

LEARNING MADE EASY

Druva Special Edition

Data Resiliency

for
dummies[®]
A Wiley Brand



Maintain a secure
data backup in the cloud

Automate your defenses
and disaster response

Recover quickly and fully
after an attack

Brought to you
by

druva 

Steve Kaelble

About Druva

Druva enables cyber, data, and operational resilience for every organization with the Data Resiliency Cloud, the industry's first and only at-scale SaaS solution. Customers can radically simplify data protection, streamline data governance, and gain data visibility and insights as they accelerate cloud adoption. Druva pioneered a SaaS-based approach to eliminate complex infrastructure and related management costs and deliver data resilience via a single platform spanning multiple geographies and clouds. Druva is trusted by thousands of enterprises, including 60 of the Fortune 500 to make data more resilient and accelerate their journey to cloud.



Data Resiliency

Druva Special Edition

by Steve Kaelble

for
dummies[®]
A Wiley Brand

Data Resiliency For Dummies®, Druva Special Edition

Published by

John Wiley & Sons, Inc.

111 River St.

Hoboken, NJ 07030-5774

www.wiley.com

Copyright © 2022 by John Wiley & Sons, Inc.

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. Druva and the Druva logo are registered trademarks of Druva. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: WHILE THE PUBLISHER AND AUTHORS HAVE USED THEIR BEST EFFORTS IN PREPARING THIS WORK, THEY MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES REPRESENTATIVES, WRITTEN SALES MATERIALS OR PROMOTIONAL STATEMENTS FOR THIS WORK. THE FACT THAT AN ORGANIZATION, WEBSITE, OR PRODUCT IS REFERRED TO IN THIS WORK AS A CITATION AND/OR POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE PUBLISHER AND AUTHORS ENDORSE THE INFORMATION OR SERVICES THE ORGANIZATION, WEBSITE, OR PRODUCT MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING PROFESSIONAL SERVICES. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR YOUR SITUATION. YOU SHOULD CONSULT WITH A SPECIALIST WHERE APPROPRIATE. FURTHER, READERS SHOULD BE AWARE THAT WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ. NEITHER THE PUBLISHER NOR AUTHORS SHALL BE LIABLE FOR ANY LOSS OF PROFIT OR ANY OTHER COMMERCIAL DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR OTHER DAMAGES.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN: 978-1-119-89422-3 (pbk); ISBN 978-1-119-89423-0 (ebk). Some blank pages in the print version may not be included in the ePDF version.

Publisher's Acknowledgments

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

Project Manager:

Carrie Burchfield-Leighton

Sr. Managing Editor: Rev Mengle

Managing Editor: Camille Graves

Acquisitions Editor: Ashley Coffey

Client Account Manager:

Jeremith Coward

Content Refinement Specialist:

Tamilmani Varadharaj

Table of Contents

| | |
|--|-----------|
| INTRODUCTION | 1 |
| About This Book | 1 |
| Icons Used in This Book..... | 2 |
| Beyond the Book..... | 2 |
| CHAPTER 1: Making Data Resilient across Multiple Clouds | 3 |
| Looking at Today's Data Environment..... | 3 |
| Facing The Risk of Data Loss..... | 4 |
| Understanding Resilience Requirements..... | 5 |
| Identifying key assets and automating data protection | 7 |
| Watching for backup vulnerabilities and threats..... | 8 |
| Responding automatically to prevent spread of threats | 9 |
| Quickly recovering clean and complete data | 9 |
| CHAPTER 2: Making the Move to SaaS..... | 11 |
| Comparing Data Protection Options | 12 |
| Storing data on-premises | 12 |
| Moving to the hybrid cloud..... | 13 |
| Choosing the hosted cloud..... | 13 |
| Opting for cloud-native | 14 |
| Managing Data, Not Infrastructure..... | 14 |
| Seeing the Advantage of Pay-for-What-You-Use Pricing..... | 15 |
| Moving to Infinite Scale | 17 |
| CHAPTER 3: Realizing the Need for Multi-Layered Cyber Defense | 19 |
| Recognizing New Variants of Cyberattacks..... | 19 |
| Looking at Current Data Protection Solutions..... | 20 |
| Upgrading Protection for Today's Threat Landscape | 21 |
| Protect..... | 21 |
| Detect | 23 |
| Respond | 23 |
| Recover..... | 24 |
| Identify | 25 |

CHAPTER 4: Making Data Resilient with Automation..... 27

- Assessing Challenges to Data Management..... 27
- Seeing Challenges in the Real World 28
 - Fleet management..... 28
 - Construction 28
 - Legal 29
 - Sports 29
 - Global energy and urban development..... 29
- Meeting Challenges through AI and ML 30
 - Monitoring access..... 30
 - Detecting anomalies..... 31
 - Orchestrating your response and quarantine resources 31
 - Recovering with confidence..... 32
 - Saving time and reducing data loss..... 33

CHAPTER 5: Realizing the Value of Your Data 35

- Using a Central Cloud Repository..... 36
 - Keeping data in one place..... 36
 - Simplifying compliance 37
- Seeing Current-Day Data Resiliency Use Cases..... 38
 - Enterprise cloud backup 38
 - Disaster recovery 39
 - Ransomware recovery..... 40
 - Enabling e-discovery..... 40
 - Compliance..... 41

CHAPTER 6: Ten Reasons to Adopt SaaS for Data Resiliency in the Cloud..... 43

Introduction

The seemingly daily headlines of ransomware attacks that shut down operations at hospitals, pipeline companies, banks, and virtually all other kinds of businesses are enough to make you turn off the TV news and cancel your newspaper subscription. For business and IT leaders, the prospects are frightening and depressing. And it's not like you can realistically build an impenetrable defense. It takes just one innocent slip by an otherwise fantastic employee to open the door to a ransomware attack.

But with data resiliency, you can change that trajectory. *Data resiliency* is the art of both strengthening your defenses and responding quickly and effectively. It means having a safe, solid data backup that you can access quickly with minimal loss of data. It also means keeping a watchful eye on your systems so you can respond early and minimize the impact of an attack. And simplifying your data management can improve other data functions such as e-discovery and long-term data retention processes.

The best defensive approach is to centralize your data resiliency solution in the cloud, through a Software-as-a-Service (SaaS) model. You can expect sophisticated automation that makes it easy to implement, manage, and maintain, and you also get consumption-based pricing that eases your budget planning and rewards your wallet.

About This Book

This book isn't a Shakespearean tragedy. It's not a horror story. In fact, it has a happy ending. This is a book about resiliency — a feel-good story about being ready for trouble and being prepared to bounce back quickly.

Data Resiliency For Dummies, Druva Special Edition, narrates this happy story for you and your organization. This book explores today's data environment, what makes it challenging and what it takes to achieve resiliency. It outlines the benefits of a SaaS approach that lets you focus on data, not infrastructure, in a way that makes infinite scaling a snap. It explores the need for a multi-layered cyber defense, discusses how automation adds

incredible power, explains why a central cloud repository is the best bet, and shows how it works for current use cases.

Icons Used in This Book

Check the margins of this book and you'll see a few lovely little icons. They're there to draw your attention to some important points. Here is what they mean:



REMEMBER

If you don't have time to read every word in the book, please don't miss the paragraphs marked with the Remember icon.



TIP

The aim is to provide actionable ideas for data resiliency. The Tip icon puts the spotlight on some of those ideas.



WARNING

This entire topic may seem like one big warning, but the Warning icon gets specific about a point you need to keep in mind when thinking about your data.



TECHNICAL
STUFF

This book helps you simplify complicated processes, but if you like the techie details, I go into greater detail with the Technical Stuff icon.

Beyond the Book

While you're reading this book, you may get an appetite for more details about data resiliency. Check out the following resources:

- » www.druva.com: Druva is a provider of SaaS data resiliency solutions and its website offers a wealth of online informational resources, white papers, and videos.
- » go.druva.com/report-idx-understanding-the-data-resilience-challenges.html: Download a free report on how to build true data resilience.
- » www.druva.com/blog/principles-of-a-data-resiliency-cloud-multi-layered-cyber-defense: Learn more about building a multi-layered, data-resilient cyber defense.

IN THIS CHAPTER

- » Understanding today's data environment
- » Grasping the risk of data loss
- » Learning what it takes to be data resilient

Chapter 1

Making Data Resilient across Multiple Clouds

Data lives everywhere, across multiple locations and clouds, distributed globally. That's both powerful and risky because it makes data loss increasingly possible. This chapter explores the realities of today's disparate data environment, discussing the risk of data loss and the challenges of preventing the unthinkable. Explore the concept of data resiliency, understanding the need for not just data protection but greater automation, visibility, analytics, and recovery.

Looking at Today's Data Environment

Not long ago, enterprises made big investments in data centers. Their operations relied increasingly on fast and easy access to ever-growing volumes of data, so it made sense to build cutting-edge castles where data could live, seemingly secure inside a well-guarded perimeter. Now, thanks to cloud services and Software-as-a-Service (SaaS), users and applications can access and store data anywhere in the world (as long as they comply with regulatory requirements). The bottom line is this: Data has expanded. It's on endpoints, in data centers, SaaS apps, and

Infrastructure-as-a-Service (IaaS) platforms such as AWS, Azure, and Google Cloud. This global expansion of new IT services and capabilities, sometimes referred to as a *cloud journey* or *migration*, has unlocked powerful new capabilities and efficiencies for companies and their customers alike.



WARNING

Along with great opportunities come great challenges of protecting fragmented data across clouds, data centers, devices, and more. Any glance at daily headlines makes clear how much the risk has also grown.

Facing The Risk of Data Loss

Today's widespread data environment creates several significant challenges — from the operational to nearly existential. With complexity comes the increased possibility of error and the chance that a misconfiguration will lead to data loss. Add to that the increased challenges related to compliance and other intricacies of data management.

Think about how much easier it is to accidentally delete data living in the cloud. The more places your data lives, the more backups are required for e-discovery and legal proceedings. And those are just the routine, operational issues. Then comes the ever-growing threat of malicious activity from both outsiders and insiders. Your data is incredibly valuable to you, your customers, and your supply chain vendors — that makes it valuable to those with bad intentions.



REMEMBER

Ransomware is, perhaps, the hottest topic of the day in this regard. It's a type of malware that encrypts and holds your critical business and customer data hostage. After an attack, you're denied access to the encrypted data until you pay a ransom to the attacker. That is costly in terms of the actual ransom and even more so for business disruption and impact to your brand. What's more, you're depending on the attacker's sense of honor to release your data — even if you pay the ransom, there's no guarantee a decryption key will be supplied in a timely manner, or that your data will not be used or released widely.

Cybersecurity Ventures predicts an attack takes place every 11 seconds, and researchers at Palo Alto Networks say the average

ransom paid rose to \$541,000 in 2021, nearly double its 2020 amount. As such, the longer your business is without data and in limbo, the more likely it will suffer permanent, expensive damage. It's estimated that nine of ten enterprises losing data for ten or more days end up filing bankruptcy within the year.

That, of course, leads to the observation that if you can just avoid disruptions, you can come out okay. If bad guys have taken your data hostage, just retrieve a backup and you're back in business, right? Absolutely, but far easier said than done. For one thing, bad actors are getting smarter all the time with malware increasingly targeting backups. Even if backups remain uninfected, quickly restoring data without significant data loss isn't exactly a piece of cake, as many customers have discovered.



WARNING

Ransomware poses unique challenges for organizations that want resilient IT and data infrastructure. Traditional recovery techniques, originally built for operational mistakes or infrastructure failures, aren't suited to enable fast recovery of both data and applications. Protecting against attacks and recovering data in the event of an attack are challenging and expensive because of the following:

- » Attacks are frequent, advanced, and expensive.
- » Identifying and restoring clean backups or files is slow.
- » Data is lost due to the inability to recover a complete data set.
- » Contamination can occur during the recovery process.
- » Coordinating and orchestrating incident response at scale is difficult, with hundreds or thousands of endpoints or servers.
- » Business downtime impacts revenue and brand reputation.
- » There's high potential for legal and regulatory fines.

Understanding Resilience Requirements

What your organization needs is true resilience to threats, *data resilience*, to which the title of this book refers. Put simply, you need to be able to steer clear of as much trouble as possible and also recover quickly when trouble arises.



REMEMBER

Fast recovery is the part that's sometimes misunderstood and happens to be the most important part. No matter how strong your team's cyber defenses, you must be ready for the possibility that these efforts aren't enough. Successful attacks are still likely.

Think about the building where you work. It has smoke detectors and sprinkler systems to protect against a fire. Those protections are smart, but they don't eliminate the need for a fire department, and they don't mean you won't have a fire someday. You'd better be ready to recover if a fire happens.

Similarly, data resilience doesn't assume that every malicious actor can be stopped. The goal is absolutely to be as secure as possible, but ultimately, a resilient system quickly detects, reacts, and recovers from breaches, minimizing the impact on data.



REMEMBER

When it comes to ransomware, if you assume an attack will happen, your best defense is reliable backup and recovery. That includes having secure snapshots of critical data and applications so you can roll back in time to a point before the infection. The best system will automate the process, finding clean data from throughout the timeframe of the attack (because not all systems are encrypted at the same time).

Secure backups are important but equally important is *where* those backups are located. In an age when operations, apps, and data live across a distributed, multi-cloud infrastructure, having centralized operations is a must.

Running multiple data protection solutions across these environments isn't sustainable. Your enterprise needs visibility from one control pane into the status of all data to minimize silos, better identify suspicious patterns, and implement effective data governance, especially in an era of increasing regulations.



REMEMBER

Here's another must in today's cloud-reliant data environments: cloud-native data protection. Adapting a legacy product for a cloud environment isn't going to get you the resilience you need. Your backups must be *air-gapped* — a network security measure ensuring your computer network is physically isolated from unsecured networks — to minimize exposure to ransomware and cyberthreats. Your solution must also be infinitely scalable and able to immediately accommodate shifts in data volumes.

In this section, I cover the series of steps you can take to secure and protect your business assets, limit the impact of ransomware, and greatly boost your data resiliency.

Identifying key assets and automating data protection

You can't adequately protect key assets unless you know what exactly they are. You must understand the full scope of the applications and the data that need protection. After you understand that scope, you'll know what must be committed to a secure backup copy. This includes

- » **End-user data:** The most likely source of a ransomware attack comes through social engineering attacks. Vulnerable data lives on desktops, laptops, mobile devices, and more. It must be protected to prevent and limit the spread of ransomware.
- » **SaaS applications:** Microsoft 365, Google Workspace, and Salesforce are popular SaaS apps containing critical enterprise data. That's evidenced by the fact that these providers have started to offer backup services. The question to consider is how many separate backup services you want to manage across your IT environment.
- » **Data center applications and data:** These systems are the high-value targets of ransomware. Loss of access to tier one or two applications (or data) in the data center can critically impact a business.
- » **Cloud workloads:** Shared responsibility models make it clear that cloud providers don't provide backup or recovery services. Cloud workloads and environments may have higher levels of availability and durability, but just like data centers, they require secure data protection and management.

Automating processes and policies for your key assets ensures you have your data when it matters most. Unified cloud backup as-a-service removes operational complexity and data silos from tools and on-premises hardware.

You can shift your focus from “how” to protect to “what and why” to protect data across your environment. The on-demand

nature of SaaS lets you turn up or down the scope of protection as you better understand what you're protecting and how much it costs.



TIP

Long-term retention (LTR) of data in the cloud can be valuable for compliance, costs, and your security posture. While simply sending copies of backup data to the cloud is the least useful way to use cloud storage, LTR or archiving is a way to leverage cheaper storage tiers and simplify the management of old or rarely accessed data.

It's one thing to move a blob of deduped data across storage tiers — a bit like signing up for a storage unit and moving rooms of old furniture, including stuff you don't need. It's quite another thing to globally deduplicate data across all cloud storage tiers.

Ask your backup/deduplication storage vendor how it can deduplicate data in the cloud and across storage tiers. Some providers offer a global, cloud-based file system with significant savings and no additional management or impact on performance.



REMEMBER

Your backups should be air-gapped and immutable. With your backup data fully isolated, ransomware won't be able to exploit its go-to security vulnerabilities. You're not data resilient if a ransomware attack or rogue insider can reach your backup data.

Watching for backup vulnerabilities and threats

Hackers gain access to your systems is through unpatched vulnerabilities. If you're running legacy solutions, you need to scan for vulnerabilities and install updates and patches. You also need to monitor for anomalies that may indicate an attack.



REMEMBER

A data resilience cloud for your data can enhance or provide early detection of threats and risks by observing activity across backup data and the environment. Of course, IT should monitor for anomalous data activity, but how do they know what patterns to look for? Ideally, this is done automatically by machine-learning (ML) algorithms that have knowledge of your historical patterns and patterns from other environments.

Catching ransomware early is vital because it can potentially sit undetected for days or weeks otherwise. The ideal state of data

resilience for your backup environment offers complete visibility into access patterns and user activity. This makes it easier to spot credentials that may have been compromised.

Responding automatically to prevent spread of threats

Quick response is essential, and that's one reason ransomware attackers often strike when your organization is less likely to respond quickly. Expect attacks on weekends or holidays, not necessarily during working hours.



TIP

Resiliency requires a watchful eye not only on primary and secondary storage systems but also on the very systems designed to help you respond and recover. Multi-layer security strategies include monitoring, detection, and response integration.

Your organization's playbook for disconnecting resources from the business network in the event of an attack needs to include your backups. An enterprise with data resiliency will have a system to immediately stop backups from infected resources and prevent restores from infected snapshots. Administrators should have the option to delete infected snapshots, preventing anyone from mistakenly restoring them, as well as wiping clean any infected devices.



REMEMBER

In a data-resilient organization, after a threat has been contained, you should

- » Quickly analyze the environment and discover the scope and nature of the infection.
- » Integrate with enterprise security systems.
- » Understand when and where an attack occurred; this is the first step to knowing what data you can recover.

Quickly recovering clean and complete data

After an enterprise-wide ransomware attack, you need to confidently recover clean data across users and workloads as fast and inexpensively as possible. Resiliency enables file and single system recovery as well as bulk recovery of multiple devices and systems. Recovery procedures may include

- » Administrator-driven restoration of end-user data
- » End-user-driven restoration of end-user data
- » Support for bulk redeployment
- » Restoration of virtual machines (VMs) to virtual private clouds (VPCs)
- » Bulk export for recovery
- » Identification of safe snapshots from which to restore



REMEMBER

Resiliency recognizes that you need multiple recovery methods and that traditional techniques are most likely inadequate against ransomware. Effective recovery from ransomware requires the ability to recover different files from different snapshots at scale, rather than manually. This requires being able to automatically identify the most recent clean copy of data across end-user and enterprise files. Doing this manually is costly and nearly impossible.

IN THIS CHAPTER

- » Exploring the options for data protection
- » Paying more attention to data itself
- » Appreciating pay-for-what-you-use pricing
- » Scaling your solutions to meet your needs

Chapter 2

Making the Move to SaaS

Where should an enterprise maintain its data? Two conflicting imperatives exist. On one hand, you may be thinking, “Time is money, so I want my data and I want it right now!” On the other hand, if you can’t keep that data safe from either accidental corruption or malicious encryption through something like a ransomware attack, you’ve got major problems.

Legacy and hybrid data protection solutions require significant IT overhead for consumption planning, system maintenance, and security monitoring. Software-as-a-Service (SaaS) solutions, on the other hand, empower IT teams to manage the data themselves, without the hassles of dealing with infrastructure. The result is pay-for-what-you-use pricing, infinite scale, and fully managed security and maintenance.

Your focus needs to be on data more than on infrastructure. The best data management and protection options have you paying only for what you need, while allowing you to scale up quickly and easily as those needs grow. This chapter explores today’s major categories of data storage and protection, including on-premises, hybrid cloud, hosted cloud, and cloud-native.

Comparing Data Protection Options

To provide the best protection for your data, you need to understand the pros and cons of each potential option so you can make the best choice for your specific needs. For starters, you may have in place a mix of different systems that have traditionally lived on-premises.

You may also have tapped into hybrid data protection solutions that can function on-premises and extend to the cloud (albeit in a way that's the same as on-premises). You may have embraced SaaS applications for many of your enterprise's IT or business services (for example, customer relationship management [CRM], enterprise resource planning [ERP], or marketing systems) and may wonder what the as-a-service potential may be for data protection.

Storing data on-premises

On-premises data centers have been the traditional venue for storing your data and keeping it safe. Your data center includes primary storage devices, usually high-speed disks, along with a backup. Your backup devices also contain disk subsystems (bring your own device [BYOD] or appliances) and perhaps a tape drive library for offsite retention. On-premises options are known for offering some of the fastest-possible access to large amounts of data.



WARNING

On-premises has lost a lot of its luster over the years as cloud options have become more prevalent. A 2022 Enterprise Strategy Group (ESG) report, *The Evolution of Intelligent Data Management*, found that the cloud provides 41 percent better security than on-premises resources. A 2021 ESG industry survey called *The Evolution of Data Protection Cloud Strategies* revealed that 46 percent of respondents view the cloud as offering better recoverability and reliability of backups. In today's data environment, the once-standard on-premises way of operating has developed a number of glaring weaknesses:

- » **The costs are high.** Maintaining and updating traditional storage and backup solutions can be expensive. Tape backups, in particular, can absorb significant IT resources and generate ever-growing maintenance costs.

- » **Scalability is a challenge.** On-premises data protection and storage solutions often fail to keep up with data growth, especially as data sources outside the data center increase (cloud workloads, SaaS applications, and new remote offices). When the volume of data outstrips the capabilities of a data protection system, you can end up with silos of data as you add new data movers and dedupe storage systems, or an expensive upgrade of hardware and software.
- » **Reliability impacts data loss risk.** If you're in a situation with failed or incomplete backups, recovery from events such as outages or ransomware attacks can be compromised.
- » **Compliance becomes more difficult.** If you extend your on-premises backup solution to different countries and public clouds (for example, archiving backup data to different public cloud zones), meeting compliance and data privacy regulations becomes operationally complex.

Moving to the hybrid cloud



REMEMBER

A hybrid cloud data protection strategy uses on-premises infrastructure for short-term access along with software to perform backups, and the cloud for offsite or long-term archiving. Tasks are divided between on-premises assets and the cloud, with frequently used data stored on-premises.

This approach can potentially reduce total cost of ownership (TCO) versus a purely on-premises implementation. On the other hand, scaling can be more complex and require a high level of IT involvement and expertise. Data management is complex with policies juggling multiple locations, opaque consumption, and billing models.

Your IT team has responsibility for on-premises servers, while the cloud provider manages the public cloud storage systems. Ultimately, though, IT is responsible for reliable data recovery, regardless of where data is stored.

Choosing the hosted cloud



REMEMBER

Hosted cloud uses traditional on-premises backup software running in the cloud — typically in a private virtual cloud. Data can be stored on-premises as well as in the cloud. You can achieve

greater availability than hybrid cloud, with guaranteed single-tenant architecture.

A hosted cloud solution doesn't fully deliver the economies of scale that the cloud has to offer, because it doesn't use cloud-native architecture. The TCO is similar to hybrid cloud or slightly better, and as with hybrid cloud, even though you may employ a third-party vendor for management, hosted cloud solutions still require considerable IT oversight.

Opting for cloud-native

Cloud-native data protection was born to live in the cloud. It is optimized for performance and scalability over the public cloud. A cloud-native solution offers centralized management of backup and recovery processes, along with consistent performance, even with petabytes of data. Plus, you'll find the TCO to be lower than either hybrid cloud or hosted cloud solutions.

The time and resource commitment from your IT operation is lower, too. A cloud-native solution requires much less day-to-day involvement or maintenance compared with hybrid or hosted solutions.

Managing Data, Not Infrastructure

What's more important to your enterprise — the distributed backup infrastructure of hardware and software instances across your on-premises and public cloud environments or the data that lives inside this infrastructure? That's not really a particularly difficult question to answer.

Sure, the miles of cables running up into the racks can be pretty in their own way; the countless blinking LEDs are rather mesmerizing, and the hum of cooling systems can generate a blanket of white noise that's downright soothing. But at the end of the day, it's all just infrastructure: a big and complicated tool that's there to run your applications and protect your data, which is the lifeblood of your operation.



WARNING

Oh, the backup infrastructure certainly demands attention. Managing that infrastructure is a huge and challenging task that never gets any easier. You absolutely must manage the

infrastructure well from storage capacity planning to patching if you're going to adequately protect your incredibly valuable data.

Here's where the cloud and the proliferation of as-a-Service models can really shake up your way of thinking. You're buying the service you need while leaving the infrastructure concerns to someone else. It can mean tremendous operational cost savings because it gives time back to you (and your team) and enables anyone in IT to manage backups, which are reasons enough to head down this path.

But just as important is the increased focus you can give to the things that really matter in your organization's life. The more you can subtract backup infrastructure management needs from your organization's overall to-do list, the more you can really pay attention to the data. And it's the data that really matters.

Seeing the Advantage of Pay-for-What-You-Use Pricing

Pay-for-what-you-use is a concept that makes sense. That's why you see insurance companies marketing policies that cover only what you need to cover. It's why more people are dropping cable television plans with hundreds of channels and opting instead for stream-only programming that meets their own unique interests. And it's why an ever-increasing number of enterprises turn to SaaS models for many of their IT-related needs.

With a SaaS model, you save money and hassle if you don't have to buy infrastructure for the next three years and can promptly ramp up or down capabilities whenever needed. The ability to pay for what you need and use is one of the factors driving cloud service adoption.

Given that, consider the SaaS solution from Druva: the Druva Data Resiliency Cloud, offering Data-Protection-as-a-Service (DPaaS) and more. DPaaS vendors run their backup systems in a cloud account and charge for the amount of service you consume. Depending on the vendor, primary licensing may be calculated in a number of different ways:

- » Some DPaaS vendors price their service per user.

- » Some charge per terabyte stored, either the source (also known as *front-end*) or post-dedupe (also known as *back-end*).
- » Others charge per virtual machine (VM).



TIP

Ask about “other charges” such as fees for restoring data, also known as *egress fees*, and bulk import/export fees. Some vendors have no charges to restore data while others offer a percentage at no charge.



REMEMBER

As you compare one DPaaS vendor with another, be sure you’re aware of all costs that each vendor may charge. You can’t compare apples to apples if there’s unknown fruit involved.

If you’re seeking a data backup product that will run in your own cloud account, you may find it harder to determine costs up front because your eventual cost will be based on the number of cloud vendor services you use. These services may include VM and storage charges, egress charges for restores, and the like.

Such variations mean you won’t be able to come up with a predictable cost per month for your entire environment. You’ll get a bill from the backup vendor, plus another bill varying from one month to the next depending on the number of backups and restores that you do.



TECHNICAL
STUFF

How a vendor implements its software in the cloud also affects your management costs. It may appear to be less expensive to buy a company’s data protection product and lift-and-shift it into your cloud account instead of using an all-inclusive service that runs in the vendor’s cloud account. But that approach doesn’t consider other possible costs, such as the following:

- » VMs
- » Block storage
- » System maintenance
- » System upgrades
- » Patching vulnerabilities
- » Protection from cyber threats
- » Backup platform and data security monitoring

Keeping all these balls in the air can be a major challenge, and you can't afford to drop even one of them.

Moving to Infinite Scale

One of the challenges confronting IT management is the fact that optimization of an organization's requires constant reassessment. The business environment is constantly changing, and IT resources must adapt if IT is to continue to provide effective service. IT teams need infrastructure models that are both simple and flexible so they can match the organization's changing needs.



TIP

Moving IT services and applications to the cloud can help provide on-demand and unlimited scale (whether for bursting or the new status quo). SaaS solutions unburden organizations from backup appliances, on-premises storage, and the need to maintain the hardware and software traditionally required for strong data protection.

Cloud backup as-a-service eliminates most backup challenges. Backing up directly to the cloud, not to be confused with pushing data from on-premises to the cloud, means managing fewer components. You don't have to deal with backup servers (physical or virtual) or appliances. Auto-scaling in the cloud eliminates the need to tune schedules, manage appliance capacity, and balance server load. And with built-in security, the need to retrofit air-gapped backups into an offsite replication model is a hassle of the past.

IN THIS CHAPTER

- » Keeping an eye on new variants of cyberattacks
- » Understanding current solutions
- » Meeting today's threats effectively

Chapter 3

Realizing the Need for Multi-Layered Cyber Defense

Many data protection solutions are no match for the sophisticated new variants in today's cyberattacks, including ones that target backup data. You need a solution that keeps your data safe, helps you bake security into every level of business operations, and automates the process of recovery in the event of an attack. This chapter focuses on upgrading your protection to a multi-layered defense.

Recognizing New Variants of Cyberattacks

If you look up the dictionary definition of “resilience,” it refers to being able to recover from difficulties — and recover quickly and successfully. That's the big reason, of course, why you maintain backup copies of your most vital data.

Reliable backup and recovery is the best defense against ransomware. Having secure backup images of critical business data and applications enables you to roll back in time to a point before the infection occurred, or even automate the process of finding clean data from throughout the timeframe of the attack. Backups are vital — but are they safe?

Unfortunately, cybercriminals have a habit of getting smarter as they go. Ransomware variants update regularly, and these new variants often directly target your enterprise's ability to be resilient.

Ransomware is a lucrative business, and it's estimated there's an attack every 11 seconds. In Chapter 1, I mention that ransoms paid are very high and only getting higher — the chance to increase the price tag offers tremendous incentive for the bad guys to get better and better at doing bad things.



WARNING

Cybercriminals are expanding the scope of their attacks. Importantly, they're recognizing that backup data is an organization's last line of defense and best chance at recovery. So, new attacks are increasingly targeting backup data, working either to encrypt or delete it. Without the backup, victims are more likely to pay the ransom.

The more data and applications you use to operate your business, the broader your “attack surface.” Do all you can to prevent attacks, but be realistic in the knowledge that an absolutely impenetrable defense is nearly impossible. That's why you must also be certain that your recovery options remain strong in the event that an attack breaches your front lines of defense.

Looking at Current Data Protection Solutions

As I mention in Chapter 2, on-premises protection has been a standard for a long time, and this approach used to make a lot of sense. If operations were fully on-premises, your main focus could be on the perimeter. If you keep the bad actors outside that perimeter defense, everything on the inside could live happily ever after.

But that thinking is nearly as dated as a rotary dial telephone. Your data and applications live all over, in many places far beyond that safe on-premises perimeter.

Not only that, on-premises solutions have their own shortcomings that can't be overlooked. Indeed, once a business network or data center has been breached, your protection solutions are subject to many of the same vulnerabilities facing the items they're there to protect. For example, even if you send backups to an alternative site, if that site is on your WAN network, a network breach by an insider could reach that site as well. Shoring up these vulnerabilities often requires significant effort and complex configuration.



REMEMBER

With that said, there's no reason to feel defeated. A solid multi-layered defense is possible, and many solutions can protect your organization. They begin with the understanding that the cloud offers exceptional options that simply can't be had on-premises.

Upgrading Protection for Today's Threat Landscape

The Druva Data Resiliency Cloud offers an example of how a multi-layered defense can take shape and boost the quality of your data protection and cyber resiliency.

Protect

Exceptional protection begins with systems that employ the most stringent security available. Security provided by AWS adds to the level of protection. Backup data managed by Druva sits amid ransomware-resistant architecture and multi-layered defenses to be sure it's inaccessible to encryption or deletion.



REMEMBER

On-premises or Windows-based solutions are vulnerable to ransomware. By contrast, Druva holds backups offsite in a separate account.



TECHNICAL
STUFF

In fact, Druva's backups are protected while in motion by the cryptographic protocol TLS 1.2, the latest Transport Layer Security that provides powerful end-to-end security as data moves across the internet. When at rest, backups are encrypted

with both envelope encryption, which provides multiple layers of protection, and an Advanced Encryption Standard (AES)-256-bit cypher, initially developed to help protect sensitive national security data. *Fun fact:* With a 256-bit cypher, the total number of different possible key combinations runs out some 78 digits, clearly no piece of cake for a hacker to decrypt.

In the Druva solution, data is also split into smaller blocks and stored separately from the metadata needed to reconstruct it. Even if ransomware entered the Druva environment, it couldn't encrypt or delete data.

Druva's encryption for data in flight and at rest offers one unique encryption key only accessible to the customer. That creates crypto-segmentation between customers, with no data leakage possible.

The Druva Data Resiliency Cloud architecture is built on zero-trust security that's applied to every aspect of backup services. Zero-trust is what it sounds like: strict verification treating every access attempt as if it's from an untrusted source.



TIP

Among users, role-based access control is recommended. That way only a small group of administrators is allowed the option of deleting backup data. Check out Figure 3-1 to see how admin roles can be customized to prevent deletion. Beyond that, there are geofencing capabilities that can block access from unknown IP addresses or embargoed countries.

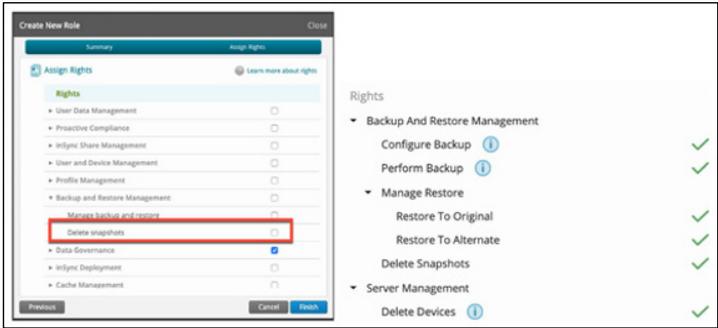


FIGURE 3-1: Customizable admin profiles.

Detect

If your backups are protected, you still need to respond when bad things happen — and recover. That’s what resiliency is all about. You need to respond to ransomware quickly and accurately, which requires anomaly detection working hand-in-hand with your primary security tools.

You can enhance detection capabilities in many ways. The question to consider when enhancing data resiliency is how your backup systems can augment detection already in place without adding complexity. Consider the following two questions:

- » What is the security posture of your backup system?
- » What type of monitoring and alerting exists for users and data across your backup systems?

If you’ve been alerted to a potential or actual breach, the first step is to quarantine infected resources in both the primary and backup environments. That includes putting a quick stop to backing up data from infected machines and keeping people from recovering any data from affected snapshots. This stops the inadvertent spread of ransomware.

Respond

Even the most common criminal knows that the odds of success are best when the potential victim isn’t paying much attention. Ransomware attackers probably aren’t going to pull the trigger during regular business hours. Prime time for them is on the weekend or national holidays.



REMEMBER

Your response activities will be most effective if they happen automatically. Built-in integrations should automate response activities like quarantining infected resources and snapshots across both primary and secondary systems. This prevents the spread of ransomware — and the fact that it’s automated maintains a strong response even if the IT team is off for the evening or weekend.

With affected resources quarantined, it’s important to understand the problem that’s occurring. If the problem occurred within the backup environment, Druva provides both access insights and unusual data activity (UDA) monitoring to help you understand what went wrong during an attack.

Druva access insights check into which users and application programming interfaces (APIs) accessed the backup environment, where the attempts came from and when, and what kinds of actions were attempted. This information can be fed into security information and event management (SIEM) apps and related tools such as Splunk Enterprise Security and Micro Focus ArcSight to accelerate and support investigation. Figure 3-2 shows how these kinds of insights appear in Druva’s solution.

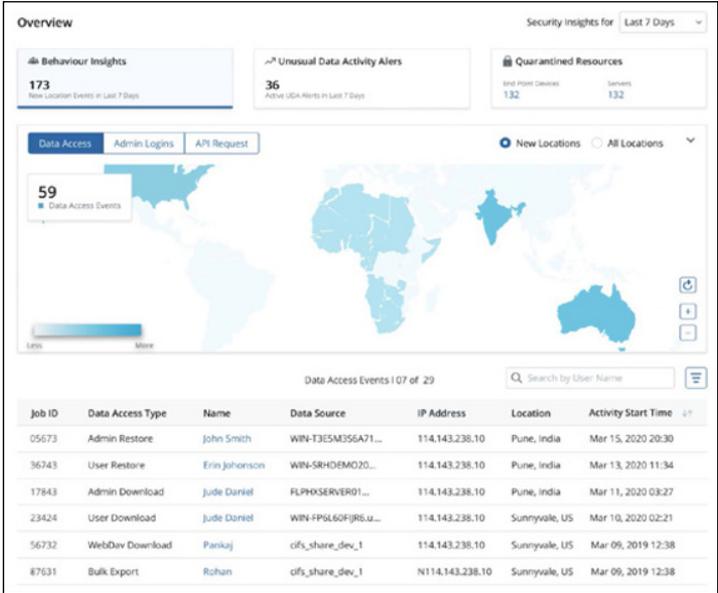


FIGURE 3-2: Gaining access insights.

UDA monitoring detects anomalies at a data level specific to each customer. Druva’s UDA feature continuously watches backup activity with machine-learning (ML) capabilities that understand what’s normal and send alerts when things are not — such as when bulk deletion or encryption are detected. See what that looks like in Figure 3-3.

Recover

After you’ve been hit with ransomware, you’ve got to be mindful of the possibility of reinfection. Your solution must be able to scan snapshots for malware or other evidence of

compromise — before recovery begins. Ensure your recovery data is clean. No point in taking a step forward if you’re going to take two steps back.

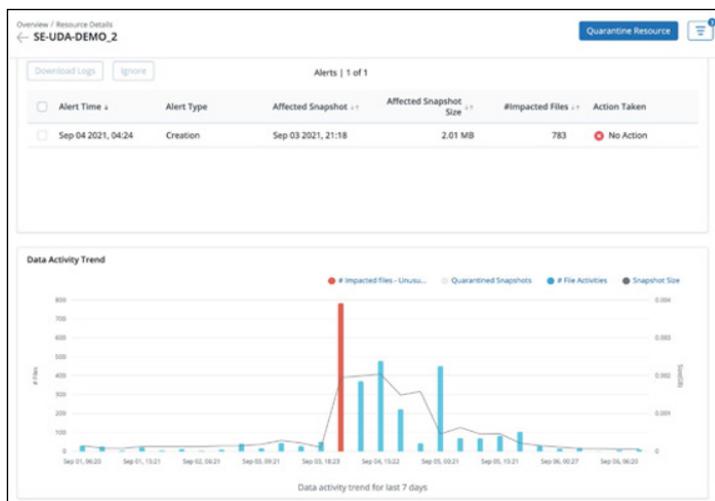


FIGURE 3-3: Tracking unusual data activity.



REMEMBER

Avoiding reinfection means having to remove infected files, and if you’re working from only one snapshot, that means you’ve lost some data. You can reduce data loss with a “golden snapshot” capability that automatically finds the most recent clean version of every file across multiple snapshots over days and weeks.

Druva’s solution is Curated Recovery, and it allows users to define the time period of the attack from initial infection to the present before automatically finding the most recent unencrypted data from that timeframe. In addition, a Recovery Scans feature can scan selected snapshots and automatically remove malicious content. This is possible by using antivirus software or customer-provided hash values from threat intel feeds or forensic investigations.

Identify

You’re back in business, but you’re not fully recovered from a ransomware attack unless you come away with a full understanding of what went wrong. Key takeaways help you improve defenses going forward and bolster your approach to resiliency.



TIP

Your solution should make it easy to access quarantined datasets for forensic investigation. Doing so can help identify the ransomware variants that caused you grief and potentially find and fill gaps in your security infrastructure.

- » Understanding the challenges of data management
- » Seeing the solutions in AI and ML

Chapter 4

Making Data Resilient with Automation

With modern data complexity and scale, it's challenging for humans to manage data, configure policies, and respond quickly when things go wrong. Artificial intelligence (AI) and machine learning (ML) can fill in for where human capabilities can't keep up, helping make data truly resilient while allowing humans to focus on what really matters to their organizations. This chapter discusses some of the challenges related to data management and explores how automated solutions can resolve those challenges.

Assessing Challenges to Data Management



REMEMBER

As your enterprise works to build data resiliency, some of the earliest steps involve assessing the challenges you currently experience with data management. Your team is best off planning for recovery now, when you're not in the middle of a ransomware crisis. Without careful preparation and the right solution in place, in particularly stressful moments, your organization may

experience factors that could impact the ability to respond and recover:

- » A lack of adequate orchestration in your incident response
- » The spread of contamination between systems
- » An inability to access targeted datasets for forensics
- » Reinfection due to contaminated recovery data
- » Data loss due to failed recovery of a complete dataset

Seeing Challenges in the Real World



REMEMBER

Sometimes the best way to recognize your own challenges is to hear the stories of the challenges that others have faced. In this section, I run through a number of real-world situations to try to help shed some light.

Fleet management

To dive right into a bad situation, one provider of fleet management solutions faced the challenge of data loss due to a ransomware attack. This company implemented a protection and recovery solution but found that the recovery wasn't as smooth as it had hoped. Data management was complex, happening across different consoles. And the time required to restore files was longer than ideal. Given that time is money for virtually any organization, speed is of the essence. This company ultimately centralized its data management and protection for greater visibility and faster response.

Construction

Take the case of a major construction company with more than four dozen remote offices. Most of those offices kept data in on-premises file servers, increasing the complexity and cost to manage data. The company was hit with a ransomware attack and found itself unable to restore many of its critical backups. What's more, many of its remote sites had network links that didn't have the bandwidth required for a complete restore.

Its solution focused on maximizing the cloud for better ransomware protection and long-term retention for its remote file servers. Going forward, centralized management and visibility have made it easier to add new sites and manage backups and restores.

Legal

Another real-world case involved a law firm. Its IT team observed the difficulty of manually finding and restoring clean profiles after ransomware. Even without ransomware, data management challenges create real hurdles to productivity.

Sports

One professional sports team had a goal of upgrading on-premises Microsoft Exchange Online infrastructure to Microsoft 365 but found its backup provider didn't have a cloud-native solution for Microsoft 365.

The last thing the IT team wanted to do was buy, provision, and manage hardware when the goal was to be cloud-native. Instead, it found a solution with a single pane of glass for managing backups, implemented role-based access controls for better compliance and security, and got up and running faster than an on-premises solution would've allowed.

Global energy and urban development

Consider the case of a global energy and urban development enterprise, challenged by a lack of a suitable Microsoft Exchange backup that left data open to attack, deletion, and corruption. The company relied on default litigation elements and email journaling for recovering lost or deleted emails. Its challenges came to a head when an executive's calendar was deleted, and it took days to get restored.

The company sought a cloud-native solution that would make such tasks faster and easier. What it found was easily scalable, with self-service capabilities and always-on Microsoft 365 data availability.

Meeting Challenges through AI and ML



After you run through the challenges and hassles that go along with data management and ransomware situations, consider how an advanced data protection solution can help ease those burdens. You can meet these challenges through automation, with the real-world magic of AI and ML.

Monitoring access

The challenge of monitoring access is a lack of insight into access attempts. Your solution should provide situational awareness of your backup infrastructure. With the right solution in place, you can expect to see the following access insights:

- » Which users and APIs accessed your backup environment
- » Where access attempts originated geographically
- » When access attempts were made
- » What actions were attempted (such as recovery or deletion)
- » Insights that come from integration with security information and event management (SIEM) and security orchestration, automation and response (SOAR) solutions

Figure 4-1 shows access monitoring on-screen with Druva's solution.

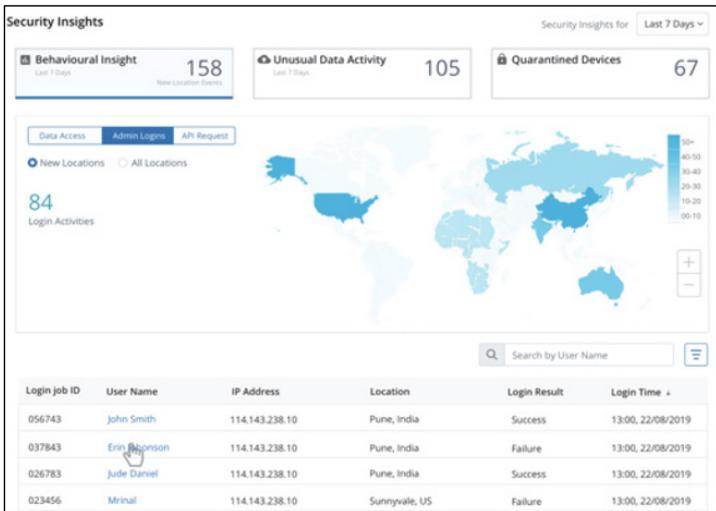


FIGURE 4-1: Keeping an eye on access.

Detecting anomalies

If you don't have data-level insights into the backup environment, you'll have a hard time spotting anomalies with enough time to do something about them. ML can offer you an added layer of accurate detection, as shown in Figure 4-2.

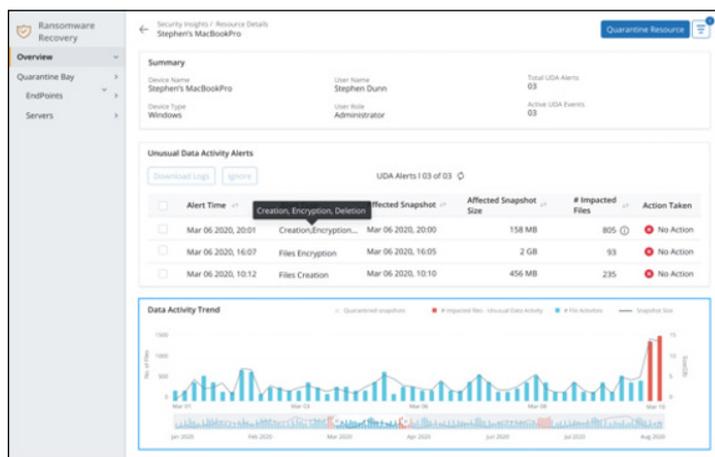


FIGURE 4-2: Spotting anomalies through automated detection.



REMEMBER

Through the power of ML, you can expect the following improvements:

- » Quick identification of affected snapshots during recovery
- » Automatic alerts for unusual data activity (learned through analysis of your specific backup environment), including bulk deletion and encryption
- » Fruitful integration with SIEM and SOAR solutions

Orchestrating your response and quarantine resources

The biggest challenge in orchestrating your response and quarantine resources is a slow response time because of a lack of coordination with your security systems. Automation can really

make a difference here in quickly isolating your resources and reducing the spread of malware through your systems.



REMEMBER

ML and AI can upgrade your automation by

- » Integration with SIEM and SOAR tools to allow automated responses based on ransomware playbooks
- » Prevention of further infections among backups or systems
- » Automated quarantining of infected snapshots in order to prevent inadvertent recovery of infected files
- » Easy inspection of quarantined snapshots for forensic investigation

Recovering with confidence

You're not recovered until you've gotten your data back and your systems running. But it's common to feel more than a little nervous that your recovery will be marred by reinfection from contaminated data as things come back online.



REMEMBER

The solution is to automatically scan snapshots for malware and other indicators of compromise (IOCs) — and do these scans *before* recovery. As Figure 4-3 illustrates, automation through AI and ML can make that a comparative snap:

- » Built-in recovery scans that make it easy to inspect snapshots
- » Scans using built-in antivirus tools to look for known malware
- » Scans using hashes from your own forensic investigations or third-party threat intel vendors

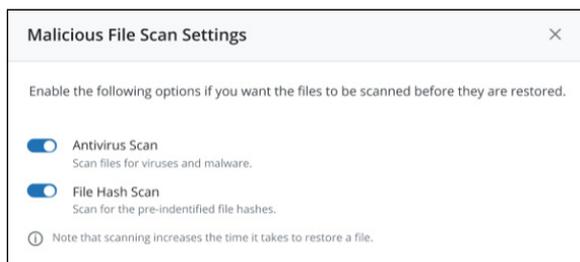


FIGURE 4-3: Preventing reinfection during recovery.

Saving time and reducing data loss



WARNING

What takes a lot of time? Doing stuff manually. And without the right solution in place, point-in-time recovery can require a lot of manual searching for clean files. If you don't undertake a meticulous search for clean files, though, you're likely to wind up with unsatisfactory levels of data loss.

Take a look at Figure 4-4 to see how problematic it can be when ransomware affects different files on different days, which unfortunately is a common situation. In this illustration, if you were to simply go back to day 123 for recovering all files, the last day that *all* files were clean, you'd lose newer data in four of the five files shown here.

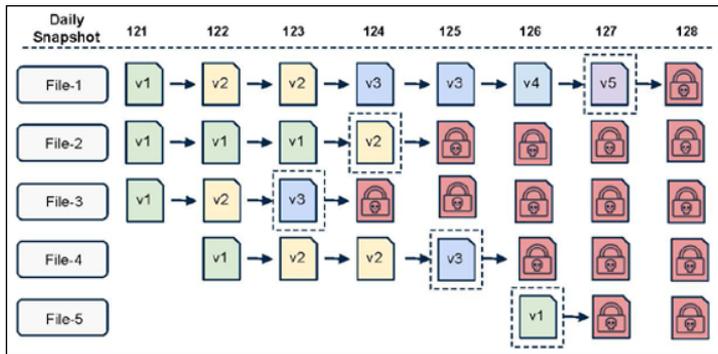


FIGURE 4-4: Searching for the most recent clean files.



REMEMBER

You could search for those newer clean files manually, but that would be a major hassle. That time-consuming drudgery can be avoided by a solution that automatically recovers the most recent clean version of each file. The solution from Druva, shown in Figure 4-5, provides an example of how that can be set up easily:

- » Set a time period and let Accelerated Ransomware Recovery automatically find the best version of every file.
- » The solution assembles clean versions into a single “golden snapshot” for recovery.

Synthetic Snapshot
 Synthetic snapshot is the snapshot with best version of files across the selected date range.

Date Range
 Date range can be maximum of 180 days

Retention Period days
 After retention period, snapshot will be auto deleted.

Exclusions

- System Files
 System, App Settings backs up settings of the following applications on a user's laptop.
- File Hashes
 File Hashes added in the Scan Setting page
- File Extensions

 Type and enter extensions here without using * or . at the beginning

FIGURE 4-5: Automatically putting together the “golden snapshot.”

- » Seeing the benefits of a central cloud repository
- » Checking the use cases for data resiliency solutions

Chapter 5

Realizing the Value of Your Data

It's time to reframe the way you look at your backup and recovery solution. It's there, of course, to protect your data from either accidental or malicious troubles and help you get back quickly in business when disruptions occur. But the reality is, it's not just there to prevent or deal with rainy days.

On the contrary, data resiliency solutions can actually make your data work for you more effectively. When you back up your data to a central cloud repository with a consolidated management system, you'll have a much easier time seeing and acting on all the information in one place. In that regard, it's not just a cost center but a way your enterprise can save money, with simpler governance and compliance and the power of deeper intelligence.

In other words, it's not just about dealing with bad news but creating truly good outcomes. This chapter offers the details on the benefits of using a central cloud repository and the various use cases for data resiliency solutions.

Using a Central Cloud Repository

There's a reason you keep your important documents in folders in a filing cabinet — or even in a bank safe deposit box — rather than strewn across your kitchen table. There's a reason you hang your extra keys on a key rack and your hammers and screwdrivers on a tool rack. When you need important stuff, you want it to be safe and handy and easy to access.



REMEMBER

Employing a central cloud repository for your backup data isn't all that different. The way many organizations have gradually developed systems, data may be stored here and there and everywhere, in multiple disparate locations, and that can be a security vulnerability. But for daily operations, it's also just downright inefficient.

Choosing a backup and recovery option with a central cloud repository offers everyday benefits that go beyond data resilience. There are numerous cost and operational benefits.

Keeping data in one place

Users are tapping into data in countless ways, through a multitude of different methods. They're accessing emails and documents through such Software-as-a-Service (SaaS) apps as Microsoft 365 and Salesforce, or tapping in from the edge via Android or MacOS. They're making use of data in the public cloud on AWS, Azure, or Google Cloud. They're connecting with data centers by way of Oracle or VMware or SQL Server.



TIP

Consider all you can accomplish on a good day with your data in one place. A wide range of data management tasks can be simplified, from long-term retention to e-discovery. Consider archiving cold storage in the cloud. Traditional backup processes can be complex and expensive if they involve moving data from on-premises or cloud-hosted solutions. It's possible to lower storage costs by up to 50 percent with automated global deduplication capabilities across storage tiers, automated storage tiering, and a lower-cost archive storage tier.

This kind of solution is fully automated and policy-driven, with centralized dashboards and predictable costs. File-level search and recovery is fast across warm and cold storage tiers. It yields storage insights to help with long-term planning. It makes the

process of e-discovery much more powerful. And it simplifies compliance with various privacy and residency requirements.

And yes, this is the way to go for rainy days, too. Data is secure and well-protected from ransomware; it's air-gapped and ready for fast recovery from malicious activity.

Simplifying compliance

Every enterprise needs to protect its data and be ready to recover quickly should anything happen. Data breaches and loss, not to mention lengthy operational downtime, can be anywhere from costly to devastating. You already knew that, of course.



REMEMBER

But that doesn't even take the matter of compliance into account. Your organization may have compliance responsibilities tied to the geographical locations where you do business (in the European Union or California, for example). If you're operating in the healthcare sphere, or financial services, or you've got publicly traded stock, or you deal with the government, you may have mountains of additional compliance considerations.

You often know them by their acronyms. The letters HIPAA strike fear in the hearts of many who deal with the Health Insurance Portability and Accountability Act of 1996. GDPR, short for General Data Protection Regulation, is a no-nonsense set of privacy rules for those doing business in the EU. Perhaps your organization deals with such things as CCPA (California Consumer Privacy Act), FIPS (Federal Information Processing Standards), SOC 2 (Service Organization Control 2), SOX (Sarbanes-Oxley Act), or FedRAMP (Federal Risk and Authorization Management Program).

Whatever the acronym and the associated anxiety it might cause you, the truth is it's a daily consideration, not just something that comes up in the midst of a ransomware attack. One CoreView study found that there has been a 250 percent increase in data regulations in the past few years, and Thomson Reuters reports that three-fifths of a compliance officer's time is spent on such things as managing regulatory implementation.



TIP

Your data resiliency solution helps to ease this anxiety every day, not just on your worst days. Take the solution from Druva as an example. It delivers third-party validation ensuring the trustworthiness of its security, not just relying on cloud service provider

certifications but also achieving its own compliance and attestations. Its certifications include

- » SOC 2 type II
- » HIPAA
- » FIPS 140-2 (GovCloud environments)
- » FedRAMP moderate ATO (InSync GovCloud environment)

What this means is your compliance tasks are simplified. Those who use the system can obtain the certifications upon request. Beyond the certifications, Druva maintains an open Vulnerability Disclosure Policy and has multiple third parties conduct ongoing penetration tests to watch for security vulnerabilities and ensure the highest level of security compliance.

Seeing Current-Day Data Resiliency Use Cases

The capabilities I discuss in this book aren't science fiction. When you read about some formerly complicated or time-consuming thing turning into a simple action, it may sound like *Star Trek*. But in this section, you discover some of the use cases for data resiliency solutions that you can acquire *today* and implement faster than you may believe.

Enterprise cloud backup

Protecting data across all cloud applications can be costly and complicated — but it doesn't have to be. The right data protection solution can accelerate projects in an environment such as AWS, for native and migrated applications alike. Druva makes disaster recovery, e-discovery, search, and compliance much simpler by giving you full visibility through a single, easy-to-use console.



REMEMBER

Druva's SaaS platform is one such example, powered by AWS and simplifying the protection of application data across all AWS workloads. That includes cloud-native core AWS services such as EC2 and RDS, containerized apps running on Kubernetes, and migrated workloads such as Oracle and SQL. As a purpose-built SaaS data protection solution, it doesn't require additional hardware, software, or maintenance.

This use case can cut total cost of ownership in half, with improved IT efficiency, no need for dedicated resources, and global source-based deduplication to cut down on storage costs by transmitting only unique data blocks.

Data protection is centralized and policy-driven, archiving is automated, and recovery is fast, granular, and application-aware. And complying with data sovereignty is simplified by the ability to replicate backup copies to multiple AWS availability zones and regions.

As another example, consider also the complexities involving protection and long-term retention of Microsoft 365 data. Microsoft is responsible for platform uptime, but your organization is responsible for data protection.

Choosing a cloud-native data resiliency option can deliver comprehensive backup and protection for your various data that's part of Microsoft 365. That includes OneDrive, Exchange Online, SharePoint, and Teams. And that brings in endpoints, data centers, SaaS applications, and cloud-native workloads.

Going down this path for data resiliency offers protection against malicious activities such as ransomware, as well as accidental data loss. It makes data compliance monitoring simpler and more consistent, and also simplifies such tasks as legal hold and e-discovery management.

Disaster recovery

With a cloud strategy, your organization can achieve remarkably simple, “one-click” disaster recovery (DR) to the cloud for on-premises virtual machines. It's possible to recover across the planet in any AWS region, and to do so quickly.



How quickly? Let's talk about RPO and RTO. RPO is short for *recovery point objective*, which is the time interval that might pass in a disruption before the quantity of data lost during that period exceeds the tolerance specified in the business continuity plan. *Recovery time objective*, or RTO, is the time duration and service level needed in a restoration before there are unacceptable consequences.

Looking for RPO of an hour or less, and RTO measured in minutes? The optimal cloud DR strategy can make that happen,

in automated fashion. You can also gain flexibility to recover Amazon EC2 and RDS resources across AWS regions or accounts.

In Druva's 100 percent SaaS solution, backup and cloud DR are consolidated in a single platform. No additional software, hardware, or managed DR sites are needed — the data footprint is reduced because you only spin up DR resources for an event. Policy-driven automation means employees are more productive.

Secure, air-gapped backups mitigate risk, and downtime is minimized with automated failover/failback execution. You can be confident of readiness through end-to-end, automated, and non-disruptive disaster recovery testing to validate RPOs and RTOs. And centralized management and visibility makes cloud DR all the simpler.

Ransomware recovery

You could say that ransomware recovery is like the poster child for why data resiliency is a worthwhile pursuit. An effective backup is the best defense against ransomware, giving you safe, unencrypted data for recovery. This happens to be the centerpiece of a resiliency solution.

The key to this use case is access to secure, air-gapped backups that can't be hit by the malware that's causing pain elsewhere in the system. Add in such tools as anomaly detection, quarantine, and malware scanning to aid in the recovery.



REMEMBER

Your backups should be strengthened by zero-trust architecture, multi-factor authorization, cloud disaster recovery, Amazon S3 multi-availability zone durability, and delayed deletion. There should be continuous monitoring and vulnerability scans, common vulnerability and exposure patching, and a backup environment with no root access.

Ultimately, your recovery depends on being able to tap into the latest clean backups possible for every file, which is the only way to keep data loss to a minimum. That's only feasible if the recovery is automated, with curated snapshots and the ability to choose the most advantageous snapshot for every file.

Enabling e-discovery

When your organization, or your lawyer, needs to locate specific data, odds are it's a need that must be fulfilled promptly. But the

process of electronic discovery, or e-discovery, is often anything but prompt. With more and more data volumes, cloud service options, mobile devices, and many more places data might live, e-discovery can often be cumbersome and maddeningly manual.

Tapping into the right vendor for managing your data and its resilience can be a real lifesaver. Imagine cloud-native technologies that can proactively collect and manage all data across all of your organization's cloud applications, mobile devices, and computers. Imagine the ability to easily put legal holds in place and store immutable information securely in the cloud.



REMEMBER

The bottom line is that by truly tapping into the value of your data, you can cut e-discovery time and costs in half. Legal teams can minimize the risk of data spoliation and select the most relevant data right at the beginning of case assessments. That's a good start down the road to victory.

Compliance

When it comes to regulation, it's vital to be proactive and not just react to violations and sky-high fines. The right data resiliency solution will include proactive compliance monitoring that not only reduces data risks but also helps your organization avoid those data regulation fines.



REMEMBER

Look for the ability to automatically monitor data compliance across all workloads, with alerts spotlighting violations across such sources as endpoints, emails, Microsoft 365, OneDrive, and Google Drive. You want it all on a single dashboard for ease of use and speed of response, and expect access to customizable templates and rules linked to such key regulations as GDPR, HIPAA, and CCPA.

Your solution should be able to deliver these kinds of compliance, at speed:

- » Defensible deletion of noncompliant data
- » "Right to be forgotten" and "right to access"
- » Limiting distribution of sensitive files

IN THIS CHAPTER

- » Simplifying operations with always up-to-date software
- » Reducing costs, time, and labor
- » Accessing data anytime, anywhere
- » Getting back to business quickly after any type of disaster

Chapter 6

Ten Reasons to Adopt SaaS for Data Resiliency in the Cloud

We're all faced with decisions where the best choice is far from the easiest option. This, however, is *not* one of those situations. Seriously, can you imagine that your organization could establish the most powerful data resiliency solution, and find out that it can cut your costs, save you time, reduce hassles, and simplify your life? And yet, when disaster strikes, you can be back in business quickly? Go ahead and suspend your disbelief and read this chapter for ten reasons your enterprise should adopt a Software-as-a-Service (SaaS) solution for data resiliency in the cloud:

- » **Simplifying pricing:** A SaaS platform for data resiliency offers all-inclusive, simple consumption-based pricing. You're paying for services consumed with an easy-to-understand model, with no additional hardware, appliances, or software to manage, and no maintenance costs.

- » **Security compliant with leading standards:** Expect security and compliance with third-party certifications — all the key certifications you need for compliance and peace of mind.
- » **Maintaining the safest backups:** You get *truly* immutable backups. All backups are copied to three locations for redundancy.
- » **Making management simple:** A SaaS solution is easy to deploy and manage through a simple web portal. And software updates are automated.
- » **Recovering quickly:** A SaaS solution provides much faster recoveries from attacks or disasters, with clean data and an approach aimed at minimizing data loss. This keeps costly business downtime to a minimum.
- » **Gaining intelligent insights:** From storage trends and recommendations to unusual data activity and enhancing sensitive data governance, a SaaS platform uses artificial intelligence (AI) and machine learning (ML) algorithms, so you don't have to deploy them yourself. This enables you to gain greater value from your data.
- » **Reducing storage costs:** Going with a SaaS platform for data resiliency reduces your storage costs while eliminating extra work for your IT team. Expect your total cost of ownership to drop by up to 50 percent.
- » **Scaling up and down:** Because it's SaaS, you can expect limitless and dynamic scaling, either up or down.
- » **Reducing data sprawl:** This kind of solution solves the challenges of data sprawl and siloed, multiple protection solutions. Backing up in one place in the cloud is safer and simpler.
- » **Tackling costly duplication:** A SaaS solution can provide global dedupe capabilities across the entire enterprise to simplify storage and save money. Through the power of deduplication, you could use less than half the storage of traditional solutions.



Your data. Always safe. Always ready.

Leave legacy backup and data protection
in the dust — make a leap to the cloud with Druva.

Experience the industry's first 100% SaaS platform
for data resiliency.

- Get up and running in minutes
- Eliminate hardware and software
- Cut costs up to 50% across all workloads

See for yourself with a free demo.

druva.com

Defend data and recover quickly from attacks

Your enterprise works hard to defend against cyber threats, but no defense is impenetrable. Data resiliency focuses not only on your defenses but also on your plan to respond and bounce back when an attack happens. A centralized cloud-based approach with a Software-as-a-Service (SaaS) model is easy to implement and manage, and it's cost-effective. This book explains how to securely back up data, keep an automated watch over your systems, respond quickly, and minimize the impact of an attack.

Inside...

- Recognizing how threats are changing
- Defining what data resiliency requires
- Building a solution in the cloud
- Appreciating the SaaS approach
- Managing data, not infrastructure
- Using AI to automate your response
- Recovering quickly and completely

druva 

Steve Kaelble is the author of many books in the *For Dummies* series, and his writing has also been published in magazines, newspapers, and corporate annual reports. When not immersed in the *For Dummies* world or writing articles, he engages in healthcare communications.

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

ISBN: 978-1-119-89422-3

Not For Resale



for
dummies
A Wiley Brand

WILEY END USER LICENSE AGREEMENT

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