



8 tips to simplify AWS backup and recovery

Why choose Druva for data protection of native AWS workloads

Overview

As organizations migrate workloads to AWS cloud, or expand their existing environments, many struggle with data protection and find it challenging to keep pace with rapid data growth. Data is severely critical to your business, and the possibility of data loss puts your organization at risk, so it's important to properly prepare for any potential failures.

The AWS shared responsibility model highlights that customers are responsible for managing their own data. Having an effective data protection strategy in place to maintain business continuity is essential. The most effective defense against data loss is to regularly back up your business-critical resources and applications. Creating a reliable strategy for backup and disaster recovery, protects against the risks of accidental deletion, malicious intent or ransomware.

A centralized approach helps cloud-focused IT teams effectively manage their data, mitigate loss, while maintaining compliance and optimizing costs. In this white paper, you'll learn the key eight tips to simplify backup and recovery for native AWS workloads, that Cloud Administrators, or Cloud Solution Architects should take into consideration when they set up a data protection strategy for applications running in AWS.

Table of contents

1. Automate backup policies	3
2. Customize retention	3
3. Copy backups to another AWS region or account	4
4. Run application consistent backups	5
5. Evaluate file-level recovery options	5
6. Implement disaster recovery plans	6
7. Determine RTOs and RPOs	7
8. Alerts, notifications and integrations	8

1. Automate backup policies

A recent Gartner press release provides, “The worldwide public cloud services market is forecast to grow 17% in 2020 to total \$266.4 billion, up from \$227.8 billion in 2019, according to Gartner, Inc.”¹. To us, this highlights that organizations are going to be managing more data and more applications in the cloud, which will increase the workload for Cloud Solution Architects and SysAdmins.

In the past, it may have been a quick fix to have these members of your IT teams write an AWS Lambda script to automate backup jobs, but as you scale your AWS environment, this becomes harder to manage, and even harder to maintain effectively. Managing data protection at scale, where many of your AWS accounts share common requirements, demands a simplified approach to managing backups.

As a result, organizations should be looking at third-party tools, who specialize in cloud-native data protection, which can allow Cloud Solutions Architects, SysAdmins and DevOps teams to automate backup policy orchestration to simplify day-to-day management, and free up their time to focus on value-add activity within their team.

Druva provides an enterprise-level backup and disaster recovery solution for AWS, designed to automate processes and save IT teams both time and money. With Druva’s global backup policies, you can easily manage thousands of resources across AWS accounts within your organization. By using the rule based selectors option, you can choose to include or exclude resources by region, account, resource ID, tags, VPC or subnet. Also, by using the resource scheduling feature you can automate custom stop, start, and reboot schedules for Amazon EC2 and Amazon RDS instances to optimize resource utilization and reduce monthly AWS bills.

Challenge:

In-house scripted solutions aren’t reliable as they come with many disadvantages such as ongoing maintenance, scale, inefficiency, and risk of staff turnover. This increases the risk of data loss.

Solution:

Choose a cloud-native solution that can automate policy orchestration, enabling your team to be more efficient and eliminate the need for manual scripting, patching or upgrades.

2. Customize retention

As organizations continue to migrate to AWS to enable their cloud-first strategy, they’re challenged by a constant increase in monthly cloud bills due to unused or underutilized resources. Since more AWS resources are being used, it can become harder to manage and monitor data protection. If you have no policy-based orchestration of snapshots, as previously mentioned, this can result in tedious manual scripting processes.

To manage resources efficiently, companies need to adapt to the growing number of applications and data retention requirements and plan accordingly. Your cloud center of excellence team will be looking to automate the creation, retention, and deletion of Amazon EBS snapshots, Amazon RDS DB snapshots, and Amazon Machine Images (AMIs). This will free up internal resources, ensure compliance, and allow for simplified management as you scale.

Third-party solutions like Druva empower you to customize snapshot retention based on a traditional grandfather, father, son (GFS) approach all within a standard backup policy. This means you can choose custom retention periods for your backups, and ensure you have weekly, monthly, and yearly copies available to meet your internal or regulatory compliance mandates.

¹ Gartner Press Release, “Gartner Forecasts Worldwide Public Cloud Revenue to Grow 17% in 2020,” 13 November 2019. <https://www.gartner.com/en/newsroom/press-releases/2019-11-13-gartner-forecasts-worldwide-public-cloud-revenue-to-grow-17-percent-in-2020>

Challenge:

The on-demand model of the cloud can easily lead to an increase in costs due to new resources, related snapshots, and AMIs that are created without retention periods for deletion. This means your monthly AWS storage bill will increase as a result of these orphaned snapshots.

Solution:

Formulate an automated backup strategy that follows predefined rules of creation, retention and deletion. This optimizes data lifecycle management and ensures lower monthly AWS storage bills.

3. Copy backups to another AWS region or account

The entire point of disaster recovery (DR) related backups is to ensure business continuity in the face of a natural or man-made disaster. As such, storing your backups in close geographic proximity to your normal production environment leaves all of your data exposed to the same threats. Hypothetically, if your business was located in the path of the storm, where would you rather have your backup data: in Virginia, a state that is also in the path of the hurricane, or in Oregon, which is located safely on the opposite side of the country?

Clearly, distance plays a key role in keeping your business online in the event of a disaster. And with AWS cross-region disaster recovery, you have the ability to store your backups in numerous locations around the world, which makes putting distance between you and your important data easier than ever. Even if an entire AWS region were to be knocked offline, your business could continue to operate with little-to-no disruption as long as a proper, geographically-diverse recovery plan is in place.

Also, after some high-profile attacks, companies have become more strategic and want to ensure their backups are isolated from the primary production account. Many organizations are now looking to ensure their data and snapshots are stored in an additional AWS account. This means your backups are 'air-gapped' and protected separately from the origin AWS account to ensure backup data is safe from any potential security breaches and or compromised accounts related to malicious intent or malware.

AWS provides out-of-the-box capabilities to copy snapshots between different regions, but a third-party solution provided by an AWS Storage Partner such as Druva, can automate this process, and allow administrators to set the regions that they want to use as a target DR site. Simultaneously, Druva also offers the same functionality for cross-account requirements, including encryption.

Aside from streamlining your backup process, Druva enables administrators to easily take additional copies of your snapshots and AMIs — and store them in other AWS regions or another AWS account within the AWS data center network. Druva's centralized view across AWS environments with file-search functionality allows users to search across all accounts and snapshots to find specific files and restore them quickly and easily.

Challenge:

Relying solely on your source region for business continuity in a disaster situation makes your data vulnerable. This can leave all of your data exposed to the threat of a natural disaster, accidental deletion, or an AWS region outage.

Solution:

Choose a third-party solution that leverages the native AWS global infrastructure. It should be able to automate cross-region and cross-account backup copies to your chosen DR site.

4. Run application consistent backups

For users who are attempting to use EBS snapshots with a Windows server instance, steps must be taken to ensure that files in use during the snapshot process are not excluded, which would lead to an incomplete backup. If an Amazon EC2 instance is in use when the backup is occurring, any data associated with that instance will most likely be excluded, which could lead to major problems. Not only would that particular backup be corrupted, but due to the incremental nature of EBS snapshots, future backups will also miss the data.

That's where Microsoft VSS comes into play. The service creates a backup of volumes even when they are in use. It does so by quickly freezing the volume and briefly suspending the write operation. It buffers any new data, so it can be added after the backup is completed. The process creates a consistent point-in-time snapshot of the volume – ensuring that data isn't missed because it is in use. As such, when used in conjunction with EBS snapshots, VSS eliminates the risk of any missing data during an incremental backup.

One of the main benefits of using Druva is the ability to easily schedule and automate application consistent backups of your data without the need for in-house scripting. For organizations that are running instances of Windows server, they have the option to utilize VSS as part of their backup process. By simply clicking the “Consistent Snapshot” checkbox when creating or editing a backup policy, Druva's system will automatically generate consistent, application aware snapshots for any Windows servers with VSS installed – and will do so without the need for any scripting from the users. As a result, worrying about missing key data because it is in use during a backup is a thing of the past.

Challenge:

When creating standard Amazon EBS snapshots you cannot turn off the instance during backup. This might exclude data that has been cached by any applications or the operating system. The result? Inconsistent data – unless you create application consistent snapshots.

Solution:

Choose a third-party solution that utilizes cloud-native technologies like VSS for EBS snapshots, and AMIs of Windows EC2 resources, so you can achieve application consistent backups without maintenance windows.

5. Evaluate file-level recovery options

One of the most useful features offered by AWS is the ability to utilize EBS snapshots for your company's backup process. Essentially, these snapshots capture incremental block-level changes to EBS volumes, eliminating the need to repeatedly perform a complete backup. That being said, EBS snapshots can still leave a lot to be desired when it comes to individual file-level recovery.

The ability to recover a single image file or database file rather than an entire volume or instance is important, especially if a user is simply trying to confirm that a volume contains the files they wish to restore. But thanks to the platform's rudimentary user interface, file-level recovery for AWS EBS snapshots can often prove to be a frustrating task.

The most basic method of file-level recovery for EBS snapshots is also the most time-consuming. It involves manually finding the snapshot containing the file you wish to recover, creating a new volume of it, and attaching that volume to a newly-created instance. Once this is completed, you should be able to find and restore the file in question. But afterwards, you'll need to backtrack and detach the volume from the instance in order to delete it.

While manually restoring an individual file from an EBS snapshot might not be the most difficult task for an IT professional, the process can be somewhat intimidating for the average user. And even for a trained professional, manual recovery can be tedious, especially if it needs to be done on a regular basis or requires multiple files from different volumes. In order to save time,

it is possible to automate file-level recovery for EBS snapshots. However, this will require in-house scripting, which can also prove to be a time-consuming process that's beyond the capabilities of an average non-IT professional.

While both manual and automated file-level recovery methods for EBS snapshots have their drawbacks, third-party providers like Druva can help by replacing the clunky AWS user interface with an easy-to-use console. Druva greatly simplifies the restore process, allowing non-IT professionals to carry out file-level recovery without the need for in-house scripting.

In addition, Druva's file-level search uses metadata as a guidepost, so users are no longer forced to search through individual EBS snapshots one-by-one. Instead, you can search for specific files across all EBS snapshots simultaneously, greatly reducing the time it takes to find and recover a specific file (or files).

Challenge:

Native tools for AWS offer limited restore functionality, which can be time consuming and cumbersome. In a recovery situation, you want to be able to restore at an instance, volume, and file level to provide greater flexibility and reliability.

Solution:

Look for a solution that provides file-level recovery as well as volume and instance-level recovery. To save time during the recovery process, make sure you choose a solution that allows for granular file search to greatly reduce the time it takes to find and recover a specific file (or files).

6. Implement disaster recovery plans

When it comes to disaster recovery, having recent backups of your company's data is essential. But it's important to understand that simply scheduling regular backups of your data stored on Amazon EC2 and the attached EBS volumes is not enough. It's only one part of a much larger process. After all, what good is your backup data if you have no means to access it when it's urgently needed?

In the end, even the best DR plans may run into trouble when faced with the unforeseen problems of the real world. So once you've settled on a disaster recovery method and have everything up and running, it's extremely important to test its effectiveness. Fortunately, the fact that cloud platforms allow users to create duplicate production environments means testing your DR plan against various scenarios has never been easier. The more you test and retest your DR plan, the better prepared you'll be in the event of an actual crisis. Remember, an ounce of prevention is worth a pound of cure.

While disaster recovery planning is often handled in-house, more and more businesses are turning to third parties to implement and maintain their DR plans. Just as insurance policies allow businesses to mitigate the risk of property damage and financial loss, treating disaster recovery as a service (DRaaS) allows companies to focus on core issues. Meanwhile, businesses can have peace of mind – knowing their data is safe. So it's no surprise that surveys have shown IT professionals who utilize DRaaS are significantly more confident in their company's disaster recovery planning than those who don't. And while utilizing a third-party vendor for DRaaS might not make sense for all businesses, it's certainly an option worth exploring.

With Druva, you can set up and automate a comprehensive DR plan on AWS cloud to allow for single-click failover of your mission-critical data. You'll have your organization up and running again within minutes.

Challenge:

By relying solely on backups in your source AWS region, could your business continue to operate if an entire AWS region went down as a result of an outage or a natural disaster?

Solution:

Choose a solution that leverages native AWS technologies to copy snapshots and AMIs to additional AWS regions or secondary AWS accounts as part of your disaster recovery strategy.

7. Determine RTOs and RPOs

Determining your company's recovery time objective (RTO) and recovery point objective (RPO) is a bit more complicated than it seems at first glance. After all, when asking someone about the amount of downtime and data loss they're willing to absorb, "as little as possible" is the natural response. But as with most things in life, there is a trade-off. And in this case, that trade-off is between data, time, and money.

When it comes to saving time, or determining your RTO, it is technically possible to implement a recovery method such as a "warm standby" or "multi-site solution" that greatly reduces or nearly eliminates the amount of downtime a company experiences in the event of a disaster. For example, if your company's website is knocked offline by a hurricane in Texas, but traffic can be immediately syphoned to an exact replica being hosted in Maine, why wouldn't you choose to have an RTO of virtually zero minutes? The answer is money.

Constantly running two identical sites simultaneously is going to cost you a lot more than simply backing up your data and restoring it as needed. In general, the faster you want to be back online after a disaster, the more you are going to have to pay to make it happen. For some companies, even a slight outage can have a major impact on the bottom line, and in this case, it makes sense to bite the bullet and set an RTO that is as short as possible. But for other companies, such a costly process makes no financial sense, and a longer RTO should be set.

The same concept applies when determining a company's RPO. For instance, a company could set a backup schedule that prevents almost all of its data from being lost in a disaster. But data backups and data storage aren't free. As with RTO, the more ambitious the RPO, the more money it's going to cost. If losing two hours worth of data can cause your company major headaches, perhaps an RPO of one hour or less is truly necessary. But if losing a day's worth of data isn't going to break the bank, then perhaps an RPO of 12 or 24 hours makes more sense.

Whatever the case may be, it's up to you to determine the right metrics to meet your company's needs. Once you've determined your company's RTO and RPO, other aspects of your DR plan should start falling into place. And thanks to Druva's easy-to-use solution, this allows organizations to validate their RTO and RPO through a tested execution of an end-to-end process.

With Druva, you can quickly and easily determine what gets saved and where, and schedule regular data backups at predetermined intervals to ensure your RPO is met, whatever it may be. And in the event of a disaster, Druva makes it easy to retrieve and restore your data, ensuring that your RTO also stays on schedule. Best of all, Druva does all this and more without the need for scripting.

Challenge:

Constantly running two identical sites simultaneously for DR is cost-prohibitive to almost all organizations. As a result, you need to consider your budget, as well as your RTO and RPO requirements for DR readiness.

Solution:

Choose a cost-effective solution that allows your organization to validate RTO and RPO through the actual execution of an end-to-end process.

8. Alerts, notifications and integrations

Whenever a critical backup job fails, having the right information at the right time, can make a significant difference between success and failure. Making the relevant people aware of the situation, whether they are customers, team members, or senior management, is essential. The sooner a potential issue is recognized, the sooner it can be resolved.

However, many organizations still rely on staff and manual processes to monitor backup jobs and distribute updates to customers and personnel. With so much information being stored in the AWS cloud, it can also be difficult to identify specific issues in a timely fashion, causing a bottleneck or delay in resolving the problem. Predefined rules, events or triggers, can vary across AWS accounts within your environment. Automatically sending a critical alert to an admin user can be of huge benefit, but what if other important employees need this information? For example, what if that business alert or notification is ignored by the recipient or the recipient is on annual leave?

Automating the notifications and alerts of the status of jobs within your backup policies can simplify management and ensure all relevant team members are kept informed. Failed backup jobs can be immediately recognized by notifying the relevant persons or teams via email or incorporating into your IT existing stack via webhooks or APIs. With 100% visibility of critical information, actions can be performed quickly and serious issues can be averted.

The Druva Cloud Platform facilitates multi-recipient capabilities so that users in different departments can automatically be presented with real-time backup job information that is critical to the business. With comprehensive reporting on backup policies, and disaster recovery plans that can be scheduled on any frequency, you'll be able to meet internal compliance mandates and external audits if required.

Challenge:

Whenever a critical backup job fails, if the relevant team members are not informed and the issues are left unresolved, this can lead to potential downtime, data loss, and missed SLAs.

Solution:

Choose an automated solution that can perform multi-recipient alerting, with the capabilities to integrate with your existing IT stack. Your internal stakeholders can be quickly notified when things go wrong and can take immediate action.

Druva for native AWS backup and recovery

Your data. Always on, always safe.

Druva gives you the powerful tools and features you need to protect your AWS resources, applications, and data. Reduce data protection complexity, and effectively manage AWS cloud backup, to meet business and regulatory backup compliance requirements.

Centralized backup and recovery for AWS environments

Built for multi-tenancy to protect multiple accounts, from a single console. Apply backup policies across multiple accounts using include/exclude rule options.


Minimize disruption and reduce risk of data loss

For extra protection against data loss, easily replicate backup copies between different regions, and accounts to ensure business continuity.

Enterprise protection, delivered “as-a-Service”

SaaS data protection solution built for AWS, designed for infinite scale, security, and flexibility. No software or hardware to install or maintain.

For more insights on AWS backup and recovery, check out druva.com/solutions/aws/



Find Druva in AWS Marketplace

Get Started



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

Americas: +1 888-248-4976

Japan: +81-3-6890-8667

Europe: +44 (0) 20-3750-9440

Singapore: +65 3158-4985

India: +91 (0) 20 6726-3300

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](https://druva.com) and follow us [@druvainc](https://twitter.com/druvainc).