data that would be captured if a full backup were to be run. They reduce the time needed to complete backup processing, but they still consume administrative resources and suffer from the same gremlins that haunt all tape processes.

The real issue is why it should be necessary to make so many copies of data that have already been made so many times before. Considering the incredible advances in computing technology over the years, it seems absurd that more intelligence could not be applied to data protection, and it highlights the fundamental weakness of tape as a portable media for off-site storage.

Restoring from tape

It would almost be comical if it weren't so vexing, but exceptions are normal where recovering from tape is concerned. Things often go wrong with backup that keeps it from completing as expected. It's never a problem until it's time to recover data and then it can suddenly become extremely important in an unpleasant sort of way. Data that was skipped during backup cannot be recovered. Even worse, tape failures during recovery prevents data from being restored.

Unpleasant surprises tend to be just the beginning of a long detour where restores are concerned. Fortunately, there may be copies from earlier backup jobs that are available to recover. Unfortunately, several weeks or months of data could be lost. When this happens, somebody has a lot of reconstruction work to do to recreate the data that couldn't be restored.

One thing to expect from disaster recovery is that more tapes will need to be used than assumed. Another is that two different administrators are likely to vary the process enough so that the tapes they use are different—as well as the time they spend before deciding the job is done, which implies the job is never completely finished. Most people who have conducted a disaster recovery would say there was unfinished business that they didn't have time to figure out. Their efforts were good enough—they passed the test—but unknown problems were still lurking.

Backing up to disk

With all the challenges of tape, there has been widespread interest in using disk instead of tape as a *backup target*. At first glance, it would seem that simply copying files to a file server could do the job, but that doesn't provide the ability to restore older versions of files. There are workarounds for this, but workarounds add complexity to something that is already complex enough.

Several disk-based backup solutions have been developed and have become popular, despite the fact that they tend to be more expensive than tape and require careful planning and administration. As replacements for tape, they rely on backup and archiving software to make a complete solution. When all the pieces of disk-based backup are put together it can get fairly complicated; however, most users of the technology believe it is well worth it as a way to avoid all the problems of tape.

Virtual tape: A step in the right direction

The desire to reduce the dependency on tape for recovery gave rise to the development of *virtual tape libraries* (VTLs) that use disk drives for storing backup data by emulating tapes and tape hardware. Off-site storage of backup data is accomplished by copying virtual tapes onto physical tapes and transporting them to an off-site facility. This backup design is called *disk-to-disk-to-tape*, or *D2D2T*—where the first disk (D) is in file server disk storage, the second disk (D) is in a virtual tape system, and tape refers to tape drives and media. Figure 2-1 shows a D2D2T backup design that uses a virtual tape system for storing backup data locally and generating tape copies for off-site storage.

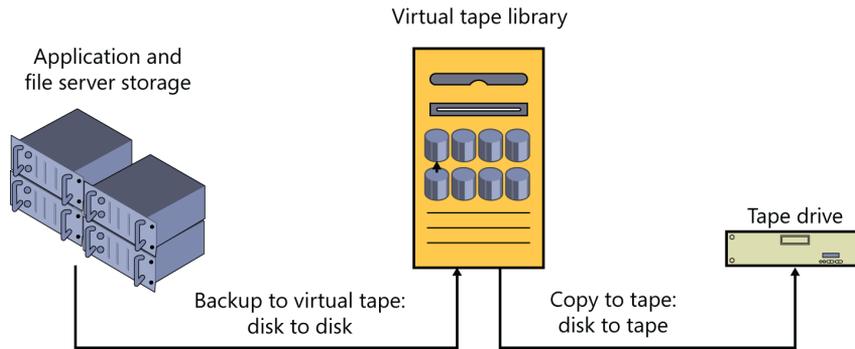


FIGURE 2-1 An illustration of the disk-to-disk-to-tape backup design.

VTLs significantly improve the automation of backup processes and provide good backup performance, but are more expensive than tape backup systems. Because the storage capacity of virtual tape products is limited, it might not be possible to backup as many servers or retain as much backup data as desired. For cost, capacity, and performance reasons, VTLs were mostly used in niche environments until dedupe technology was integrated with them and made them more widely applicable.

Incremental-only backup

The *incremental-only* approach to backup makes a single full backup copy and thereafter makes incremental backup copies to capture newly written data. If synthetic full tapes are not made, this approach leads to horrendously long and troublesome restores because every tape that was ever made might be needed for recovery. This implies copies need to be made of every tape in case they fail and also requires them to be stored in different locations, which means it might be necessary to have multiple copies at each location to account for media failures and so on and so forth. (It's funny what disaster paranoia will lead you to think about.)

That's why backup vendors developed disk-based, incremental-only backup systems that automatically copy backup data from a backup system at one site to another system at a remote location. When a disaster happens at the primary site, a full recovery can be made at the remote site from backup data in the remote system.

Incremental-only backup solutions integrate database, replication, and backup software along with the redundant hardware systems and facilities overhead at the remote site. Like other disk-based backup systems, they have capacity limitations that restrict the amount of backup data that can be kept, requiring management diligence and planning. Incremental-only backup systems are effective for solving backup problems, but, if the IT team also wants to reduce the cost of storage, incremental-only systems probably don't fit the bill.

Dedupe makes a big difference

A breakthrough in virtual tape technology came when dedupe technology was integrated with VTLs. Like previous-generation VTLs, dedupe VTLs require backup software products to generate backup data, but the dedupe function eliminates redundant data from backup streams. This translates directly into backup storage capacity savings and makes them much more cost-competitive with tape systems. Not only that, but dedupe VTLs improve backup performance by simultaneously backing up a larger number of servers and by keeping more backup copies readily available online. Many organizations happily replaced their tape backup systems with dedupe VTLs.

While dedupe VTLs have transformed backup for many IT teams, it has done relatively little to make disaster recovery easier. In most cases, tape copies still need to be made for off-site protection and the challenges of restoring data from tape are the same whether they were generated by tape drives or a dedupe VTL. However, like incremental-only backup solutions, some dedupe VTLs can also replicate data off site to another remote dedupe VTL, eliminating the need to make off-site tape copies—with the familiar caveats that remote replication adds additional systems and facilities costs as well as being more complicated to manage.

Dedupe variations: source and primary dedupe

Due to the success of dedupe backup systems, most people associate dedupe technology with target-side backup protection, but the technology can be successfully implemented other ways as well. *Source dedupe* implements dedupe technology before sending it over the network to be backed up. The main advantage of source dedupe is that it consumes far less bandwidth to transfer data and the main disadvantage is that it takes more processing resources on the server where the dedupe process runs.

Primary dedupe is the application of dedupe technology for primary production data, as opposed to being limited to backup data. The main advantage of primary dedupe is that it reduces the amount of capacity consumed on primary storage—which tends to be the most expensive storage in the data center. The main disadvantage of primary dedupe is the performance impact of running dedupe on production data.

For the love of snapshots

Snapshot technology is an alternative to backup that was first made popular by NetApp in their storage systems. Snapshots are a system of pointers to internal storage locations that maintain access to older versions of data. Snapshots are commonly described as making *point-in-time* copies of data. With snapshots, storage administrators are able to recreate data as it existed at various times in the past.

Snapshot technology is widely appreciated by IT teams everywhere for having saved them innumerable hours that they would have spent restoring data from backup tapes. It's no wonder that snapshot technology has become a key element of storage infrastructures and is one of the most heavily utilized features on most business-class storage systems.

While IT teams have largely replaced backups with snapshots for restoring historical versions of data, the two technologies are often used together in backup scenarios. Snapshots are used to capture updates to data and then backup processes capture the data from the snapshot. This keeps backups from interfering with active production applications and their data.

Continuous and near-continuous data protection

A technology that is closely related to snapshots is called *continuous data protection*, or *CDP*. CDP solutions are typically third-party software that monitors disk writes from servers and replicates them locally to a CDP server. The CDP server keeps the data and tags it with time stamps that help the IT team identify all the various versions that were saved. CDP solutions provide much greater granularity of data versions.

A variation of CDP is *near-CDP*, where the system doesn't quite stay up to speed, but is probably close enough for affordability's sake. *Data Protection Manager* from Microsoft is an example of a near-CDP solution that integrates tightly with Microsoft server products and allows users to pick and choose from numerous copies of the data on which they were working.

One problem with snapshots is that they consume additional storage capacity on primary storage that has to be planned for. The amount of snapshot data depends on the breadth of changed data and the frequency of snapshots. As data growth consumes more and more capacity the amount of snapshot data also tends to increase and IT teams may be surprised to discover they are running out of primary storage capacity. A remedy for this is deleting snapshot data, but that means fewer versions of data are available to restore than expected. In many cases, that may not be a huge problem, but there could be times when not being able to restore previous versions of data could cause problems for the IT team. Otherwise, the ease that snapshot capacity can be returned to free space depends on the storage system and may not be as simple as expected.

A big breakthrough: Cloud snapshots

The Microsoft HCS solution incorporates elements from backup, dedupe, and snapshot technologies to create a highly automated data protection system based on *cloud snapshots*. A cloud snapshot is like a storage snapshot but where the snapshot data is stored in Windows

Azure Storage instead of in a storage array. Cloud snapshots provide system administrators with a tool they already know and love—snapshots—and extend them across the hybrid cloud boundary.

Fingerprints in the cloud

The data objects that are stored as snapshots in the cloud are called *fingerprints*. Fingerprints are logical data containers that are created early in the data lifecycle when data is moved out of the input queue in the CiS system. While CiS systems store and serve block data to servers, they manage the data internally as fingerprints. Figure 2-2 illustrates how data written to the CiS system is converted to fingerprints.

The Microsoft HCS solution is not a backup target

One of the main differences between disk-based backup solutions that emulate tape, such as VTLs, and the Microsoft HCS solution is that the HCS is not a *backup target*, does not emulate tape equipment and does not need backup software to generate backup data. Instead, it is primary storage where servers store data and use cloud snapshots for protecting data in the cloud.

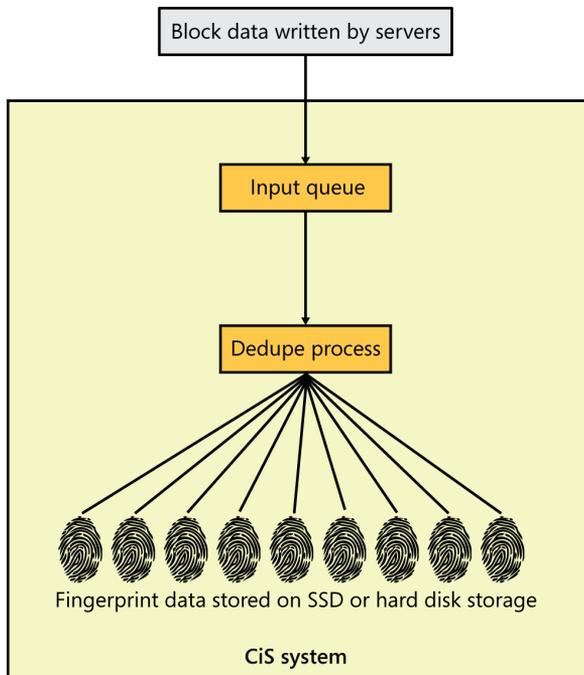


FIGURE 2-2 Block data is converted to fingerprints in the CiS system.

Just as backup processes work by copying newly written data to tapes or disk, cloud snapshots work by copying newly made fingerprints to Windows Azure Storage. One of the biggest differences between backup and cloud snapshots is that backup transforms the data by copying it into a different data format, whereas cloud snapshots copy fingerprints as-is without changing the data format. This means that fingerprints in Windows Azure Storage can be directly accessed by the CiS system and used for any storage management purpose.

Cloud snapshots work like incremental-only backups insofar that fingerprints only need to be uploaded once to Windows Azure Storage. Replication services in Windows Azure Storage makes multiple copies of the data as protection against failures. With most backup systems, there are many different backup data sets that need to be tracked and managed, but with cloud snapshots, there is only a single repository of fingerprints. In addition, there is no need to create synthetic full tapes because all the fingerprints needed to be recovered are located in the same Windows Azure Storage bucket.

Scheduling cloud snapshots

IT teams can flexibly configure their CiS systems to perform automated cloud snapshots to meet a broad range of requirements. Unlike tape backup systems that necessarily tie data expiration to tape rotation schedules, cloud snapshots can be assigned any expiration period. For instance, if the IT team decides they want to keep all cloud snapshot data for a minimum of three months, they can do it without having to worry about which tapes to use. Also, if the IT team wants to upload data more frequently, they can run cloud snapshots several times a day.

Efficiency improvements with cloud snapshots

Cloud snapshots eliminate tape problems and operator errors because there are no tapes to manage, lose, or go bad. No tapes need to be loaded for the next backup operation, no tapes are transferred off site, there are no tape names and labels to worry about, and no courier services need to be engaged. The arcane best practices that were developed for tape backup no longer apply to cloud snapshots. This is an enormous time saver for the IT team and removes them from the drudgery of managing tapes, tape equipment, and backup processes.

Data protection with cloud snapshots also eliminates the need to make full or synthetic full tapes. The incremental-only approach of cloud snapshots means that a minimal amount of data is copied and transferred. In addition, the fact that data is deduped on-premises before it is snapshotted means the amount of data that is uploaded is minimized.

Comparing cloud snapshots

The biggest difference between cloud snapshots with the Microsoft HCS solution and other backup products is the integration with Windows Azure Storage. Cloud snapshots improve data protection in three important ways:

1. Off-site automation. Cloud snapshots automatically copy data off site to Windows Azure Storage.

2. Access to off-site data. Cloud snapshot data stored off site is quickly accessed on premises.
3. Unlimited data storage and retention. The amount of backup data that can be retained on Windows Azure Storage is virtually unlimited.

Remote replication can be used to enhance disk-based backup and snapshot solutions by automating off-site data protection. The biggest difference between cloud snapshots and replication-empowered solutions is that replication has the added expense of remote systems and facilities overhead, including the cost of managing disk capacities and replication links.

Table 2-1 lists the differences in off-site automation, access to off-site data from primary storage, and data retention limits of various data protection options.

TABLE 2-1 A Comparison of Popular Data Protection Technologies

	Automates off-site storage	Access off-site data from primary storage	Data retention limits
Tape Backup	No	No	None
Incremental-only backup	Uses replication	No	Disk capacity
Dedupe VTL	Requires replication	No	Disk capacity
Snapshot	Requires replication	No	Disk capacity
Cloud snapshot	Yes	Yes	None

Remote office data protection

Cloud snapshots are also effective for automating data protection in remote and branch offices (ROBOs). These locations often do not have skilled IT team members on site to manage backup, and as a result, it is common for companies with many ROBOs to have significant gaps in their data protection.

Installing the Microsoft HCS solution in ROBO locations allows the IT team to completely automate data protection in Windows Azure Storage. This highlights another important architectural advantage—the many to one (N:1) relationship of on-premises locations to cloud storage. This design makes it possible for a Microsoft HCS solution at a corporate data center to access data from any of the ROBO locations. In addition, alerts from CiS systems running in the ROBOs can be sent to the IT team so they can remotely troubleshoot any problems that arise.

The role of local snapshots

CiS systems also provide *local snapshots* that are stored on the CiS system. Although local and cloud snapshots are managed independently, the first step in performing a cloud snapshot is running a local snapshot. In other words, all the data that is snapped to the cloud is also snapped locally first. The IT team can schedule local snapshots to run on a regular schedule—many times a day and on demand.

Looking beyond disaster protection

Snapshot technology is based on a system of pointers that provide access to all the versions of data stored by the system. The Microsoft HCS solution has pointers that provides access to all the fingerprints stored on the CiS system and in Windows Azure Storage.

The fingerprints and pointers in a Microsoft HCS solution are useful for much more than disaster protection and accessing point-in-time copies of data. Together they form a hybrid data management system that spans the hybrid cloud boundary. A set of pointers accompanies every cloud snapshot that is uploaded to Windows Azure Storage, referencing the fingerprints that are stored there. The system of pointers and fingerprints in the cloud is a portable data volume that uses Windows Azure Storage for both protection and portability.

This hybrid data management system enables additional data and storage management functions beyond backup and disaster recovery. For example, data tiering and archiving both take advantage of it to manage data growth and drive storage efficiency. Figure 2-3 illustrates the hybrid data management system spanning the hybrid cloud boundary, enabling data tiering and archiving, by storing and tracking fingerprints in Windows Azure Storage.

Cloud storage is for blocks too, not just files

One of the misconceptions about cloud storage is that it is only useful for storing files. This assumption comes from the popularity of file-sharing services and the fact that it's difficult to conceive of data objects that aren't files because that's what people use.

But Windows Azure Storage works very well for storing fingerprints and fingerprints definitely are not files—they are logical packages of contiguous *blocks* of data. Blocks are an addressing mechanism that operating systems use to calculate where to put data to maintain system performance. CiS systems exchange blocks, not files, with servers. On the cloud side, CiS systems exchange objects with Windows Azure Storage—and fingerprints, chock full of blocks, are the objects used by the Microsoft HCS solution. Just as the operating system translates files into blocks in order to store them on storage devices, CiS translates blocks into fingerprints so they can be stored in the cloud.

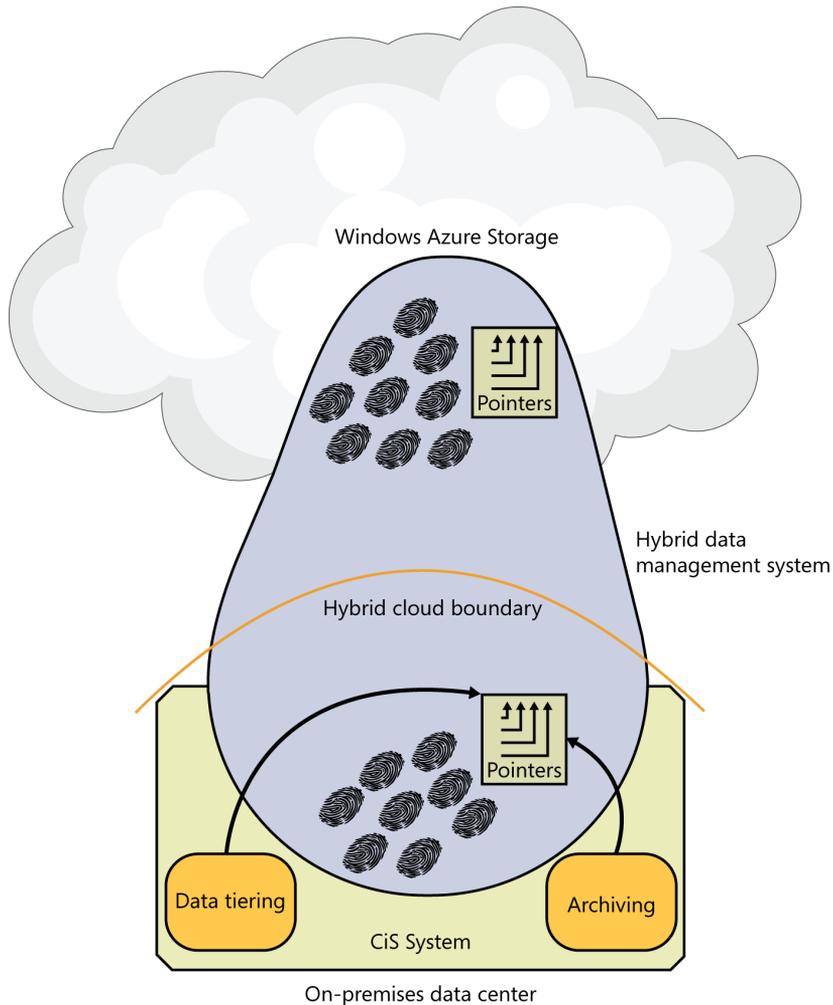


FIGURE 2-3 The Microsoft HCS solution unifies data management across the hybrid cloud boundary.

Summary

Backing up data, in preparation for recovering from a disaster, has been a problem for IT teams for many years due to problems with tape technology and the time-consuming manual processes that it requires. New technologies for disk-based backup, including virtual tape libraries and data deduplication, have been instrumental in helping organizations reduce or eliminate the use of tape. Meanwhile, snapshot technology has become very popular with IT teams by making it easier to restore point-in-time copies of data. The growing use of remote

replication with dedupe backup systems and snapshot solutions indicates the importance IT teams place on automated off-site backup storage. Nonetheless, data protection has continued to consume more financial and human resources than IT leaders want to devote to it.

The Microsoft HCS solution from Microsoft replaces traditional backup processes with a new technology—cloud snapshots, that automate off-site data protection. The integration of data protection for primary storage with Windows Azure Storage transforms the error-prone tedium of managing backups into short daily routines that ensure nothing unexpected occurred. More than just a backup replacement, cloud snapshots are also used to quickly access and restore historical versions of data that were uploaded to the cloud.

One of the key technology elements of the Microsoft HCS solution are granular data objects called fingerprints, which are created by the integrated dedupe process. The Microsoft HCS solution tracks all fingerprints on premises and in the cloud with a system of pointers that provide disaster recovery capabilities as well as the ability to recover point-in-time copies of data. The system of fingerprints and metadata pointers in the Microsoft HCS solution forms a hybrid cloud management system that is leveraged to provide functions that supersede those of backup systems grounded in tape technologies and processes. The next chapter, “Accelerating and broadening disaster recovery protection,” continues the discussion by showing how the hybrid data management system enables deterministic, thin full recoveries of data from Windows Azure Storage.

Accelerating and broadening disaster recovery protection

IT teams are constantly looking for ways to improve the processes, equipment, and services they use for disaster recovery (DR). The pressures of data growth and the importance of data to their organizations mean they need to cover more data with better DR technology at lower costs. The problem is that disaster preparation, like insurance, is a cost that returns nothing to the organization until something bad happens. DR solutions that can be leveraged for other purposes give a much better return on their investment.

Reducing downtime and data loss are important elements of any DR strategy. Many IT teams have DR strategies that focus most of the attention on a small number of mission-critical applications and largely ignore everything else. Everybody involved knows that this is unacceptable, which is why they are looking for better solutions to reduce downtime and data loss for all their applications, not just their top-tier applications.

This chapter begins by examining the requirements for DR, including recovery planning and testing before discussing remote replication. The Microsoft hybrid cloud storage (HCS) solution is introduced as a new, more flexible, and simpler approach to solving DR problems by virtue of being designed with the hybrid data management model.

Minimizing business interruptions

The goal of disaster preparation is to reduce disruptions to business operations. The ultimate goal is to avoid any downtime whatsoever. This can happen when the IT team has adequate time to prepare for an oncoming disaster and possesses the technology to shift production operations to an unaffected secondary site. For example, a company in the path of a hurricane may be able to execute a smooth transition of certain key applications from the primary site to a secondary site in a different geography before the storm arrives. Unfortunately, the disruption caused by disasters is usually unavoidable and unpredictable. That's when simple designs, reliable technologies, and practiced processes are most valuable.

Planning for the unexpected

DR plans are customized documents that identify the roles, processes, technologies, and data that are used to recover systems and applications. The best plans are comprehensive in scope, identifying the most important applications that need to be recovered first as well as providing contingency plans for the inevitable obstacles that arise from the chaos of a disaster. DR plans should be regularly updated to address changing application priorities.

Data growth complicates things by forcing changes to servers and storage as capacity is filled and workloads are redistributed. Hypervisors that allow applications and storage to be relocated help IT teams respond quickly to changing requirements, but those changes are somewhat unlikely to be reflected in the DR plan. That doesn't mean the application and data can't be restored, it simply means that the IT team could discover the plan isn't working as expected and have to rely on memory and wits. The Microsoft HCS solution accommodates data growth without having to relocate data onto different storage systems. The advantage of knowing where your data is and where it should be recovered to after a disaster cannot be emphasized enough.

You can't believe everything, even though it's mostly true

Statistics are often quoted for the high failure rate of businesses that do not have a DR plan when a disaster strikes. These so-called statistics probably are fictitious because there is no way of knowing if a business had a DR plan or if it moved to a new location, changed its name, or been out of business temporarily while facilities were being rebuilt. It's difficult to put calipers on survival when it can take so many forms.

However, it is also obvious that some businesses do indeed fail after a disaster and that businesses struggle to return to pre-disaster business levels. The loss of business records, customer information, and business operations contributes heavily to the eventual failure of a business. There are many things that are weakened by a disaster and losing data certainly doesn't help.

Practicing is a best practice

Testing DR plans and simulating recovery situations helps the IT team become familiar with products and procedures and identifies things that don't work as anticipated. It's far better to have an unpleasant surprise during a test run than during an actual recovery.

Unfortunately, many IT teams are unable to test their DR plans due to the disruption it would cause to normal business operations. Team members may need to travel, systems, servers, storage and applications may need to change their mode of operation or be temporarily taken off-line, workloads may need to be adjusted or moved, and any number of

logistical details can create problems for everyday production. The IT team can spend many days simply planning for the tests. It is somewhat ironic that a process intended to improve business continuity can cause interruptions to the business.

System virtualization technology has been instrumental in helping IT teams practice DR by making it simple to create temporary recovery environments. Unfortunately, verifying recovery processes may require restoring a large amount of data, which can take a long time. With compound data growth, this situation is only getting worse. The Microsoft HCS solution uses a recovery model called *deterministic recovery* that significantly reduces the amount of data that needs to be restored. It is discussed later in this chapter in the section titled “Deterministic, thin recoveries.”

Recovery metrics: Recovery time and recovery point

There are two metrics used to measure the effectiveness of disaster recovery: *recovery time* and *recovery point*. Recovery time is equated with downtime after a disaster and expresses how long it takes to get systems back online after a disaster. Recovery point is equated with data loss and expresses the point in the past when new data written to storage was not copied by the data protection system. For instance, if the last successful backup started at 1:00 AM the previous night, then 1:00 AM would be the presumed recovery point.

A good visualization for a recovery time and recovery point is the timeline shown in Figure 3-1. The disaster occurs at the spot marked by the X. The time it takes to get the applications running again at a future time is the recovery time. The time in the past when the last data protection operation copied data is the recovery point.

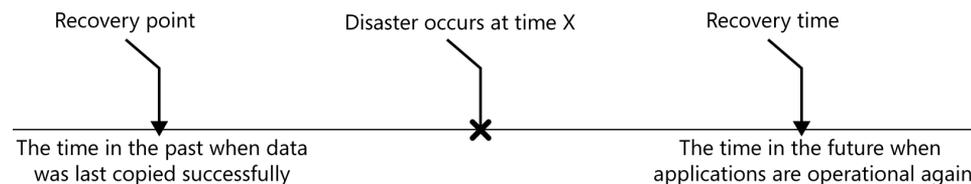


FIGURE 3-1 The timeline of a recovery includes the disaster, recovery point, and recovery time.

The timeline shown in Figure 3-1 will likely have different dimensions and scales depending on the importance of the application. The IT team sets recovery point objectives (RPOs) and recovery time objectives (RTOs) for applications based on their importance to the organization and the data protection technology they are using for recovery. The highest priority applications typically are protected by technologies providing the shortest RPOs and RTOs while the lowest priority applications are protected by tape backup technology, which provides the longest RPOs and RTOs.

Shortening RPOs and RTOs with remote replication

Remote replication, or simply *replication*, works by sending copies of newly written data to a remote site for the purpose of minimizing downtime after a disaster. In the best case, when a disaster strikes an application at the primary site, it can continue running at a secondary site without a loss of service. This can happen when the sites are relatively close to each other and are part of a remote cluster configuration or when special technologies and processes are used to *failover* the application to a different set of servers and storage. Replication allows the IT team to establish the shortest RPOs and RTOs.

Unfortunately, replication alone is insufficient as a data protection technology because it does not protect against threats like data corruptions or virus attacks. Data that is corrupted prior to being replicated will be corrupted on the secondary site too. In that case, the IT team will have to restore non-corrupted data from backup tapes.

Remote replication has been implemented different ways with various degrees of effectiveness and a range of costs, as discussed in the following sections.

STORAGE-BASED REPLICATION

Replication between storage systems is a proven method for providing excellent RPOs and RTOs. With storage-based replication, applications running on servers at a secondary site can read data that had been written to storage at the primary site only a few moments earlier.

Storage-based replication can easily multiply data center costs by adding the costs of duplicate storage and server systems, backup systems and software at both sites, low-latency network equipment, and the management, maintenance, and facilities overhead associated with running dual sets of equipment.

SERVER SOFTWARE REPLICATION

Replication solutions are also available through server software. For example, both Microsoft Exchange Server 2013 and Microsoft SQL Server 2012 have remote replication features for DR purposes. There are also several server software products from a number of vendors that replicate VMs and their data to remote sites. In general, server software replication is used for smaller, less active data sets than storage-based replication, but there is a great deal of overlap in the range of applications and scenarios where they are used.

Server software replication tends to be less expensive than storage-based replication, but still requires storage capacity at both sites, although the storage systems do not need to be similar. Other costs include the cost of servers, backup systems at both sites, and maintenance and management of the equipment and the facilities overhead of two sites.

DEDUPE VTL REPLICATION

Some dedupe VTLs feature remote replication in order to provide automated off-site copies of backup data. The amount of data that is transferred and stored by dedupe VTL replication is reduced because the data is deduped before it is copied. It is also likely that more applications may be protected because dedupe backup VTLs tend to protect a broader set of applications than storage or server software replication. When you consider the breadth of application coverage and the fact that a fully featured backup system is provided with the dedupe VTL, it can be argued that dedupe VTL replication is a more complete solution than storage-based and server software replication.

RPOs and RTOs with dedupe VTL replication are longer than they are in storage or server software replication because the replicated data is in a backup format and must be restored before it can be used by applications. RPOs with dedupe VTL replication are determined by when the last backup operation finished. The cost of dedupe VTL replication includes duplicate VTL systems, servers and storage at both sites (no need for them to be from the same vendor), backup software at both sites, and maintenance for the equipment and facilities overhead associated with operating two sites.

Replicating data growth problems is a problem

For all its strengths as a disaster recovery tool, remote replication has one very serious flaw: it doubles the amount of data that is stored. In the context of high data-growth rates, it's clear that replication should be used with discretion to avoid making the problem of managing data growth even worse.

Pragmatic IT teams know there are limited resources available to recover, which limits the number of applications that can be restored immediately following a disaster. That's why prioritizing applications for recovery is so important—mission critical applications need to be recovered and made operational before other lower priority applications are brought online. There is no point in jeopardizing higher priority applications by complicating replication with lower priority applications that aren't needed until later.

Dedupe has its advantages

Dedupe VTL replication is more efficient than storage-based or server software replication because it replicates data after deduping it on the primary site. Even though the RTOs and RPOs with dedupe VTL replication might extend further into the future and past, reducing the capacity needed for DR storage capacity is an advantage. However, once the data is restored at the recovery site, it will consume the same amount of capacity as at the primary site because the dedupe process is in the VTL and not primary storage. Now, if primary storage was also deduped, as it is with the Microsoft HCS solution, then the capacity efficiencies of deduplication are carried over after recovery.

Unpredictable RPOs and RTOs with tape

The problems encountered when recovering from tape were discussed in the section “Restoring from tape” in Chapter 2. IT teams struggle with setting RPOs and RTOs when tape is the data protection technology used for restoring data. RTOs established with tape are usually based on a best case scenario, something that rarely happens with tape DR scenarios. RPOs usually assume that backups finish successfully—an assumption that is, unfortunately, too often wrong. Considering the nature of backup failures and tape rotation mechanisms, IT teams can discover the actual recovery point changes from one day to one or two weeks in the past. That starts a completely different set of involved and thorny management problems.

In general, having overly aggressive RPOs and RTOs for tape restores sets expectations for the organization that might not be realistic, creating additional pressure on the IT team that may contribute to errors that lengthen the recovery process.

Disaster recovery with the Microsoft HCS solution

IT teams are looking for DR solutions that are less expensive and more comprehensive than remote replication and more reliable and faster than tape. Using cloud storage for data protection can be a solution, but slow download speeds must be overcome to achieve RTOs that can compete with tape.

The intelligent hybrid data management system in the Microsoft HCS solution combines excellent RPOs and RTOs with cost-competitive Windows Azure Storage, without adding to data growth problems. It is an excellent example of how using cloud resources to manage the IT infrastructure can improve existing data center practices.

The concept of DR with the Microsoft HCS solution is simple: fingerprints that were uploaded by cloud snapshots to Windows Azure Storage are downloaded again during a recovery process that is driven by a CiS system at a recovery site. Figure 3-2 illustrates the data flow for recovering data with the Microsoft HCS solution.

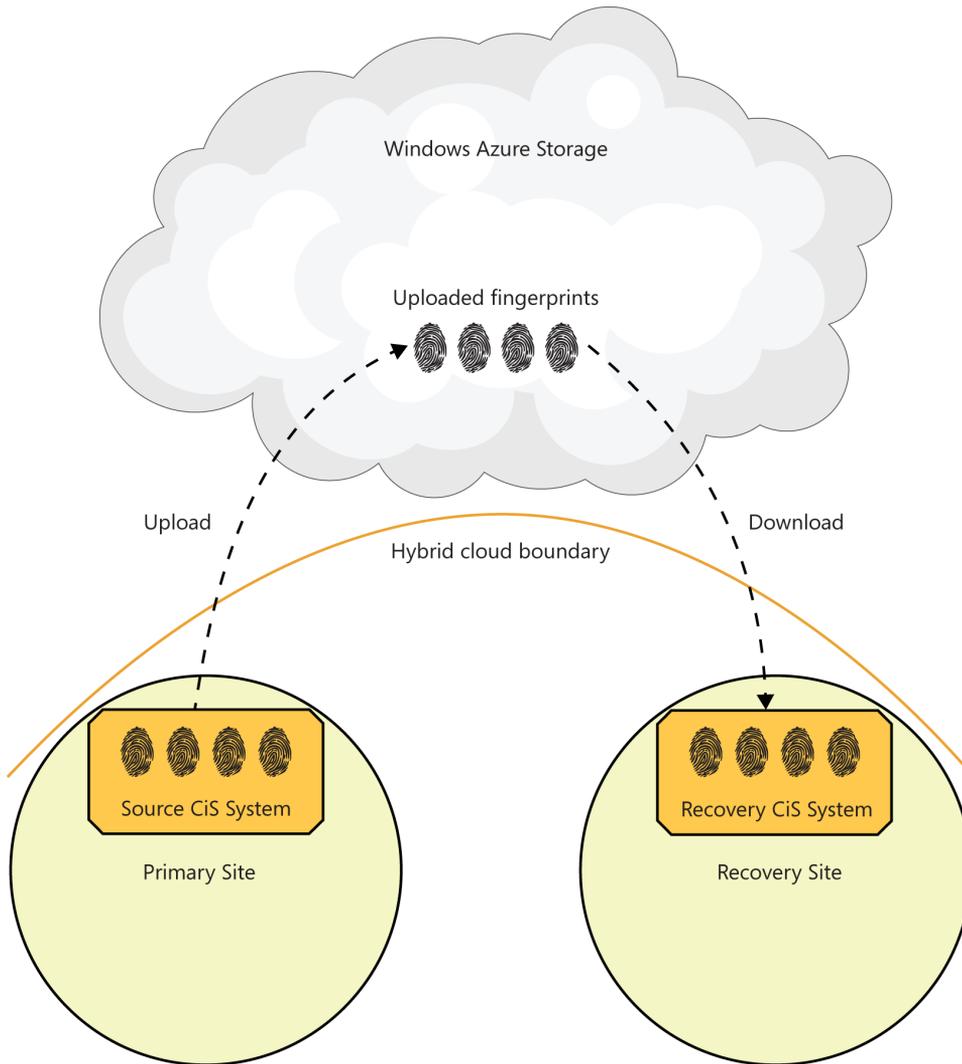


FIGURE 3-2 The data flow for recovering data with the Microsoft HCS solution.

Introducing the metadata map

The section titled “Looking beyond disaster protection” in Chapter 2 described the hybrid data management system of fingerprints, pointers and cloud snapshots that spans on-premises and Windows Azure Storage. One of the key elements of this system is the *metadata map*, a special object containing the pointers to all the fingerprints stored in the cloud. Every cloud snapshot operation uploads an updated version of the metadata map

as a discrete, stored object. When the process ends, the Windows Azure Storage *bucket* (cloud storage container) has an updated collection of fingerprints and a new metadata map with pointers to the locations of all fingerprints in the bucket. An individual metadata map consumes less than 0.3 percent of the capacity consumed by fingerprints.

A bucket by any other name

A *bucket* is the generic word for a storage container that holds data objects in the cloud. They are sometimes compared to large disk drives, but it is more useful to think of them as specialized servers that store data objects. They are accessed and managed using cloud APIs. A storage *volume* is the generic word for a storage container for data in on-premises storage systems. It is more frequently used for block data than for file shares, but it is sometimes used to refer to the container where a file share is.

In the hybrid cloud storage model, the contents of a volume are protected by uploading them as fingerprints to a Windows Azure Storage bucket. A Windows Azure Storage bucket typically stores fingerprints from multiple volumes on the CiS system. In fact, it is not unusual for a single Azure storage bucket to store the fingerprints for all the volumes in a CiS system.

Disaster recovery operations begin by selecting a cloud snapshot date and time and downloading the metadata map from its bucket to a *recovery CiS system*. When the map is loaded, servers and VMs at the recovery site can mount the storage volumes that had previously been on a *source CiS system*, and then users and applications can browse and open files. The fingerprints from the source CiS system are still on the other side of the hybrid cloud boundary, but can now be accessed and downloaded in a way that is similar to a remote file share.

Figure 3-3 shows the relationship between Windows Azure Storage, source and recovery CiS systems, and illustrates how the metadata map is uploaded, stored, and downloaded.

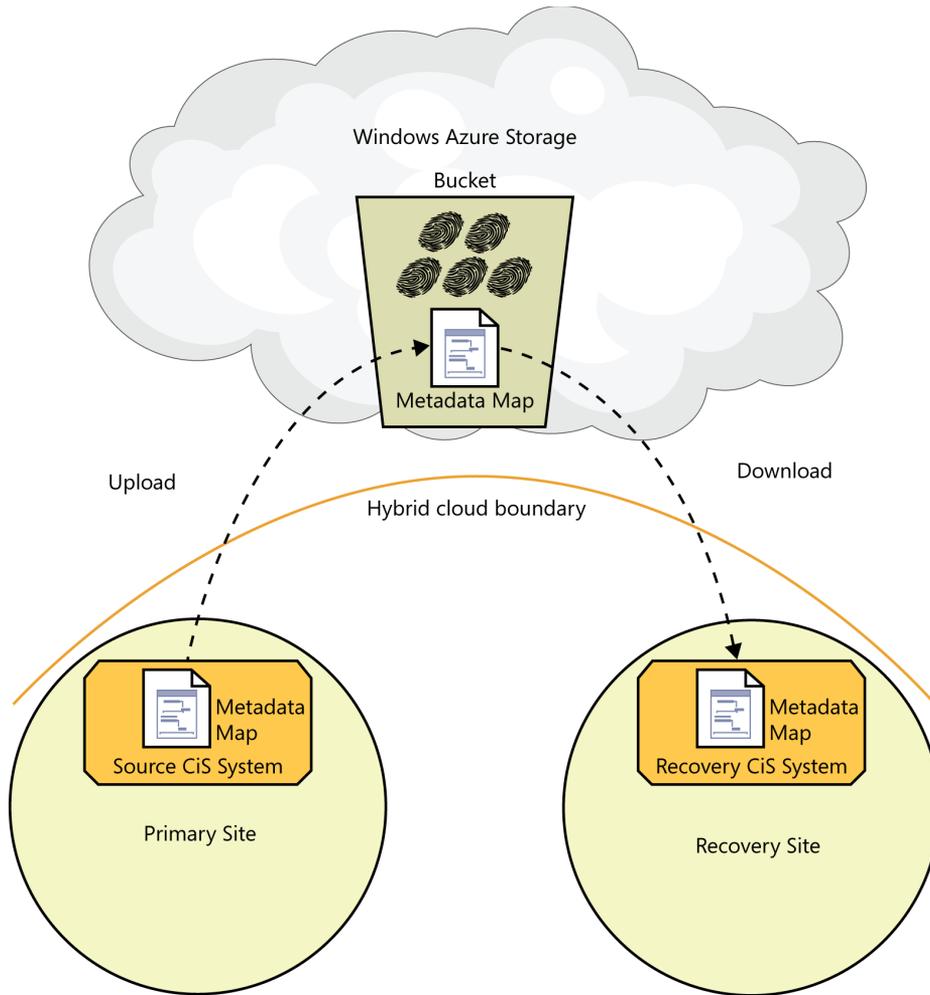


FIGURE 3-3 The metadata map that was uploaded by the source CiS system and stored in the Windows Azure Storage bucket is downloaded by the recovery CiS system.

Recovery times with the Microsoft HCS solution

As applications and users access data stored in the cloud the CiS system downloads their fingerprints and stores them on internal storage. It then sends the data to the requesting application. The time it takes to download data depends on its size and the available bandwidth, but office automation files can typically be downloaded in a few seconds.

As Figure 3-1 illustrates, recovery time is determined by the amount of time it takes for applications to resume operations. It follows that recovery times with the Microsoft HCS solution are determined by the time needed to access the Windows Azure Storage bucket

and download the metadata map. Download times for the data used by an application will impact the application's performance, but after the download completes, normal application performance will resume.

Your mileage will vary

There are many variables that can influence download performance and individual results will vary, nonetheless, readers will want some idea of download times for the metadata map. In a hypothetical example of 10 TB of data stored in the cloud over an unencumbered DS3 (44.7 Mbps) internet connection, the metadata map would likely be downloaded in less than 2 hours.

Deterministic, thin recoveries

One of the best DR practices is prioritizing the applications that will be brought online following a disaster. The IT team works to ensure the most important applications are brought back online first so the business can resume its operations. Prioritizing applications for recovery with the Microsoft HCS solution is a matter of bringing the applications online and connecting them to the CiS system in a prioritized order (or as prioritized groups). This ensures the most important applications get all the bandwidth they need to complete their data downloads before downloading lower priority applications.

Application-driven data recoveries with the Microsoft HCS solution are *deterministic* because every fingerprint that is downloaded is determined explicitly when an application accesses its data. They are also considered *thin* because data that is not needed is not downloaded. Deterministic, thin recoveries download far less data and consume far less network bandwidth than backup systems that are not driven by application behavior.

Deterministic, thin recoveries have critical efficiency benefits for the IT team. They don't require as much storage capacity at the secondary site as remote replication and backup solutions and they are also much easier to use for DR testing because they are far less intrusive. IT teams that have been unable to test their DR plans will appreciate the relative simplicity of DR testing with the Microsoft HCS solution.

In contrast, recoveries driven by backup software, including tape and dedupe VTL solutions, are *opportunistic* and restore as much data as possible without regard for application priorities. Data is read sequentially and application data is restored as it is encountered. While data transfer rates for tape and VTL systems are usually fairly fast, they recover everything, which takes a lot more capacity and resources at the secondary site.

Comparing recovery times with cloud storage as virtual tape

Just as disk drives in VTLs are used in place of tape in disk-to-disk-to-tape (D2D2T) designs, cloud storage can replace tape technology through the use of cloud storage gateways that emulate tape equipment. Instead of storing backup data on disk drives, they upload backup data to cloud storage.

Cloud-storage-as-virtual-tape automatically transfers data to off-site storage while avoiding the problems of physical tapes. However, performance is constrained by the bandwidth of the cloud connection, which tends to be several orders of magnitude slower than on-premises tape connections. This means that every operation done with cloud-storage-as-virtual-tape is very slow compared to physical tape drives and media. In other words, backup jobs or tape-to-tape copies that were designed with assumptions for high performance can take a very long time.

Unlike the Microsoft HCS solution, cloud-storage-as-virtual-tape is managed independently of primary storage by backup software. Recovery operations first download virtual tape images from the cloud before restoring data to primary storage arrays. Also, it is highly likely that multiple tape images will have to be downloaded to restore all the data needed by applications. The opportunistic restore model that tape backup uses, wastes a lot of time with cloud-storage-as-virtual-tape. The IT team needs to be aware of this when formulating their RTOs.

Your mileage will vary, part 2

Let's take the hypothetical example of 10 TB of data stored in the cloud with a DS3 Internet connection and estimate the difference in recovery times between the Microsoft HCS solution and using cloud-storage-as-virtual-tape with backup software. In the previous sidebar, we estimated the time to download the metadata map to be less than 2 hours. From then on, applications can access their data. With virtual tape, however, all the tape images would be downloaded first, which would probably take over 3 weeks. Using dedupe with a cloud-storage-as-virtual-tape would improve download performance considerably, but recovery times would likely be slower by an order of magnitude compared to the Microsoft HCS solution. Clearly, there are big differences in the way that cloud storage is used.

The working set

The fingerprints that are downloaded by applications during DR operations constitute a special instance of what is called *the working set*. Under normal circumstances, the working set is the data that users and applications access during daily application processing. During

recovery operations, applications and users determine the working set when they open files. After the CiS system returns to normal production operations, the working set becomes a dynamic entity that changes as new data is created and old data is accessed less frequently.

The Microsoft HCS solution was designed with the concept of placing the working set data on-premises and dormant data in the cloud. It provides applications and users access to dormant data in Windows Azure Storage whenever it is needed. This not only provides powerful management for data growth, but also has big implications for recovery.

Application coverage and data protection continuity

With the Microsoft HCS solution, every application, regardless of its priority, is recovered efficiently with deterministic, thin restores. The result is applications resume operations with their working sets at the recovery site while the data that is not needed remains on Windows Azure Storage.

Continuing data protection for all applications running at the recovery site is an important step that can be easily overlooked after all the excitement of a restore. IT team members can quickly and easily configure a new set of cloud snapshots on the recovery CiS system so that new data can continue being uploaded to Windows Azure Storage.

More cloud snapshots = more recovery points

Recovery points are determined by the cloud snapshot schedule and the data retention policies configured by the IT team. Typically, cloud snapshots are taken once in a 24-hour period—usually at night. However, cloud snapshots can be scheduled more frequently than once a day. IT teams that want three or four recovery points during the workday can easily set up a schedule for it.

The length of time that fingerprints are stored in Windows Azure Storage by the Microsoft HCS solution is determined by the data retention policy assigned to the cloud snapshot. IT teams typically set retention periods that match the tape rotation schedules they are familiar with, including weekly, monthly, quarterly, yearly and multi-year data retention. This subject is explored further in Chapter 5, “Archiving data with the hybrid cloud.”

Recoveries with spare and active CiS systems

The Microsoft HCS solution has an N:N architecture for recovering data. Some examples demonstrating the flexibility of this architecture are discussed in the following paragraphs.

A single, spare CiS recovery system can be installed at one of the sites operated by the IT team using an N:1 relationship to protect other data centers or ROBO locations. If a disaster occurs at any site, the spare could be used to recover data and resume operations.

This design works well except when disaster strikes the location housing the spare. In that case, other production CiS systems in other data centers can act as the recovery system. In the simplest example, a pair of CiS systems running in different data centers can be used to recover data for each other. The unaffected CiS system would serve data to the servers it normally does and would also add applications and workloads from the disaster site. This sort of 1:1 relationship is similar to one where two storage systems remotely replicate data to each other, however in this case, the two CiS systems do not communicate directly with each other.

In more interesting cases, one or more active CiS systems can be used to recover for disasters that strike multiple CiS systems. The general purpose N:N recovery architecture of the Microsoft HCS solution shows the power of using hybrid cloud management for DR by locating all recovery data in a centralized location and enabling recoveries to be conducted wherever there are sufficient resources to do so.

The cost advantages of an N:N architecture are appealing to IT teams that want to distribute the cost of DR equipment across multiple sites. Not only do they get flexible DR capabilities, but they also reduce their investment in capital equipment and the fully burdened cost of managing and operating that equipment.

Recoveries and cloud storage buckets

The ability of an active CiS system to download the metadata map for another CiS system highlights the fact that CiS systems are designed to work with multiple Windows Azure Storage buckets simultaneously.

Metadata maps are associated with a particular bucket and all the fingerprints stored in it. If a source CiS system is uploading data to two different buckets, it follows that there are also two metadata maps to download in a DR scenario. Furthermore, two different recovery CiS systems can be used to recover the data from the source system, each working with a different bucket. This allows the recovery operation to be done in parallel. Figure 3-4 illustrates a recovery operation where the data from a source CiS system that uses two storage buckets is being recovered on two different recovery CiS systems.

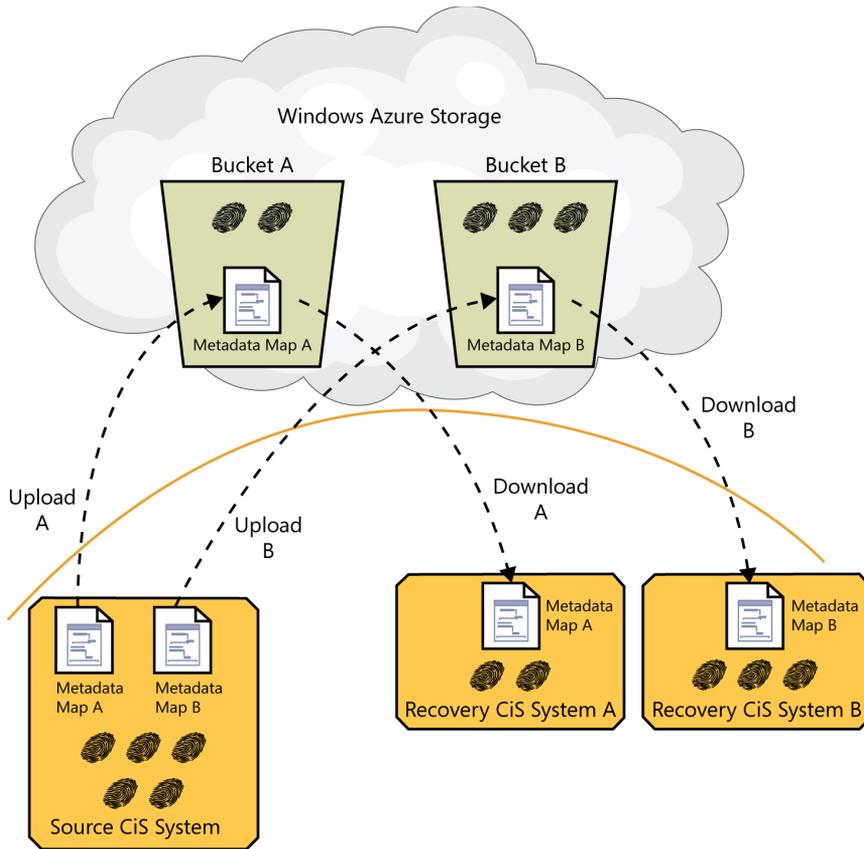


FIGURE 3-4 The process of parallel recovery from data stored in two Windows Azure Storage buckets to two different recovery CiS systems.

Windows Azure Storage as a recovery service

Windows Azure Storage provides granular scalability and built-in data protection for the Microsoft HCS solution. The following sections describe the recovery roles that Windows Azure Storage takes in the solution.

Disaster recovery services

Long before there were cloud services, organizations engaged DR service companies to help them prepare for disaster recoveries. These companies provide a number of valuable services, which might include a facility to recover in, storage and tape equipment, server systems, networking equipment, system software installation, recovery planning, and disaster simulation exercises to test the readiness of an IT team. Unfortunately, recovery services tend to be expensive and are not affordable options for many application scenarios.

Windows Azure Storage does not offer the same types of services, but instead provides affordable and reliable storage with built-in data protection features that IT teams can rely on to recover from a disaster. Rather than consulting on how to recover, Windows Azure Storage services is part of the actual recovery process.

Redundancy as a service: local and geo-replication

Windows Azure Storage has built-in data replication services that make redundant copies of data that has been uploaded to the cloud. When data is first uploaded, Windows Azure Storage makes three copies within the same (local) Windows Azure data center. Each copy is written to a separate fault domain within the Windows Azure data center so that a device or system failure will not result in data loss.

In addition to local replication, Windows Azure Storage also offers a service called *geo-replication*. Geo-replication replicates data asynchronously from one Windows Azure data center to a remote Windows Azure data center. As the replicated data is ingested at the remote Windows Azure data center, the local replication service there makes three copies of it.

Location-independent recovery

Through cloud snapshots, the Microsoft HCS solution uploads all the data needed for recovery into one or more Windows Azure Storage buckets. The portability of fingerprints and the metadata map makes it possible for one or more recovery CiS systems to access those buckets from virtually any location with a suitable Internet connection.

An organization does not have to operate multiple data centers in order to take advantage of location-independent recovery. An example would be a business with a primary data center that has the ability to quickly setup VMs and a spare CiS system in a local colocation facility. Location-independent recovery gives the IT team many options for developing a DR strategy that fits their operations and their budgets.

ROBO protection and recovery

As mentioned in Chapter 2, the Microsoft HCS solution can be effectively used to protect data at remote and branch office (ROBO) sites. With the N:N recovery architecture, each ROBO location uploads its fingerprints and metadata map to Windows Azure Storage, where it can be recovered to a CiS system in another ROBO location or a corporate data center.

Summary

Disaster recovery is a fundamental best practice for all IT teams, yet many of them struggle with the technologies, tools, and processes they have. The combination of data growth, the difficulty writing, updating, and testing DR plans, and the need to make DR more cost-effective is making it very difficult for IT teams to do the job the way they know it needs to be done. Solutions like remote replication work well to reduce RPOs and RTOs for a limited number of mission-critical applications, but the expense of owning and operating dual environments for replication means that a lot of data does not get the DR coverage that the organization needs.

The Microsoft HCS solution is based on the hybrid management model where deduped fingerprints on a source CiS system are uploaded to Windows Azure Storage where they can be downloaded to another recovery CiS system for DR purposes. The recovery data that is stored in the cloud does not consume floor space, power, or cooling costs in any of the organization's data centers. Fingerprints in Windows Azure Storage are protected in the cloud by replication and geo-replication services. One of the key management elements is an object called the metadata map, which contains pointers to all the fingerprints that were uploaded by the source CiS system. The combination of the fingerprints and the metadata map creates a portable, deduped data volume that can be downloaded to another CiS system during recovery operations.

In a recovery operation, the metadata map is downloaded first and then all the data that had been uploaded becomes visible to applications and users. Thereafter, the download process is driven by applications as they access their data. This deterministic, application-driven recovery process limits the data that is downloaded to only the deduped working set, leaving all the data that is not needed in the cloud. The thin, fast recovery capabilities of the Microsoft HCS solution enable IT teams to test their DR plans without disrupting their production operations. Recovery times with deterministic restores are short. Recovery points can be reduced by taking cloud snapshots several times a day.

The hybrid cloud management model enables a number of flexible, cost-reducing data recovery architectures. A single CiS system can be a spare for other CiS systems in a N:1 topology, or one or more CiS systems can be used to recover data for one or more disaster-stricken CiS systems in a N:N topology. There is no need to duplicate a data center environment for DR with the Microsoft HCS solution.

The flexibility and leverage gained through the hybrid cloud management model does not end with DR scenarios, but extends to other aspects of storage management as well. Chapter 4, “Taming the capacity monster,” continues the exploration by showing how the same fingerprints that were uploaded to Windows Azure Storage and used for DR purposes are also used to extend the capacity of on-premises CiS systems.

Taming the capacity monster

Data growth appears to be an unstoppable force in our digital world. There is virtually no cost to create and copy data, but it can be relatively expensive to store and manage. The challenge for IT teams is finding solutions that mitigate the storage and management costs of all this growing data and give them a way to accommodate data growth more flexibly in a more orderly, measured fashion. This chapter discusses some of the existing technologies IT teams use to manage their storage capacity and how the hybrid data management architecture of the Microsoft hybrid cloud storage (HCS) solution gives them powerful, new automated tools for managing data growth.

The need for flexible storage

File systems and databases depend on having static, fixed-size volumes to place data in so they can meet performance expectations and identify when storage capacity is running low. From the perspective of the IT team, the combination of static, fixed-sized storage volumes and high data growth rates is an unfortunate mismatch. When an application's data grows unexpectedly, storage volumes can run out of available capacity and the application may not be able to write data, causing systems to grind to a halt and crash.

The IT team does what it can to prevent this from happening by monitoring capacity levels and setting alerts for capacity thresholds. Despite these warnings, the IT team still has limited time to make difficult decisions about how to best respond, including emergency capacity purchases that are not in the budget, archiving and removing data that might be needed by users, and rebalancing application storage and workloads. Frustrated by chronic interruptions to their work, IT teams seek solutions that allow them to manage storage capacity with more flexibly and less stress and expense.

Migrating data with server virtualization technology

Server virtualization technology enables flexible storage capacity management through *storage migration* technology, such as the Storage VMotion (SVMotion) feature of VMware's ESX hypervisors and the Storage Live Migration feature of Windows Server 2012 and Hyper-V.

When a storage system approaches a capacity-full condition, some of the VMs storing their data there can have their data migrated transparently to another storage system.

Here's how it works. VM data is formatted and stored as a portable storage object called a virtual machine disk (VMDK) in VMware ESX environments or a virtual hard disk (VHD) in Windows Server 2012 Hyper-V environments. A VMDK/VHD on a source storage system can be synchronized with a copy running on a destination storage system. When all the data on the source is synchronized at the destination, the VM switches to using the destination VMDK/VHD. Figure 4-1 illustrates the process of storage migration.

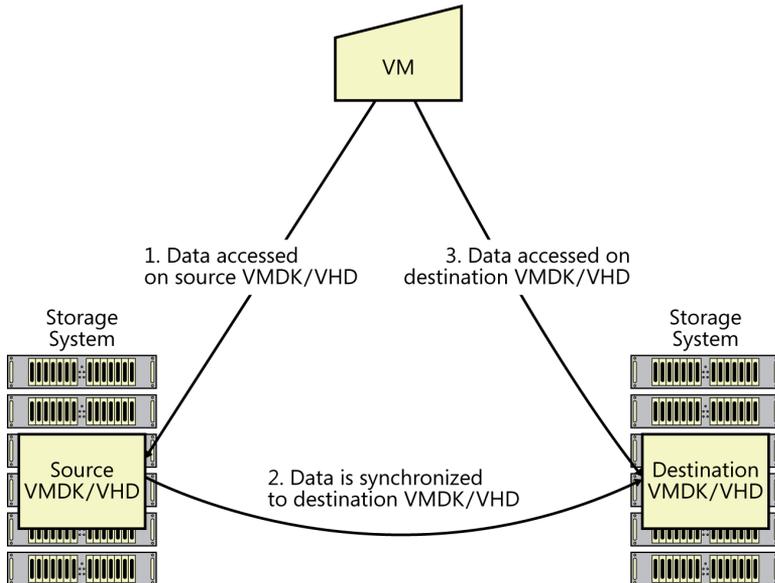


FIGURE 4-1 In storage migration, a VM moves its VHD from one storage system to another.

After the synchronization process has completed and the VM is accessing data from the destination VMDK/VHD, the source VMDK/VHD and all its data can be deleted and its capacity *reclaimed* and used for other purposes.

Storage migration is a powerful tool for managing data on a storage system that is approaching its capacity limits. Without it, the IT team would have to stop all the applications on the source volume, copy them to the destination volume, redirect the VM to access the volume on the destination, and then bring up the applications. This is a lot of work that involves application downtime and planning.

Storage migration depends on having a suitable destination storage system with sufficient capacity and performance capabilities to accommodate the new workload. Finding a suitable destination storage system becomes more difficult as utilization levels increase across available storage resources.

Solving VM sprawl with StorSimple

The Microsoft HCS solution works with SVMotion and Storage Live Migration in either the source or destination roles. Its scalability characteristics make it a valuable tool for managing environments suffering from VM sprawl.

VM sprawl occurs when there are so many VMs created that it is virtually impossible to manage them all. Over time, many of these “zombie” VMs are no longer used, but the data that was stored for them continues to consume storage capacity. IT teams use the Microsoft HCS solution as a destination storage system for migrated zombie data from other storage systems. The bottomless capacity of the Microsoft HCS solution with Windows Azure Storage really comes in handy when you are fighting zombies.

Thin provisioning brings relief

Thin provisioning technology was briefly discussed in Chapter 1, “Rethinking enterprise storage,” as a means to provide capacity to applications on a first-come, first-served basis. Instead of pre-allocating storage capacity as multiple, fixed-size volumes, thin provisioning aggregates storage capacity into a large, shareable resource and then parcels it out incrementally to elastic, virtual volumes as they need it.

With static provisioning, storage volumes are discrete resources, each with its own bounded address range. As new data is written, the capacity of the specific volume is consumed. With thin provisioning, each volume is expandable, using the shared free space as new data is written. Figure 4-2 compares how capacity is consumed using static provisioning, as opposed to thin provisioning.

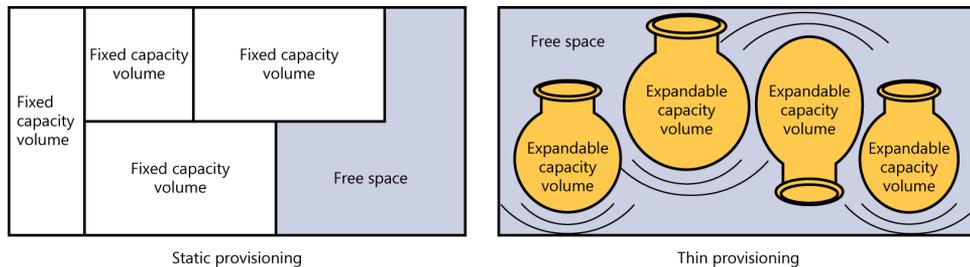


FIGURE 4-2 A comparison of static provisioning with thin provisioning.

Thin provisioning is widely appreciated by IT teams as a way to flexibly manage storage capacity in response to data growth. As a result, most enterprise storage systems today have thin provisioning as a feature. The Cloud-integrated Storage (CiS) system in the Microsoft HCS solution uses thin provisioning to allocate its internal storage resources to volumes. Some enterprise server software products, such as VMware vSphere and Windows Server 2012, also provide thin provisioning features.

The problem with thin provisioning is that most storage systems have fixed maximum capacities that can become exhausted. As the shared capacity becomes full, applications will start slowing and eventually crash if the IT team does not take action to resolve the situation. Thin provisioning is not a particularly dangerous tool, but it needs to be understood and managed according to best practices. There can be significant differences in how it is implemented. For example, the Microsoft HCS solution does not have a fixed maximum capacity by virtue of its ability to use Windows Azure Storage as a tier for data. The expandable free-space resources of Windows Azure Storage are used by the Microsoft HCS solution to change the dynamics of thin provisioning, as illustrated in Figure 4-3.

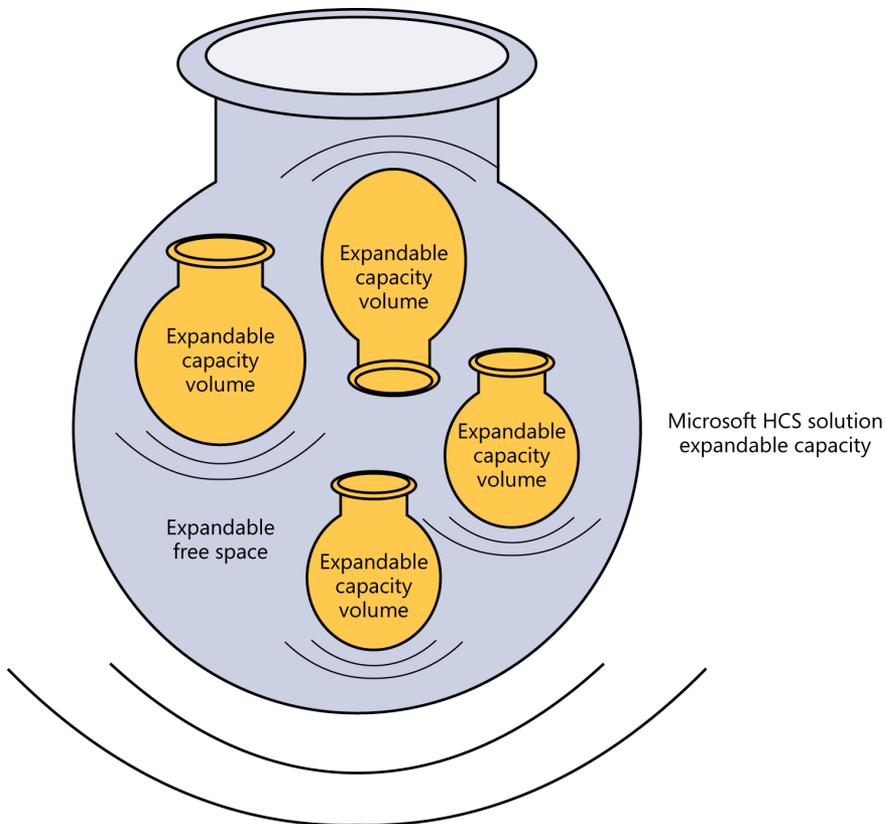


FIGURE 4-3 The Microsoft HCS solution offers expandable capacity for thin provisioning.

Storage on credit?

Thin provisioning solves serious capacity management problems for IT teams, but some people advise against using it to avoid running out of capacity unexpectedly and creating a self-inflicted disaster. An analogy that is sometimes used is thin provisioning is like buying capacity with unlimited credit backed by limited funds. The biggest problem with this line of thinking is that capacity problems are inevitable anyway, so why not use a power tool that helps deal with them?

The concept of over-provisioning is familiar to most IT team members through the use of virtual networking and virtual server technologies. Thin provisioning is just the over-provisioning concept applied to storage capacity. Thin provisioning solutions can—and usually do—provide ample warnings as capacity is consumed and exceeds certain thresholds. If the warnings are ignored, serious problems can occur.

Server software is also evolving to accommodate thin provisioning in storage arrays so that software can signal to storage what capacity can be reclaimed. For instance, a relatively recent standards-based command, UNMAP, allows thin provisioning storage systems to be used for more use cases, including storage with high turnover rates—where a large number of files are created and deleted again on a regular basis.

Storage architectures: Scale-up, scale-out, and scale-across with cloud storage as a tier

Storage systems are typically designed so that customers can purchase some amount of capacity initially with the option of adding more as the existing capacity fills up with data. Obviously, increasing storage capacity is an important task for keeping up with data growth. Scaling architectures determine how much capacity can be added, how quickly it can be added, the impact on application availability, and how much work and planning is involved.

Scale-up and scale-out storage

Storage systems that are not integrated with cloud storage increase their capacity in one of two ways: by adding storage components to an existing system, or by adding storage systems to an existing group of systems. An architecture that adds capacity by adding storage components to an existing system is called a *scale-up* architecture, and products designed with it are generally classified as *monolithic storage*.

Scale-up storage systems tend to be the least flexible and involve the most planning and longest service interruptions when adding capacity. The cost of the added capacity depends on the price of the devices, which can be relatively high with some storage systems. High availability with scale-up storage is achieved through component redundancy that eliminates single points of failure, and by software upgrade processes that take several steps to complete, allowing system updates to be rolled back, if necessary.

An architecture that adds capacity by adding individual storage systems, or nodes, to a group of systems is called a *scale-out* architecture. Products designed with this architecture are often classified as *distributed storage* or *clustered storage*. The multiple nodes are typically connected by a network of some sort and work together to provide storage resources to applications.

In general, adding network storage nodes requires less planning and work than adding devices to scale-up storage, and it allows capacity to be added without suffering service interruptions. The cost of the added capacity depends on the price of an individual *storage node*—something that includes power and packaging costs, in addition to the cost of the storage devices in the node.

Scale-up and scale-out storage architectures have hardware, facilities, and maintenance costs. In addition, the lifespan of storage products built on these architectures is typically four years or less. This is the reason IT organizations tend to spend such a high percentage of their budget on storage every year.

Scale-across storage

The Microsoft HCS solution uses a different method to increase capacity—by adding storage from Windows Azure Storage in a design that *scales across* the hybrid cloud boundary.

Scale-across storage adds capacity by incrementally allocating it from a Windows Azure Storage bucket. This allocation is an automated process that needs no planning and requires no work from the IT team. The cost of additional capacity with this solution is the going rate for cloud storage and does not include facilities and maintenance costs. However, using cloud storage this way involves much higher latencies than on-premises storage, which means it should be used for application data that can be separated by latency requirements—in other words, data that is active and data that is dormant.

Figure 4-4 compares the three scaling architectures.

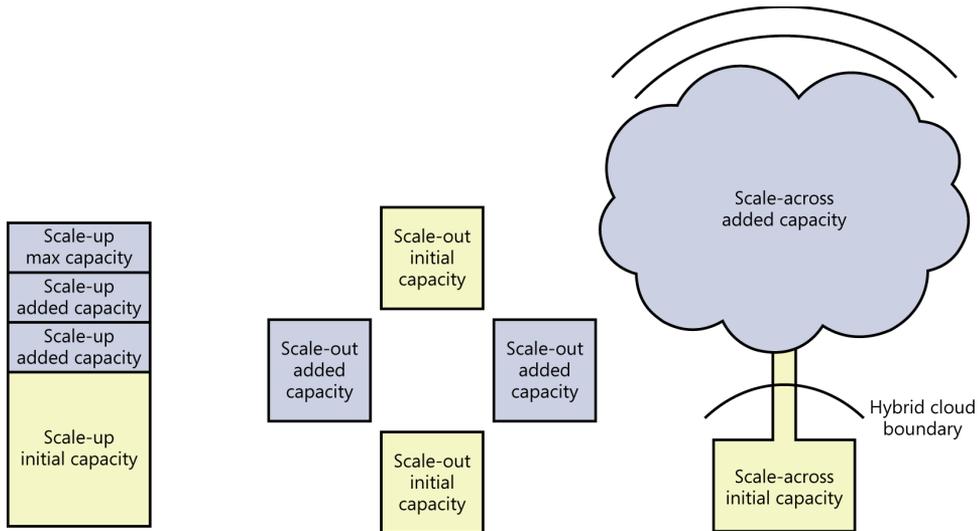


FIGURE 4-4 A comparison of storage scaling architectures.

Buying storage by the sip instead of the gulp

One of the advantages of purchasing storage capacity for the Microsoft HCS solution is that it can be done in very small increments. This removes the sometimes unpleasant process of purchasing capacity in large increments, the way IT teams are accustomed to buying capacity for scale-up or scale-out storage products. Another advantage for some IT teams is that the incremental capacity in the cloud can be acquired with operating budget money as opposed to spending from a capital budget.

Separating dormant data from active data with cloud-as-a-tier

Storage tiering is a feature of many storage systems. It is most commonly implemented as a way to place latency-sensitive data on high-performance, low-latency storage such as solid state disks (SSDs) and to place the rest of the data on disk drives. Storage tiering is also often described as a way to generate a high number of I/O operations per second (IOPs) from the optimal number of storage devices—SSDs generate most of the IOPs, while disk drives are responsible for most of the capacity requirements. *Cloud-as-a-tier* takes this idea one step further, by placing the data that has few to no IOPs on cloud storage.

Chapter 3, “Accelerating and broadening disaster recovery protection,” explains how deterministic, thin recoveries download only the working set and leaving the rest of the data in the cloud. The data that is left in the cloud stays there until it is accessed later, if ever. Viewed from the perspective of storage tiering, this *dormant data* has no IOPs requirement.

However, CiS systems do not start their existence as recovery targets, but as storage systems that servers use to store their application’s data. In these cases the working set is determined by the frequency that data is accessed—and not by deterministic, thin recoveries. To get a better understanding of how this works, it is important to understand data *life cycles*.

The life cycles of fingerprints

The data life cycle is an analysis of usage patterns for data as it ages. A high percentage of data on corporate file servers is *unstructured* data that is accessed most frequently in the days and weeks after it is created, and then is accessed less in the months that follow, until it is rarely, if ever, accessed. A large source of unstructured data is user-generated documents, including spreadsheets, word processing files, and presentations. Figure 4-5 illustrates the typical life cycle of corporate data.

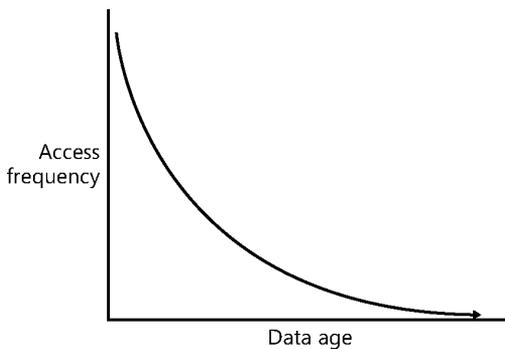


FIGURE 4-5 The typical data life cycle with access frequency as a function of data age.

The access frequency of data depends on the application, and there are certainly other life cycles that don’t follow the curve of Figure 4-5. For instance, accounting applications that process data on monthly or quarterly schedules will have a different curve. Nonetheless, a lot of data does follow this asymptotic curve, where access declines over time.

Data in the Microsoft HCS solution is encapsulated as fingerprints. As each fingerprint ages, it moves down the life cycle curve from top left to right bottom and at some point the capacity management algorithms in the CiS system determine that data needs to be relocated to the cloud. There is no specific point along the curve that determines when this happens because the process is driven by the available free capacity in the CiS system.

Remapping pointers for dormant data and accessing it in the cloud

When a dormant fingerprint is relocated, it would be logical to assume it is copied to cloud storage. However, in most cases, it was already copied there during the first cloud snapshot operation after it was created. Not only that, but every copy of the metadata map that was uploaded since that time has pointers to it and is also stored in the same cloud storage bucket. The pointers in the CiS system on-premises also reference its location in the cloud so it can be restored, if necessary. Everything is in place to locate the fingerprint in the cloud simply by changing its pointers—and that’s how tiering data to the cloud is accomplished, by changing pointers.

After the fingerprint’s pointers are changed, the fingerprint’s capacity in the CiS system is reclaimed so it can be used by new fingerprints. However, the data is still accessible and viewable on-premises as if nothing happened. From that point, if the data is accessed, the CiS system downloads its fingerprints from Windows Azure Storage transparently without any action required by IT team members.

The fingerprint twins, separated by the hybrid cloud boundary

The life cycle of a fingerprint in the Microsoft HCS solution includes the creation and management of twin copies that are on both sides of the hybrid cloud boundary. Both copies are managed as part of the hybrid cloud data management system. This is different than cloud backup or archiving solutions where the copies in the cloud are managed by a different process.

Because there is a single management system that spans on-premises and cloud storage, the processes for deleting on-premises copies and downloading cloud copies are transparent and automated. Transparency combined with automation is a hallmark of hybrid cloud storage.

Cumulative results of data growth and data life cycles

Another analysis compares all the accumulated data that is stored to the data that is currently active. If data is actively accessed for only several weeks, it follows that most of the data that has been stored over an extended period of time is dormant. Figure 4-6 shows a graph of active versus dormant data over time.

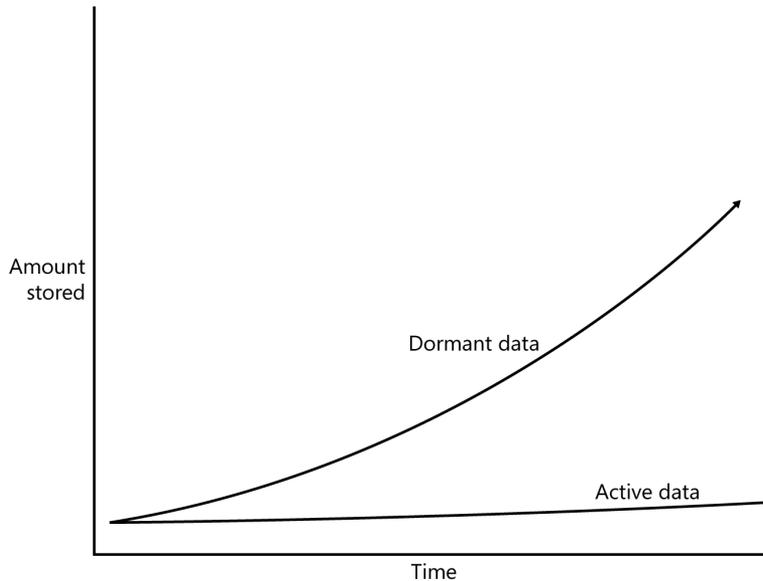


FIGURE 4-6 A graph comparing dormant data to active data over time.

The active data in Figure 4-6 is the working set stored in a CiS system. It follows that the dormant data in the figure is stored in cloud storage. As the graphic indicates, the size of the working set grows much more slowly over time, even though the overall data growth—and the growth of dormant data—is much faster. The reason for the difference in the two curves is because the working set does not accumulate data; instead, it continuously expels aging active data to dormant storage at a rate that tends to match the ingestion rate of new data.

Windows Azure Storage as a “bottomless” tier for dormant data

All this dormant data amassing in the cloud raises concerns about the scalability of cloud storage. One of the advantages of the Microsoft HCS solution is the fact that Windows Azure Storage is an enormous, endlessly-scaling object storage service that expands without restraints. The IT team does not need to provision or pre-allocate storage or thin provisioning storage pools, nor do they have to segment their expanding volumes or move them from one container to another. The storage volumes in the solution are thinly-provisioned on the CiS system and incrementally provisioned on Windows Azure Storage.

That said, the solution has capacity limitations based on the storage resources of the CiS system. All storage solutions have maximum capacities that should be understood and monitored. Capacity planning is something that should periodically be done as a best practice, to ensure there are not any unexpected problems looming.

CiS designs for efficient working set storage

Cloud-as-a-tier provides smooth, transparent scalability that removes most of the work of managing storage capacity problems. That said, there are additional capacity management technologies used by the Microsoft HCS solution. The rest of this chapter looks at these other data reduction technologies and how they minimize the impact of data growth.

Data reduction and tiering within the CiS system

As mentioned in the Chapter 1 section “Data tiering,” CiS systems also tier data across their SSD and HDD storage. The SSD tier contains data that has been deduped and the HDD tier contains data that has been both deduped and compressed. There is also an NV-RAM tier in the system that stores metadata but it does not contain application data.

Deduping primary storage

The high IOPs performance of the NV-RAM and SSD storage in the CiS system allow its dedupe process to run with a minimal impact on overall system performance. The end result is that all data stored on SSDs, HDDs, and Windows Azure Storage by the Microsoft HCS solution has been converted to fingerprints through the dedupe process. IOPs performance in the deduped/SSD tier is consistent with expectations for SSDs. The overhead from the dedupe process adds some latency, but this is much less than the physical latencies of disk drives. The small decrease in SSD performance is compensated for by the ability to store much more data on them.

Dedupe metadata spans volumes that are using the same Windows Azure Storage bucket. Deduping across volumes increases its effectiveness by increasing the potential of identifying duplicate data. Data that is sent to different Windows Azure Storage buckets belong to different dedupe domains.

Deduping data in the SSD tier of the CiS system has important implications for the capacity data consumes throughout its entire life cycle, increasing the amount of data that can be stored in the CiS system and reducing the capacity consumed in Windows Azure Storage. When you consider that the cost of cloud storage depends on the amount of capacity consumed, deduping it on-premises is a cost-saving benefit.

Dedupe differences matter

Backup dedupe ratios are typically much higher than dedupe ratios for primary storage. Dedupe processes are I/O intensive, which can create performance problems for other concurrent storage processes. Backup dedupe doesn't have any other storage processes to interfere with, but primary dedupe designs have to consider the performance of primary storage. Two ways to reduce the performance impact of dedupe is to put the dedupe data on low-latency storage and use a less aggressive dedupe process.

Although the dedupe ratios for primary storage are less than they are for backup storage, they can be leveraged in more significant ways. For starters, deduping primary storage directly counteracts the effects of data growth—which means the IT team does not have to purchase storage capacity as frequently. Deduping primary storage can also act like source-side dedupe, which reduces the amount of data that needs to be copied during data protection operations, for example, nightly cloud snapshots will copy less data to Windows Azure Storage.

Compression completes the reduction

The HDD data tier in a CiS system contains data that has been both deduped and compressed. As fingerprints age further and are accessed less frequently in the SSD tier, they are moved from SSDs, compressed to further reduce their capacity footprint, and written to HDDs. Data that is not likely to benefit from compression, such as media data that has already been compressed, is identified and skipped by the compression stage. Fingerprints that become more active while they are in the HDD tier are automatically moved back to the SSD tier.

As with dedupe, the data reduction benefits of compression are realized, not only by making more efficient use of resources on the CiS system, but by reducing the amount of data that is copied during cloud snapshots and reducing the capacity needed for them on Windows Azure Storage.

Summary

IT teams looking for ways to manage data growth have many things to consider, including how often they will need to upgrade the capacity of their storage systems and what the associated management overhead is. Automated storage and data management tools that provide flexible ways to structure storage and lower its costs are highly valued.

Virtual server tools, such as SVMotion and Windows Live Migration, that migrate VM data from one storage system to another are an effective way to manage capacity crises when

they occur. The Microsoft HCS solution works in these environments and provides a great destination for migrated VM data due to the elastic nature of its hybrid storage architecture.

Thin provisioning has proven itself as one of the best tools for dealing with data growth by allocating capacity dynamically to expandable storage volumes, instead of parceling it in advance into fixed-sized volumes. Nonetheless, the IT team needs to carefully manage the overall capacity of a thin provisioning system to avoid hitting the proverbial storage wall. The Microsoft HCS solution is a thin provisioning system that uses cloud-as-a-tier with Windows Azure Storage to create an expandable amount of storage that is shared among multiple expandable volumes.

Scale-up and scale-out storage architectures that were developed for enterprise storage have both hardware and facilities costs that companies want to control. The scale-across architecture of the Microsoft HCS solution is a completely new approach that automatically expands storage into Windows Azure Storage and significantly reduces the on-premises storage and management costs. As data grows, additional capacity is allocated from the cloud instead of consuming data center resources.

Built-in data reduction processes also contribute heavily to reducing the impact of data growth. Primary deduplication and compression that occurs early in the data life cycle increase the efficiency of storage wherever data is stored, whether it's on low-latency SSDs, capacity-oriented HDDs or low-IOPS Windows Azure Storage. The next chapter, "Archiving data with the hybrid cloud," talks about how that efficiency applies to archived data in Windows Azure Storage.

Archiving data with the hybrid cloud

Organizations depend on their IT teams to find and restore data that was *archived* for historical purposes in order to recall the conditions, discussions, decisions, and results of the past. Digital archiving is required in many industries to comply with government regulations for storing financial, customer, and patient information. For example, the health care industry is required to keep patient records for many years in order to inform future health care providers of a patient's history of conditions, diagnosis, and treatment. Many businesses that do not have explicit regulations governing digital archiving have defined internal policies and best practices that archive data for legal reasons because courts expect companies to produce internal records when they are requested. As businesses, governments, societies, and individuals increase their dependence on data, archiving it becomes more important.

This chapter discusses the technologies used for digital archiving and describes how the Microsoft hybrid cloud solution (HCS) can be used to archive data to Windows Azure Storage.

NOTE The word *archive* is sometimes used instead of “backup,” a topic that was discussed in Chapter 2, “Leapfrogging backup with cloud snapshots.” Sometimes the word *archive* is used to refer to migrating data from primary storage to free disk capacity, a topic that was discussed in Chapter 4, “Taming the capacity monster.” In this chapter the word *archive* refers to storing digital business records for an extended period of time.

Digital archiving and electronic discovery

There are two different use cases for digital archiving. The first is to create a repository of data that has intrinsic value and that people are interested in accessing. Libraries are excellent examples of archiving repositories that contain all sorts of information, including research data or documents that students and scientists may need to reference as part of their work. This form of digital archiving has become one of the most

important elements of library science, with specialized requirements for very long-term data storage (think millennia) and methods which are beyond the scope of this book.

The other use case for digital archiving is for business purposes and is the subject of this chapter. Digital archiving in the business context is one of the most challenging management practices in all of IT because it attempts to apply legal and compliance requirements over a large and growing amount of unstructured data. Decisions have to be made about what data to archive, how long to keep the data that is archived, how to dispose of archived data that is no longer needed, what performance or access goals are needed, and where and how to store it all.

Like disaster recovery (DR), archiving for legal and compliance purposes is a cost without revenue potential. For that reason, companies tend to limit their expenditures on archiving technology without hindering their ability to produce documents when asked for them. There are other reasons to archive business data but, in general, business archiving is closely tied to compliance and legal agendas.

The ability to find and access archived data tends to be a big problem. Storing dormant data safely, securely, and affordably for long periods of time is at odds with being able to quickly find specific files and records that are pertinent to unanticipated future queries. The selection of the storage technology used to store archived data has big implications for the long-term cost of archiving and the service levels the IT team will be able to provide.

Compliance and legal requirements have driven the development of electronic discovery (*eDiscovery*) software that is used to quickly search for data that may be relevant to an inquiry or case. Courts expect organizations to comply with orders to produce documents and have not shown much tolerance for technology-related delays. Due to the costs incurred in court cases, *eDiscovery* search and retrieval requirements are often given higher priority than storage management requirements. In other words, despite the desire to limit the costs of archiving, in some businesses, the cost of archival storage is relatively high, especially when one considers that the best case scenario is one where archived data is never accessed again.

Complete archiving solutions often combine long-term, archival storage with *eDiscovery* software, but there is a great deal of variation in the ways archiving is implemented. Many companies shun *eDiscovery* software because they can't find a solution that fits their needs or they don't want to pay for it. Unlike backup, where best practices are fairly common across all types of organizations, archiving best practices depend on applicable regulations and the experiences and opinions of an organization's business leadership and legal team. Digital archiving is a technology area that is likely to change significantly with the development of cloud-storage archiving tools in the years to come.

Protecting privacy and ensuring integrity and availability

Despite the relative importance of eDiscovery, it is not always the most important consideration in a digital archiving solution—protecting the privacy of individuals is. Privacy concerns extend to all forms of data storage, but archives have been explicitly addressed in regulations such as the Health Insurance Portability and Accountability Act (HIPAA) in the United States. In Europe, the Data Protection Directive covers all forms of stored data, including archives.

Encryption technology is commonly used to protect private, archived data from theft when it is in-flight or stored in the public cloud. That means IT teams need to consider the encryption features of their digital archiving solutions, including how encryption keys are managed in their storage solutions.

The integrity of archived data also needs to be ensured. Archival systems should be able to determine if the data being read is the same as the data that was written. The fact that data has been stored for an extended period of time makes it more susceptible to errors introduced by the physical degradation of stored data, sometimes called *bit rot*.

Some regulations for archiving require that redundant copies of archived data are stored in geographically remote locations to protect them from disasters. This can be done by making copies of archive tapes, taking regular backups of archive systems, replicating archive data to a remote site, or using cloud snapshots with Microsoft HCS. All remote copies need to meet the requirements for privacy and integrity.

Policies for managing data archives

One of the most common ways to manage data archives is by implementing policies and rules for data. Policies can be used to determine what data is selected for archiving and how long or where the archived data will be retained. For example, a policy could be established to retain all the data placed in a special archive folder for a minimum of ten years.

Most IT teams enforce data archiving policies through automated tools in their backup, archiving and eDiscovery solutions. Automation removes most of the human errors that could result in the loss of archived data and establishes the intent to comply with regulations. Corporate IT auditors will typically look to see that policies exist and are implemented to comply with applicable regulations.

Storage options for data archives

IT teams have chosen between tape and disk storage for storing archived data based on the cost and ease that archived data can be located and accessed. Cloud storage is now being looked at as an alternative. The sections that follow compare these three options.

Archiving to tape

IT teams often choose tape for storing digital archives because it is the least expensive option. Tapes can be stored for long periods of time with minimal operating costs, although they should be stored in low-humidity, air-conditioned facilities and be maintained by periodically rewinding them and checking their error rates, which may necessitate making new copies.

While data transfer speeds for tape are very good, the time it takes to retrieve tapes from an off-site facility is not. In addition, if no disk-resident archive *index* exists to identify the files or messages with pertinent data and map them to individual tapes, the process of finding archived information, that may have been stored across a large number of tapes, can take days or weeks. In this case, the best approach may be to restore the contents of archive tapes to a temporary disk storage volume and then search the data there. Depending on the amount of disk capacity available, this process may need to be repeated multiple times, clearing out the capacity of the temporary search volume each time and refilling it with data from different archive tapes. The money saved using tapes for archiving can be offset by lengthy data searches through the archives. Courts have not shown much patience in these matters and have levied expensive penalties to businesses that have not been able to produce data in a timely manner.

Archiving to disk

High-capacity disk systems are also used for archiving, even though they are much more expensive to operate than tape. A feature, referred to as drive *spindown*, has been incorporated into some disk systems to reduce power and cooling costs by selectively stopping individual disk drives in the storage system. When data is needed on drives that are spun down, the system starts them again and reads the data. The problem with spindown technology is that disk drives are generally not made to be powered on and off and sometimes they do not respond as expected. Application performance can also be erratic.

There is no question that disk is superior to tape for searching with eDiscovery solutions. The immediate access to files and the ability to search both production data and online archived data on disk saves everybody involved a great deal of time—which is a big deal to corporate legal teams. However, disk-based archiving still requires some form of disaster protection, which is usually tape, and all the overhead related to data protection, including administrative time, equipment, media, and facilities costs.

Archiving to cloud storage

Cloud storage is another option for archiving data that will evolve in the years to come with new cloud storage services. There are many variables to consider with cloud-based data archiving, including the type of interfaces used and the level of integration with on-premises storage. For example, cloud storage for archiving could be achieved by making it look like virtual tape in the cloud for backup, as discussed in the section titled “Comparing recovery times with cloud storage as virtual tape” in Chapter 3, “Accelerating and broadening disaster recovery protection.” In general, disk-resident indices that are accessed on-premises to

identify data objects in the cloud should be used for the same reasons they should be used with tape.

Another important consideration for using cloud storage for archival and compliance purposes is the documentation that is required by regulations. A cloud solution for corporate compliance needs to meet fairly strict guidelines to be a valid solution.

The remainder of this chapter examines Microsoft HCS as storage for long-term data archives using Windows Azure Storage.

Archiving with the Microsoft HCS solution

IT professionals are accustomed to thinking about archiving and backup as two related but different tasks and practices. The data and media used are often managed and maintained separately. Confusion over tapes for backup and archive can result in archive tapes being overwritten by backup processes and unexpected problems during recoveries. The media and equipment for tape backup and tape archiving might be similar, but the practices for both are decidedly different.

In contrast, Microsoft HCS automates both archiving and backup using cloud snapshots to upload fingerprints to Windows Azure Storage. Cloud snapshots used for backups typically expire in a few days to a few months, but cloud snapshots used for archiving may expire many years in the future, depending on compliance and governance requirements for archives.

Data archiving with Windows Azure Storage

Windows Azure Storage is being used successfully for data archiving. One of the most obvious advantages is that data stored there is off-site, but online, combining remote protection against site disasters with immediate access. Combined with the geo-replication service, Windows Azure Storage makes it significantly easier for IT teams to comply with regulations that mandate multisite disaster protection for archived data.

Archived data can be uploaded or downloaded from multiple corporate locations, enabling IT teams to flexibly design archiving workflows while simultaneously providing a centralized repository for accessing and exchanging data archives. Consolidating archives in Windows Azure Storage simplifies management of archived data by reducing the number of variables involved, including security management and encryption keys for all stored data.

As mentioned previously in the section “Archiving to cloud storage,” it is recommended that data archived to the cloud for long-term storage be searchable through on-premises, disk-resident indices.

Compliance advantages of Windows Azure Storage

Compliance with regulations can be complicated, especially the documentation that is required by auditors. Windows Azure has a website called the Trust Center that has information about compliance topics related to Windows Azure, including the HIPAA Business Associate Agreement (BAA), ISO/IEC 27001:2005 certification, and SSAE 16 / ISAE 3402 attestation. Windows Azure Storage services are named features for these compliance documents. The URL for this site is: <http://www.windowsazure.com/en-us/support/trust-center/compliance/>.

Integrated archiving with the Microsoft HCS solution

Long-term storage for archived data is an integrated feature of the Microsoft HCS solution. Data can be kept for an extended period of time on Windows Azure Storage simply by configuring cloud snapshot operations for that purpose. The next section, “A closer look at data retention policies with the Microsoft HCS solution,” describes how this is done.

An important advantage of archiving with the Microsoft HCS solution is that archived data on Windows Azure Storage is viewable on-premises by scanning folders or mounting cloud snapshots. The details of how this works depends on how archiving is implemented, either by archiving data in-place or by copying data to archive folders.

Archiving data in place provides default archival storage for the contents of a storage volume. It is essentially the same as backing up data with cloud snapshots, but with extended data retention policies for storing data in the cloud. IT team members can view data that was archived in place and later deleted from primary storage by mounting cloud snapshots that were taken before the data was deleted. For instance, a cloud snapshot with a data retention policy of five years could be mounted to look for archived data that was deleted from primary storage any time in the last five years.

Often IT teams want to create special folders for archive data. Containerizing archives this way may be required by archive software or best practices designed to enforce special treatment of archived data. For instance, an archive volume could be established with long-term data retention policies so that data written to it would be protected, long-term, by the next cloud snapshot process that runs.

A closer look at data retention policies with the Microsoft HCS solution

To illustrate how cloud snapshots are used for both backup and archiving, we'll follow the processes of using cloud snapshots to retain data for different lengths of time. Let's consider two cloud snapshot operations, referred to as CSbackup and CSarchive. CSbackup is configured to retain data for one week and CSarchive is configured to keep data for one year.

Assume a new fingerprint, FP1, is created before CSbackup runs. CSbackup then uploads a copy of FP1 to Windows Azure Storage and, at the end of its process, uploads an updated

version of the metadata map, MM1. MM1 contains pointers to all the fingerprints that were on premises when CSbackup ran, including FP1. The hybrid data management system is also updated, including a reference linking FP1 with CSbackup. For the next week, MM1 could be used to restore the CiS system as it existed when CSbackup ran.

Now, let's assume a second new fingerprint, FP2, is created before CSarchive runs. When CSarchive runs, it uploads a copy of FP2 to Windows Azure Storage and then uploads an updated version of the metadata map, MM2. This version of the metadata, MM2, has pointers to all the fingerprints that were in the system at the time, including FP1 and FP2. The hybrid data management system is also updated, linking FP1 and FP2 with CSarchive.

Next, assume that the retention period for CSbackup expires. At this point, metadata map MM1 is deleted from Windows Azure Storage and the hybrid cloud management system is updated to remove all the references linking fingerprints to CSbackup. However, MM2 and all its fingerprint links still exist, including those for FP1 and FP2. MM2 and all the fingerprints linked to it will be retained in Windows Azure Storage for the next year and can be used to restore the data that was on-premises when CSarchive ran.

Fingerprint expiration and managing temporary data

Fingerprints are deleted from cloud storage when there are no longer any metadata maps that reference them. Let's assume a very simple scenario where a volume is protected by a single cloud snapshot that expires in one week. If no other cloud snapshots are run, eight days later the hybrid cloud data management system will remove all links to that cloud snapshot and the metadata map and the expired fingerprints will be deleted. If another cloud snapshot is run again on the volume, the fingerprints and a new metadata map would be uploaded again.

Policies and practices for protecting temporary volumes are easily established so that unimportant data does not consume cloud storage capacity longer than necessary. For example, a member of the IT team could be working on a project that requires testing a process or application where data is stored for a week and then is discarded. Snapshots for this volume can be setup to make daily copies that expire in two days. At the end of testing, the volume and its data can be removed from the CiS system, and two days later, all its data in the cloud will expire and be removed from cloud storage.

Cloud snapshot policies for long-term digital archives

Every cloud snapshot has a data-retention policy that determines when it will expire. Each volume in a CiS system can be covered by multiple cloud snapshots with different policies and retention periods.

To illustrate, a hypothetical mix of data-retention policies can be created to keep weekday cloud snapshots for 15 days, weekend cloud snapshots for 8 weeks, beginning-of-the-month cloud snapshots for 36 months, and beginning-of-the-quarter cloud snapshots for 10 years. Although it would be highly unlikely for this sequence of snapshots to be fully realized

(10 years is a long time to keep doing anything the same way), if it did there would be a total of 99 cloud snapshots for this volume (15 + 8 + 36 + 40). The IT team would be able to restore the volume as it existed for each of the last 15 days, the last 8 weekends, the first day of the month for the last 36 months, or the first day of the quarter for the last 40 business quarters.

Snapshot granularity and aligning volumes with snapshot policies

All snapshots in the Microsoft HCS solution are done at a volume level. There is no way to separate the data from different applications into different snapshots if they are in the same volume. They will all be combined together in one snapshot.

The IT team can work with business managers in their organizations to determine what the data-retention policies should be for the data they are creating and using. It may be advantageous to align application data into different volumes based on their retention requirements. For example, a department that needs to retain a small amount of data for a long time should probably be kept in a different volume from a department that needs to store a lot of data but retain it only for DR purposes. Otherwise, a lot of data will be kept for a long time that does not require that level of protection.

Managing policies for migrated volumes

The section “Migrating data with server virtualization technology” in Chapter 4, “Taming the capacity monster,” discussed storage migrations in VM environments. It is important for IT teams to remember that data protection policies will likely change when the VMDK or VHD is moved to another storage system. IT teams are advised to align migrated VMs with policies and practices that maintain continuity with previous archives.

Using the Microsoft HCS solution as secondary storage for enterprise data archiving

As discussed previously in the section “Integrated archiving with the Microsoft HCS solution,” special volumes can be established with extended data retention policies for storing archived data. This is useful for IT teams that want to use the Microsoft HCS solution for offloading archive data from an enterprise storage system that is running low on capacity. This tends to work best when the archive software is able to migrate data from their existing on-premises storage system to the Microsoft HCS solution, but other techniques such as migrating virtual machine disks can be used. Unlike most volumes that have cloud snapshots configured for both short-term and long term data retention, archive volumes could be configured with only long-term data retention policies.

When the Microsoft HCS solution is being used this way for storing migrated data from another enterprise storage system, it is referred to as *secondary storage*. The transparent

and incremental capacity expansion characteristics also apply to archive volumes and enable IT teams to archive large amounts of data without worrying about overfilling their primary storage capacity, all while automatically protecting it on Windows Azure Storage.

Figure 5-1 shows the Microsoft HCS solution being used as secondary storage for an enterprise archiving system. The CiS system component of the solution is connected to a file server that provides LAN access to the archive server that is migrating the data.

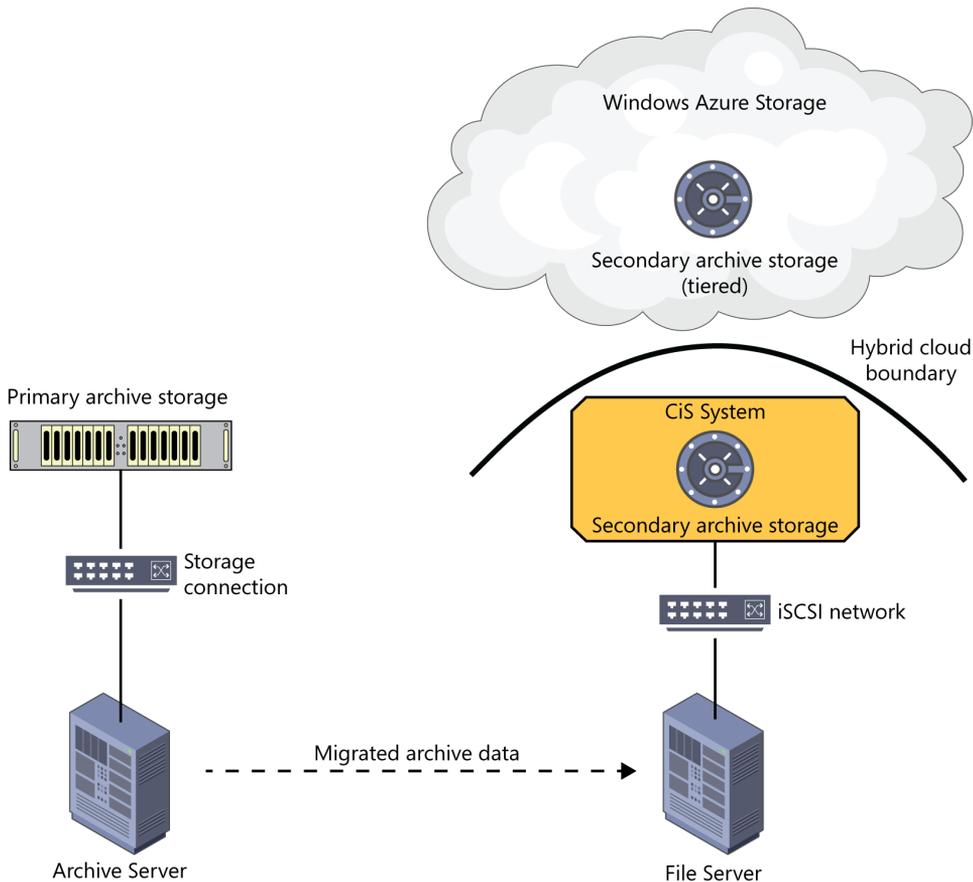


FIGURE 5-1 Archive data is migrated from primary archive storage connected to an archive server to a file server connected to Microsoft HCS.

Meeting regulatory requirements for privacy, data integrity, and availability

The Microsoft HCS solution protects the privacy of individuals by encrypting data in-flight between the CiS system and Windows Azure Storage and at rest in the cloud. In other words, all data uploaded to the cloud is encrypted from the time it leaves the CiS system.

Data integrity is enforced through the use of hashing algorithms that are calculated for every fingerprint in the system. Every time a fingerprint is read, its hash value is compared to ensure its data integrity.

Replication services in Windows Azure Storage provide availability by making three copies of uploaded fingerprints and storing them in different fault domains. The geo-replication service of Windows Azure Storage provides geographically remote protection to another Azure data center where another three copies of the data are kept.

Archiving data from ROBO locations

ROBO locations pose problems for the IT team due to the lack of IT skills in those locations. However, that does not mean they are exempted from complying with regulations. The Microsoft HCS solution can be deployed in ROBO locations to automatically create long-term archival copies of data generated and captured there. The IT team can manage ROBO archive data remotely, ensuring the privacy of data archived in Windows Azure Storage through the strong encryption provided by the Microsoft HCS solution.

Summary

Digital archiving is a growing concern for most IT teams due to the necessity and challenges of complying with regulations and policies. IT teams are looking for solutions that help them comply with regulations and reduce the cost of archiving by reducing the cost of archival storage and by automating the process.

The Microsoft HCS solution features Windows Azure Storage and is capable of storing data for many years in compliance with government regulations. Windows Azure has achieved several important compliance milestones including the HIPAA Business Associate Agreement (BAA), ISO/IEC 27001:2005 certification, and SSAE 16 / ISAE 3402 attestation.

Cloud snapshot policies determine the length of time that data is retained in the cloud and are easily customized to fit a wide variety of data retention requirements. IT teams can choose to archive data in place or in special archive volumes, or choose to use the Microsoft HCS solution as secondary storage for offloading existing enterprise storage systems. Encryption technology protects privacy, hashing provides data integrity checks, and Windows Azure Storage replication services provides availability and DR protection.

The next chapter, “Putting all the pieces together,” looks at the broader system capabilities of the Microsoft HCS solution and discusses a number of use cases where IT teams can successfully deploy it.

Putting all the pieces together

Chapters 2-5 discussed how the Microsoft hybrid cloud storage (HCS) solution protects data, provides disaster recovery (DR) protection, manages storage capacity, and archives data. This chapter explains how it does all of this by leveraging fingerprints as data objects that are accessed directly on Windows Azure Storage and on-premises.

It's important to have realistic performance expectations for enterprise storage and the Microsoft HCS solution is no exception. Its scale-across architecture delivers excellent performance while accommodating the enormously high latencies that cloud data transfers have. But, like all other storage solutions, it has strengths and weaknesses that need to be understood. The performance discussion in this chapter will give readers a way to understand the types of workloads that work best for the solution. To close the chapter, a number of use cases are discussed that show the versatility of the Microsoft HCS solution.

The complete picture of hybrid cloud storage

The comprehensive storage functionality of the Microsoft HCS solution is enabled by a data structure that follows the hybrid cloud management model discussed in Chapter 1. The Cloud-integrated Storage (CiS) system on-premises is much more than an edge device or gateway that transfers data between two dissimilar environments; it divides common storage tasks between on-premises and cloud storage. The CiS system and Windows Azure Storage assume complementary roles as they exchange data and management information across the hybrid cloud boundary. The CiS system manages the elements and operations that are pertinent to on-premises storage and Windows Azure Storage provides the application programming interfaces (APIs) and services for storing and protecting data off-site in the cloud. To get a better understanding of how this works, we return to an examination of fingerprints.

The system of fingerprints and pointers

When data is initially written to the CiS system it is stored as block data and placed in an input queue, waiting to be deduped. This input queue is also referred to as the linear tier. Data in the linear tier is a very small percentage of all the data in the system, and like fingerprints, it is protected by cloud snapshots.

When data exits the linear tier, it is run through the dedupe process and if there is a match to an existing fingerprint, the data is assigned to that fingerprint's pointer. If the dedupe process does not find a matching fingerprint, a new fingerprint is generated, and a new pointer is created mapping the incoming data to the new fingerprint. It follows that the block storage address spaces that are exported externally to servers are mapped internally by pointers to fingerprints. When a server accesses data stored by the Microsoft HCS solution, the CiS system looks up the fingerprints needed to service the request using these pointers.

Fingerprints can be stored on any tier, determined by how recently they were accessed. Every tier is managed as a queue where least recently accessed fingerprints are moved to lower tiers as part of capacity management.

The pointers that map block storage to fingerprints are sometimes called the *metadata*. Metadata is also tiered by the system based on least recently used information but the tiering of fingerprints and metadata are independent of each other. In general, metadata is located on a tier with less latency than where the fingerprints are placed in order to minimize the performance impact of looking up fingerprints.

Fingerprints provide a data abstraction across all tiers, including the cloud tier in Windows Azure Storage. They are identified uniquely, by name, regardless of which tier they are in. Fingerprints stored in Windows Azure Storage are not encapsulated in other data formats, such as tape formats, and are managed as part of a single hybrid data management system. This allows fingerprints to be immediately accessed by all storage functions, without having to convert and load them into primary storage before using them. Fingerprints stored in the cloud tier are encrypted, however, so they must be unencrypted before they can be used again on-premises.

Because the Microsoft HCS solution uses the same names for fingerprints on-premises as is the cloud, tiering is done by updating metadata to use a fingerprint's cloud copy. Tiering this way is much faster and more efficient than uploading data and helps explain how capacity can be reclaimed so quickly when facing a sudden influx of new data.

Recovery operations also leverage the fingerprint naming scheme. Using the pointers downloaded with the metadata map, a recovery CiS system can access and download the fingerprints needed to fulfill IO requests from servers. As the fingerprints are downloaded to the CiS system they are placed in the SSD tier and their pointers are updated to reflect their new location. The name of the fingerprint does not change.

Following the fingerprints

Managing fingerprints across the hybrid cloud boundary transforms standard storage functions by eliminating certain constraints that have long been a part of storage operations. Unlike the capacity limitations of thin provisioning in other storage systems, thin provisioning in the Microsoft HCS solution is continuously expandable through cloud-as-a-tier functionality. Unlike snapshots in other storage systems that consume on-premises storage capacity, cloud snapshots with the Microsoft HCS solution are stored in Windows Azure Storage where they are also automatically protected from disasters. Unlike backups that require additional types of storage and different management practices, automated cloud snapshots with the Microsoft HCS solution are just another part of the hybrid data management system. Unlike remote replication in other storage systems that requires duplicate equipment, software, management and facilities, storing fingerprints in Windows Azure Storage requires no such additional capital investment, management or facilities overhead. Unlike tape-based DR that is highly error prone and nearly impossible to test and prepare for, DR with the Microsoft HCS solution is deterministic and can be tested relatively quickly with far fewer interruptions.

While fingerprints can be stored across all tiers, different storage functions direct fingerprints to be stored either on-premises or in Windows Azure Storage. Figure 6-1 shows the storage locations for primary storage, cloud snapshots, local snapshots, secondary archive storage, and tiered primary and secondary storage.

There is a data flow implied in Figure 6-1, starting at the bottom when data is written to the CiS system by the servers using it. Three storage functions store and access data in the CiS system: primary storage—where data is stored for everyday operations; local snapshot storage—for local point-in-time copies of file versions; and secondary archive storage—for storing long-term digital archives that were generated independently with archiving software and migrated to a CiS system.

There are three storage functions shown in the Windows Azure Storage cloud: primary storage with data that was tiered to the cloud, cloud snapshot storage for off-site backup and DR protection, and tiered secondary archive storage.

The remote Windows Azure Storage data center cloud in the upper right corner of Figure 6-1 is storing fingerprint copies that were transferred to it by the Windows Azure geo-replication service. To clarify, Windows Azure Storage is subdivided into multiple fault domains and data stored in Windows Azure Storage is already replicated three times within each data center. Geo-replication provides protection in case a site disaster strikes a Windows Azure data center.

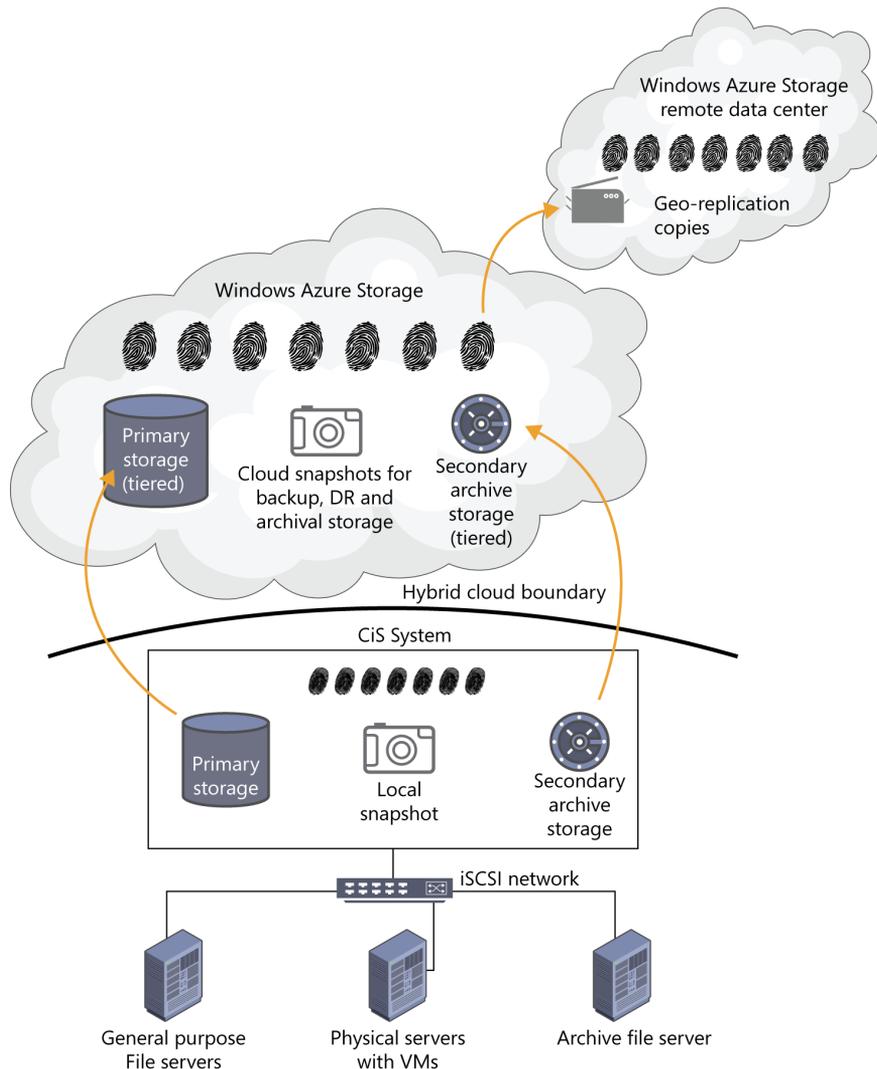


FIGURE 6-1 Servers access the Microsoft HCS solution through an iSCSI network. Primary storage, secondary archives storage and local snapshots are located on the CiS system. Cloud snapshots DR and tiered primary and secondary archive storage are located in Windows Azure Storage. The data on Windows Azure Storage is also stored in a remote Windows Azure data center through the use of geo-replication services in Windows Azure.

Reprising the metadata map

Chapter 3, “Accelerating and broadening disaster recovery protection,” discussed the metadata map and its role in disaster recovery operations. It is an excellent example of management information in the Microsoft HCS solution that crosses the hybrid cloud boundary as part of disaster recovery preparation and operations.

Understanding hybrid cloud storage performance

Implementing any storage solution requires a solid understanding of its performance capabilities, strengths, and limitations. The Microsoft HCS solution is no different and its scale-across design results in some unique performance characteristics.

Even though there are high-performance components, such as NV-RAM and flash SSDs in the CiS system, the Microsoft HCS solution was designed primarily to be a solution for data growth problems. The cloud-as-a-tier storage function emphasizes this design goal by putting dormant data in the cloud where access times are many orders of magnitude slower than in the flash SSD tier. In addition to storing data, the CiS system uses NV-RAM and SSDs for its dedupe processes, leveraging performance components to create a capacity benefit.

The old capacity versus performance trade-off

Performance and capacity utilization typically have an inverse relationship in storage products. In order to get the best performance out of many disk-based storage systems, capacity utilization is sacrificed by using techniques such as short stroking disk drives, which intentionally limits the range of motion of disk arms in the disk drives, thereby improving performance and reducing available capacity.

Another technique is wide striping, which uses many disk drives in order to increase the available input/output per second (IOPS) in the system. The more disk drives there are in a system, the more IOPS can be generated to support transaction processing applications. Performance is achieved by adding disk drives—and capacity—with far less interest in optimizing capacity than in boosting performance. Capacity utilization can be increased in wide-striping disk systems by filling them up with data from more applications, but at the risk of creating contention for disk resources, which adversely impacts performance.

Storage systems that use flash SSDs are turning the tables on the old “performance first” designs of disk-based storage. The IOPS generating capabilities of SSDs is great enough that their performance can be used to increase capacity through techniques like dedupe—which are too IOPS-intensive to be practical on disk-based storage systems. Increasing the effective capacity through dedupe allows more data to be stored in flash SSDs, which improves the overall performance of the system.

The best performance with the Microsoft HCS solution is achieved when the least amount of data needs to be downloaded from cloud storage. In other words, the working set of the data easily fits within the capacity resources of the CiS system.

Establishing the working set

In normal (non-DR) scenarios, the CiS system selects the working set primarily on how recently fingerprints have been used. Fingerprints that are accessed more frequently than others remain in the working set and those that are accessed less frequently are tiered to cloud storage. Other variables that impact working set selection are the availability of free storage resources in the CiS system and weighting factors for different data types. As fingerprints are tiered to the cloud, the amount of available CiS storage capacity on-premises increases. Figure 6-2 illustrates how tiering fingerprints is determined by the relationship between their access frequency and the available storage in the CiS system.

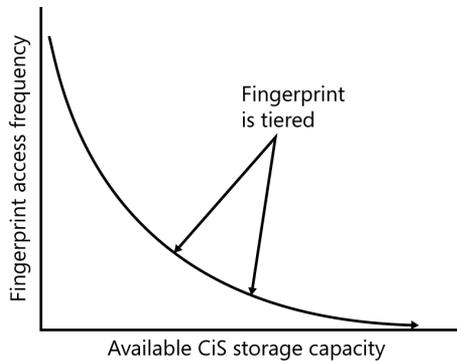


FIGURE 6-2 Fingerprints are tiered to cloud storage based on their access frequency and the available capacity in the CiS system.

The working set in a CiS system is a dynamic entity that changes as new data is written to the CiS system and as applications and users change the data they are actively working with. When fingerprints that were previously tiered to the cloud are accessed and downloaded from the cloud, they are added back to the working set.

Application volatility matters

Performance with the Microsoft HCS solution is not an exercise in perfection but in probabilities. Applications may constantly change the data they access, but many are reasonably predictable. Occasionally the CiS system will need to download data from Windows Azure Storage that it had previously tiered there. That's expected and is accommodated by the design of the system.

The applications that tend to cause problems with the Microsoft HCS solution are those that have high data volatility—where they access a high percentage of large data sets in a relatively short period of time. These applications do not follow the traditional data life cycle that was discussed in Chapter 4, “Taming the capacity monster,” but access data randomly at unpredictable intervals. Examples of applications that have access profiles like this are transaction processing applications and many database systems. Microsoft Exchange is another application that can access data unpredictably.

Avoid defrag!

Disk defragmentation utilities are considered best practices by some IT teams because they help maintain reasonable performance from their disk storage systems. However, their impact on the Microsoft HCS solution has the exact opposite effect by accessing all the data in the volume several times, ruining access histories for the data, and causing any data that had been tiered to the cloud to be downloaded. If the CiS system had been nicely over-provisioned and the working sets working well, defrag can spoil everything.

Just don't do it. Don't defrag. Not only does it lack the ability to improve the performance of the SSD tier in a CiS system, it also contributes to them wearing out faster. Other server utilities that access large amounts of data should likewise be avoided.

Migrating data to a CiS system

The Microsoft HCS solution works on the principal of over-subscription, where more data is stored by the combined hybrid system than fits in the capacity of the CiS system on-premises. The order that data enters the system is important because the working set is determined by access history. When data is first written to a CiS system, such as when copying data as part of a data migration, there is no established working set. If the amount of incoming data is less than the capacity of the CiS system, there is little to think about because the system will form a working set and tier data later.

However, if the amount of incoming data is greater than the capacity of the CiS system, it makes sense to first copy the data that has been least recently accessed on its previous storage location. With no established working set, the most recently written data stays on the CiS system and data that was copied first is tiered to the cloud.

While there may be a temptation to fill up a CiS system all at once with data to see how it works, it's much more satisfying to take a more gradual approach and not over-run the system, forcing it to tier data to the cloud. There are a few practices that can be used to achieve a smoother migration with the hybrid cloud storage solution from Microsoft. The first is to migrate data incrementally, a server or several servers at a time, over the course of several days. The second is to migrate data in an inverse order to application priority—in other words, plan to migrate the most important servers and applications last. The third is to start running applications using the migrated data as soon as the IT team is ready in order to establish access history for their data.

The geek-out way to migrate data into a CiS system

The optimal practice for migrating data to a new CiS system is to copy data based on its access dates in its pre-CiS storage. One way to do this would be to first copy half the data from each volume that has been least recently accessed. Then copy the other half of the data that was most recently accessed. That way the CiS system will have instant access history that roughly resembles the access characteristics of the data before it was migrated. True geeks will think of all kinds of interesting ways to do this and, as usual, their mileage will vary.

Deployment scenarios for the Microsoft HCS solution

The versatility of the Microsoft HCS solution allows IT teams to use it many different ways. The sections that follow describe a number of them.

Refreshing file-server storage technology

A common scenario for implementing the Microsoft HCS solution occurs when an IT team is refreshing file-server storage that is running out of capacity. These older storage systems may be supporting legacy applications that are still needed by the organization, but are not heavily used and have diminishing value. They often contain a lot of dormant data but the IT team doesn't have the time to find it and archive it. Sometimes this storage is comprised of internal server disks and sometimes it is part of an older SAN architecture. The IT team knows they need to improve their ability to manage storage, but they want to reduce their storage costs instead of increasing them.

Instead of buying a traditional storage system that has more capacity than they need, the IT team can deploy a Microsoft HCS solution and migrate the storage volumes from their existing file server storage systems onto the CiS system. Migrations can be accomplished by copying files from one volume to another using server software utilities, or by using SVMotion in VMware environments or Storage Live Migration in Hyper-V environments.

Replacing backup technologies and processes

Chapter 2, "Leapfrogging backup with cloud snapshots," described the problems of backup at length and how the Microsoft HCS solution solves them. To recap, many IT teams are looking for ways to streamline their backup processes and reduce the amount of time administrators spend on backup processing every week. It's not unusual for members of the

IT team to spend a considerable amount of their time on repetitive backup tasks and media management. They are looking for new solutions that permanently change the backup equation from something that is rife with problems to something that is automated and error-free.

Off-site transportation of backup data is also problematic. The recovery point (time in the past when data is recoverable) can be several days if tapes have not been taken off-site. Recovery times (how soon recovery can be finished) can also take longer than desired due to the time it takes to retrieve data from off-site storage. Tapes can be damaged during transport, lost, or stolen, which can create disastrous privacy and security exposures to the organization.

The Microsoft HCS solution circumvents tape backup technologies and processes by protecting data with automated cloud snapshots to Windows Azure Storage. Cloud snapshots allow the IT team to get back almost all the time they used to spend on backup and spend it on other projects instead.

Establishing DR competency

Chapter 3, “Accelerating and broadening disaster recovery protection,” was largely devoted to the problems IT teams have with DR and how the Microsoft HCS solution improves the situation. To summarize briefly, data growth is making an already bad situation worse because the amount of data that needs to be protected and later restored is too large using existing technologies and processes. Backup processing can’t be counted on to complete within the backup window, which means data might not be restorable. Many IT teams are struggling with how to protect a swelling volume of data across their systems and applications.

IT teams recognize that their ability to recover is impaired. In many cases, they cannot test their recovery plans because the disruption to production operations would be too great, or if they do test them, they encounter too many problems and can’t finish. As a result, they can only guess what the recovery time objectives (RTOs) or recovery point objectives (RPOs) might be for their applications, and although they find this situation unacceptable, they don’t know what to do about it.

The Microsoft HCS solution gives the IT team a new, efficient DR tool. Cloud snapshots are much more efficient than other data protection technologies, completing the task in much less time and requiring far less administrative effort. As importantly, successful recovery tests can be conducted with a relatively small amount of hardware and minimal interruptions to production operations. The CiS system used for the test requires an adequate Internet connection to access fingerprints and from there on, deterministic restores ensure that only the metadata map and working sets are downloaded to establish realistic RPOs and RTOs.

Externalizing BLOBs on large SharePoint installations

Microsoft SharePoint is used in many organizations as a way to share and exchange files. SharePoint files are stored in a Microsoft SQL Server database that becomes larger as more employees share data this way. Eventually, the database may become so large that backup and recovery of SharePoint data can become a problem for the IT team.

The Microsoft HCS solution addresses SharePoint backup and recovery by externalizing binary large object (BLOB) storage to the CiS system. This means the data files in the SharePoint database are relocated to the CiS system, where they are referenced by a link using a SharePoint API. Emptied of large data objects, backup and recovery is much faster.

Another benefit of externalizing BLOB storage with CiS systems is that BLOBs that become dormant will be tiered to Windows Azure Storage. Tiered BLOBs can be accessed again quickly and transparently the next time somebody wants to use them, but they no longer contribute to data growth problems on-premises.

Controlling capacity for enterprise document management

Enterprise document management software helps organizations manage projects by automating workflows and organizing large numbers of related documents on file servers. Document management repositories have a tendency to become quite large and require a great deal of storage capacity. Aging project data is rarely accessed, but the organization may be required to keep it in online to comply with contract terms and government regulations. Off-site copies may also be required, which are made using time-consuming manual processes or, more expensively, with storage replication. The storage costs for document management can be high for a function that is largely historical record keeping.

Document management is a great example of a management function that can benefit from primary storage dedupe. Many of the files are derivatives of other files with slight modifications made for different aspects of the project. With so much commonality in the data, dedupe ratios can be very good. The Microsoft HCS solution dedupes primary storage and provides significant capacity relief for storing documents. The cloud storage-as-a-tier feature of the Microsoft HCS solution allows capacity to be expanded transparently to Windows Azure Storage. The scale-across architecture of the Microsoft HCS solution is an excellent match for the capacity needs of enterprise document management.

Automating compliance coverage

As discussed in Chapter 5, “Archiving data with the hybrid cloud,” all data stored on CiS systems can be configured to have long-term copies made periodically by cloud snapshots for archiving purposes. This establishes a blanket of coverage that the IT team can automate for all data stored in the Microsoft HCS solution.

The caveat for this is that data needs to be on the CiS system when a long-term cloud snapshot is run. Data that is deleted before then would only be held in Windows Azure Storage until the expiration of any cloud snapshots that reference it. The IT team might want

to implement best practices that ensure data is not deleted prior to the next long-term cloud snapshot operation. For example, if the IT team sets up long-term cloud snapshots to run on the last day of the month, they would want to implement policies and take measures to keep data available for archiving purposes for a minimum of 31 days.

Migrating archived data to secondary storage

Archiving software stores data for historical purposes and creates metadata and indices for locating data that might be needed in the future. The archived data is typically stored on primary storage that is protected by backup or remote replication technology. Some IT teams want to continue to use their existing primary storage for recently archived data, but want to migrate older archives to secondary storage that costs less to own and operate.

The Microsoft HCS solution provides additional online capacity for archived data through its cloud-as-a-tier feature. Just as document management files tend to be derivatives of other files, archived data can also contain a high number of derivative files, which means primary dedupe can effectively reduce the capacity consumed. Archived data that is migrated to the Microsoft HCS solution can be protected by cloud snapshots with long-term data retention periods that ensure availability on Windows Azure Storage for many years into the future. Archived data stored on Windows Azure Storage is replicated three times across different fault domains in a Windows Azure data center and the option of using geo-replication exists for additional protection.

Managing VM sprawl storage

Organizations that develop technology tend to make extensive use of server virtualization technology. VMs allow developers to experiment with different ideas and to quickly and inexpensively test ideas. While VMs are virtually free to setup, they can have a very real cost in storage capacity long after they are no longer being used. VMs that are not in use consume no processor resources, but if their storage volumes were not deleted, they continue to consume storage capacity. The IT team that supports the development team might not know if dormant volumes are important or not. Problems can occur when the primary storage they are on nears its capacity limits and the IT team has to make uninformed decisions about which ones to remove.

The Microsoft HCS solution can be used to support this type of development environment, either as primary storage for the VMs or as secondary storage that the IT team migrates VM volumes to. Either way, there will likely be good dedupe ratios due to the commonality between operating environments and applications across VMs. If the IT team ever wanted to move a VM volume back to its original or another storage system, they could use SVMotion or Storage Live Migration to accomplish the task.

ROBO data protection compliance

The Microsoft HCS solution helps IT teams comply with best practices for data protection as well as government regulations. Putting CiS systems in remote offices automates all aspects of data protection and gives the IT team access to the data that was uploaded to Windows Azure Storage by the remote offices. That means data from remote and branch offices (ROBO) is available and can be quickly downloaded and recovered at a centralized location by the IT staff.

The solution also provides a way to securely distribute data to remote sites. For example, the IT team can use a cloud snapshot to upload marketing and sales information to a bucket in cloud storage and then access it by downloading the contents of the bucket (including the metadata map) to a new volume at each remote location.

Cloud storage for infrastructure managers

A lot of IT teams are looking for ways to exchange corporate data similar to how they use consumer cloud storage services, such as SkyDrive from Microsoft, but with more security and control. The Microsoft HCS solution provides a way for IT teams to do this that is completely under their control, secured by encryption, and with data integrity ensured.

Summary

The scale-across architecture of the Microsoft HCS solution is unique in the industry. Designed for the problems of data growth, it does more than simply transfer data between on-premises storage and cloud storage—it keeps data online and instantly accessible using the same names and data format regardless of where it is stored. This means data never has to be copied by additional storage products and processes, such as tape or dedupe backup equipment. It also explains how a fingerprint uploaded to Windows Azure Storage by a cloud snapshot during a backup operation can be used by the cloud-as-a-tier feature months later.

The fact that data spans on-premises and Windows Azure Storage has important implications for performance. The key to performance with the Microsoft HCS solution is maintaining a reasonably stable working set and knowing that least recently accessed fingerprints are tiered first when the system needs to expand capacity. Applications that do not have a high degree of volatility form the most predictable working sets.

There are a number of ways IT teams can use the Microsoft HCS solution to solve their storage problems. In addition to the backup, disaster recovery, capacity growth, and archiving examples detailed in Chapters 2-5, it also can be used very effectively for enterprise document management and for large Microsoft SharePoint environments.

The Microsoft HCS solution is still a relatively new development in enterprise storage. It will be interesting to see what new capabilities it assumes in the years to come and what influence it might have on the rest of the industry. For a hypothetical discussion on that topic, turn to the next and last chapter, “Imagining the possibilities with hybrid cloud storage.”

Imagining the possibilities with hybrid cloud storage

Cloud computing has enormous potential as an infrastructure technology, but as the new kid on the IT block, it has a lot of catching up to do to match the enormous legacy of on-premises technologies and practices. Therein lies the importance of hybrid cloud computing—the integration of on-premises and cloud infrastructures to achieve the best of both worlds. Over the coming decade, hybrid cloud computing will surprise people in the directions it takes and the solutions that sprout from it. This chapter takes a hypothetical look at the future of hybrid cloud storage and what role the Microsoft hybrid cloud storage (HCS) solution may play in it.

NOTE This chapter is a piece of technology fiction that is based on the state of hybrid cloud storage technology today and projecting how it might develop in the years to come. Hybrid cloud storage is certain to have an interesting future that will be shaped by unforeseen events, technical inventions, and business changes. This chapter is bound to get some things wrong but, with some luck, may get a few things right. Read at your own risk.

Thanks to VMs, everything done in data centers today can be done in the cloud tomorrow

Servers, storage, networking, and management applications can all be provided in Infrastructure-as-a-Service (IaaS) offerings from cloud service providers (CSPs). The services that CSPs offer are continually evolving, with much of the progress coming in the form of instantly available virtual environments. For example, Windows Azure enables IT teams to easily and quickly create and manage VMs, storage, and virtual network connections using Windows PowerShell scripts or the Windows Azure browser management portal.

MORE INFO An excellent blog post on setting up Windows Azure environments can be found here: http://sqlblog.com/blogs/buck_woody/archive/2013/04/17/creating-a-windows-azure-virtual-machine-the-right-way.aspx

VMs have become the granular infrastructure building blocks of corporate data centers. With VMs everywhere on premises and VMs everywhere in the cloud, it follows that effective *VM portability* across the hybrid cloud boundary will be an important enabler to hybrid infrastructures. IT teams want to copy successful VM implementations between their data centers and the cloud where they can be run with different goals and circumstances. VM portability provides the flexibility to change how processing is done and is a guard against being locked in by any one CSP. System Center 2012 App Controller is an example of a management tool that automates the process of uploading and installing VMs in Windows Azure, and it is an excellent example of the progress being made to integrate cloud and on-premises data centers.

There will always be differences between the things that on-premises and cloud data centers do best. Certain types of applications and data are likely going to stay on premises, while others can only be justified economically in the cloud. Then there will be everything else that probably could be run either on premises or in the cloud. The final decision on those will be made based on cost, reliability, and security.

Infrastructure virtualization

Abstracting physical systems as VMs started a revolution in computing that continues today with cloud computing. The initial breakthrough technology for VMs was VMware's ESX *hypervisor*, a special operating system that allowed other *guest operating systems* to run on it as discrete, fully-functioning systems. IT teams used ESX to consolidate many server instances onto a single physical server, dramatically reducing the number of physical servers and their associated footprint, power, and cooling overhead. There are several hypervisors in use today, including VMware's ESXi and Hyper-V, which is part of both Microsoft Server 2012 and Windows Azure.

But virtualization technologies have been around for much longer than ESX. The technology was first invented for mainframes, and both virtual networking and virtual storage were well-established when the first VMs from VMware were introduced. Virtualization is one of the most important technologies in the history of computing and will continue to be.

In addition to VMs, virtual switches (v-switches) and virtual storage appliances (VSAs) were also developed to run on hypervisors in server systems. Of these three virtualized infrastructure technologies, VSAs have been the least successful at imitating the functionality of their corresponding hardware systems. This is not so hard to understand considering the performance challenges of running a storage system developed for specialized hardware on a PC-based hypervisor.

However, hypervisors in the cloud are different, simply by virtue of where they run, and are much more likely to attract the interest of storage vendors. The most successful infrastructure transitions tend to be those that require the least amount of change and storage vendors will want to sell VSA versions of their on-premises products to ensure that customers making the transition to the cloud will continue to use their technologies—whether they are products or services.

Orchestrating clouds

IT teams are always looking for ways to manage their infrastructures more efficiently. Installation wizards simplify deployment by defining resources and initiating operations and connections between them. But wizards are usually limited in scope to automating the deployment of a single product or service. The next level of automation that coordinates the deployment of multiple products and services is called *orchestration*. Orchestration automates multiple complex setup and configuration tasks for technologies that work together to form a solution.

Orchestration is an excellent example of a relatively new technology area with enormous potential for managing hybrid cloud computing environments. As orchestration matures, the breadth and depth of the automation it provides will expand. Eventually, orchestration may be able to create complete virtual data centers by defining all the VMs, storage volumes, virtual networks, data management processes, and policies on both sides of the hybrid cloud boundary.

For example, orchestration could be used to create a group of VM servers for several departments, the VLANs they use to access a CiS system on-premises, the storage volumes on the CiS system, the Windows Azure storage bucket to expand the capacity for these file servers, and the cloud snapshot policies for daily data protection and end of month data archiving.

The rise of the pocket data center?

Automated orchestration of virtual resources in hybrid cloud environments will transform IT. Free from the time-consuming drudgery of installation planning and change management, IT teams will turn their focus to managing applications, data, devices, and people.

Just as there is role-based administration of systems today there will be role-based orchestration in the future. Informed and guided by policies and templates, role-based orchestration will enable IT team members to create miniature virtual data centers (VDCs) for specific business missions. For example, role-based orchestration could allow an IT team member to create a virtual data center that provides public information and forums about a new product while simultaneously conducting data analysis from internal and external sources to help generate the sales forecasts used by manufacturing.

Managing data growth in a hybrid cloud

Organizations using hybrid cloud designs will expect to limit the costs of running their own corporate data centers. Considering the cost of storage and the rate of data growth, it follows that most of the data growth needs to be absorbed by cloud storage while maintaining steady storage capacity levels on premises. The Microsoft HCS solution provides an excellent way to limit capacity growth on-premises by deduplicating primary storage and using the cloud as a tier for low-priority data. There is nothing particularly futuristic about that however, because the solution does that already.

Another way to limit the storage footprint on-premises is to migrate applications to the cloud. Just as the Microsoft HCS solution migrates lower-priority data to the cloud, the applications that are migrated from the on-premises data center to the cloud could also have lower priorities, and less sensitivity to the effects of migration.

Data portability in the hybrid cloud

Hybrid cloud designs will accommodate diverse and changing workloads by providing computing resources for unique and temporary projects, as well increasing capacity for expanding applications on premises. In theory, the flexibility of hybrid cloud computing will enable IT teams to move applications and data according to changes in business priorities.

For these things to transpire, it must be possible to transfer both the applications and their data across the hybrid cloud boundary. Data portability is an aspect of hybrid cloud technology that will likely see a lot of development in the years to come. The sections that follow discuss aspects of data portability in hybrid clouds and how the Microsoft HCS solution could provide it.

Migrating applications and copying data

Data and the applications that use it must be located in the same data center for performance reasons. Because the amount of data is usually much larger than the application software, the migration process will be effectively gated by the time it takes to complete the data migration.

Migration time is the amount of time it takes between shutting an application down in one location and starting it in another. Unlike disaster recovery (DR), where some of the most recently changed data might be lost, migrations are expected to completely copy all data. In other words, there is no recovery point where migrations are concerned and there will always be some amount of data to copy for an application migration.

Copying a lot of data takes a long time, even when there is a lot of network bandwidth available. So, if the goal is to minimize migration time, it follows that minimizing data transfer times is paramount. Data protection technologies that copy only recently changed data to a remote site, such as replication, continuous data protection (CDP), and cloud snapshots, may

be useful. There will undoubtedly be several different approaches developed to decrease migration time in the coming years.

The Microsoft HCS solution could be used for application/data migrations someday. Here's a quick overview of how it might work: After pausing the application, the IT team would take a cloud snapshot, which would upload the most recently written data and the most recent metadata map. After the cloud snapshot completes, they would start a CiS VSA in the cloud that would access and read the metadata map and copy the data to where it would be accessed by the application running in the cloud. This is not much of a departure from the way deterministic recoveries are done today with the Microsoft HCS solution.

Can you get there from here?

Cloud data centers can do many of the same things that on-premises data centers do, but they are distinctly different. One way they differ is the method used to access, write, and read data. The most popular cloud storage service is object storage, which provides buckets for storing data. The CiS system in the Microsoft HCS solution accesses data stored in Windows Azure storage buckets using the Windows Azure Storage API.

By comparison, data center storage is typically accessed through remote file systems and block storage interfaces. This means VMs and VSAs in the cloud might need to use a different data access method than they do on premises. This is not an insurmountable challenge, but it's not necessarily trivial either, and its solution has to be built into VMs, VSAs, or cloud services.

You call that an infrastructure?

Cloud data centers have the basic server, storage, and network elements of an infrastructure, but from the perspective of the IT team, they are missing a lot of the elements they work with on a regular basis, including NICs, HBAs, device drivers, backup systems, tapes, cables, and power. Avoiding all these infrastructure bumps in the road is the point of cloud computing, after all. The idea of infrastructure as a *service* should be to simplify the work of providing resources so more attention can be paid to applications, data, devices, and people.

Like anything else in IT, there are big differences between rookie-level and guru-level plays. The best hybrid clouds will be designed by people that understand the nuances and details of virtual appliances and how to make them operate most effectively in various cloud platforms. In the years to come, organizations will likely be looking for cloud administrators who know their stuff. The combination of cloud and virtualization skills will be highly desired by IT teams looking for individuals that can get the cloud job done.

Virtual disks as a porting medium

One solution to data portability is to copy the VM's VMDK or VHD file across the hybrid cloud boundary. The hypervisors managing those VMDK/VHD files could also manage this task or special import/export processes could be used. That said, storage developers have always found ways to add value where data transfers are concerned and will likely find ways to offload this work to participate in hybrid cloud data migration scenarios too.

Emulating on-premises storage methods as a service

A different approach to solving the problem of different data access methods on-premises and in the cloud is to provide cloud storage services that emulate on-premises storage. In other words, a service that allows data to be stored and accessed in the cloud by applications using the same storage methods as servers running on-premises. Amazon's Elastic Block Storage (EBS) and Windows Azure Drive are examples of this type of storage emulation.

One potential application for block storage access in the cloud is storage-based replication. As discussed in Chapter 3, "Accelerating and broadening disaster recovery," storage-based replication copies new storage blocks from the primary site to the secondary site, where servers at the secondary site can access it. Presumably, replication to the cloud would use a vendor's VSA in the cloud to receive updated blocks from their on-premises storage system. Storing the replicated blocks might work best with emulated block storage.

Recovery in the cloud

Considering the difficulties many IT teams have with DR, it's not surprising that one of the most compelling aspects of cloud computing is the potential for conducting DR in the cloud. The ability to have an inexpensive, ready-made recovery site on demand is very appealing. Recovering data in the cloud allows the IT team to work in a green field environment and accommodates a certain amount of mistakes and retries. There is little for the IT team to not to be excited about if DR has been broken and deemed unfixable.

Predicting recovery time objectives (RTOs) and recovery point objectives (RPOs) for hybrid cloud computing will likely become much more important than they are today because they will be factored into service level agreements (SLAs) for cloud services. The good news is that practicing recoveries will be much easier in the cloud due to the instant availability of resources to test with and the ability to isolate those tests so they don't interfere with production operations.

Recovering with a CiS system VSA in the cloud

DR in the cloud with a CiS system would be similar to the migration scenario discussed previously in this chapter, except there would be no opportunity to quiesce and synchronize the most recently updated data. A CiS VSA running in the cloud, would need to be configured with the encryption key from the original CiS system, as well as the URL to access the bucket and the account userID and password.

Once the bucket is accessed, the CiS VSA would copy the most recent metadata map. Notice that this would be a local copy operation in the cloud, as opposed to a remote download, which means the performance would be many times faster. After loading the metadata map in the CiS VSA, a couple different scenarios could unfold. One is that the data would be exported to cloud storage where applications running in cloud-based VMs would access it. Another scenario is that VMs could use an emulated block storage service, if one existed, to access the CiS VSA. Other solutions may also be possible and time will tell if and how the solution evolves. The main difference is that recovery times with a CiS VSA will be much faster than they are today on-premises—and they are already very good.

Recovering with backup software running on a VM in the cloud

Fast data copies within the cloud would apply to backup data too. Assuming the backups were written using a tape format, the contents of the tape images would be read by a VM running a cloud-enabled version of the backup program that would either export them to cloud storage or restore them to block storage, where it would be accessed using a suitable block storage service.

Other solutions are also possible and will likely be developed. Restores would be fairly fast, and there would not be any of the media errors or handling problems that plague tape. However, there would probably be many tape images to read before VMs could access the data. Compared to downloading all those tape images to an on-premises backup server, restoring in the cloud would be a major improvement.

Recovering with server software and services in the cloud

New forms of data protection for hybrid cloud environments are being developed. Examples include the Hyper-V Replica feature in Microsoft Windows Server 2012, which was discussed briefly in Chapter 3, as well as Windows Azure Backup and Windows Azure Hyper-V Recovery Manager services. As solutions and services for hybrid cloud recovery continue to evolve, it is likely that the use of virtualization technology and virtual disk files will be key features.

Recovering with storage system VSAs and replication in the cloud

Replication was discussed in Chapter 3 as one of the most effective ways to protect data against a disaster. Storage-based replication could be done using a VSA in the cloud that functions as the remote, secondary system. The VSA in the cloud would almost certainly have to be from the same vendor, just as on-premises replication solutions require a matched pair. Any hardware-centric replication solution would need to address the differences between on-premises and cloud data access. That said, most of the storage system vendors have solved tough problems before and customers who use their solutions will likely want to continue using them.

A cost consideration for storage-based replication is that the VMs or VSAs in the cloud may need to be running at all times in order to receive the data that is being replicated.

In contrast, other hybrid cloud data protection products could write data to cloud storage services without requiring a VM or VSA.

Big Data and discovery in the cloud

Data analytics, or Big Data, is a method of finding indirect dependencies or relationships between information that would otherwise appear not to be related. It works by combing through large amounts of unstructured data using map-reduce algorithms that were initially pioneered at Google and have been popularized by an open source implementation called Hadoop.

Hadoop, by nature, is a cloud application, involving a large number of systems and data sets that are processed by many parallel operations done in sequence. One of the challenges for Big Data is getting all that unstructured data uploaded into the cloud where it can be formulated for Hadoop processing.

The Microsoft HCS solution regularly uploads unstructured data to the cloud as part of regular cloud snapshots. That means data that could be processed by a Big Data application could be synchronized in the cloud once a day. If there was a VSA version of the CiS system in the cloud this would be possible. As described previously, a CiS VSA could access and read the metadata map and copy the data to where it would be accessed by the Hadoop process.

Discovering new discovery techniques in the cloud

Related to the Hadoop scenario, another possible use for corporate data in the cloud is to perform eDiscovery searches on it. EDiscovery from VMs in the cloud could access copies of the data there with minimal impact to IT operations. The two things needed would be a way to get all this data into the cloud initially and then to make it accessible in the cloud for the search operation to work on.

The core functionality of the Microsoft HCS solution fits this scenario very well: it uploads data to the cloud every day. All that would be needed is a VSA version of the CiS system that would access and load the metadata map and then create a copy of the data to be used by search tools.

Will discovery be hadoopified some day?

Searches of data archives are best done with data that has been indexed by archiving software, however, it might not always be possible to run everything through the archiving system. Map-reduce processes could be applied to comb through unstructured data that has not been indexed—that's more or less what it was originally invented for when Google was first trying to figure out Internet search.

Summary

Hybrid cloud computing is expected to be an important new technology, but with all the various on-premises data center technologies and processes that IT teams are using today, there is going to be a lot of development to integrate on-premises with cloud infrastructures over the coming years. Hybrid cloud storage, as a foundation for this change, will have an important role to play.

Migrating data between on-premises data centers and the cloud will be an important function to enable hybrid cloud applications. There are many possible ways to do this, including the potential to use cloud snapshots uploaded by an on-premises Microsoft HCS solution and accessed by a hypothetical hybrid cloud storage VSA running in Windows Azure that can access and make the data available to VMs running there.

One of the most interesting possibilities for hybrid cloud storage is its potential for DR. The latency and performance problems with downloading huge amounts of data would be alleviated if that data could be accessed locally in the cloud by VSAs running there. This would apply to many different forms of data protection including tape backup and cloud snapshots. A VSA that could load the metadata maps uploaded by the Microsoft HCS solution could also probably use deterministic restores to minimize the amount of data needed to be processed by the recovery process. Hybrid cloud storage could very likely be the technology to finally solve the broken state of affairs for DR in many organizations.

Big Data is another new, important technology that could benefit from the potential data synchronization capabilities of the Microsoft HCS solution. In general, any application that might be too expensive and too disruptive to run on premises might be able to use cloud snapshots to synchronize data to VMs running in Windows Azure.

This chapter tried to look down the road at the future of hybrid clouds and hybrid cloud storage. It's going to be a very interesting road with unexpected twists and turns and a few long, beautiful views every now and then. I hope it has been enjoyable and entertaining.

Considerations and recommendations for networking, privacy, and data protection

This appendix discusses a variety of topics that IT team members should consider when implementing the Microsoft hybrid cloud storage (HCS) solution.

A bifurcated data path from application to cloud

The data path for most storage I/Os starts at the server and traverses a network or bus on its way to storage. The hybrid cloud storage solution from Microsoft has a data path that is bifurcated at the Cloud-integrated Storage (CiS) system. Writes—and most reads—follow a typical data path over an iSCSI SAN to the CiS system, but reads that download fingerprints from cloud storage have the additional network component of an Internet connection between the CiS system and cloud storage. Writes to the cloud are not part of an I/O operation but are typically made as part of the data protection process during cloud snapshots.

There are minimal requirements for operating both networks in the data path, as described below.

iSCSI considerations

Setting up an iSCSI SAN to connect servers to the CiS system is relatively straightforward and simple. Most server operating systems have device drivers to establish iSCSI sessions between them and storage. iSCSI provides robust communications for storage I/O traffic over Ethernet networks. The CiS system has multiple Ethernet ports for high availability. iSCSI SANs should be segregated from LAN traffic using subnets, VLANs, or separate physical networks. In general, the larger the SAN and LAN, the greater the need for segregation.

Internet connection considerations

Unlike an iSCSI network, where it is easy to segregate SAN and LAN traffic with subnets or VLANs, the connection to the Internet is almost always shared with other Internet traffic at the site. The minimal dedicated bandwidth recommendation for the Internet link between the CiS system and Windows Azure Storage is 20 Mb/second. Obviously connections with more bandwidth will provide faster uploads and downloads. The CiS system can have its bandwidth throttled during production hours to keep it from interfering too much with other work. This typically doesn't create problems for the Microsoft HCS solution because most of the time its Internet traffic is generated at night during cloud snapshots.

Privacy and data integrity technologies

The hybrid cloud storage solution from Microsoft protects the privacy of data copied to the cloud by encrypting all fingerprints prior to uploading them with AES-256 encryption. Data is protected by encryption when it is in-flight and at-rest in the cloud. For this reason, it is necessary to input the encryption key in a recovery CiS system that is accessing fingerprints that were uploaded by a different CiS system that was lost in a disaster.

Encryption keys are managed by the IT team and are not stored in cloud storage. The IT team and those responsible for security in the organization may have policies and measures in place for managing encryption technology and keys. This important aspect of operating the hybrid cloud storage solution from Microsoft requires consideration and planning.

Data integrity is ensured for each fingerprint by keeping a SHA-256 hash value for its contents in the fingerprint's metadata. Every time a fingerprint is opened by the CiS system, its hash value is generated and compared with the embedded hash value. If they do not match, the data is known to have been corrupted.

Data protection considerations

There are no default snapshot schedules that automatically start protecting data when the Microsoft HCS solution is powered on. The IT team is responsible for establishing all snapshot schedules, including daily cloud snapshots and cloud snapshots intended to retain data for longer periods of time. The scheduling mechanism in the system is designed to adapt to a wide variety of policies and best practices.

Snapshots are done on the volume level. When this book was written, there was a maximum of 256 snapshots per volume, including both local and cloud snapshots.

Glossary

Active data Data that is expected to be accessed again relatively soon or periodically

Archiving A storage process that preserves data for an extended period of time

At-rest An IT resource that has a stable state and is not being copied

Backup target A storage device or system that backup software writes data to when performing backups

Backup A data protection method that was developed to work with tape and usually combines periodic full copies of data with incremental copies of new data

Best practices IT management informed by advanced knowledge and experience

Big Data Vernacular term for data analytics, associated with, but not restricted to, Hadoop technology and methods

BLOB Binary large object, often a file

Block storage A storage environment characterized by devices and protocols that are designed to consume storage based on the granular element, blocks

Bucket A storage container provided by a cloud storage service

CDP (continuous data protection) A method of data protection that makes copies of all changes made to data

CiS (Cloud-integrated Storage) An on-premises storage system that stores data for on-premises systems

and incorporates cloud storage services as a resource for storing on-premises data

Cloud computing Scalable computing services provided on a short- or long-term basis by a large number of systems

Cloud snapshot A data protection method that stores point-in-time copies of data for on-premises systems in cloud storage

Cloud storage-as-a-tier Scale-across storage, CiS that works with cloud storage to provide a single, scalable storage system

Cloud storage Scalable, object-based storage capacity provided as a service on a short- or long-term basis

Cloud A data center providing scalable computing and storage services, characterized by a large number of systems that can be accessed for long-term or short-term projects

Clustered storage Scale-out storage, a tightly coupled group of storage systems that function as a single, scalable storage system

Data analytics Computing processes looking for patterns or correlating factors in large amounts of data

Data reduction Processes that reduce the amount of storage capacity consumed for a given amount of information

Data tiering A storage management process that determines the performance and storage requirements for data and locates it on a cost-effective storage resource

Data volatility An indication of the percentage of an application's data that may be accessed in day-to-day operations

Dedupe (Deduplication) A process that identifies duplicate copies of data and eliminates them by linking to reference copies

Deterministic Precisely specified

Discovery The process for finding data that may be needed for legal or regulatory reasons

Dormant data Data that is accessed very rarely, if ever

Downtime The amount of time systems and data are not available for processing, usually associated with a disaster or failure, but also maintenance operations

DR Disaster recovery, the process of resuming operations following a disaster that causes the unexpected stoppage or failure of computing equipment and processes

Enterprise A business or government entity of substantial size

Fingerprint The granular data structure that is managed in the Microsoft hybrid cloud storage solution comprised of data and metadata

Geo-replication A cloud storage service that copies data from a cloud data center to a remote cloud data center

Hash A numerical value calculated from processing a data string that uniquely identifies the string

High availability A system design designed to continue operating after the loss or failure of system components or entire systems

Hybrid cloud boundary The distance, time, or technology barrier that separates an on-premises data center from a cloud data center

Hybrid cloud storage Data storage formed by the combination of on-premises storage and cloud storage services

Hybrid cloud A computing service that combines public compute services with private compute services

Hypervisor A software program that provides an operating environment for virtual machines

IaaS (infrastructure-as-a-service) A cloud service offering the use of virtual computer, storage, and network systems

Index A way of condensing or collating data electronically that facilitates searching

In-flight An IT resource that is being copied or transmitted from one location to another

IOPS The total sum of read and write operations per second; input/output operations per second

iSCSI (Internet Small Computer System Interface) An Ethernet protocol for exchanging storage commands and data between computer systems and storage systems

IT team Employees and contractors that plan, acquire, manage, and operate IT

IT Information technology, the profession and industry of developing, manufacturing, selling, implementing, and operating data processing and communications products and services

Local snapshot A point-in-time copy of data or pointers to data stored on a storage system's own disk drives

Metadata Data that describes data or attributes of data, such as a hash value of its contents

Migration time The time that an application is offline while it and its data are being relocated from one data center to another

Monolithic storage Scale-up storage, a single system storage design that scales by adding components

Near CDP A method of data protection that makes copies of most changes made to data

NV-RAM (non-volatile random access memory) Fast memory storage that retains data even after the loss of power

On-premises A facility owned and operated by an organization such as a business or government

Orchestration An intelligent installation process that manages multiple related technologies to create a solution

Portability The ability to relocate compute resources and processes from one location or set of resources to another

Primary site A data center where production operations run and where replication copies originate

Primary storage Storage where applications read and write data during normal processing

Private-cloud A computing service that is restricted to a specific set of users, often implemented at a corporate-owned facility

Public-cloud A multi-tenant computing service that is provided openly over the Internet

Recovery point The time in the past when the last data was captured prior to a disaster event

Recovery site A data center where DR operations are conducted

Recovery time The amount of time it takes to return a system to full functionality after a disaster event

Replication A data protection method that copies written data from one location to another

ROBO Remote office / branch office

SAN Storage area network

Scale-across A storage design that scales by adding resources from cloud storage

Scale-out A storage design that scales by adding additional systems to a group of systems

Scale-up A storage design that scales by adding components to a single system

Secondary site A data center where replicated data is copied to

Secondary storage Storage that is used for data protection or archiving

Snapshot A method of data protection that uses a system of pointers to make point-in-time copies of data

Spindown A process of stopping the rotation of disk drives for dormant data to reduce the power costs of storing it

SSD (solid state disk) Amassed memory technology that functions like a disk drive

Tape rotation The schedule that backup software creates for naming, using, and cycling tapes

Thin provisioning A method of allocating storage capacity from a common resource to individual storage volumes on a first-come, first-served basis.

Thin A storage process designed to minimize resource consumption

Virtual machine (VM) The functionality of a physical computer provided by software operating as a logical system image

Virtual storage A storage resource that is comprised of elements of other storage resources

Virtual tape The use of a disk storage system to replace tapes used for backup

Virtualization The process of using software to mimic the functionality of physical equipment

VM sprawl A phenomenon where the number of virtual machines in an organization scales beyond the ability of the IT team to manage

VM Virtual disk A virtual disk managed by a hypervisor for storing the data for a VM, typically stored as a file.

VSA (virtual storage appliance) The equivalent of a VM, but for a storage system

VTL (virtual tape library) A storage system used to backup data to disk

Working set The data normally accessed in regular daily processing

Index

A

- active data, 52, 93
- Amazon, 86
- App Controller, 82
- applications
 - data volatility and, 72–73
 - migration of, 84
 - prioritizing, 34
 - recovery of, 36
- archiving
 - data in place, 62
 - definition of, 93
 - electronic discovery and, 57–58, 88
 - encryption and, 59
 - importance of, 57–58
 - in CIS system, 22
 - management policies for, 59, 62–64
 - to cloud storage, 60–61
 - to disk storage, 60–61
 - to magnetic tapes, 60
 - Windows Azure Storage, 61–66
- at-rest, 65, 92–93
- automation, 4, 21, 51, 83

B

- backup
 - best practices, 7
 - definition of, 93
 - disk-based, 15
 - incremental-only, 16, 20
 - magnetic tape, 12–14
 - problems of, 11–12, 75
 - synthetic full, 14, 16, 20
 - using cloud snapshots, 20

- virtual tape, 15–17, 87
- backup targets, 15, 19, 93
- best practices, 7, 93
- Big Data, 88, 93
- binary large objects (BLOBs), 76, 93
- block data, 19–20, 22, 68
- block storage, 68, 86, 93
- buckets, 31–33, 37–38, 48, 53, 93

C

- capacity
 - management of growth, 47–52, 84
 - performance and, 71
 - reduction of, 5–6
 - requirements of dedupe VTL replication, 29
 - storage arrays and, 7
- cloud, 5, 40, 86–88, 93
- cloud computing, 3, 5, 81–82, 93
- cloud service providers (CSPs), 81
- cloud snapshots, 18–21, 88–89, 93
 - data protection and, 20–21, 36, 92
 - retention, 62–65
 - storage of, 69–70
- cloud storage, 1, 60–61, 93
- Cloud-integrated Storage (CiS)
 - data tiering, 9, 53–54
 - definition of, 93
 - fingerprints, 19–22, 68–69
 - Internet connection and, 92
 - migration, 73–74
 - overview, 8–9, 67
 - performance, 71–73
 - recovery, 30–33, 36–38, 86–87
 - retention, 62–65
 - snapshots, 9, 21

- thin provisioning, 9, 45–46
 - working set, 72
- cloud-storage-as-a-tier, 49–53, 68, 76, 93
- clustered storage, 48, 93
- compliance, 58, 61–62, 77–78
- compression, data, 54
- continuous data protection (CDP), 18, 93
- cost considerations, 4, 7, 17, 37, 87

D

- data. *See also* working sets
 - access to, 72–73
 - archived, 62
 - order of incoming, 73–74
 - unstructured, 50, 58, 88
- data analytics, 88, 93
- data availability, 4, 48, 59, 66
- data centers
 - cloud, 8, 82, 85
 - on-premises, 1, 53, 82, 84, 94
 - virtual, 83
- data growth, 3–4, 26, 29, 84
- data integrity, 59, 66, 92
- data life cycles, 49–52
- data protection
 - at ROBO sites, 40
 - cloud snapshots and, 20–21, 36, 92
 - continuous, 18, 93
- Data Protection Manager, 18
- data reduction, 5–6, 53–54, 93
- data tiering, 9, 22, 53–54, 93
- data volatility, 47, 72–73, 94
- dedupe ratios, 54, 76, 78
- deduplication (dedupe), 6, 17, 29, 53–54, 94.
 - See also* primary dedupe; source dedupe
- defragmentation, 73
- deterministic recovery, 27, 34, 89
- deterministic, definition of, 94
- digital archiving. *See* archiving
- disaster recovery (DR)
 - as a best practice, 7
 - definition of, 94
 - in the cloud, 86–88
 - problems of, 12
 - strategies of, 25–30
 - with hybrid cloud storage, 30–39, 68–70, 75–76

- discovery, 57–59, 88, 94. *See also* eDiscovery
- disk storage, 15, 60–61
- disk-to-disk-to-tape (D2D2T), 15–16
- documentation, 61–62
- dormant data, 36, 49–52, 58, 94
- download performance, 34
- downtime, 25, 27, 94

E

- eDiscovery, 58–60, 88
- Elastic Block Storage (EBS), 86
- electronic discovery. *See* eDiscovery
- encryption, 59, 61, 92
- enterprises, 3, 94
- erasure coding, 3

F

- fingerprints
 - block data and, 22
 - data integrity, 92
 - data life cycle of, 50–52
 - definition of, 94
 - expiration of, 63
 - in CiS system, 68
 - overview, 19–20
 - storing in Windows Azure Storage, 31–32, 69
 - working set of, 35–36

G

- geo-replication, 39–40, 66, 69, 94
- Google, 88

H

- Hadoop, 88
- hard disk drives (HDDs), 8–9, 53–54
- hash, 94
- hashing algorithms, 6, 66
- high availability, 48, 94
- HIPAA Business Associate Agreement (BAA), 62

hybrid cloud, 94
 hybrid cloud boundary, 2, 22–23, 82, 84, 94
 hybrid cloud storage
 architecture, 1
 definition of, 94
 disaster recovery (DR) with, 30–39, 68–70, 75–76
 management model, 1
 performance, 71–72
 storage volumes in, 32
 Windows Azure Storage in, 8
 Hyper-V, 4–5, 74, 82
 Hyper-V Recovery Manager, 2, 87
 Hyper-V Replica, 87
 hypervisors, 4, 26, 82, 86, 94

I

IaaS (Infrastructure-as-a-Service), 81, 94
 incremental-only backup, 16, 20
 index, 60, 94
 in-flight resources, 65, 94
 Internet connection and CiS system, 34–35, 92
 IOPS (input/output per second), 6, 49–50, 53, 71, 94
 iSCSI (Internet Small Computer System Interface),
 8, 70, 91, 94
 ISO/IEC 27001 2005 certification, 62
 IT, 94
 IT managers, 1, 7
 IT team, 94

J

Joyner, John, 2

L

life cycles, 50–52
 linear tiers, 68
 local replication, 39
 local snapshots, 21, 69–70, 94

M

magnetic tapes, 12–14, 60
 metadata, 68, 77, 94
 metadata maps, 31–34, 37, 70, 87
 Microsoft HCS
 benefits of incremental storage, 49
 cost advantages of, 37
 data growth, 26
 defragmentation with, 73
 deployment scenarios, 74–78
 differences from other storage systems, 69
 Microsoft Sharepoint, 76
 Microsoft System Center 2012, 82
 migration, 43–44, 64, 73–74, 77, 84
 migration time, 84–85, 94
 monolithic storage, 47, 94

N

near-CDP solutions, 18, 94
 nodes, storage, 48
 NV-RAM (non-volatile random access memory), 53, 94

O

object storage, 85
 on-premises data centers, 1, 53, 84, 94
 orchestration, 83, 94
 overprovisioning, 47
 over-subscription, 73

P

performance
 capacity and, 71
 cloud-storage-as-virtual-tape, 35
 download, 34
 hybrid cloud storage, 71–72
 primary storage and, 6
 solid state disks (SSDs), 53

pointers, 17, 22, 40, 51, 68
 portability, 5, 82, 84, 94
 primary dedupe, 17, 22, 29, 54, 77
 primary site, 16, 28–29, 95
 primary storage, 21, 54, 95

- archived data in, 62, 77
- capacity management, 17–18, 64
- data protection in, 19, 21
- dedupe ratios, 54, 78
- deduping, 53–54
- performance and, 6
- storage location in CiS system, 69–70

 privacy, protection of, 59, 65, 92
 private cloud, 95
 public cloud, 59, 95

R

recovery. *See also* disaster recovery (DR)

- deterministic, 27, 34, 89
- location-independent, 39
- metrics, 27
- opportunistic, 34–35

 recovery CiS system, 32–33, 36–38
 recovery point objectives (RPOs), 27–30, 75–76, 86
 recovery points, 27, 36, 75, 84, 95
 recovery site, 29–30, 32, 86, 95
 recovery time, 27, 33–35, 75, 87, 95
 recovery time objectives (RTOs), 27–30, 35, 75–76, 86
 redundancy, 39–40
 remote and branch offices (ROBOs), 21, 40, 66, 78, 95
 replication

- dedupe VTL, 29
- definition of, 95
- local, 39
- remote, 21, 28–29
- server-software, 28
- storage-based, 28, 86–87

 retention, 21, 36, 62–65
 ROBOs (remote and branch offices), 21, 40, 66, 78, 95

S

SAN, 9, 91, 95
 scale-across storage design, 48–49, 55, 67, 76, 78, 95

scale-out storage design, 47–49, 55, 95
 scale-up storage design, 47–49, 55, 95
 scheduling, 20–21, 92
 secondary archive storage, 69–70, 77
 secondary site, 25, 28, 34, 86, 95
 secondary storage, 64–66, 69–70, 77, 95
 server virtualization technology, 43, 77
 service level agreements, 86
 short-stroking technique, 71
 snapshots, 9, 17–18, 92, 95.

- See also* cloud snapshots; local snapshots

 solid state disks (SSDs), 6, 9, 53–54, 71, 95
 source CiS system, 32–33, 37–38
 source dedupe, 17
 spindown, 60, 95
 SSAE 16 / ISAE 3402 attestation, 62
 SSDs (solid state disks), 6, 9, 53–54, 71, 95
 storage. *See also* cloud snapshots;

- hybrid cloud storage; primary storage
 - cloud, 1, 3, 60–61, 93
 - clustered, 48, 93
 - distributed, 48
 - monolithic, 47, 94
 - need for flexible, 43
 - secondary, 64–66, 69–70, 77
 - secondary archive, 69–70, 77
- storage arrays and capacity, 7
- storage design
 - scale-across, 48–49, 55, 67, 76, 78, 95
 - scale-out, 47–49, 55, 95
 - scale-up, 47–49, 55, 95
- Storage Live Migration, 45, 74, 78
- storage migration, 43–44, 64, 77
- storage tiering, 49–50
- Storage VMotion, 45, 54, 74, 78
- storage volumes
 - archived data in, 62–65
 - dedupe and, 53
 - in hybrid cloud storage, 32
 - snapshots and, 64
 - thin provisioning, 45–47
- SVMotion, 45, 54, 74, 78
- synthetic full backup, 14, 16, 20
- System Center 2012, 82
- system virtualization, 26–27

T

tape rotation, 13–15, 95

thin, 34, 95

thin provisioning, 9, 45–47, 69, 95

transparency, 51

Trust Center, 62

V

virtual disks. *See* VM virtual disks

virtual hard disks (VHDs), 4, 43–44, 64, 86

virtual machine disks (VMDKs), 4, 43–44, 64, 86

virtual machines. *See* VMs (virtual machines)

virtual storage, 82, 95

virtual storage appliance (VSA), 82, 85–88, 95

virtual switches (v-switches), 82

virtual tape, 15, 35, 95

virtual tape libraries (VTLs), 15–17, 29, 34, 95

virtualization, 4, 26–27, 43, 82, 95

VM sprawl, 45, 77–78, 95

VM virtual disks, 86–87, 95

VMs (virtual machines), 4, 43–44, 81, 87–88, 95

VMware, 4, 74, 82

volumes. *See* storage volumes

W

wide striping technique, 71

Windows Azure, 81–82

Windows Azure Storage

block data in, 19–20, 22

buckets, 31–32, 37

cloud snapshots in, 20–22

compliance in, 62

data archiving using, 61–66

deduplication in, 53

disaster recovery in, 39

dormant data storage in, 52, 76

fingerprint storage in, 31–32, 69

hybrid cloud storage using, 8

recovery at ROBO sites, 40

redundancy in, 39–40

Windows Live Migration, 54

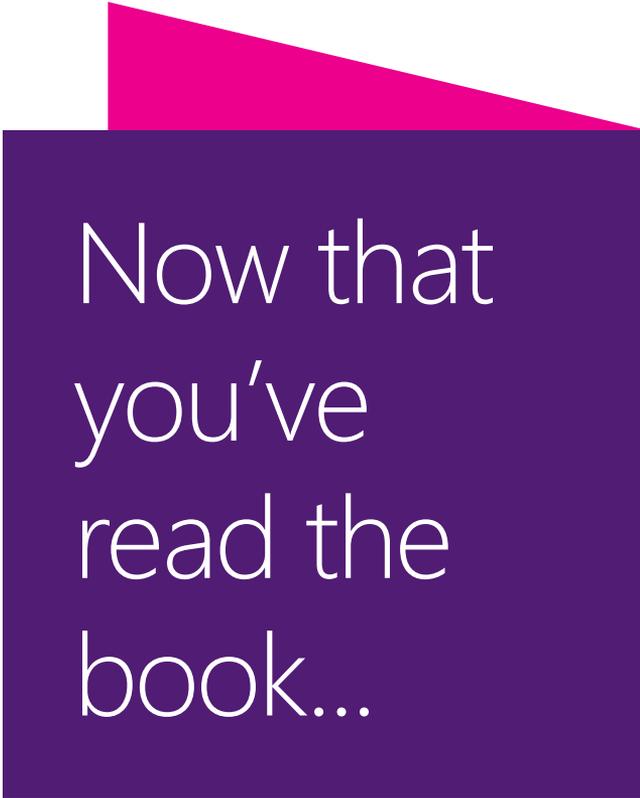
working sets, 35–36, 52–53, 72, 95

About the author



MARC FARLEY is a senior product marketing manager at Microsoft working on hybrid cloud storage solutions. *Rethinking Enterprise Storage: A Hybrid Cloud Model* is his fourth book on network storage; his previous books are the two editions of *Building Storage Networks* (McGraw-Hill, 2001 and 2002) and *Storage Networking Fundamentals* (Cisco Press, 2004). In addition to writing books about storage, Marc has blogged about storage while working for EqualLogic, Dell, 3PAR, HP, StorSimple, and now Microsoft.

When he is not working, Marc likes to ride bicycles, listen to music, dote on his cats, and spend time with his family.



Now that
you've
read the
book...

Tell us what you think!

Was it useful?

Did it teach you what you wanted to learn?

Was there room for improvement?

Let us know at <http://aka.ms/tellpress>

Your feedback goes directly to the staff at Microsoft Press,
and we read every one of your responses. Thanks in advance!

