

EXECUTIVE SERIES

**THE  
GORILLA  
GUIDE TO...<sup>®</sup>  
EXPRESS EDITION**



# Modern Backup and Recovery

David Chapa

---

## INSIDE THE GUIDE:

- Understand Why Legacy Backup Isn't Good Enough Anymore
- Find Out Why You Need a 'Backup Bill of Rights'
- Learn From Real-life Backup Success Stories

**TAKE A QUICK WALK  
THROUGH THE IT JUNGLE!**

Compliments of

**COHESITY**

**THE GORILLA GUIDE TO...**

# Modern Backup and Recovery

**Express Edition**

## **AUTHOR**

David Chapa

Copyright © 2019 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](http://www.actualtechmedia.com)

# TABLE OF CONTENTS

<b>Introduction: A New Era Requires New Thinking.....</b>	<b>4</b>
<b>Chapter 1: Why Legacy Solutions Leave You Wanting More.....</b>	<b>8</b>
Backup Should Not Be Hard—So Why Is It?.....	12
<b>Chapter 2: The Backup Bill of Rights.....</b>	<b>15</b>
Examining the Bill.....	17
<b>Chapter 3: Why Modernization Needs Transformation.....</b>	<b>22</b>
Defining ‘Mission-Critical’.....	23
Service-Level Agreements.....	25
<b>Chapter 4: When It’s Time to Change.....</b>	<b>27</b>
Success Stories.....	28
Enter the Modern Era.....	33

# INTRODUCTION

## A New Era Requires New Thinking

Welcome to this Gorilla Guide Express to Modern Backup and Recovery! If you're reading this, it's likely that you're an IT practitioner or executive with a traditional backup environment in your data center. You've probably also been thinking about how to take advantage of current technology to improve your backup situation.

That's smart. With the growth of public cloud, more options are available than ever; choices that you just didn't have a decade ago. Your backup data may be mirrored offsite somewhere, costing you thousands of dollars and not providing any business advantage at all. It may be sitting in a vault under a mountain somewhere, only to be unearthed if the worst should happen. Then, if it does, you have to wait for the trucks with your tapes to arrive, and start the laborious process of rebuilding your data infrastructure.

Meanwhile, you've been reading or hearing about companies that have moved past that, organizations that

have fast backup and—perhaps more important—fast restore. These businesses have moved ahead of yours, and you don't want to fall any further behind.

## Big Numbers: the Data Explosion

According to estimates,<sup>1</sup> 2.5 quintillion bytes of data are being created every day. And it's a pace that's increasing with the growth of the Internet of Things (IoT).



The same Forbes article also states that, “over the last two years alone 90 percent of the data in the world was generated.”

Those are staggering figures. They're also an indication of the problem organizations face with traditional backup methods and infrastructure. There's simply too much data being produced for outdated systems to keep up. Where do you put it all? And with so much data being backed up, how can you possibly get to it in a timely manner should the need arise?

The clear lesson here is that you need to start the process of rethinking backup and recovery *now*, if you haven't already.

<sup>1</sup> <https://www.forbes.com/sites/bernardmarr/2018/05/21/how-much-data-do-we-create-every-day-the-mind-blowing-stats-everyone-should-read>

You're right to be concerned. Overcoming that fear is the focus of this short book. We'll show you how your current legacy backup solution falls short, and point you in the right direction for modernizing your backup operations. It may seem like an overwhelming task, but it really isn't; you just need to be methodical, ask the right people the right questions, and proceed in a step-by-step fashion.

We start with a discussion of the limitations of legacy backup, and why it's no longer optimal. At one time, it was fine, but no longer. Much of this will likely apply to your situation.

The next section deals with what you *should* be expecting with your backup. It discusses what your backup needs to be providing in the current hybrid cloud-era, and those things you have a right to demand. Can your current system—whether you do it yourself or work with a partner—meet those criteria? If not, it's time to think about a change.

Following that is a discussion of some of the nuts-and-bolts of how to start the transformation to modern backup and recovery. A good place to start is to define, in concrete terms, what is truly mission critical in your operations.

Finally, we take a look at some real-world examples of how large companies were able to transform from the legacy method of backup and recovery to the modern method, and how they benefited from it. We'll also introduce you to the solution responsible for those upgrades, and give you some food for thought.

So if you're ready, let's start exploring ways to get more from your backup and recovery!

## CHAPTER 1

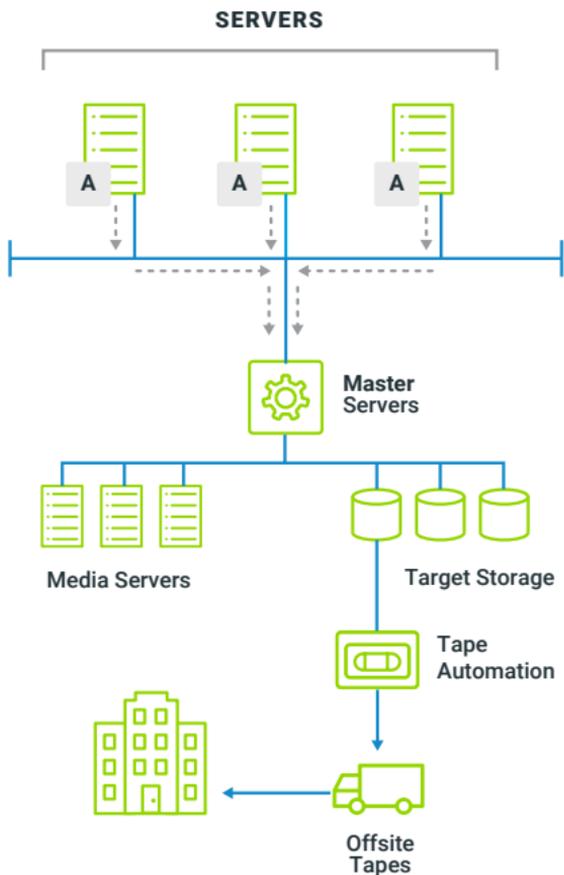
# Why Legacy Solutions Leave You Wanting More

Backup and recovery has changed considerably over the years. Or it should have, based on how organizations function today. About 30 years ago, it was typically quite simple: a backup window opened, scheduled backups ran, databases were shut down through pre- and post-scripts to perform a clean backup. At the end of all the scheduled backups, the databases were restarted and life was good. **Figure 1** shows an example of a typical backup workflow.

That is not the world we find ourselves in today. Gone are the days of a backup window. Companies are 24x7x365, which means the lights must stay on all the time to serve customers. Organizations are global today: their data and applications exist in multiple locations, across multiple clouds. This, combined with a growing data footprint, raises some unique challenges for IT.

Data continues to grow, yes, but it's also sprawling out, becoming fragmented across legacy backup infra-

## Legacy Backup



**Figure 1:** How a legacy backup system typically worked.

structure, locations, and functions. This leads to lack of visibility into and control of data.

That brings up important questions: How is the data being protected and managed? Is the data covered by a backup policy? Is there a plan for recovery?

It gets worse: add to this the constant threat of cyber-attacks, including ransomware and data breaches. So not only is IT struggling with data growth, data sprawl, fragmentation, and data protection, it must also be agile enough to respond quickly and effectively when the bad guys threaten their environments.

These issues continue to plague IT leaders every day—yet, sadly, some of the backup and recovery tools haven't adapted or kept up with current trends in data centers. Although we're living in a modern IT world, many of the backup tools still seem to be stuck in the 1980s and 1990s.

Face facts: these legacy tools are holding you back. As mentioned earlier, chances are your company is always open for business, in a digital sense. That's why business leaders are consistently looking to IT for ways to be as innovative with solutions as the business is in driving revenue. This means you need tools that will help the business continue to innovate and move ahead.

The problem is that older tools haven't kept pace with modern workloads and infrastructure trends, including data and devices everywhere, application mobility, edge computing, and the cloud. These backup tools, to give just one example, are complete resource hogs.

There's also the question of the backup success rate. It's common to measure backup success with percentages: the higher the number, the better you feel. But even with a 99 percent success rate, there's a critical question that must be answered:

What data was *not* protected?

You may not even know! If the data was scheduled for backup, it must have value to your business. So, which server, volume, folder, file, database—whatever—wasn't protected? Losing something today could severely impact operations if you have to use the previous backup for the restore.



**Missing one backup** may not sound too terrible, but think about the domino effect if other servers were dependent on the database for the non-backed-up (or improperly backed up) server. You'd then have to restore all those other servers to align with the broken one. And on and on... you can see where this is going.

The problem can quickly compound, if any amount of data loss—even seemingly tiny bits—has been a regular occurrence with your backup jobs. In addition, you may

not know which job is at issue; it could be the same job or different jobs. And if the problem goes back a long time, consider the potential damage to your environment if you have a real disaster recovery scenario.

## **Backup Should Not Be Hard—So Why Is It?**

Let's face it: backup is slow. There's no getting around that. And missing a backup job creates exposure and risk. If something happens and you need to restore, you may lose much more data than anticipated. What if some of that is mission critical?

Another consideration is data creation and modification. Are you backing up or protecting data as soon as it's touched or created? Chances are your legacy solution isn't providing this kind of protection. This is another level of exposure and risk you open yourself up to, and something you should be thinking about as you look to modernize your backup and recovery solution and strategy.

And backup is only part of the equation. Backup may seem like the most important aspect of managing our data, because it's what we do every day to protect it. But often, very little thought is given to recovery and recoverability. Recovery is like using the insurance you hoped you'd never have to use. Consider what would happen if

your house gets flooded, and you find, to your horror, that you never bought proper flood insurance.

It's the same with recoverability, which is the ability to successfully recover from incidents and achieve Service-Level Agreements (SLAs). When you have to rely on it, you may be in for a shock if you haven't taken the right steps to ensure your backed-up data is ready to go. Contrary to common belief, recoverability is *not* the same as availability.

Those steps include confirming that your backups are available and recoverable, and measuring your infrastructure's ability to quickly and efficiently restore a file, volume, or entire system. You should be using the right tools and metrics to measure recoverability, then determine real-world recovery time.

The emphasis in the modern data center is moving beyond just backup. Now, you must be able to quickly, confidently, even instantly, be able to restore your company's lifeblood—your data.

The bottom line here is that backup shouldn't be this hard to manage, but it is. Is there a better way? Yes, there is. Imagine a future in which you could say with

confidence that you know what's being protected, and how, when, and what your absolute success rate is.

In that ideal environment you would also eliminate the backup overhead within your ecosystem, freeing up resources such as CPU, memory, and network bandwidth for other jobs as you operate your always-open digital business.

Believe it or not, it's achievable today. Read on to find out how.

## CHAPTER 2

# The Backup Bill of Rights

The Backup Bill of Rights is a concept. It states that you should have certain basic expectations of your systems: A fundamental level of protection that should be met. These are things, in other words, you should require of your backup.

### **The Backup Bill of Rights**



1. You have the right for more operational efficiencies from your solution, eliminating the “fragmented silo” approach
2. You have the right for better outcomes through faster backup performance
3. You have the right to recovery that’s not compromised
4. You have the right to a solution to protect both traditional and modern workloads

- 5.** You have the right to a better way to manage and leverage backup data sets
- 6.** You have the right to a solution supporting a hybrid cloud strategy for:
  - a.** Long-term retention (LTR)
  - b.** Protecting modern workloads in the cloud
  - c.** Seamless user experience between on-premises and cloud
- 7.** You have the right to make your backup data more productive:
  - a.** Use backup data for dev/test in the cloud without impacting production
  - b.** Compliance testing against protected data
- 8.** You have the right to derive value from your backup data:
  - a.** Gain insights from your backup data through analytics
- 9.** You have the right to defend your organization against cyber threats such as ransomware



You should have the right to a solution that will meet your needs and expectations, and not one that you have to settle for. That solution should also not need so much customization that it requires additional care and feeding on the part of your IT staff; that defeats the purpose, if it takes them away from higher-priority tasks.

This Backup Bill of Rights lays the groundwork for how you should be looking at taking a fresh, modern approach to backup solutions.

## **Examining the Bill**

Let's explore some of these rights and expand on them a bit more to provide a deeper, richer look at a modern approach to backup.

One major drawback of a legacy backup approach is the inability to derive value from your backup data. In some cases that data sits in a proprietary format, and can only be accessed by the application itself or through some scripting method.

Wouldn't it be great if you could double-click on the data you're protecting and crack open a wealth of treasure stored in it through analytics? What kind of actionable insights could you provide business decision makers, while at the same time not impacting your production workflow?

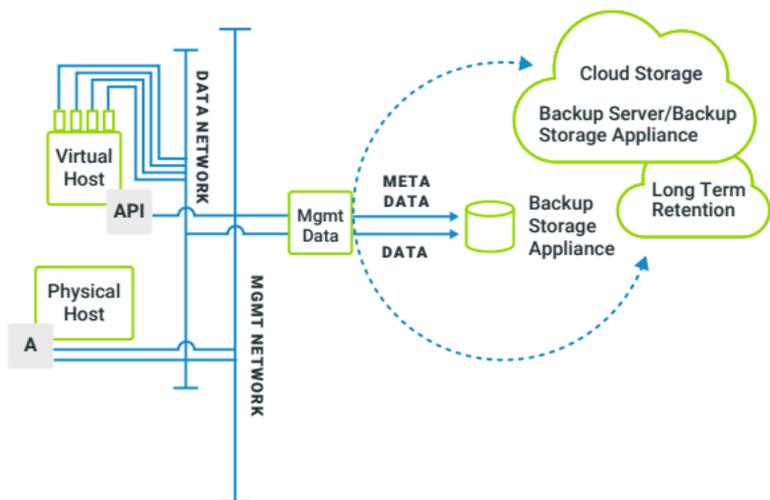
Imagine the benefits of IT not being seen as just a cost center, but instead as an innovator that can help drive better business outcomes. You aren't just protecting the data anymore; you're making it much more agile for your core business, improving your company's bottom line.

That innovation can be applied to the cloud as well. Cloud computing is a key component in enabling digital transformation, and organizations have been looking at ways to best leverage its advantages to provide better products and services to its customers.

More and more, this has meant adopting a hybrid cloud approach. In a hybrid environment, some data, processes, and resources live in your corporate data center, while others live in a public cloud or clouds. The key to a good user experience in this setting is consistency across the board, regardless of where the data, processes, or resources live. Any solution you're considering today that doesn't support a hybrid cloud strategy should be eliminated from the list. **Figure 2** shows an example hybrid cloud architecture.

Given current trends and the adoption of the cloud, your hybrid cloud strategy should be front-and-center as you go through your selection process. Part of that strategy

## Modern Hybrid Cloud Data Protection



**Figure 2:** A typical example of backup and recovery in a hybrid environment.

has to include a few critical components like disaster recovery, archive, and/or long term retention (LTR).

The cloud has its drawbacks, however, and one of them is that it contributes to the problem of data sprawl mentioned before. To be considered a modern backup solution, your data needs to be accessible immediately, from anywhere. You can't wait days anymore to get your data back; that's the kind of antiquated thinking that's holding you back. You need to be able to recover your data—whether it's individual files, hundreds or thousands of virtual machines (VMs), application objects, and so on—instantly. As in, Right Now.



**Given current trends and the adoption of the cloud, your hybrid cloud strategy should be front-and-center as you go through your selection process.**

That's just the starting point, too. Having reliable backups and instant restores is fast becoming a requirement, but the cloud isn't just about storing data or transferring data for backup or archive—it's also being used for compute. Many companies today are creating cloud-native applications running on VMs. So not only should the solutions you're considering support disaster recovery, archive, and/or LTR, they should also provide the same type of protection that your on-premises workloads enjoy for those new modern workloads in the cloud.

The developers creating those cloud-native applications need data to work with and model, which provides another opportunity for IT to innovate: You can accelerate the agile iterations of your development teams by providing access to production data clones within the context of the backup data. In terms of agility, that's a great way to help increase productivity and efficiency.

Finally, the solution you select should be filling the backup gaps you have—not creating more gaps. Take a

close look at the Backup Bill of Rights and identify the gaps most important to you, using it to narrow down the selection of solution providers.

Fill those current gaps by employing a solution that puts your backup data to work for you: gaining insights through analytics, conducting compliance validation testing and application testing with protected data, and so on. The list is long, and the possibilities are endless.

## CHAPTER 3

# Why Modernization Needs Transformation

Now we turn to your needs as an organization. There are some non-negotiables when it comes to backup and recovery, and this chapter will provide some insights to help you make better decisions. So let's look under the hood and begin looking at how to build a strategy you can use to support your Backup Bill of Rights. (This isn't meant to be an exhaustive list; rather, it should give you a good framework to use within your organization.)

Let's first talk about the most important part of backup: recovery. You can't start to build a strategy without first talking about the end game. You need to know what your expectations are for recovery before you can begin to build a strategy for data protection. And you can't talk about data protection without first addressing those items you need to back up. That starts with a deceptively simple question: What things needing protection support the mission of your business?

## Defining ‘Mission-Critical’

That question is deceptive for this reason: Before you can answer, you need to fully understand the mission of the business. If you’re not 100 percent sure of this answer, don’t go any further: you need to invite business unit leaders to the table to help you better understand the core business and mission.

(This has an ancillary benefit: Once you get business leaders involved and identify an executive sponsor, it will raise others’ awareness of you and your team, increase collaboration across business units, and strengthen the validity of the findings.)

Back to the mission. Let’s say, for example, you work for an airline. The primary mission of the business is to keep planes safely in the air and get them safely on the ground at their destination. Knowing this, determining what supports the business’s mission from a technology perspective is where you begin to build out your plan. Here are some of the questions to ask your team, and cross-functional teams, to get at the heart of your mission and thus develop a solid, modern backup and recovery plan:

## Defining Your Mission

---

- 1. What systems, applications, and/or platforms support the mission?**
  - a. Where do these systems reside?
  - b. Is the data in a single location or distributed?
- 2. What's the recovery strategy for each of the primary applications/systems?**
  - a. Where is recovery taking place?
    - Original location
    - Alternate location
    - Public cloud
    - On-premises
  - b. How does the backup plan work in conjunction with the recovery strategy?
    - Are the RTOs acceptable to the business?
- 3. What key elements are missing in our existing backup solution?**
  - a. Can we improve operational efficiencies?
    - Slow performance (backup/recovery)
    - Fragmentation (data sprawled across multiple silos)
    - Management too complex
    - Extensive and complex to operate and upgrade

- b. Is it hybrid-cloud ready?
  - Consistent user experience between on-premises and cloud
  - Easily move data between on-premises and cloud for dev/test
  - Archive for LTR data
  - Support for disaster recovery/recovery to cloud
- c. Are we making use of protected data?
  - Dev/test agility
  - Insights through analytics
  - Scanning for compliance or vulnerabilities
- d. How do we respond to cyber threats?
  - Ransomware
  - Malware

Those are sure to spark a good series of conversations and help focus in on what's needed to help better protect your data and serve your end users.

## Service-Level Agreements

One increasingly-popular way to deal with these questions is through the creation of SLAs. In years past, very few IT organizations had documented SLAs, and when

a breakdown happened, it was often because IT made a decision for the business that had broad ramifications. Having your business leaders engaged in this discussion at the start makes it much easier to clearly identify the right level of service for the various business units within an organization.

Here are the steps for properly defining the scope of SLAs for your company:

- a. Understand the “why” behind the systems of each object you’re protecting
  - Why it matters to the business, and why it should matter to IT
  - What SLAs does the object require?
  - Which systems have the highest priorities?
- b. Will legacy backup meet these objectives? If not, why not?
- c. How many backup applications are running today?
  - What’s the cost of protection?
  - What’s the cost of recovery?
- d. Do you have to make tradeoffs?
- e. Cost of downtime and data loss to the business
- f. Does this fragmentation across backup silos limit visibility and insight?

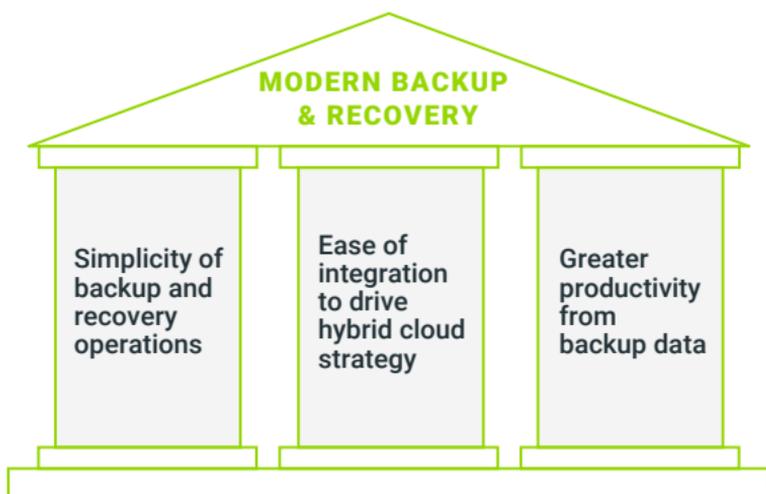
## CHAPTER 4

# When It's Time to Change

So far, we've been discussing the problems with legacy backup, including the time and expense you're incurring because of your outdated infrastructure, and how it's impairing your ability to modernize operations.

Now, we'll talk about solutions. Specifically, the solution Cohesity offers for meeting the evolving demands of today's evolving businesses.

Cohesity's transformative modern backup and recovery approach has three pillars:



**Figure 3:** The foundational supports of a modern backup and recovery solution.

Cohesity offers customers a single, easy-to-navigate user interface, backed by enterprise-grade features to mitigate the massive data sprawl that's fragmented across multiple data silos. It supports the blended model of traditional and modern workloads, as well as providing the insights into the data you're already protecting, to achieve better business outcomes.

Since real-world experience speaks loudest, we'll take a look at a couple of actual customers who chose Cohesity, what primary challenges they faced, and how Cohesity helped change their business outcomes.

## **Success Stories**

### **WestLotto**

WestLotto is Germany's largest lottery provider and the control center for the transnational lottery, Eurojackpot, which includes 18 European countries. The company has been in business for more than 60 years and employs more than 350 people.

WestLotto was faced with a dilemma: the vast sets of sensitive data it takes in on a daily basis place a premium on privacy. The IT team identified three primary goals:

1. Faster time to market with new solutions/products

2. Be capable of cloning complex and interconnected IT systems when needed by dev/test
3. Reduce backup and restore times while ensuring better service and availability

This sounds exactly like the needs that most organizations have when they start to modernize, doesn't it? WestLotto had to provide its customers and users the highest level of security, privacy, and discretion.



**The results speak volumes.** For example, it previously required a day to recover a mission-critical gaming transaction server. That task now takes about 15 minutes to accomplish.

All this had to be done while the business continued to grow, making scale a top concern. In the end, WestLotto modernized and simplified their backup and recovery infrastructure with a software-defined solution from Cohesity. This helped to eliminate multiple products, and consolidate data onto a single, web-scale platform.

The results speak volumes. For example, it previously required a day to recover a mission-critical gaming

transaction server. That task now takes about 15 minutes to accomplish.

On the dev/test side, Cohesity helped IT accelerate the development of new products and games. One key reason was Cohesity's zero-cost backup cloning feature, which allowed IT to quickly and easily provision the systems required by the dev/test teams—without any overhead.

Another issue was the explosion of data, a critical problem most organizations are dealing with in the modern era. The amount of data being written to tape libraries by WestLotto's legacy solution resulted in several new source terabytes every day, making their existing solution untenable. With Cohesity's deduplication, the data volume has been drastically reduced, down to a less than 100GB per day.

## **AutoNation**

AutoNation is another customer. They've been around since 1996 and are publicly traded on the New York Stock Exchange. They're the largest automotive retailer in the United States, selling both pre-owned and new vehicles, managing collision centers, and running a sophisticated auto parts business.

AutoNation has a massive IT environment with its primary data center in Denver Colorado, and a secondary

one in Plano, Texas. It manages more than 1,300 VMs, 200 physical servers, and a growing data workload in excess of 1.2PB that supports the retail operations.

You can imagine the number of backup jobs it takes to protect that much data. AutoNation was dissatisfied with its legacy solution and backup success rates. It suffered a host of issues, but the backup job failures were the most painful—there were nearly 6,000 backup failures per month.

This was an alarming situation, with backup costs skyrocketing. The company was also in desperate need of a way to support its application development effort in Amazon Web Services (AWS).

The IT team put together its requirements, which consisted of:

1. Web-scale architecture, with a predictable cost model
2. The ability to protect, store, and manage data with next-gen technology
3. Capabilities beyond modern backup and recovery
4. A true single pane of glass for management with a hybrid, multi-cloud approach

After an extensive proof-of-concept (POC), the team selected Cohesity. During the POC, they immediately

noticed a drastic decline in the job failure rate; and any failures that did occur were easily traced and remedied. They were also able to take corrective actions to ensure these failures wouldn't repeat themselves again. This degree of troubleshooting insight wasn't available previously.

In the end, AutoNation deployed Cohesity across their environment. They consolidated backup software, media and master servers, and target storage, all onto a single platform. The “pay as you grow” model means the IT team can easily expand by adding new nodes into the cluster to address growing requirements without having to perform forklift upgrades.



**AutoNation was dissatisfied with its legacy solution** and backup success rates.

It suffered a host of issues, but the backup job failures were the most painful—there were nearly 6,000 backup failures per month.

The result? The numbers speak for themselves: AutoNation saved 60 percent per-terabyte compared to a traditional SAN, and reduced their annual maintenance costs by 50 percent.

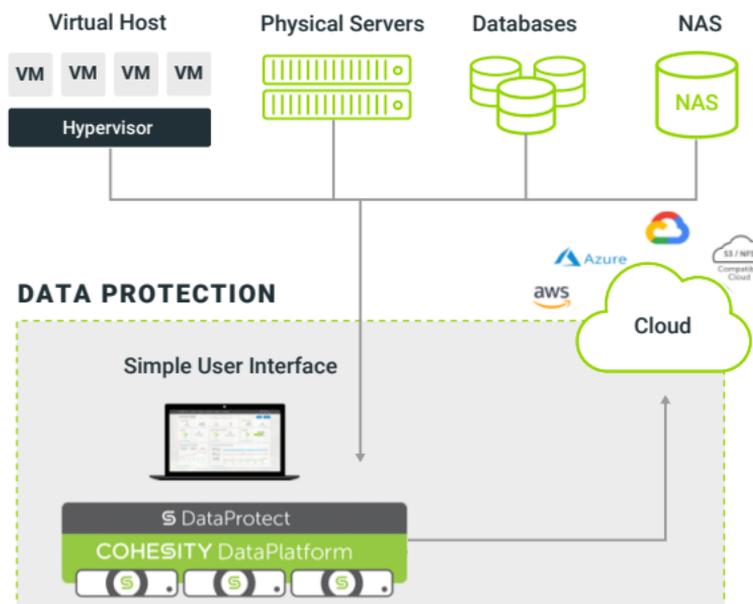
AutoNation now has a disaster recovery plan in place with Cohesity, and they're able to target recovery for the 75 percent of applications that run in their on-premises data center in Colorado, as well as the remaining 25 percent that live in the AWS cloud. Cohesity's ability to seamlessly integrate with AWS from a single pane of glass management interface was a big win for AutoNation, because they can now easily take advantage of AWS Glacier for their LTR data.

## **Enter the Modern Era**

Changing large parts of your business can be scary, to be sure. The fear ramps up even more when it comes to messing with your data, the crown jewels of your organization.

The other side of that equation, however, is the cost you pay by keeping your feet planted firmly in the past. It should be obvious by now that if you think about backup and recovery the same way you always have, you're hurting your company's ability to compete in the modern IT era.

This era has data and workloads everywhere. It erases the traditional data center lines, putting physical and virtual infrastructure wherever the need is, based upon business requirements. It's also required that systems



**Figure 4:** An example of a modern backup and recovery system.

be available at all times and accessible from all places; there's no more room for backup windows that can take a whole night to complete. (Figure 4 shows what a newer, modern backup system might look like.)

It's an era with new threats proliferating everywhere, threatening your precious data from every conceivable direction. The question is no longer if a breach will occur, but when. That's a reality every organization must grapple with and be best prepared to recover from.

It's also an era of great opportunity—the chance to mine your data for insights and information that can drive revenue, rather than just sitting, unused, on a bunch of aging tapes. An opportunity to gain an edge on your competitors, those who persist in living in the past.



**This era has data and workloads everywhere.** It erases the traditional data center lines, putting physical and virtual infrastructure wherever the need is, based upon business requirements. It's also required that systems be available at all times and accessible from all places. There's no more room for backup windows that can take a whole night to complete.

This wonderful, terrifying, increasingly virtualized modern era holds a lot of promise, but also a lot of danger for those wedded to the past. You can join in by adopting a modern approach to backup and recovery, or you can remain stuck in an era that no longer exists, using outdated tools that address none of your current or future needs.

Which path will you take?