

WHITE PAPER

Building an innovative cloud disaster recovery plan

Best practices, tools, and templates

Introduction

The rise of virtualization as a business tool has revolutionized the way companies operate today. By decoupling data from the underlying physical hardware, businesses are freed from the limitations imposed by a need to keep their data within arm's reach — it can now be stored anywhere in the world that it makes operational and regulatory sense to do so. However, while data is more mobile than ever, many organizations do not have a comprehensive plan in place for how to recover their data in case of natural disaster, ransomware, and other unexpected events.

Cybersecurity Ventures predicts¹ cybercrime damages will cost the world \$6 trillion annually by 2021, up from \$3 trillion in 2015. They also project cybercrime “will be more profitable than the global trade of all major illegal drugs combined.” In such a climate, an attack is not a matter of if, but when. The secret to recovering from a ransomware breach or other disaster lies in preparation. It is no longer sufficient, or even wise, for an IT manager to cross their fingers and hope that their company won't be attacked — they must take active, purposeful steps to prepare for the inevitable. In today's 24x7x365, always-on world, a solid disaster recovery (DR) plan is no longer a “nice to have” — it should be a critical part of every company's data management strategy.

What disaster recovery is...and what it isn't

Don't be fooled: Simple backup is not the same as disaster recovery. Many backup vendors claim to provide DR as part of their solution, but unless they are optimized to provide the fast recovery of all enterprise data and applications, they aren't doing enough. A simple backup app may only update during off-peak times (typically at night) and may store data on legacy media such as tapes in slow-to-access, remote locations. If you needed quick access to mission-critical data, such as following a ransomware attack, it could be weeks before you got it.

“ Only 22% of businesses have high confidence in their DR plan in the event of an emergency. ”

— TechTarget Research

On the other hand, a backup/restore solution that includes effective DR stores data in one or more separate locations. Most importantly, the enterprise can restore operations quickly and easily, including spinning up VMs off-premises or in the cloud to run applications in order to maintain business continuity. With a comprehensive DR strategy in place, companies can achieve a recovery point objective (RPO) of less

than 24 hours and a recovery time objective (RTO) of mere minutes, restoring any amount of data — from an individual file to a complete virtual machine.

Emerging methods for implementing a DR strategy

Historically, on-premises infrastructure was used for backup and disaster recovery, but the cost limited DR to only the most mission-critical applications. With the transition to the cloud, these costs have been dramatically reduced, allowing more applications to be included in the DR planning. Effective DR, however, requires remote capabilities, typically leveraging the cloud, and there are a number of ways to make this happen:

- Duplicate a primary data center precisely, right down to the last cable. However, managing a secondary data center to ensure proper failover is both time-consuming and costly, particularly for VMware environments. If the secondary location is near the primary data center, it may be affected by the same disaster (for example, a power outage or earthquake).
- Use an appliance hosted in the cloud such as an Amazon Snow Family device that can be physically shipped from an Amazon facility for on-premises recovery. This kind of solution is typically marketed for extremely data-intensive analytics use cases.

- Manage self-service recovery using enterprise-owned or -leased infrastructure hosted in the cloud that enables restoration of data, applications, or both.
- Engage a disaster recovery as a service (DRaaS) provider. This resembles self-service recovery, but the DRaaS provider eliminates any IT overhead and infrastructure costs. It backs up and restores your company's operations via a cloud service provider such as AWS.

The advantages of cloud-native backup and DRaaS

A comprehensive answer to the challenges of merging backup with DR is a solution that is architected with the cloud in mind, leveraging the advantages of the

Essential requirements for effective DRaaS:

- SaaS-based
- Cloud-native
- Source-side deduplication
- Separation of data and metadata
- Unified dashboard for management
- Encrypted architecture

public cloud in terms of instant availability and long-term retention, while being optimized for bandwidth reduction and minimizing the impact on end users. Effective DRaaS uses technologies such as global deduplication of virtual data, ensuring that only one copy of each file is maintained. This can allow for bandwidth savings of up to 80 percent and ensures that even remote office locations, potentially with suboptimal WAN speeds, can still be effectively protected. Leveraging the efficiencies of public cloud vendors such as AWS lets companies take advantage of tiered storage, with data sorted into hot, warm, and cold storage depending on retention and recovery needs. This provides for long-term storage at an affordable price.

Designing an effective DR plan

While it should be obvious that a comprehensive DR plan is an essential requirement for any modern company, the path toward achieving that goal may not be so clear. To help prepare a DR plan for your company, here are the four essential steps in the process, as well as the tools for helping you get it right:

Step 1: Perform a business impact analysis (BIA)

Any comprehensive backup and DR planning process must begin with an accurate assessment of your current virtualized environment. How much data are you currently managing? Where is it located? How critical is it to your business operations? Once you have completed this step, the vital question becomes: How

would a disruption of this data access impact your business? Think in terms of business opportunities lost, time spent restoring files and rebuilding databases, and so on. This step is crucial to the process because it will inform every decision you make from here on, including how much you will budget for the solution. Obviously, it makes sense to invest more to protect the data that is vital to your company's ongoing success.

Step 2: Perform a risk assessment

A BIA (business impact analysis) is essential for looking inward at your business-critical data and the impact on your business of any disruption to it. A risk assessment, on the other hand, is focused on potential external situations that could negatively impact your business, and the likelihood of such situations occurring. These could include natural disasters (such as tornadoes and floods) as well as manmade events (such as power outages and terrorist acts). This will allow you to gauge the probability that your DR plan will one day need to be activated. When you're preparing a risk assessment, be sure to leverage all available records to assess the threat of disaster.

Such sources might include (but are not limited to):

- Company records of disruptive events
- Employee recollection of disruptive events
- Local and national media records
- Local libraries
- First-responder organizations
- National Weather Service historical data
- U.S. Geological Survey maps and other documentation

- Experience of key stakeholder organizations
- Experience of vendors doing business with the firm
- Government agencies such as the Federal Emergency Management Agency (FEMA), Department of Homeland Security, and U.S. Department of Energy

Use this *comprehensive guide* for preparing a risk assessment for your company, as well as *FEMA's Risk Assessment Table*.

Step 3: Design a risk-management strategy

Once you have identified the critical elements of your virtualized landscape, the business impact of any disruption to it, and the likelihood of disaster, the question becomes: What can I do to mitigate the damage? This is when you need to decide upon a specific solution for backup and DR of your business-critical data. Although there are sure to be multiple, possibly contradictory, demands, a few elements that will inform your decision might include the following:

- RPO (how much data you can afford to lose)
- RTO (how quickly your business needs to be back in operation)
- Data residency laws (where your data can legally be stored)
- Budget for implementation

These considerations will allow you to calculate the ROI of competing vendors and select the one that best fits your organization's requirements. As mentioned earlier, utilizing the public cloud for DR can provide savings in all of these areas.

Step 4: Configure and test (and keep testing!)

It should be obvious, but you need to know whether your backup and DR solution is configured correctly before you actually have to use it. The only way to achieve that is by regularly testing your DR. A cloud-native backup and DR solution allows you to immediately spin up your virtual machines in the cloud for development testing (dev-test) purposes. Ensure that your VMs operate as expected and that data has been backed up in compliance with the RPO you have set.

Bear in mind that testing is not a “one and done” affair — it should be a regular, and ongoing, part of your work. Set a cadence that makes sense for your organization and stick to it.

The key takeaways

Disaster will strike

Ensuring that your company has a comprehensive DR plan in place is every bit as important as ensuring that your production environment is operational. There can be no doubt that organizations today are under threat of data loss like never before — and even with the best precautions, it is only a matter of time before an attack happens.

Be prepared

The solution lies in dutiful preparation and planning. Companies that have foreseen the inevitable and, far from running scared, have confronted the dangers and designed a strategy for overcoming them, will prove to be far more resilient, and more profitable, organizations in the long term.

Learn how to execute your DR plan in the cloud at www.druva.com/cloud-disaster-recovery/.

Druva™ delivers data protection and management for the cloud era. Druva Cloud Platform is built on AWS and offered as-a-service; customers drive down costs by up to 50 percent by freeing themselves from the burden of unnecessary hardware, capacity planning, and software management. Druva is trusted worldwide by over 4,000 companies at the forefront of embracing the cloud. Druva is a privately held company headquartered in Sunnyvale, California, and is funded by Sequoia Capital, Tenaya Capital, Riverwood Capital, Viking Global Investors, and Nexus Partners.

Visit [Druva](#) and follow us [@druvainc](#). Find Druva in AWS Marketplace.

GET STARTED

