

EXECUTIVE SERIES

**THE
GORILLA
GUIDE TO...**®
EXPRESS EDITION



Converged Disaster Recovery and Backup

David A. Chapa

INSIDE THE GUIDE:

- Why Your Old Method of DR/Backup Won't Cut It Anymore
- A Comprehensive Data Protection Strategy Checklist
- How To Achieve True IT Resilience

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

Compliments of

Zerto

THE GORILLA GUIDE TO...

Converged Disaster Recovery and Backup

Express Edition

AUTHOR

David A. 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

TABLE OF CONTENTS

Chapter 1: When ‘Files’ Were Pieces of Paper.....	4
End-to-End Coverage.....	5
Increased Requirements for Backups.....	7
The Evolution of Backup and DR.....	10
Chapter 2: The Limitations of Legacy Backup and DR.....	14
Cloud-Ready	18
Viruses, Malware, and Ransomware.....	18
Orchestration and Automation.....	19
Chapter 3: Availability or Resilience? Or Can You Have Both?.....	21
Chapter 4: Protecting Your Organization.....	25
Chapter 5: Dare to Challenge the Status Quo.....	32
The Foundations of Zerto’s Converged Disaster Recovery and Backup.....	33
Elements of a Modern IT Resilience Platform.....	36
The Choice Is Clear.....	38

CHAPTER 1

When ‘Files’ Were Pieces of Paper

As long as we’ve had data we’ve had a need to back it up, protect it, save it from loss. If you’re as “experienced” (aka *old*) as some of us, you may remember big companies with rooms dedicated to nothing but storing paper files. It was typically called — big surprise — the file room; in it, you’d be able to find nearly every company document dealing with clients, vendors, employees, and business operations.

Remember getting a copy of an invoice that you signed? Some of those forms were in triplicate: The original, one for you, and one for the archives.

Those days are largely gone, but as we’ve progressed into the digital age, we didn’t let those basic principles of protecting data veer too far from our original practices. Digitally, we’d have an original copy on disk, then we’d either back up a copy to another disk or a tape device; this copy would then be stored in the archive, or at another office location in the event our original was lost or corrupted and we didn’t have access to our second copy for some other reason.

In the 1960s, 1970s, and even the early 1980s, protecting data was easier: Our critical data existed, for the most part, only in the data center. Very rarely was critical data out in the wild, so our backup strategies and plans were quite basic: We'd send out a system message informing users that the system would be going down at a certain time for backup, then we'd log everyone off, and lock down login capabilities while we performed the backup.

Once that was complete, the system was available for users to access again. On the backend, we began the process of making copies of our copies for offsite purposes. Backup was very simple and straightforward, albeit very time- and resource-consuming.

End-to-End Coverage

In those days, the focus was on backup more than disaster recovery (DR). But these days, DR has become the focus. This isn't to discount the crucial aspect of backup: After all, there's nothing to recover if your backup systems fail. But DR has changed mightily, and its value is more and more related to the concept of speed. In other words, how fast can your business get back on its feet again after an unplanned – or, increasingly, planned – outage?



When we talked about DR “back in the day,” it meant we were sending a full copy of our data offsite in the event of corruption or data loss. Many times, it was a copy that was anywhere from one day to six days old, and we just hoped, in the event of an actual disaster recovery scenario, that it was good enough.

Before we get there, let’s talk a little bit about what constitutes a disaster. It really comes down to one primary factor and two distinct descriptions of events that occur in your environment. The first primary factor is the amount of time your organization can tolerate the outage before it hits a critical mass. The only way to determine this is by conducting due diligence for each of the particular business systems that are core to your operation.

Once you’ve performed due diligence and have defined a *pre-determined* time before you declare a disaster and kick off DR operations, you can begin to look at events in your environment in two distinct categories: Business interruptions and disasters.

A business interruption is something that interferes with normal business operations. It can be anything from a deleted file, corruption, system crash, or network outage, to planned downtime for upgrades, patches, and other necessary events.

If this interruption occurs on a system critical to a core business function, the clock starts ticking. If you reach the pre-determined time outlined in your plan, you've reached the point where you declare a disaster.

When we talked about DR “back in the day,” it meant we were sending a full copy of our data offsite in the event of corruption or data loss. Many times, it was a copy that was anywhere from one day to six days old, and we just hoped, in the event of an actual DR scenario, that it was good enough.

Increased Requirements for Backups

The more connected we became technologically, the more we needed to find better ways to protect our data and institute new DR strategies to improve our recovery point. This ushered in the era of a multi-tiered, server-based approach to backup with full file agent support.

This architecture allowed much more flexibility over how and when we protected our data, and, based on our design, how fast we could back it up. The tiered approach meant that the more servers we deployed to drive data to the tape devices, the faster (potentially) our backups could complete.

In this model, there would be a single server that would be the primary traffic cop; that server would direct the other servers what to do and where to go to retrieve the data. Once these servers established a connection with the clients, they would offload some of the work to the client agent to begin the file collection process, packaging process, and, ultimately, sending that backup package to the backup server responsible for writing it to the backup media.

This worked well for the environments at the time, but as the speed of business increased, so did the requirements for protecting the data being created.

We also found that our DR strategies were still quite limited to very much what we did in the past. While backup was done more quickly because of the multiple servers we had deployed, copies created to go offsite began to present a problem, or an “issue.”

As mentioned before, deploying multiple backup servers may potentially increase the performance of the

backup. But the problem wasn't necessarily with driving the data from the servers; it was with the backup medium of choice at the time, tape. To keep tape spinning and not "shoe-shining," many IT departments opted for an interleaving strategy – multiplexing multiple backup jobs – so the tape drive wouldn't have to pause and wait for its buffer to be filled.




If you've ever had to recover a full volume or server from a backup tape that had backup jobs interleaved on it, you know how painfully slow this process is.

This strategy worked great to meet the backup windows, but was horrendous for DR. If you've ever had to recover a full volume or server from a backup tape that had backup jobs interleaved on it, you know how painfully slow this process is.

DR had its limitations due to the way the data was protected to begin with; that meant that what took mere hours to backup, could take many more hours to recover. We didn't do ourselves any favors by implementing a strategy such as this in our IT ecosystems. As **Figure 1** reveals, most admins still aren't happy with their backup solutions.

Top Areas of Concern Regarding Backup



Gartner

Figure 1: Backups still fall far short of what IT admins need.

The Evolution of Backup and DR

Remember the good old days of the backup window? That's when we had an expectation that no work would be done on the systems or applications between certain evening hours so we could do backups without affecting production. Well, as was mentioned earlier, the speed of business has really become the primary factor in the evolution of backup and DR.

We're more connected globally than ever before, and as such we have to look at alternatives to augment our backup and DR strategies. There's no longer this idea of a backup window per se; in a globally connected world, customers and users expect 7/24/365 availability and access to the systems and data they require.

We looked at faster and faster solutions for backup. Fortunately, solutions such as storage snapshots entered in and provided us a way to at least grab a crash-consistent point-in-time; even that didn't solve the problem completely because we still needed to move that off the primary disk.

That meant we'd look at replicating that snapshot to a like storage solution, preferably somewhere offsite, in the event of a disaster. It was costly and cumbersome, but it was still an improvement over previous paradigms. It wasn't ideal, though, as it placed IT in the precarious position of being physically locked in to a particular storage vendor's hardware solution for DR.

Tape backup was becoming less and less of a primary backup repository in favor of the faster, but more costly and finite, disk-based deduplication appliances. While deduplication provided a much-needed compact data footprint, it also presented some performance issues when recovering larger data sets or creating copies for offsite purposes.

Even cloud solutions gave us promises of a better tomorrow. But some of our architectural decisions in the primary on-premises data center made DR to the cloud less practical, because not all of our systems could simply be "lifted and shifted" to the cloud.



Continuous data protection finally addressed the biggest issue and challenge for IT: the data risk exposure. In other words, when we would employ a backup window or even regularly scheduled snapshots, we were still left exposed between the last backup or the last snapshot.

Another solution that didn't get the spotlight as much as it probably should have was continuous data protection (CDP). CDP finally addressed the biggest issue and challenge for IT: The data risk exposure. In other words, when we would employ a backup window or even regularly scheduled snapshots, we were still left exposed between the last backup or the last snapshot.

Those solutions weren't capturing every data modification, creation, or deletion, but CDP does. The problem with CDP adoption is it may have been a bit before its time, but that was back in the 90s when it was first introduced by Peter Malcolm. Today, the speed of business is pushing us toward the stated requirement of a continuous protection model.

Backup and DR must change: Not necessarily *what* the function describes, but *how* we achieve the results.

Some may argue that DR as a function can and should replace backup, while others would say you must keep the two separate.

Whatever camp you reside in, we all can agree it needs to be reviewed and enhanced, and if that means that DR is the “new backup,” or that backup becomes the new DR, so be it. After all, when time is money, losing hours or days of data can have a real financial impact for every business.

The reality is that your organization needs to find a solution that will meet your business requirements and needs for recovery and business continuity of operations. This is the true evolution; it’s the mindset of how we approach this in IT today vs. in the past.

CHAPTER 2

The Limitations of Legacy Backup and DR

We just finished looking back at the evolution of backup; throughout that chapter, the term “the speed of business” was used to describe the force that ultimately caused newer technologies to emerge and improve the experience with backup and DR. Now we turn to the present, and the challenges with legacy approaches to backup and DR.

Organizations are operating more globally than ever before, which means 7/24/365 is becoming standard operating procedure. To adapt to the ever-changing business landscape, organizations are constantly looking at IT for new application solutions that will help drive core business. This means newer applications, cloud solutions, data mobility for increased collaboration, and so on.

This is just the tip of the iceberg. Sadly, we tend to focus on what we see, and not what’s below the surface. But the reality is what exists below the surface truly makes it non-negotiable. If you don’t have a modern solution for backup and DR that has kept pace with the

newer technologies, you may find your organization is at risk. And those risks can be deadly.

Organizations today must be “always-on” to remain competitive and provide the best customer experience possible. You owe it to your organization to look intently into the needs of your always-on business and whether or not your solution is able to address those requirements not only today, but three years or even five years from now.

You may be thinking that your solution is solid because it’s based on a snapshot approach. However, snapshots

Why Legacy Backup Doesn't Work Anymore




Figure 2: Old methods of backing up data and systems don't work in the modern era.

still aren't backups, and they do have additional impact on resources and production. Copy-on-write snapshots are still a common technology in many modern storage arrays, and they increase the burden on your data center. This is especially true when employing high snapshot activity to avoid the protection gap.

Legacy solutions are typically still aligned with a focus on backup (**Figure 2**). There's nothing wrong with that; backup is important, but the solutions you look for should be equally aligned with both DR and backup. This should include a virtual machine (VM), application, and cloud-first approach to provide the optimal level of protection and continuity for your organization.

The “just back up” approach is slow and dated. It can also introduce new problems: Suppose the files are corrupted, or contain malware, or carry a virus.



When we discuss cloud-ready applications, the first thing that should come to mind is data locality — where the data lives should not present a challenge to your data protection solution.

The backup could be just as dangerous as not having a backup.

In addition, today's modern environments include multiple hypervisors, cloud connectivity, and so on. Many backup solutions may claim to provide you the broadest support for these features, but you need to take a closer look and ask questions:

1. Does it give you an application-consistent backup across VMs?
2. Is it agent-based or fully integrated into the platform?
3. Does it sequentially protect the VMs?
4. Is it really and truly "cloud-ready"?
5. Does it provide viruses/malware/ransomware protection?
6. Does it support all hypervisors, a limited set of hypervisors, or perhaps just one?
7. Does it provide automation and orchestration?

Let's look at a few of these in more detail.

Cloud-Ready

When we discuss cloud-ready applications, the first thing that should come to mind is data locality — where the data lives should not present a challenge to your data protection solution. Yet some of the legacy solutions, because of the core engine and how it was originally developed, struggled in this area simply because they weren't cloud-native solutions.

Viruses, Malware, and Ransomware

Viruses and malware may seem like smaller, more manageable issues in your environment, but when you introduce ransomware into the equation, it changes the entire conversation. Are you prepared for the consequences of a ransomware attack? In some cases, ransomware can cause hours and hours of downtime if you're not prepared.



Downtime not only affects the potential loss of revenue, but in reality that could be the least of your concerns; the greater impact can be to customer loyalty and trust.

Ransomware is an ugly reminder that we still have malicious people out there, and they're getting more and more creative in their tactics. Downtime not only affects the potential loss of revenue, but in reality that could be the least of your concerns; the greater impact can be to customer loyalty and trust.

With any organization, your customers rely on you to keep your systems secure and safe. When they can't access their data or information for some undisclosed reason, it begins to raise red flags and warnings to these customers.

Orchestration and Automation

In the midst of a disaster, the last thing you want is a mistake to occur while you're executing your DR plan.

But the reality is that we're all human, and mistakes can happen when things are heating up and the pressure's on to move your workloads to your recovery zone. In the past, you might have scripts that would run to execute some of these steps, but scripts can only be so intelligent. Often, they'd be outdated due to changes in the environment, hardware, and software.

When looking for a truly modern approach, solutions with built-in orchestration and automation make the process of recovery simpler and more straightforward.

The last thing you want to do is to have to “think through” your DR execution strategy. Intelligently moving your production workloads in the event of a real disaster or as part of a DR test is the goal for data center admins.

CHAPTER 3

Availability or Resilience? Or Can You Have Both?

What does it mean to have an available and resilient IT infrastructure? It's not just about having the most expensive technology or the newest technology, but having the right technology to best serve the core business and financial or operational objectives.

An available infrastructure is just that: One that will remain up and running in the event of a system or application failure. This harkens back to the day of the very expensive high availability (HA) solutions. While costly and complex to configure, these HA solutions did the job: If something went awry on system A, it would fail over to system B. There wasn't necessarily a lot of intelligence built-in, but it provided a more or less scripted outcome. Once you remedied the issues on system A, you would then manually fail back from system B to system A.

An IT resilient infrastructure is one that will seamlessly adapt to change, while protecting your business and customers from expensive and reputation-damaging disruptions. Whether planned or unplanned, an

Converged Disaster Recovery and Backup




Figure 3: A holistic solution is what's needed for modern disaster recovery and backup.

IT resilient architecture allows you to be ready for any scenario and allow you to mitigate the risks of downtime that may otherwise affect your core business. **Figure 3** illustrates the best method for obtaining that resilience.

So, what are some of the technologies you can look at to help you shape the IT resilience architecture you're striving to achieve?

1. Snapshots
2. CDP
3. Journal-based recovery
4. Replication
5. Cloning

We've already talked about snapshots and CDP, but have yet to touch on journal-based recovery. Journal-based recovery allows you to not only recover entire volumes, but provides seconds of granularity to the file level.

The journal maintains all of the checkpoints it's been capturing during the IO stream while the VM was active during the protection. This journaled approach lets you select the point at which you'd like to recover the file or files. This was lacking in the original CDP back in the 1990s, and it means you can pinpoint your recovery, almost like rewinding to a particular scene you're watching from your DVR. You can return to the exact moment the ransomware, file corruption/deletion, or other disaster occurred.



Journal-based recovery allows you to not only recover entire volumes, but provides seconds of granularity to the file level.

Replication and cloning are good solutions, too, especially when talking about DR. However, just plain old replication and cloning won't necessarily cut it; you

need to find a solution that will have intelligence built-in for automatic failover and failback.

While this may sound more and more like the old HA story, it's more about continuous availability and IT resilience when you include the level of intelligence needed to achieve this type of outcome. Having the ability to adapt and shift as the environment requires and dictates is critical.



TRY ZERTO FOR FREE

<https://www.zerto.com/page/zerto-free-trial/>

If you're ready to give Zerto a whirl, click on the link and download the free trial. It only takes a few minutes to get started on your IT Resilience journey.

CHAPTER 4

Protecting Your Organization

A data protection strategy is much more than backup. The complexities come in when you introduce not only on-premises data to protect and recover, but, increasingly, the cloud and data locality. The level of effort to build a strategy grows exponentially when you begin to factor in hybrid environments, born-in-the-cloud applications, and data mobility.

One of the biggest challenges to modernizing backup and DR is getting management approval. Building a solid case for a continuously available and IT resilient data protection strategy typically involves three core elements:

1. Business impact use case
2. Technology solutions supporting the key objectives and mission of the organization
3. Gaining executive sponsorship and buy-in

The first step is to communicate to the business leaders in the language they can understand and comprehend.

This means to speak in business terms, not technology terms; that means you focus on the mission of your organization and use that to help drive the conversation with your business leaders. How will this new continuously available and IT resilient data protection strategy help the business meet the mission? That's where you start, and you build from there.



SEE HOW ZERTO WORKS

Get a demo for yourself and see how Zerto protects you from downtime through planned and unplanned outages, ensuring business continuity. Click on the link to schedule.

<https://www.zerto.com/page/get-zerto-demo/>

The second consideration is technology. The question here: Based on the organization's mission and key objectives, what technology solutions will best fill the gap as expressed by the business leaders? A further consideration here is how the technology solution will scale with business growth.



A solid data protection strategy provides the foundation for both availability and IT resilience, but to do this you must know what threats you're protecting against.

Once you identify these solutions, you'll need to make your case to executive management. You'll want to make sure you're using the original business use case as the guide to illustrate how the solution will help meet or exceed the requirements.

Your Data Protection Strategy Checklist

Here's a sample working outline you can use in your organization to build out your data protection strategy. You can also weigh your current strategy against this checklist.



- 1. Know what threats you're protecting against**
 - a. Why does it matter to the business, and why should it matter to IT?
 - b. What type of service-level agreements (SLAs) does the business require?
 - c. Which systems have the highest priorities?

2. Follow the cost impact

- a. What's the cost of protection?
- b. What's the cost of recovery?
- c. Do you have to make trade-offs?
- d. Cost of downtime to the business?
- e. Systemic risk?

3. A successful data protection strategy begins with the recovery in the event of a disaster. Start with the end in mind:

- a. Full disaster declaration
- b. Server disaster
- c. System/application
- d. File
- e. Security disaster
 - i. Malware, ransomware, virus, and so on

4. Agility is ability. Chances are you've determined your organization cannot live with a four- to six-hour SLA, which means the technology you more than likely will be looking to leverage is the continuous data protection model:

- i. Protects against random malware or ransomware attacks by minimizing the impact of data loss
- ii. Protects against accidental corruption by a DevOps team inadvertently introducing code that may negatively impact customer services
- iii. Protects against logical failures, not just disasters
- iv. Provides a sliding window for your recovery strategies

1. Seconds, minutes, hours, or more (i.e., Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO))

5. Planned vs. unplanned downtime

- a. What determines a disaster?
 - i. Unplanned downtime doesn't always mean a disaster declaration, as long as you and the business unit leaders are in agreement of the following:
 1. What is the pre-determined amount of time an outage or disruption (downtime) can last before a disaster is declared?
 2. What core application failure/disruption would constitute a declaration?
 3. Which core systems/platforms failure/disruption would constitute a declaration?
- b. Orchestration and Automation
 - i. You need a solution to help mitigate the challenges your team will face when declaring a full-on disaster
 - ii. This is equally as important with planned downtime. Automagically moving your workloads between your production and recovery sites with ease and confidence should be a critical selection criteria
- c. Planned downtime is just as an important topic as the unplanned
 - i. Planning for downtime may seem counter-intuitive, but having the opportunity to have production in an offline state is something many organizations do require at times. In the past this often meant

“kicking everyone off,” but in today’s always-on digital world, that’s no longer an option. But the idea of having a planned effort to maintain your production environment is as critical as being able to recover from a disaster.

1. Every organization must stay on top of security patches, OS updates/upgrades, hardware refreshes, and so on.
2. Taking down systems to perform these requisite tasks is unacceptable in an “always-on” digital world.
3. How do you test patches? Upgrades?
4. Has test/dev ever required production systems and data?

This is an abbreviated list, but covers the bases as related to protecting the business from loss. Once your strategy’s in place and the systems built, modifications, reviews, and so on become easier.

The final key component to this is executive sponsorship. You may have the best plan of attack to help the business, but without executive sponsorship and buy-in, it could take longer for it to be adopted. Once the business leaders see the value of this type of solution to support the overall mission of the organization, you're on your way!

A solid data protection strategy provides the foundation for both availability and IT resilience, but to do this you must know what threats you're protecting against. You should have this information from your prep work, but if not, make sure you ask the business leaders why the systems, data, and platforms are important to the business. Having a clear line of sight to the *why* will help you develop your *how* and define the *what* of your deliverable.

CHAPTER 5

Dare to Challenge the Status Quo

“The most dangerous phrase in the language is, ‘We’ve always done it this way.’” — Admiral Grace Murray Hopper

That statement is the most effective way to challenge existing resistance to change. Be careful at the start, though: If you’re in a situation where your technology and solution are working for your organization and your environment, you may not feel pressure to modernize.



You owe it to yourself to take a good look at your requirements, key objectives, and overall mission and identify if you’re able to achieve better availability and resilience for the business.

But even then, you should challenge whether or not you’re able to do things better. You owe it to yourself to take a good look at your requirements, key objectives, and overall mission and identify if you’re able to achieve better availability and resilience for the

business. This chapter will look at the solution Zerto offers to help IT meet those objectives that the business may be demanding.

Earlier chapters covered a great deal about backup and DR and how they've traditionally been separate functions. We also discussed how DR may be the new backup, or even if backup would be the the new DR in some cases. Well, Zerto decided to stop playing word games and just converge the two. Here are the key pillars of its converged DR and backup platform.

The Foundations of Zerto's Converged Disaster Recovery and Backup

Zerto's IT Resilience Platform™ converges DR, backup, and workload mobility whether on-premises or to, from, and between hybrid and multi-cloud platforms. CDP is the foundation of the platform, with built-in orchestration and automation. **Figure 4** provides an overview.

It provides IT leaders with simplicity, enterprise scale, and agile data protection to save time, resources, and costs. Analytics, with intelligent dashboards and live reports, provides complete visibility across multi-site and multi-cloud environments, giving companies the confidence to know business SLAs and compliance needs are met.

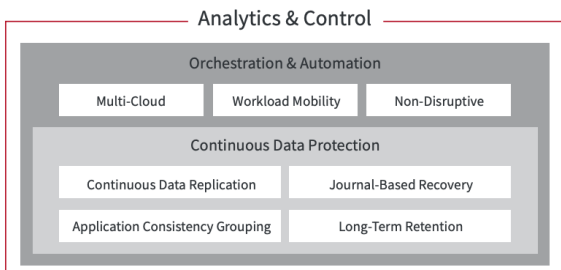


Figure 4: The core of Zerto's IT Resilience Platform.

At its core is CDP, which provides numerous advantages:

- a. **Replication.** Zerto's CDP delivers recovery point objectives (RPOs) of seconds by replicating every change being generated in near-real time. This is performed at the platform level, which enables the continuous capability by removing any production impact.
- b. **Scalability.** With a software-based platform, scaling the infrastructure to support DR processes is simple. As a new virtual host is added, simply install a new virtual appliance. Although Zerto scales to support very large environments, it provides the same granularity needed in environments of all sizes, with the same capabilities and no production impact.

c. **Application Consistency.** There are very few scenarios today where an application is run on a single VM. Instead, most applications have multiple VM dependencies. Using the traditional methods of protecting VMs individually results in some significant challenges to recovering your complete applications quickly.

Zerto differs with its Virtual Protection Group (VPG) capability. VPGs allow you to protect one or more VMs together in a consistent fashion, ensuring every checkpoint (point-in-time) inserted into the Zerto Journal is from the same point-in-time for all VMs within the VPG. This allows the consistent recovery of an entire application, and all its VM dependencies, to an in-sync moment. This is critical in the DR/business continuity and mobility use cases, but also unique within the backup space.

d. **Journal-Based Recovery.** All replicated changes are stored in a journal for up to 30 days, providing incredible recovery granularity through checkpoints inserted every few seconds. This reduces data loss to just seconds by enabling recovery of files, VMs, applications, or entire sites. They can be recovered either to the latest point-in-time, or, for example, when the VM is attacked by a virus or ransomware, to a point-in-time before the attack.

e. **Long-Term Retention.** Zerto's elastic journal concept really changes the game in data protection, merging both granular journal-based recovery with long-term repositories; this allows a continuous stream of recovery points, from seconds to years. Compliance standards often require you to keep, and ultimately recover, data for longer than 30 days. Long-term retention utilizes your existing journal to store data from any point for days, weeks, months, or even years, with no production impact because the journal lives on the secondary/DR site.

Elements of a Modern IT Resilience Platform

Modern solutions for modern architecture require agility and intelligence. Zerto takes the approach of combining both short-term journaling technology with long-term retention repositories, allowing rollback to any point-in-time across your on-premises or cloud environments with minimal data loss and downtime.

Non-disruptive recovery that's orders of magnitude faster than other solutions sounds like a utopia for any backup admin, but it's achievable with Zerto. On top of that is a rich intelligent index and search engine to make finding and restoring your files a breeze. This is

a key element that makes any modern solution truly functional across sites.

Something that makes backup and DR so difficult is the orchestration or execution of the recovery. Zerto's automated protection manages the recovery workflows you've created via their orchestration engine. This is an important feature, especially when you're faced with disasters and may not be thinking clearly. Letting Zerto do the work means you're less likely to make errors that may actually exacerbate the situation, rather than improving it.



WHAT IS 'IT RESILIENCE'?

It's a term you may not have heard before, but it has crucial implications for your business. Learn what it is, and how it will benefit your operations, by clicking on the link.

<https://www.zerto.com/it-resilience/>

Another feature of a modern solution is that it also plays well with others. Zerto has integrations with many of the existing backup targets today, including Exagrid, HPE StoreOnce, EMC Data Domain, and

others. This is crucial, because it means you don't have to buy anything new for your backup target, and you can repurpose your existing hardware solutions.

Because Zerto is a software company, it won't try to sell you an appliance. You run its product on your hardware. That makes life easy and eliminates another potential source of worry — no rip-and-replace required here.

The Choice Is Clear

The old methods of backup and DR were fine for the situations that existed in data centers in those days: The days when the amount of data created was more limited, and when most infrastructures were self-contained, and before things like ransomware existed. They were simpler days, hence the backup and DR needs were simpler.



We're now living in an always-on, completely connected world of explosive data growth, the proliferation of public, private, and hybrid clouds, and ever-increasing complexity inside and outside of your on-premises data center.

Say Goodbye To...



Figure 5: Zerto simplifies your disaster recovery and backup nightmares by eliminating multiple pain points.

Those days are gone. We're now living in an always-on, completely connected world of explosive data growth, the proliferation of public, private, and hybrid clouds, and ever-increasing complexity inside and outside of your on-premises data center.

It all adds up to the need to move on from old-school thinking about how you protect your valuable data and systems. The old DIY mindset may have been adequate back in the prehistoric era of computing, although even then it could be incredibly time-consuming.



Only by moving away from the “we’ve always done it that way” mentality can you achieve true IT resilience.

What's needed now is a comprehensive view of your environment and its protection. That vision goes beyond simple backup and DR; now, it's about IT resilience, a holistic response to today's protection challenges.

That's what Zerto provides: The Zerto IT Resilience Platform is a modern solution for a modern architecture, one that runs at the new speed of business. Only by moving away from the “we've always done it that way” mentality can you achieve true IT resilience. **Figure 5** provides a glimpse of the problems Zerto eliminates.

Consider also the advantages of going with a single vendor for DR and backup. It means “one throat to choke” to start with; no hassling with multiple vendors, each pointing the finger at the others. Then there's the headache of juggling multiple support contracts, along with the licensing, installation and management of a host of niche products (not to mention the learning curve involved for all those offerings).

It's about simplifying your environment, and letting someone else do the heavy lifting so you can get back to driving innovation and bottom-line success for your organization.

This whitepaper on the changing nature of backups is a good place to start: <https://www.zerto.com/page/future-of-backup-from-periodic-to-continuous/>