10 Tips for developing an **AWS disaster recovery plan**

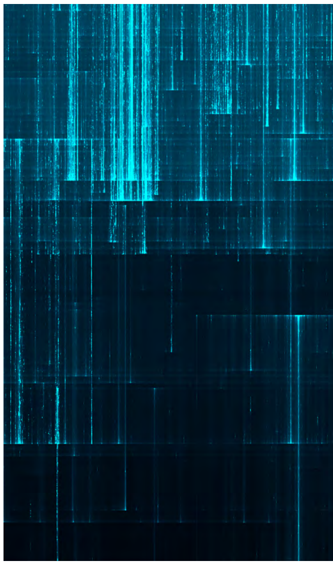

Introduction

Amazon Web Services' (AWS) cloud-based platforms provide users with a flexible, easily scalable solution to meet their computing needs. It also eliminates the need for users to invest in a costly physical computer infrastructure of their own. So, when it comes to disaster recovery, this versatile, decentralized system is an ideal solution for your business.

However, while AWS might provide the necessary features for disaster recovery, these tools will prove useless in the absence of a comprehensive strategy. With that in mind, here are 10 tips for developing an AWS Disaster Recovery Plan, along with some information about how third-party vendors like Druva CloudRanger can help with the plan's implementation.

Checklist

The following checklist describes 10 simple steps that you can — and should — take today to set up and automate a single-click AWS Disaster Recovery plan for your organization.

1. Don't equate backups with disaster recovery

While backing up your data at regularly scheduled intervals is essential, it's important to not equate the backup process with disaster recovery. As software engineer Scott Ross explains, "Disaster Recovery is not backups," but rather, it is "the process, policies, and procedures … related to preparing for recovery or continuation of technology infrastructure" in the event of a crisis. In other words, simply backing up your data won't be of much help unless you have a process in place to quickly retrieve and put it to use.

2. Prioritize downtime costs vs. backup/recovery costs

As with any successful strategy, an AWS Disaster Recovery Plan must be tailored to meet your company's specific needs. As such, choices will have to be made between the amount of money spent on backup and restoration of data versus the amount of money that might be lost during downtime. If your company can withstand a lengthy outage without hemorrhaging cash, a slower, less expensive backup and recovery option might make sense. But if even the slightest amount of downtime wreaks havoc on your bottom line, more expensive methods such as an AWS-based duplicate production environment might be required.

3. Determine your RTO

Before you can weigh downtime costs against the cost of data backup and recovery, you must determine your company's recovery time objective (RTO). Essentially, this is the maximum length of time your recovery process can take to get everything back online without inflicting unacceptable losses for your business.

4. Determine your RPO

Another factor that must be determined is your company's recovery point objective (RPO). This is the maximum amount of data loss your company is willing to accept as measured in time. For example, if a disaster strikes two hours after your company's last backup, any changes to your data that occurred within that two-hour window could be lost. If that is acceptable, then an RPO of two hours is a good fit. But if losing two hours' worth of data would cause major headaches, your RPO must be narrowed by scheduling more frequent backups.

5. Choose the right backup strategy

As mentioned above, regular backups are only one part an effective AWS disaster recovery plan. Nonetheless, they are an extremely important component. That's why choosing the right backup recovery plan for your business is vital. Even though you've already settled on a cloud-based solution, you will have to choose between various backup options such as using Amazon Machine Images (AMI) or Amazon EBS snapshots.

6. Choose the right backup management system

You'll also have to choose between using in-house scripting to create a centralized backup management system from scratch or using a third-party service like Druva CloudRanger to streamline the process. To learn more about which options are right for you, check out our blog post on choosing the best Amazon EC2 backup strategy to meet your needs.

7. Identify mission-critical applications

After determining your company's RTO, RPO, and preferred a backup strategy, it's time to choose which type of AWS disaster recovery plan is right for you. And depending on which option you ultimately choose; it may also be necessary to identify and prioritize mission-critical applications. Some of the most common methods include:

Backup and restore: A simple, cost-effective method that utilizes services such as Amazon S3 to backup and restore data.

Pilot light: This method keeps critical applications and data at the ready so that it can be quickly fired up should disaster strike.

Warm standby: This method always keeps a duplicate version of your business' core elements running, resulting in a nearly seamless transition with very little downtime.

Multi-site solution: Also known as a Hot Standby, this configuration leaves almost nothing to chance by fully replicating your data/applications between two or more active locations and splitting traffic/ usage between them. In the event of a disaster, traffic is simply routed to the unaffected location, resulting in no downtime.

8. Implement cross-region backup flexibility

As with traditional methods of backup and recovery, geographic diversification of your data is essential for your AWS disaster recovery plan. If a natural disaster or man-made catastrophe brings down your primary production environment, having a backup stored in the same building, or even the same region, makes little sense. Luckily, the global reach of AWS makes geographic diversification a breeze. If your primary AWS services are knocked offline, you can rest assured that your DR plan can be implemented using backup data that's been safely stored a world away (literally).

☑ 9. Test and retest your AWS disaster recovery plan

The best-laid plans of mice and men often go awry. Even the most detail-oriented AWS disaster recovery plan has the potential to fail when put into actual practice. That's why it's important to constantly test and retest your plan for flaws. Thanks to AWS' ability to create a duplicate environment, you can test your plan using real-world scenarios without jeopardising your actual production environment.

☑ 10. Consider a third-party service for your AWS disaster recovery

Once you've created your company's AWS disaster recovery plan, its implementation can seem like a daunting task. But with the help of a third-party service like Druva CloudRanger, it doesn't have to be. Aside from streamlining your backup process, Druva CloudRanger can also simplify disaster recovery by copying snapshots and AMIs across AWS regions for DR purposes, and quickly restore them to unattached volumes, attach them to existing instances or file-level restore.

Conclusion

Unfortunately, getting hit with a disaster event like ransomware has become a case of when, not if. Storing any significant amount of business data without a DRaaS solution in place is akin to a game of Russian roulette: the longer you play, the greater your odds of calamity become until it is eventually all but certain. By following the steps in this checklist, you can set up and automate a comprehensive DR plan on AWS cloud to allow for a single-click failover of your mission-critical data and have your organization up and running again within minutes.

Discover our capabilities with a FREE 14 day trial.

Sales: +1 888-248-4976 | sales@druva.com

Americas: +1 888-248-4976
Europe: +44 (0) 20-3750-9440
India: +91 (0) 20 6726-3300

Japan: +81-3-6890-8667
Singapore: +65 3158-4985
Australia: +61 1300-312-729

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 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.