

COHESITY PRESENTS

THE GORILLA GUIDE TO...[®]



Hyperconverged Secondary Storage

James Green, ActualTech Media

INSIDE THE GUIDE:

- How the proliferation of data infrastructure silos is hurting your business
- Consolidate your secondary data and apps at web scale
- Gain greater efficiency and control through hyperconvergence

**HELPING YOU NAVIGATE
THE TECHNOLOGY JUNGLE!**



ActualTech Media
www.actualtechmedia.com

In Partnership With

COHESITY

THE GORILLA GUIDE TO...

Hyperconverged Secondary Storage

AUTHOR

James Green, ActualTech Media

EDITOR

Keith Ward, ActualTech Media

LAYOUT AND DESIGN

Olivia Thomson, ActualTech Media

Copyright © 2018 by ActualTech Media

All rights reserved. This book or any portion thereof may not be reproduced or used in any manner whatsoever without the express written permission of the publisher except for the use of brief quotations in a book review.

Printed in the United States of America.

ACTUALTECH MEDIA

Okatie Village Ste 103-157

Bluffton, SC 29909

www.actualtechmedia.com

ENTERING THE JUNGLE

Chapter 1	6
Too Many Points of Management	6
What Is ‘Mission-Critical’?	8
Primary Storage	9
Secondary Storage	10
Silos. Lots of Silos.	11
Dark Data	11
Cloud Computing	13
Modern Technology Problems	14
Chapter 2	16
Hyperconverged Secondary Storage	16
What is Hyperconvergence	16
Hyperconverged Secondary Storage	19
Data Protection	20
File Services	22
Test and Development	23
Analytics	25
Chapter 3	27
The Role of Public Cloud	27
Own vs. Rent	27
Cloud Storage and Data Management Challenges	29
Requirements for a Cloud-First HCSS Platform	30

Chapter 4	33
The Cohesity Platform and SpanFS	33
The SpanFS Distributed File System	33
Replication and Cloud	35
Global Deduplication	36
Security and Quality of Service	37
WAN Optimized	37
Multi-Lingual Storage	38
Strict Consistency	39
Test and Development	40
Native Tape and Cloud Archival	41
Analytics Workbench	41
Chapter 5	43
Secondary Storage Workloads	43
Backup and Archive	43
Data Protection in Depth	45
Files and Objects	47
Test and Development	48
Analytics	49
Chapter 6	51
Real Savings, Real Efficiency	51


CALLOUTS USED IN THIS BOOK



The Gorilla is the professorial sort that enjoys helping people learn. In the School House callout, you'll gain insight into topics that may be outside the main subject but that are still important.



This is a special place where you can learn a bit more about ancillary topics presented in the book.



When we have a great thought, we express them through a series of grunts in the Bright Idea section.



Takes you into the deep, dark depths of a particular topic.



Discusses items of strategic interest to business leaders.

ICONS USED IN THIS BOOK



DEFINITION

Defines a word, phrase, or concept.



KNOWLEDGE CHECK

Tests your knowledge of what you've read.



PAY ATTENTION

We want to make sure you see this!



GPS

We'll help you navigate your knowledge to the right place.



WATCH OUT!

Make sure you read this so you don't make a critical error!

CHAPTER 1

Too Many Points of Management

If we look back through the years, there have been a handful of milestones that mark pivotal shifts in the way we do things in IT. The introduction of Active Directory in Windows 2000, VMware taking hold of the server infrastructure in 2008, and the evolution of flash storage over the last 10 years are prime examples of this. Each IT milestone changed how the next generation of IT solutions operated.

These sea changes have contributed to the spawning of entire ecosystems of companies and products. Over the past two decades, a growing sprawl of purpose-built products and services has resulted in fragmentation and silos. This has in turn added to the management burden of IT teams, along with increasing the number of solutions end users must interact with. This has led to increased complexity and cost for organizations of all sizes.

A comparable consumer technology analogy would be audio-visual equipment. Over the years, we've used many or most of these consumer technologies that were cutting-edge at the time: vinyl records, 8-tracks, Betamax, VHS, cassettes, CDs, Laserdiscs, Minidiscs, DVDs, MP3s, Blu-rays, and more. Many of us own the same movie (or album) on multiple formats.

The evolution of IT is similar, but the timeframe is compressed. Whereas the evolution of audio-visual technologies above took three or four generations, IT can easily go through as many protocols, formats or standards in just two or three decades.



In IT, the word “silo” means an infrastructure component or application which is isolated, and must be managed separately. In storage, specifically, this is frequently used to refer to storage separated by protocol type (file, block, object, cloud), as well as storage that’s segmented by the fragmented management of the devices providing the service.

One example of storage silos is when an organization has multiple storage units serving file-based network-attached storage, but each unit must be managed separately, and accessed separately by end users. This contrasts with a storage fabric, in which multiple devices may exist, but are managed as a single solution, and present a single point of access to end users.

The rapid pace of IT advancement complicates what would otherwise be a fairly normal, and readily understood, technological progression. As a rule, organizations don’t stop using one technology just because it’s been supplanted by another.

Multiple generations of technologies persist, each enabling applications and services that the organization is not ready to, or capable of, abandoning. It’s not uncommon, for example, to find an organization in 2018 still utterly reliant on a DOS-based application first written in the 1980s, and updated a few times in the 1990s. Twenty- and even 30-year-old applications can be found in organizations of all sizes, despite the fact no modern IT solutions rely on the same underlying Infrastructure.

What Is 'Mission-Critical'?

The majority of IT budgets are spent on roughly the top 10% of mission-critical IT, while the remainder gets the scraps; or it simply isn't maintained at all. This can create neglect cascades as IT teams deprioritize software updates, licenses, or replacement hardware for less prominent workloads. This neglect often results in lower priority workloads experiencing system performance and stability lags.

The generally-accepted definition of mission-critical IT is "IT solutions whose failure would represent a serious impact to an organization." *Serious impact* is typically defined in the following ways:

- It will cost the organization a lot of money
- It will lose the organization a lot of customers
- It will cause the organization to cease existing
- It will result in loss of life/loss of mission hardware (such as a spacecraft).

Each organization, however, defines "mission-critical" differently. In many organizations, workloads which have negligible business impact, but which have significant internal political impact, are considered mission-critical. In other organizations, mission-critical is used to refer to workloads with the highest performance demands, not necessarily applications with the highest organizational impact.


Significantly dated applications, such as the previously discussed 30-year-old DOS-based application, are often not given a mission-critical designation, despite being workloads whose interruption could significantly impact the organization. The seemingly arbitrary nature of criticality designations has real-world consequences when one takes into account the difference in IT spend between infrastructure deemed mission critical, and everything else.

To further understand one of the most pressing problems, let’s look at a few terms you may or may not be familiar with.

Primary Storage

Primary storage can be loosely defined as the storage hosting your mission-critical infrastructure, and systems required for your company to do business. Primary storage likely consumes a vast majority of your IT budget annually, gets the lion’s share of care and feeding efforts, and requires the most specialization in your staff.

Here one would expect to find all-flash arrays, high end SANs, and expensive, feature-rich filers. It’s not uncommon to see storage solutions dedicated to a single IT service, or to find that different service owners have engaged different vendors.



The **high cost of downtime** is widely considered a justification for the high price tags. As a result, there are numerous vendors targeting this space, increasing the pressure to fragment the installed primary storage base, ultimately resulting in storage silos.

Secondary Storage

Secondary storage underpins workloads *not* designated as mission critical. It also provides supporting, but non-mission-critical, storage for mission-critical workloads. Supporting, but non-mission-critical storage examples vary greatly. Storage for backups, logging, or warm archives are all examples.

Consider, for example, Facebook. Facebook is arguably the largest photo album on the planet, housing hundreds of billions of photos from its more than 2 billion users. Only a tiny fraction of these photos are regularly accessed, with the rest shifted to a warm archival storage solution.

The frequently-accessed images – which strongly correlates with the newest images – may be kept on expensive, highly-performant primary storage. The overwhelming majority of the images, however, will be moved to significantly less expensive secondary storage.

This secondary storage is less performant, and would produce much higher latency under the same loading as the primary storage solution. The secondary storage, however, isn't subjected to anywhere near the same stresses as primary storage; it's designed chiefly to have a low cost per unit of storage.

The predominant use case for secondary storage at Facebook is warm data; however, secondary storage has many use cases. These include serving as the active data solution for non-mission-critical workloads. Secondary storage can have a number of different performance requirements, depending on the specific use case.

While Facebook custom designed its secondary storage solution for the purpose, most organizations don't undertake this effort. This is why secondary storage solutions frequently include hand-me-downs from primary storage. These are often out of date, out of support, or demonstrate inadequate performance.

In addition, because secondary storage is typically the majority of an organization's storage, it also contains a collection of "least-expensive-at-the-time" storage solutions. The result is secondary storage that's under-funded and under-supported, and broken into increasingly difficult to manage storage silos.

Silos. Lots of Silos.

One unfortunate side effect of the standard IT lifecycle is that it's a never-ending refresh project. Legacy hardware and software continues to perpetuate silos throughout the data center, further segmenting infrastructure (and the teams that support it). Upgrades can also cause IT infrastructure to be broken into additional silos to facilitate the upgrade.

To further complicate the situation, cloud platforms have now been added to the mix. Cloud adoption has resulted in even more silos to manage and maintain, with very little integration across the board.

The result is a host of issues. One is slow and error-prone provisioning of new systems due to fragmented locations and scattered data. Another is low global visibility into both where data is located, and exactly what that data is.

Dark Data

Not only is a lack of data visibility an operational hindrance, but it can be a serious compliance risk; most audits require you to be able to locate and recover data at any given time, to ensure it's being properly protected. If an organization can't tell what data is stored

where, it's likely there will be an awkward day at the office when the auditor asks for data samples that exhibit compliance.

Typically, this data opacity is a side effect of a “keep everything” mentality. Organizations of all sizes often lack a plan to weed out old and irrelevant data; they don't leverage any sort of tiering strategy to segment data based on privacy, importance, or business value, either. It's a problem exacerbated by an overabundance of management points necessitated by the proliferation of silos.

To clarify the issue at hand, consider a real-world example. Imagine that you work for a publicly-traded healthcare company. Your company has a master database with Personally-Identifiable Information (PII) and Patient Health Information (PHI). This data has been collected over the last 20 years, and includes Social Security Numbers as well as the private healthcare records of millions of people.

A standard backup schedule for such a database might be hourly, with everything else in the data center backed up daily. This would be standard operating procedure in an organization like this, and likely includes replicating snapshots of the primary storage systems to an offsite data center.

One day, auditors show up and ask for 10 test restores of backups from random dates spanning the last five years. Many organizations struggle to identify where the exact data requested is stored.

Within a relatively recent window, information about the location of data will be in the backup and replication logs. For data from five years ago, however, that's highly unlikely. Backup methods, naming conventions, and destinations change often. Data is moved as new storage devices are added, and data center refreshes occur.

Without some sort of indexing and search capability, it becomes difficult to track data over time. This increases the burden of responding to audit requests, and may result in failing certain

audits. The worst-case scenario of these sorts of audit failures is that they can turn into resume-generating events.

This problem compounds exponentially over time as new infrastructure combines with staff turnover. Storage silos increase as data fragmentation works its way through the data center.



Remember: Always back up your data. If your data doesn't exist in at least two places, then it simply doesn't exist. In addition: always test your backups. If you cannot extract your data from a backup, then that data doesn't exist.

Cloud Computing

Though cloud computing is often touted as a solution to all ills, use of cloud computing is not a license to disengage one's brain. A cloud is simply automated IT infrastructure. When a cloud is managed by a vendor, using that cloud means outsourcing the busywork of IT: replacing failed hard drives, patching servers and so on. Fundamentally, however, a cloud is still IT, not much different from what one might find in one's own data center.

Early on, the primary use case for cloud computing was inexpensive storage. In the modern era, clouds are no longer just places to archive data. They've become a true extension of the on-premises data center, and are being leveraged for multiple use cases such as tiered storage, application development, Software-as-a-Service (SaaS), and data protection.

While there are tremendous benefits to this increasing adoption of cloud storage, both economically and operationally, a downside is that it adds to the complexity of an organization's IT infrastructure.

Hypothetically, an organization taking a cloud-only approach might realize silo reduction. In reality, the majority of organizations are adopting a hybrid or multi-cloud approach, resulting in increased complexity and *more* storage silos to manage.

Software developers have traditionally been the most vocal champions of cloud computing. Many initially started using public clouds in a pay-as-you-go model, to avoid having to interact with IT operations and security teams. Public clouds allow developers to get access to needed resources quickly, with little or no red tape. This “shadow IT” proved highly attractive, and in many organizations quickly progressed from dev and test to production usage.

The widespread adoption of cloud computing coincided with the rise of container technology, and with a movement in application development that saw applications breaking down to the microservices level. The result was even more storage complexity, and a requirement for even more control over where company data was going and being used.

In addition to all this, cloud-based services are still used for their popular original purpose: to store cold data. This data, such as aged backups and infrequently accessed files, had traditionally been stored on tape media; it now predominantly lives in the cloud.

Modern Technology Problems

A birds-eye view of modern IT infrastructure reveals that it’s swimming in an endless sea of proprietary hardware from too many vendors. In addition, admins have to deal with too many licenses, software updates, different user interfaces, and support programs. This sprawl has made management way too complex; and with so many offerings on the market, confusion has set in on what to run, and where.

Even in the new age of cloud, there's an entirely new generation of data silos and sprawl going on across multiple clouds; this is due to a lack of defined standards, and uncoordinated adoption happening in IT departments everywhere.

Does this sound familiar and hit a little too close to home? Are you ready to address the growing data and workload needs in a modern, cloud-first world? If so, read on...

CHAPTER 2

Hyperconverged Secondary Storage

In Chapter 1, we went over the problem of “too many” when dealing with everything an IT team needs to support their organization. In this chapter, we’ll explore the concept of Hyperconverged Secondary Storage (HCSS), and how some of the challenges described earlier are being mitigated or eliminated by organizations leveraging HCSS.

What is Hyperconvergence

Within the IT industry, the terms “convergence” and “hyperconvergence” are used almost exclusively to refer to IT solutions which combine storage and compute into a single product. A popular converged solution would be [Dell EMC’s vBlock](#), while well-known Hyperconverged Infrastructure (HCI) solutions include [Nutanix](#), [VMware’s VSAN](#), and [Scale Computing](#).

What are commonly called converged infrastructure solutions sell discrete storage and compute as a single SKU, but still require separate management for both the storage and the compute. Solutions typically referred to as hyperconverged, on the other hand, integrate storage, networking and compute into a single product, with a single management pane. These HCI systems are highly efficient, leveraging software-defined technologies that allow for very granular control of resources, the ability to scale out effortlessly, and run on commodity components, leading to predictability of cost.

Based on the above, it's reasonable to conclude that, in its purest form, HCI is the consolidation of distinct silos of data center infrastructure into a solution with a single management environment.

A smartphone is a great example of HCI in the consumer market. It consolidates a phone, modem, computer, music player, camera, GPS, calculator and even a flashlight into a single device. The user no longer has to set up and administer what were previously individual devices.

As with other instances of HCI, a smartphone is more than the sum of its parts. Using the features of the common integrated platform, powerful new apps can be created that wouldn't be possible were these solutions restricted to separate devices.

Consider an app which finds the best retail price for a product within walking distance. This app may identify the product based on a barcode scanned via the camera; determine the user's location through the GPS; contact a database located on the Internet using the browser; and, finally, combine all this data using the compute capability. The simplicity and efficiency of this approach completely changed the world of consumer electronics.

In the data center world, HCI delivers similar consolidation benefits. It rejects the previous multi-tiered model of IT infrastructure. This model involved networked storage that existed as independent layers of infrastructure, namely storage-area networking (SAN) and network-attached storage (NAS).

Instead of managing compute, network and storage as distinct resource silos, HCI allows organizations to provision IT infrastructure as a completely integrated solution. Just like a smartphone, HCI consolidates multiple functions into a single system, and creates a solution that's more than the sum of its parts.

Mythbusting Hyperconvergence¹



There is a common misconception that HCI requires use of a hypervisor (perhaps because they share a common prefix?). This is actually not the case. While it's true that a hypervisor is often used to provide workload isolation, there are other ways to accomplish the same objective (e.g., workload QoS or containers).

But the hypervisor is a design choice, not a fundamental property of HCI. Hyperconverged architectures can and do use a variety of alternative techniques (which include, but are not limited to, hypervisor technology) to logically isolate workloads on a shared infrastructure.

Another myth is that HCI doesn't allow the administrator to scale compute and storage independently, preventing optimization for different workload combinations. On the contrary, a well-designed HCI solution can support a variety of nodes that provide a different balance of compute and storage.

If you require greater compute performance, you can add CPU-intensive nodes. If you require greater storage capacity, you can scale with nodes that have a greater proportion of storage capacity to CPU. Your hyperconverged environment can be optimized with the correct balance of compute to storage to meet your collective workload requirements. Even more importantly, it can be changed dynamically to meet your changing workload needs over time.

Finally, there's a common misperception that HCI only applies to business-critical workloads. The value of HCI certainly applies to primary storage and servers, but that same value can also extend into

the arena of non-mission-critical secondary workloads such as backup, file shares, app development, object stores and analytics.

Today, secondary workloads are highly fractured, fragmented and inefficient. Furthermore, they typically account for a large percentage of the total infrastructure within a modern data center. This is a huge unaddressed area of opportunity, ripe for consolidation. Applying the principle of HCI, you can collapse these secondary workload silos into a single, intelligent, web-scale platform. Doing so allows you to reap the benefits of global data deduplication, copy data management, in-place search and analytics across a merged infrastructure. The efficiency gains are enormous.

¹ This aside is an adaptation of a blog post by Cohesity's CEO, Mohit Aron, which first appeared on the Cohesity blog in late 2016. At the time of printing, the original article can be found at <https://www.cohesity.com/blog/what-is-hyperconvergence/>

Hyperconverged Secondary Storage

HCI consolidates hardware silos of compute, networking, storage, and data services. When managed, maintained, scoped and upgraded independently, these infrastructure elements can become islands of inefficiency, creating bottlenecks that impact IT services for the entire organization.

This inefficiency is seen when individual components of the infrastructure are kept separate, run and managed by separate teams in non-standard ways, and when the systems aren't truly optimized to work together. HCI delivers the triple benefit of eliminating workload silos, hardware silos, and operations silos within the data center. The result is vast simplification and increased efficiency for your entire infrastructure.

Secondary storage is comprised of (but not limited to) at least four major areas of infrastructure:

- Data protection
- File services
- Test and development
- Analytics

Hyperconverged Secondary Storage (HCSS) converges all secondary storage into a single solution, integrating the management of what had been individual silos of storage. As a unified solution, HCSS can achieve capacity efficiency that's impossible when storage is broken out into isolated solutions.

HCSS provides a single platform for secondary storage, regardless of the use case or storage protocol required. This allows HCSS to solve the performance problems that commonly plague secondary storage. HCSS redistributes storage as needed throughout the storage fabric, ensuring that workloads get the performance they need.

To deliver the same level of performance, isolated storage solutions would have to be massively over-provisioned, and the majority of them would sit idle for most of the day. HCSS's consolidation of secondary storage results in economic benefits similar to those delivered by server virtualization. Those benefits have been well documented over the years, so much so that it transformed the entire IT industry; it makes sense to apply the same model to secondary storage.

Data Protection

Arguably the most pervasive use case for consolidation within the realm of HCSS is data protection. Data protection is also a use case in which the silo problem – including storage, media servers,

and various software, licensing, and user interfaces – makes itself known.

Buying storage from one vendor, backup software from another, and off-site storage services from yet another just scratches the surface of the challenges associated with data protection. The “whack-a-mole” approach to solving one problem with one solution can lead to extreme fragmentation, inefficiency, and low visibility.

As a result, data protection is often considered difficult, slow, expensive, unreliable, and risky. A [recent Dell EMC survey](#) found that only 18% of respondents believe that their current data protection solution will meet all future business challenges, highlighting the scope of the problem.

Difficult, slow, expensive, unreliable and risky are all descriptors that no organization can afford to have associated with their data protection solutions; especially those tasked with protecting mission-critical workloads. By definition, the failure of a mission-critical workload could result in the wholesale shutdown of an organization’s operations.

Data protection challenges are complicated by the growing pressure on IT budgets, which are being increasingly squeezed. At the same time, the number of workloads to manage is constantly growing. In many organizations, this is compounded by the ongoing “right-sizing” of IT teams.

In short, IT is asked to do more – often much more – with less. Given this situation, it’s easy to see how managing, maintaining and upgrading data protection solutions can and does end up chronically delayed. Overworked and understaffed IT teams are under constant pressure to devote efforts to higher-priority mission-critical workloads.

A candid (and more than slightly ashamed) presentation by the IT director of an organization that’s gone through an extreme failure

event is a regular feature of most IT conferences. Hindsight is 20/20, and nearly all of them will admit that they shouldn't have neglected their data protection systems. Lesson learned, they now oversee data protection as diligently as they do their mission-critical workloads.

Organizations seeking to get a handle on their data protection quickly find that consolidating all their data protection infrastructure into one cohesive platform is worth its weight in gold. But data protection is just the beginning of the secondary storage story.

File Services

Secondary storage also includes everyday file storage. This file storage includes users' home drives, departmental network shares, and general purpose long-term storage of data not required for everyday use.

While mission-critical workloads get fancy flash storage, organizations often relegate file services to secondary storage, even when those file services support mission-critical workloads. This is often done because secondary storage is perceived as cheaper, and file storage is seen as less important than the block storage that underpins the workloads themselves. This practice, however, can have slow-moving but severely negative economic repercussions.

In many organizations, secondary storage is used to offer file services to mission-critical workloads using hand-me-down storage systems. This is done in large part because storage solutions originally specified to act as highly performant primary storage solutions have the performance required to keep up with mission-critical workloads.

Unfortunately, the older a storage system gets, the higher its overall support costs tend to become. Vendors charge more for support on

older storage systems, in part as a form of negative reinforcement to pressure organizations into refreshing their storage regularly.

As a result, maintaining legacy infrastructure has a sneaky way of consuming an organization's IT budget, sometimes more so than upgrading to current product offerings.

File services not dedicated to supporting mission-critical workloads have their own challenges, too. The current *"keep all data forever"* strategy many IT organizations employ is not only unrealistic, but spreading, as data growth rates continue to increase.

Remember the earlier example of compliance in a healthcare setting? It's also relevant to mention here: keeping data an organization *should not keep* has legal repercussions. Modern compliance and privacy regulations, such as the EU's General Data Protection Regulation (GDPR), require that organizations only keep the data that they absolutely *must* keep, with the rest being securely destroyed on a regular basis.

Due to the factors listed above, when companies do decide to store data, it's important to store that data in the right place. Unstructured data, such as file services, also needs to be indexed, with a chain of custody that covers the entire lifetime of that data. At the same time, it needs to offer the performance required to support mission-critical workloads at prices that are rational for archival storage.

This is why "storage administrator" has been a well-paid and highly-respected position within IT teams for decades.

Test and Development

Test/dev is one of the functional areas of IT that's seen the most disruption recently. With the introduction of cloud computing as an extension of the data center, companies are able to extend available resources beyond what their on-premises infrastructure offers. But

to understand test/dev workloads with relation to the data center, one must understand where the current state came from.

Developers have long felt that they lived in a world where resources were held hostage by operations teams. As part of their change management processes, operations teams required developers to jump through hoops to get access to additional resources. These resources are required by developers to write and test their code, making the change management processes frustrating and seemingly arbitrary.

To get access to resources, developers would have to submit ticket requests. IT would then make copies of production data, and perform data masking in order to hide any PII/PHI. Operations teams would then go through a series of provisioning operations on various pieces of siloed infrastructure before finally giving developers what they needed.

Most organizations don't consider dev/test to be mission-critical. As a result, provisioned development workloads were usually relegated to old infrastructure that hadn't been properly maintained, and was often different than the production systems.

Responsibility for maintenance of these dev/test workloads often rested with developers. Developers were required to clean up and delete their test/dev environments when done with their test runs, even though most test/dev environments lived on indefinitely. This turned developers into junior sysadmins, furthering frustration and internal IT team strife.

With the introduction of cloud-based services such as Amazon Web Services (AWS), everything changed. Developers could simply enter a credit card number and get immediate access to a seemingly infinite pool of resources that weren't subject to the restrictions and timelines of corporate, on-premises IT.

While this “shadow IT” did present some new challenges for the business, the idea of “agility” was experienced in a new and profound way. Leveraging cloud resources has since become something that forward-thinking IT departments have accepted as the new normal.

These organizations have taken it as their goal to be able to safely and quickly serve up snapshots of running production systems without intervention from the operations team. Policy-based IT automation can then handle workload cleanup automatically.

For cloud-only organizations, running cloud-native applications, with experienced cloud developers, and IT operations teams with extensive IT automation experience, the above works great. And while chasing unicorns, your humble scribe would also like a teleportation app for his iPhone.

Non-mythical organizations, however, have to deal with reality. In reality, organizations need to make dev/test cloud-native easy, both on-premises and off. This level of ease of use needs to be available for legacy applications as well as those on the bleeding edge, with consistent performance, without costing as much as the infrastructure used for mission-critical production workloads.

Analytics

Analytics is not exclusively a secondary storage workload, but it has an outsized impact on secondary storage. When discussing analytics, the conversation often revolves around the process of shaping and forming raw datasets into actionable information. Analytics workloads often work on data found on both primary and secondary storage, accessing and processing enormous volumes of data, and straining the performance capabilities of all storage systems within an organization.

Analytics became an important class of workload when organizations started to realize that they were sitting on digital mountains of

data. The canny ones began to wonder how they could begin to make better use of this to drive business decisions; the next step was to create algorithms to mine their data, looking for actionable insights.

The high resource demands of analytics workloads spawned an entire generation of new software systems, IT infrastructure platforms, and entire specializations, such as the role of data scientist.

Analytics has reshaped the face of the earth itself. Without modern analytics, logistics companies such as FedEx would operate at a fraction of the efficiency they do today, altering international trade on a global scale.

Analytics drives everything from business to forestry management, mining operations, nature conservation and agriculture. Decisions both large and small – where to locate a hot dog truck, how to engage maximum foot traffic, how millions of square miles of the earth’s surface are put to use – are all now utterly dependent on data derived from analytics.

As organizations become more dependent on analytics, the importance of real-time, up-to-date analytics rises. Delivering data to analytics solutions requires that entire copies of datasets be made, and that these copies are created without impacting running workloads.

This means that ponderous and expensive “Big Data” systems may not be the answer; while powerful, they don’t provide the lightweight footprint needed for maximum analytics agility.

This puts secondary data in the spotlight, as the historic data contained therein is an important source of business intelligence for many companies. Real-time insight derivation requires leaving behind the antiquated methods of batch processing and post-process report generation, and entering into a new age of Artificial Intelligence (AI) and Machine Learning (ML), all of which can put a significant strain on isolated silos of secondary storage.

CHAPTER 3

The Role of Public Cloud

It's important to understand the role of cloud computing in IT, specifically as it relates to managing secondary storage. The public cloud in particular has offered redundant, resilient storage at prices that organizations struggle to match with on-premises solutions.

While the public cloud may struggle to deliver adequate performance for certain mission-critical workloads, it has become the standard against which all secondary storage is measured. Organizations with proliferating secondary storage silos regularly find themselves unable to deliver the performance that workloads using secondary storage demand, while at the same time seeing spare capacity stranded in various silos, resulting in economic inefficiency.

Against this backdrop, cloud storage is difficult to argue with.

Own vs. Rent

Owned infrastructure vs. pay-as-you-go cloud services provokes discussions similar to the “rent vs. buy” of home ownership. The decision to rent or own takes into account a number of factors, and operates on a continuum of choice.

There are times when renting makes sense. Shorter-term inhabitations, such as when one only plans to stay in a given location for a year, vacations, or when one doesn't have the financial wherewithal to acquire a home are examples.

Renting has its upsides with regard to flexibility and responsibility, but you pay a premium for those things. Over the long term, owning a home can be cheaper.

But as most home owners will tell you, owning a home can also be expensive. Maintenance, updates, and long-term care and feeding of a house can be taxing, both financially and emotionally.

The expenses involved in home ownership occur irregularly, and sometimes unpredictably. In addition, home ownership lacks the flexibility of renting. It's usually much, much harder, both financially and logistically, to just pack up and move.

This analogy translates fairly well to the decision IT leaders make about engaging cloud resources vs. owning infrastructure. The right answer for many organizations moving forward is usually somewhere in between.

Traditionally, the “in-between” solution had been to lease on-premises or co-located infrastructure. While this still occurs today, organizations are increasingly turning to a mix of on-premises and cloud infrastructure. But as with all things in IT, your mileage may vary, and a lot of questions you have will be met with “it depends” kind of answers.

It's clear why organizations may want to take advantage of public cloud resources, but the public cloud isn't a panacea. There's still an inherent need for IT resources to connect with, and seamlessly span, private data centers and public clouds.

This is often referred to as “hybrid cloud:” presenting a combination of on-premises and cloud-based resources. Today, organizations need cloud solutions that go one step further.

Cloud resources should be invisible to end users, seamlessly operating as an extension of an organization's private data center, spanning both platforms smoothly and transparently. Without this seamless integration, organizations are simply perpetuating the silo problem covered in Chapter 1, and creating yet another silo of management and cost.

Cloud Storage and Data Management Challenges

In the modern era of Application Programming Interfaces (APIs), vendors and software developers must work together to produce the digital connective tissue to unify multiple infrastructures. Today's organizations aren't content to be restricted to on-premises IT, nor are they comfortable betting everything on a single cloud provider. Hybrid and multi-cloud is the future of IT, and this requires the seamless integration of both organization-controlled and outsourced IT infrastructures.

Without tight integration, use of public cloud storage, for example, simply creates another storage asset that needs specialized staff and manual interaction to manage. In other words, another silo: and more silos are bad, especially when the auditor pays a visit.

Cloud computing operational challenges worth mentioning include privacy and data sovereignty concerns, which in turn can drive data locality issues. For example, institutions in the financial services and healthcare industries fall under strict regulations about things like where data can be geolocated.

Since many cloud providers replicate worldwide, it's important to investigate which legal jurisdictions an organization's data may be subjected to in order to gain a true understanding of ownership concerns for data stored in the public cloud. Fortunately, cloud providers are increasingly empowering organizations with tools that allow them to manage geolocation.

This makes it imperative that HCSS not only unify on-premises and cloud storage, but address data locality, encryption, and other cloud storage concerns. Public clouds offer many features, but this only makes control over one's data more important.

Requirements for a Cloud-First HCSS Platform

If cloud storage is something organizations can't do without, and HCSS can get an organization's secondary storage silo problem under control, can HCSS be applied to all of an organization's secondary storage, regardless of where it exists?

Given the complexity of cloud storage, there's an argument to be made that any HCSS platform attempting to consolidate all of an organization's storage must be designed from the ground up to do so. The HCSS platform in question must provide the ability to store, move, and manage data across locations. It should do this between both on-premises data centers and across multiple public clouds.

Furthermore, the HCSS solution must hide the complexities of underlying cloud storage silos while providing a single data management layer. Add in Service Level Agreement (SLA)-driven automation, and the HCSS solution can be used to support organizations adopting a cloud-first approach to storage.

The ideal cloud-first HCSS platform should meet the following requirements:

- **Simplify and converge storage silos.** A cloud-first HCSS platform must be able to consolidate all of an organization's traditional storage silos, including backups, archives, files, objects, test/dev, and analytics. It should provide data access through common storage protocols (NFS, SMB, object) to support the requirements of all users and applications. The HCSS solution should provide best-in-class storage services such as deduplication, snapshots, replication, and encryption. It also needs to be highly scalable to support rapidly-increasing data volumes.

- **Seamlessly span private and public clouds.** The cloud-first HCSS platform must span both on-premises data centers and public clouds. It should consolidate storage silos both in the private data center and in the public cloud, and it must provide a common set of access protocols and management interfaces across any location.
- **Efficient data transfers across hybrid clouds.** The cloud-first HCSS platform should provide a network-efficient, WAN-optimized solution to transfer data between locations. Ideally, the data will be compressed and deduplicated to minimize the load on the network.
- **Minimize costs of data storage.** The cloud-first HCSS platform must minimize the cost of storing data in the cloud. It should eliminate the need to store redundant copies of data across storage silos, and tier the data to the most cost-effective storage media based on workload profiles. It should also assist with cost control by integrating cost visibility into the management interface.
- **Maximize performance.** The cloud-first HCSS platform must maximize performance, both on-premises and in the cloud. It should provide Quality of Service (QoS) controls to minimize resource contention between users, and to allocate resources to the highest-priority workloads. The HCSS solution should be able to move data throughout the storage fabric as needed to meet performance requirements.
- **Built-in security.** The cloud-first HCSS platform should manage data security requirements in hybrid cloud environments. That includes providing encryption for data at-rest and in-flight, managing key rotations, and supporting external key management systems. Role-based access control (RBAC) must be implemented to prevent unauthorized access to data. And

data should be protected from unauthorized modifications (such as ransomware attacks).

- **True software-defined platform.** As a software-defined platform, the control layer is separated from the data layer, so all data placement and access can be managed by policy.

CHAPTER 4

The Cohesity Platform and SpanFS

The previous three chapters have outlined the overall problems and drivers behind the need for storage consolidation, as well as the benefits of doing so.


Cohesity set out to solve these problems by creating a data platform that addresses these issues, specifically for the secondary storage market. The Cohesity platform combines the hardware resources required for today's intense data management operations with software designed from the ground up to facilitate storage consolidation.

Cohesity has created a single, distributed platform that allows IT departments to take advantage of the key benefits of consolidation, including cloud integration. The result is an overall reduction of the storage footprint (and associated costs), both on-premises and in the cloud.

The SpanFS Distributed File System

To enable organizations to control and manage their secondary data at scale, Cohesity has built a completely new file system: SpanFS. It's designed to effectively consolidate and manage all secondary data: backups, files, objects, test/dev, analytics data and more. SpanFS enables organizations to achieve HCSS on a highly-scalable platform that can operate inside an organization's on-premises data center as well as multiple clouds.

Powered by **SpanFS™**



SpanFS is designed to maximize scalability. Everything in the file system is completely distributed across all the nodes in a cluster.

Distributing the file system's metadata throughout the storage fabric means that there's no master node, eliminating a crucial bottleneck that's traditionally plagued scale-out storage solutions. With SpanFS, data and IO are dynamically balanced across all the nodes, and individual nodes can be added or removed to adjust capacity or performance – with no downtime.

By distributing both data copies and metadata throughout the storage fabric, Cohesity provides always-on availability. The data remains available, even in the event of multiple node failures. Software updates are completely non-disruptive and done with rotating node updates.

Here's an example of the kind of resilience Cohesity provides. They tested the scalability of SpanFS by running Cohesity DataPlatform Cloud Edition in Microsoft Azure. The cluster was scaled from 8 to 256 nodes, and as the graphs in **Figure 1** show, IO throughput scaled almost linearly for both sequential and random IO.




Figure 1: Random and sequential read/write performance of SpanFS scaling from 8 to 256 nodes in Microsoft Azure

SpanFS is designed to run anywhere: across an organization's data center, multiple public clouds, and the edge. Since SpanFS is software-defined, it's capable of unlimited scale.

Replication and Cloud

SpanFS can replicate data to another Cohesity cluster for disaster recovery, and archive data to third-party storage like tape libraries and NFS volumes. SpanFS seamlessly integrates with the leading hyperscale public clouds like Amazon AWS, Microsoft Azure, and Google Cloud.

SpanFS makes it simple to use the cloud in four different ways:

1. **CloudArchive** enables long-term archival backup to the cloud, providing a more manageable alternative to tape.

2. **CloudTier** supports data bursting to the cloud. Cold chunks of data are automatically stored in the cloud, and can be tiered back to the Cohesity cluster once they become hot.
3. **CloudReplicate** provides replication to a Cohesity Cloud Edition cluster running in the cloud. The Cohesity cluster in the cloud manages the data to provide instant access for disaster recovery, test/dev, and analytics use cases.
4. **CloudSpin** gives organizations instant access to their backup data in the cloud, helping speed test/dev and cut time to market.

Global Deduplication

SpanFS maximizes space efficiency through the use of global deduplication. This deduplicates data across all nodes, meaning that data only exists in the number of copies configured for resiliency.

Cohesity nodes can have deduplication configured to be inline, post-process or turned off completely. And, since Cohesity is a software solution, it can be easily configured and deployed in data centers or clouds.

Deduplication flexibility has practical uses. For example, when hosting backup jobs, a volume would likely have its deduplication set to “inline,” deduping as the backup streams to the destination. Backups are typically quite large, and the space savings from inline deduplication could be significant.

When using a volume to host files, it’s likely that deduplication would be set to a post-process setting. File updates are typically small, and it’s unlikely to be worth deduplicating files during peak usage.

In contrast, test/dev workloads are likely to be used then quickly destroyed. There’s no need to waste CPU cycles deduping this data, so deduplication on test/dev volumes would be set to off.

While some storage systems out there limit the amount of deduplication you can perform across tiers of storage – and even the logical constructs within a storage system – SpanFS removes these barriers and deduplicates data globally, across all tiers of data, and throughout the entire storage fabric.

Security and Quality of Service

SpanFS enables encryption of data at-rest and in-flight, and provides RBAC that enables mapping of Active Directory or local users to roles with segregation of duty. SpanFS also has built-in Quality-of-Service (QoS) support, data isolation, separate encryption keys, and integrates with enterprise key management solutions.

QoS is designed into every component of the system. As data is processed by the IO engine, metadata store or data store, each operation is prioritized based on QoS.

High-priority requests are moved ahead in subsystem queues, and are given priority placement on the SSD tier. This can have a profound impact on the performance of hot data as it's accessed.

By moving hot data to the SSD tier, Cohesity is able to ensure optimization of resources for any given SLA without having to worry about manual provisioning, latency, and wasted capacity. SpanFS is able to manage all of this automatically, based on the QoS setting the administrator provides.

WAN Optimized

SpanFS optimizes the end-user experience by incorporating a number of technologies aimed at reducing the amount of data required to traverse the network. The end result is efficient use of WAN and Internet bandwidth, enabling data access from anywhere in the world as quickly as possible.

- Effective data deduplication optimizes the bytes transferred over the network. Metadata ensures that as data changes, only the changed elements are sent over the network. Data that is sent once doesn't need to be sent again.
- The size of the data ensures that it takes maximum advantage of the *put* API calls from the cloud vendors to enable larger upload sizes. This also limits the number of API requests to be made.
- In case of network timeouts and errors, periodic checkpointing ensures that the process doesn't restart from scratch.

Multi-Lingual Storage

SpanFS exposes industry-standard NFS, SMB and S3 protocols. Any number of volumes or object buckets can be configured simultaneously on a single Cohesity cluster.

The volumes are completely distributed, with no single choke point, and the data is spread out across all nodes in the cluster. Volumes are accessed through a virtual IP mount point, and user access and IO are distributed across the nodes using the virtual IP address.

Each of these volumes benefit from all the patented SpanFS implementations of enterprise storage services such as global deduplication, encryption, replication, unlimited snapshots, and file/object level indexing and search. While these enterprise storage services aren't unique to SpanFS, the way they're managed in the underlying subsystems is.

Strict Consistency

SpanFS implements a distributed NoSQL store that stores the metadata on the SSD tier. This metadata store is optimized for fast IO operations, provides data resiliency across nodes, and is continually balanced across all nodes.

Like many NoSQL implementations, the key-value store used by itself provides only “eventual consistency.”



To achieve strict consistency, the NoSQL store is complemented with Paxos algorithms. With Paxos, the NoSQL store provides strictly consistent access to the value associated with each key. The sophisticated nature of Cohesity’s solution can be seen in this, and by the fact that it applies strict consistency.

Distributed Backup Software

Traditional data protection is complex, slow and inefficient. It requires multiple infrastructure silos for backups, target storage, and long-term data retention. Backups are typically only done once per day, and recovery can take multiple hours; that’s a far slower window than required by most mission-critical applications.

In addition, backups must be taken anywhere that workloads operate or data is stored. For organizations with multiple sites, or which have implemented multiple infrastructures, this can make data protection expensive and complex. Either the storage solution in use must be distributed across all locations where backups are to be collected, backups must be streamed to a central location, or some combination of both must be used.

It should thus come as no surprise that many Cohesity customers first choose to adopt Cohesity to simplify their data protection environments. Cohesity’s HCSS solution eliminates data protection silos by converging secondary storage, and providing target storage, backup, replication, and cloud integration.

With Cohesity, backups can be taken as frequently as every five minutes, and recovery times reduced to a few seconds. Cohesity delivers this along with native public cloud integration, and at a fraction of the cost of traditional solutions.

Test and Development

Today's data center administrators must constantly balance the needs of developers and the ever-growing volume of data they create. Every time a bug needs troubleshooting or an application needs a revision, the admin needs to make a copy of production data to either reproduce a problem or act as a real-world environment to code against.

With the passive nature of traditional secondary storage solutions, admins often require a separate infrastructure stack to perform their testing and development, further increasing data sprawl across the enterprise.

Cohesity's HCSS platform allows virtualization administrators to leverage secondary storage infrastructure, such as data protection and management, to run test and development environments. This drastically reduces their overall storage footprint. Cohesity's SpanFS empowers developers to clone the latest backup of their production application stack and run it directly off the Cohesity DataPlatform, providing a unified foundation for copy data management.

Native Tape and Cloud Archival

In addition to an organization's typical backup workflows, Cohesity can automate the long-term archival process to a tape device or to the cloud.

While the use of tape and long-term archival is often scoffed at, many organizations, such as those involved in the financial and healthcare sectors, continue to use it, as they're required to keep

up to 10 years of data available. When handled properly, tape has proven itself as a long-term storage mechanism.



While tape backups are a proven long-term archival technology, it's important to retain and maintain the tape drives themselves. Older tapes may require older drives to successfully read data, and if you can't extract data from your backups, then those backups don't exist.

In the past, legacy backup routines were a complicated process of juggling tapes, keeping legacy devices in storage, and maintaining staff awareness of the processes involved. With Cohesity, this is as simple as just adding an additional step to one's backup and allowing the system to manage it.

Analytics Workbench

Further extending the native analytics capabilities of the Cohesity DataPlatform, Analytics Workbench (AWB) opens up the platform for customized analytics that can be tailored to fit unique business needs. AWB allows users to upload and execute custom code to efficiently and quickly process large datasets across a Cohesity Cluster.

AWB jobs run in the background, and can be prioritized in conjunction with other concurrent workloads using Cohesity's built-in QoS capabilities. AWB jobs can be scheduled to run on pre-defined intervals or on a one-time-only basis.

With AWB, organizations have the ability to leverage their secondary storage solutions to become information hubs. AWB enables businesses across different industries to build tailored apps that provide the reporting and analysis that fits their individual needs,

and distribute those apps throughout the organization in an app store-like fashion.

Examples of how AWB is being used across different verticals include:

- **eDiscovery.** Rapid content analysis to find relevant case information for legal requests or holds.
- **Compliance.** Ensure compliance with PII requirements, including cluster-wide content scans for names, phone numbers, and credit card information that may have been stored in clear text.
- **Threat analysis.** Log correlation across disparate packet capture solutions, to identify potential threats or locate the origin of a security breach.

As you can see, the Cohesity DataPlatform is a flexible, scalable, and efficient storage system. SpanFS is the powerful and intelligent workhorse that enables Cohesity to consolidate numerous use-cases and workloads. In the next chapter, we'll dive a little deeper into examples of the diversity of workloads Cohesity can support.

CHAPTER 5

Secondary Storage Workloads


It's important to understand that Cohesity isn't just providing "yet another" storage or backup system, but a complete toolbox of software. Cohesity incorporates software that organizations would typically have to buy additional software to manage.

This is at the core of HCSS. It isn't just about buying additional storage capacity; HCSS represents a fundamental shift in the way secondary systems are managed.

Backup and Archive

Cohesity DataProtect is an end-to-end data protection solution that's fully integrated into the Cohesity DataPlatform. DataProtect simplifies your data protection environment with a single unified solution for backup, recovery, replication, disaster recovery, target storage, and multi-cloud integration.

Data Protect provides sub-5-minute Recovery-Point Objectives (RPOs), instantaneous recoveries, and can reduce the cost of data protection by 50% or more when compared to leading data protection solutions.



Cohesity DataProtect and Cohesity DataPlatform

DataProtect natively provides consistent, managed backups for VMware vSphere, Microsoft Hyper-V, Nutanix AHV, KVM, Microsoft Windows Server, Linux, Microsoft SQL Server, Oracle RMAN, Pure Storage FlashArray, and NAS, including Pure Storage FlashBlade, NetApp, Dell EMC Isilon, and any generic NAS. DataProtect includes a number of features worth exploring:

- Integrated backup and recovery
- Recovery points of minutes, instead of hours
- Instantaneous recovery times
- Short backup windows
- Support for all major hypervisors
- App-consistent backups
- Native protection of primary storage
- Policy-based management
- Instant file-level search
- VM, file, and object-level recovery
- Remote replication for disaster recovery and migrations

- CloudTier
- CloudArchive
- CloudReplicate
- CloudSpin
- Tape archival
- Encryption of data at-rest and in-flight
- Role-based access control
- Ransomware protection

Data Protection in Depth

That's quite a list of features, but what does it all mean?

The short version is that Cohesity includes all the backup and archival software most organizations are likely to need. Incorporating this software directly into an organization's secondary storage solution eliminates the need for separate backup software, proxy servers, media servers, replication software, disaster recovery software and target storage.

The scale-out nature of Cohesity allows backup jobs to be parallelized, allowing for high-ingest throughput. When combined with SnapTree™ for unlimited snapshots and clones, this allows organizations to reduce their RPO to 15 minutes or less. (SnapTree is discussed later.)

Workloads can be recovered instantly by creating a clone of the backup VM and running that clone directly from the Cohesity platform. Workloads can then be returned to primary storage using Storage vMotion, if needed.

Cohesity includes the ability to back up VMs running on all the leading hypervisors, and uses hypervisor-specific APIs to perform

Changed Block Tracking (CBT). Application-consistent backups are made possible with application adapters for physical Windows, Linux, and Microsoft SQL Server (with support for Windows failover clustering and SQL AlwaysOn Availability Groups (AAG)).

These application adapters allow administrators to provision test/dev copies of SQL Server by automating clone, copy and attach of SQL databases to any SQL Server. DataProtect provides fully-managed Oracle RMAN-based data protection with source-side dedupe, support for Oracle RAC and ASM, and log backups for “any-point-in-time” recovery.

Cohesity integrates with primary storage solutions such as Pure Storage, allowing these solutions to automatically tier snapshots down to Cohesity. Cohesity can protect Pure Storage FlashBlade, NetApp, and Dell EMC Isilon with volume snapshots, or any generic NAS by mounting the live volume. Parallel tracking of changed data and multi-stream data transfers provide high-performance backups.

Administrators can create policies that specify application SLA requirements, including RPO, retention policies, offsite replication and cloud archival. These policies can be assigned on a per-VM basis, based on application SLA requirements.

Cohesity’s integrated search capabilities allow administrators and end users alike to instantly find VM and file data. Google-like wildcard search is supported on both VMs and individual files to accelerate recovery times.

Restoration methods include recovery of individual VMs, restoring files to source VMs, and recovering individual application objects for Exchange, SQL, and SharePoint.

Cohesity enables organizations to leverage flexible replication topologies, including site-to-site, one-to-many sites, cascaded, and replication to the cloud. This allows organizations to protect

their data offsite, and enables disaster recovery and migrations to remote sites.

Native cloud integration enables data tiering to the cloud, and offers the ability to replace both tape archive and offsite data protection, with long-term archival to the leading hyperscale cloud providers. Once in the cloud, data can be used for disaster recovery, test/dev, and analytics.

Cohesity supports software-based encryption using the AES-256 standard, with optional FIPS certification. Data is encrypted at-rest on the platform, and in-flight when replicated or archived to the cloud. Keys are automatically rotated and managed by either an external key management system or the Cohesity cluster.

Cohesity integrates with Active Directory to support role-based permissions. Permission levels can be customized by type of user (admin vs. end user) and by source of protected data.

Cohesity also provides comprehensive protection against ransomware attacks by enabling frequent backups (as often as every five minutes), and enabling quick recovery of data on primary storage systems with minimal data loss. In addition, Cohesity backups are stored on immutable snaps to prevent malicious alterations.

Cohesity DataProtect is one of the most integrated, powerful, and complete data protection solutions on the market today. But it's only one part of the overall secondary storage market.

Files and Objects

Cohesity DataPlatform Views provide the ability to create file shares that can be accessed via NFS or SMB protocols. Views are members of a DataPlatform storage domain. Storage domains are logical data pools with defined storage policies for efficiency (deduplication and compression), replication factor and/or erasure coding, encryption, and cloud tiering.

Data security is maintained using protocol-specific permission and access controls. Active Directory and Kerberos authentication integration provide user and group directory and credential management.

QoS policies can be created that prioritize workloads across the cluster. Advanced, integrated data protection can be enabled to protect NFS and SMB data.

Cohesity guarantees data resiliency at scale with strict consistency and data efficiency, with global variable block length deduplication and compression between different workloads like VMs, physical machines, databases, NAS, file share, etc.

Test and Development

Cloning operations used for rapidly spinning up test/dev environments are powered by Cohesity SnapTree snapshot capabilities. SnapTree is the Cohesity-patented, tree-like data structure that enables a virtual, fully-hydrated image of every snapshot created.

Since SnapTree clones are designed with performance in mind, access time is identical, regardless of the number of snapshots or clones created. This is radically different from legacy snapshot architectures, which rely on a chain-linked data structure. With SnapTree, read operations don't have to traverse every link in the chain to find the requested piece of data, making snapshots and clones instantaneous.

Consider an organization with a VMware-based virtual infrastructure. In the case of test/dev, an additional snapshot is taken when the clone is made, and presented back to VMware as a new unique virtual machine.

All reads requested to old data are read directly from the older snapshots by following the short hops in the tree, while all new or changed blocks are written to the new clone instance. This provides

a full clone of live data with zero overhead that can be given to developers to perform their testing.

To facilitate rapid test/dev, organizations have full access to the exposed REST API. This offers the ability to refresh the test/dev environments automatically, either on a scheduled basis or on-demand.

Analytics

Rather than maintaining expensive secondary storage solutions to be used as an insurance policy, Analytics Workbench (AWB) can unlock the vast potential of this data and provide meaningful information to the organization.

AWB is powered by Cohesity's Indexing Engine, which runs against all data stored on the platform. This index can immediately provide some very interesting high-level storage information, such as:

- **Storage metrics.** Detailed storage-level information, including storage utilization, available capacity, and data growth trend analysis.
- **File-level metrics.** Comprehensive breakdown of file-level information, such as file-type and user access history, to understand how secondary storage is being used in a particular environment.
- **VM-level metrics.** Dashboards show storage consumption by application and data change rates, to better anticipate future storage needs.

The Cohesity DataPlatform also has natively-integrated MapReduce, to enable the efficient and quick processing of large amounts of data. Leveraging the power of Cohesity's Indexing Engine, finding a needle in a haystack of data becomes simple.

Cohesity is able to generate detailed storage, file-level, and VM-level reports on the fly. This allows organizations to remove the unnecessary complexity of having to leverage an external business intelligence (BI) tool to extract information from data. Queries run in the background as data is streamed onto the cluster, delivering real-time insights into the data that's being collected and stored.

CHAPTER 6

Real Savings, Real Efficiency

Today's organizations face significant challenges related to their secondary storage implementations. In nearly every organization, secondary storage is fragmented into silos, stranding capacity and forcing organizations to over-build to meet performance requirements.

The integration of cloud computing into an organization's IT mix does little to simplify matters; in many cases, it merely adds a new set of silos located on infrastructures that IT teams don't directly control. Developers and business units don't want to know about the underlying infrastructure issues; they just want IT to work as easily as their cloud computing instances, regardless of where those workloads exist.

All of this data must be protected, secured, managed, and audited. IT teams often feel like they're drowning under the management burden, even as their budgets are slashed and expectations rise.

It likely feels like an unmanageable chore to even begin to imagine overhauling the way secondary storage is managed within your organization. Fortunately, you're not alone; Cohesity is here help.

Hyperconverged Secondary Storage (HCSS) is real. It offers real savings, real efficiency, and opens the doors to other potential solutions, such as streamlined disaster recovery, tight cloud integration, test/dev and the ability to build an analytics app store within Analytics Workbench.

If you were to design and build a brand new data center from scratch, with today's best-of-breed technologies, the result would be a data center designed to behave like a cloud. To accomplish this, HCI would be heavily leveraged. HCI primary workloads and HCSS for secondary storage would form the building blocks of an easy-to-manage, simple-to-maintain data center.

HCI makes infrastructure simple, allowing organizations to handle all their data seamlessly. IT staffs that don't have to focus on keeping the lights on can instead focus generating value for the business, whether it's through further IT automation initiatives or extracting actionable insights from stored data.

Fortunately, organizations don't have to build a completely new data center from scratch to transform how they manage their data. Existing data centers can realize the benefits of storage consolidation, cloud integration, and beyond with hyperconverged secondary storage from Cohesity.