ZONES
First Choice for IT™

Microsoft Azure VERITAS

MOVE YOUR DATA

OUT OF THE DARK AGES

Microsoft Azure VERITAS

In this day and age, most companies have at least a basic understanding of the need for a disaster recovery (DR) framework.

IT leaders know that to keep business humming along without interruption, they need to have backup plans in place for when access to their data is impeded. This way, even when something goes wrong, employees can stay productive and revenue can keep flowing.

But in a lot of cases, IT decision-makers have only a basic understanding of DR best practices. Many think it's enough simply to duplicate their employees' data, back it up on another server, and call it a day. In reality, companies today need to do more.

In this guide, we'll talk about how. Let's go through the pain points you may encounter in the DR planning process and discuss strategies you can use to combat them.

How traditional DR misses the mark

Traditionally, companies have opted for a DR process that's fairly formulaic – rarely do they take into account their specific organizational needs or the nuanced ways in which DR best practices are evolving. They end up with strategies that have a wide range of shortcomings:

- **They're too complex.** Failover requires using a plethora of applications, servers, networks, and storage systems in disparate locations.
- **They're too slow.** Recovery time objectives (RTOs) can range from hours to days. Data backup processes are often disjointed, leading to unnecessary friction and delays.
- **They're unreliable.** Basic processes are subject to change when they're under sloppily designed DR frameworks, and compliance checks are often infrequent and inconsistent.
- **They're inefficient.** Does your business know how much of its data is valuable and is actually being used? If you don't, you might be wasting large amounts of money on backing up useless data.

That last point may be the most important one, as most companies have vast stockpiles of data that they're using for nothing at all. According to data compiled by IDG Research Services, only about 28% of companies' data has any real value to it. This means, in all likelihood, that the majority of the data you're storing and maintaining is going unused, simply taking up space. This is called "dark data," and it's a major issue for companies today.

Being aware of hidden DR costs

In the past, investing in disaster recovery has been an incredibly costly aspect of managing IT – and decision-makers have done little to address the issue. Often, they simply write off an expensive DR infrastructure as a cost of doing business, rather than stopping to consider ways they can trim the fat.

It turns out, their legacy DR frameworks often have a host of hidden costs involved. Consider the following:

- **Extra maintenance:** The typical DR server has to be running at all times, passively waiting to be called upon to serve traffic. This effectively means enterprises have to maintain twice the infrastructure for the same amount of operational activity.
- **Extra staffing:** Businesses must keep IT staff members ready to deploy in both locations – their primary data storage site and their backup. They must have team members prepared both to complete maintenance on the primary server to make the passive DR server active.
- **Downtime costs:** Even with a DR approach that works relatively well, it can take hours to recover your information and restore business operations. The Ponemon Institute has estimated that the typical enterprise loses revenue during this time at a rate of $7,900 per minute.
- **Lost business:** As your IT screeches to a halt, you're going to have servers that are over capacity, leading to delays in all different aspects of your business. The customer experience is going to be affected, and that will lead some people to walk away.

When you try to run a full-capacity redundant DR system in a secondary site, the costs can be numerous. Some are obvious; others are more subtle. It's important to find strategies for keeping these costs under control for your organization.

Making cloud DR work for your organization

Moving your DR framework to the cloud is a great way to make the process more efficient and eliminate many hidden costs. It's not rocket science, either – it's simply a matter of taking a close look at your operations and determining a path forward that fits your organization's needs. In many cases, you don't even need to rewrite DR applications when moving from on-premises to the cloud. You can make the move with minimal friction.

The objective should be to move your existing apps and data with the least amount of refactoring work or process change as possible. Some companies look to revolutionize apps and processes and hire a lot of new people as their first resort; the truth is, this is rarely necessary. You don't want the cloud transition to become a multi-year process that drains your organization's time and money. In fact, that would defeat the purpose, since the whole goal here is greater efficiency.

A simple baby step here is to use the cloud as a DR site first. Rather than change your primary data storage site and disrupt your existing business, you can use the cloud for backup, getting a small taste of how it works without introducing any major changes to your employees' basic work processes.

Choosing the right cloud for you

Of course, it's not enough simply to throw your data into the cloud – any cloud. Not every cloud environment works perfectly for DR, and you'll need to choose one that's a good fit for your organization and its needs. Understanding the nuances is important here. For example, a cloud-native management stack will behave differently from a vSphere management stack.

The ideal cloud DR solution should be able to tackle a wide range of objectives for your business, such as:

- Ingesting data from vSphere environments that you have on-premises.
- Storing your data with backup-style data reduction and retention on low-cost media.
- Avoiding VM conversion and management stack differences by leveraging software-defined data centers (SDDCs).
- Booting directly and bringing up workloads from S3 backups onto SDDCs as live data stores.

Don't fret if you don't fully understand these objectives, or you don't know how a cloud DR system can help achieve them. Working with the right business partner can put you on the fast track to success.

ZONES

What we do at Zones

Remember: It's not enough simply to have a backup server and keep your data stashed there, just in case a disaster scenario comes up. The goal here is to have a DR strategy that's smart, cost-effective, and tailored to the needs of your business.

At Zones, we work every day to help clients put such systems in place. We have a proven track record of designing and implementing reliable DR solutions that help increase recovery speed and cut costs for organizations everywhere. We also have strong partnerships with a wide range of tech OEMs, which means we can always source the hardware and software you need to run your system right.

It's time to move your DR out of the dark ages. There's a better way to keep your data secure and maintain continuity for your business. With our help, you can find it.

First Choice for IT.™

Visit our website below to learn more ways Veritas, Microsoft, and Zones can benefit your business.

ZONES
First Choice for IT™

Visit **zones.com** or call **800.408.ZONES** today.