Foreword by Martin Glassborow, aka Storagebod, storage industry expert



# Rethinking Enterprise Storage A Hybrid Cloud Model

Marc Farley

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

	Foreword	ix
	Introduction	xi
	Next steps	XV
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
	Index	97

# Contents

	Foreword	ix
	Introduction	xi
	Next steps	xv
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? Doing things the same old way doesn't solve	7
	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	
	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	
	Introducing the metadata map	31
	Recovery times with the hybrid cloud storage solution	33
	Windows Azure Storage as a recovery service	
	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	
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

Appendix	91
Glossary	93
Index	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:

#### 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

#### CHAPTER 3

# Accelerating and broadening disaster recovery protection

T 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

S tatistics 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

et'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**

ong 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.

# 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

# С

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

# Η

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

# Μ

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

# Ν

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

# 0

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

# Ρ

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

# Т

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 ape, 15, 35, 95 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