

Modernizing Backup & Recovery with Commvault through AWS Marketplace

Introduction

Cloud adoption is happening faster than ever as many organizations look to move operations out of the data center and improve business agility. Amazon Web Services (AWS) is one of the market leaders in this space, providing reliable, scalable, secure cloud services that are available to customers to consume on-demand. As your organization plans the migration of data and workloads into AWS; however, it's important to consider how the data being used and created by these services is protected and managed. Without a comprehensive data management strategy that spans all your data – whether on-premises or in the cloud – you risk your data becoming fragmented or siloed. When that happens, it opens your organization up to potential risks, including:

- Data being unprotected or unrecoverable in the event of a loss, accidental deletion, or outage
- Data being at risk of being noncompliant with industry or governmental regulations
- A greater surface and opportunity for potential attacks like ransomware to strike

To avoid these issues, you need to a truly modern data protection solution. One that provides complete control of your data regardless of where it lives. One that can natively manage, protect, and ensure recoverability of data across your entire data estate – including on-premises applications, databases, VMs, containers, IaaS, PaaS, SaaS, etc. – and do it from a single interface with an automated, policy-based approach. A solution like this minimizes the risks described above and provides you the agility to use your data to innovate and drive business value.

Key benefits of Commvault Backup & Recovery on AWS

The Commvault Intelligent Data Services platform provides the tools to address all the concerns outlined above and enables your organization to innovate, grow, and thrive. With Commvault Backup & Recovery, available on AWS Marketplace, you'll enjoy all the benefits of the Commvault solution in a simple, easy-to-deploy instance that runs on Amazon EC2. These benefits include:

Efficient management of all your data

Commvault combines its best-in-class support for applications, databases, hypervisors, and other traditional on-premises technologies with unparalleled support for AWS services to enable SLA-based, automated protection across your organization's data estate. This efficiency means reduced complexity – fewer management tools and interfaces, fewer resources (and less time) needed to manage them – and allows you to focus more of your effort on utilizing this data to drive business value.

For customers that already leverage native AWS snapshots (for Amazon EBS, RDS, etc.) as part of their cloud protection strategy, Commvault can help drive efficiency there, as well. Commvault orchestrates and manages AWS snapshots, meaning you can utilize the native capabilities provided by Amazon to not only protect the data within these AWS services, but automate how these snapshots are created to fit in with your defined SLAs/policies. Additionally, Commvault will store these snapshots deduplicated and automate the lifecycle management to ensure cost optimization around how these point-in-time copies are stored. Commvault can even manage **cross-account** and **cross-region** replication of snapshots, providing another layer of customization to allow you to meet your AWS protection needs.

Reduced cloud costs

Continuing with the theme of driving cost-efficiency in protecting and managing AWS data, Commvault has native capabilities that reduce compute, storage, and network egress fees while protecting AWS services. These features can provide significant cost savings when compared to other solutions that don't offer the same level of integration.



Automated Data Tiering & Lifecycle Management

Commvault automates the tiering and deletion of backup data within Amazon Simple Storage Service (Amazon S3), ensuring that you meet your SLAs for data retention but aren't keeping data beyond its required lifecycle.

Commvault also has a feature that combines storage classes within Amazon S3, S3 Glacier, and S3 Glacier Deep Archive and intelligently moves data between those classes based on usage patterns. These combined storage classes optimize recovery times, enabling you to recover your data quicker, and reduce the cost of network egress fees for those recoveries.



Data Deduplication and Compression

Deduplication of backup data (including AWS snapshot data, as described above) can drastically reduce the amount of storage consumed, reducing your costs for Amazon S3. Customers have reported that Commvault can achieve deduplication rates of up to 99%!

Additionally, source-side compression can show significant savings on network egress charges when doing large scale data migrations or replicating data between regions.



EBS Direct Read (Changed Block Tracking)

Commvault integrates with Amazon Elastic Block Storage (Amazon EBS) Direct APIs, allowing them to streamline backup operations for EBS volumes. By utilizing the EBS Direct Read (which some may know as Changed Block Tracking), Commvault understands which blocks within a volume have changed since the last backup and can access those blocks directly.

Testing with EBS Direct APIs have shown up to an 85% improvement in backup times for EBS volumes² over the traditional method of creating and mounting a new EBS volume and doing a block-by-block comparison. This reduces backup windows and saves on both storage costs (for this secondary EBS volume) and compute costs (for the EC2 instance mounting that EBS volume).



Power Management and Auto-Scaling Compute Nodes

When running in AWS, Commvault has multiple infrastructure components required for operation that run as Amazon EC2 instances. Built-in power management and auto-scaling functionality help control software instances that reside in the cloud. Using this feature, Commvault automatically turns off the EC2 instances for their MediaAgents when not in use. Similarly, it will create and scale EC2 instances used as Access Nodes (to access AWS resources) as needed and delete those instances when they're no longer in use. This helps keep the costs associated with running the backup infrastructure down, in contrast to some competitive products that require multi-node clusters running 24x7.

Modern workload management

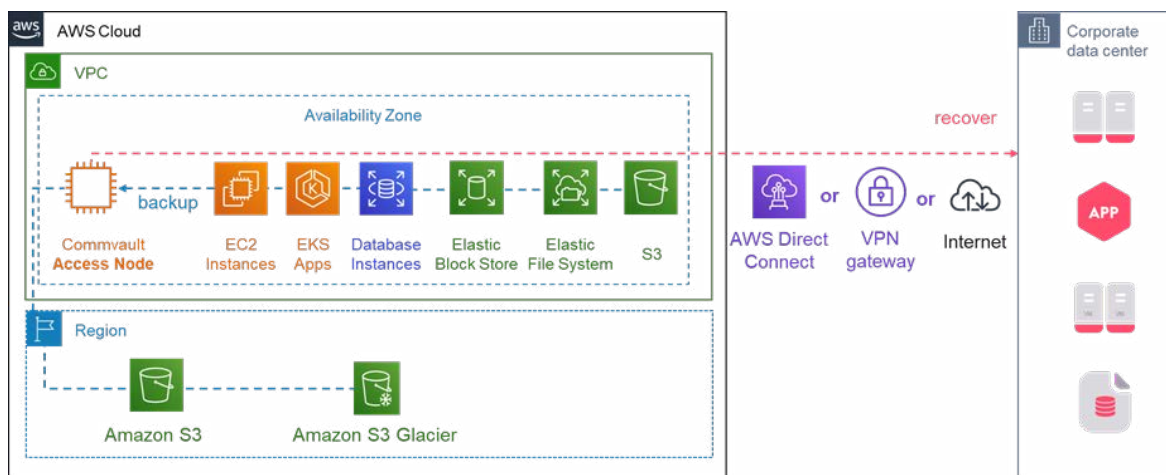
As we've discussed previously, a modern backup & recovery solution needs to be able to protect and manage data wherever your workloads may reside. Commvault provides the industry's broadest workload coverage, from databases to virtual machines, to SaaS applications and containers. Gartner agrees, rating Commvault 1st in all three use cases – Physical Environments, Virtual Environments, and Public Cloud Environments – in their Critical Capabilities for Data Center Backup and Recovery Solutions report for both 2020 and 2021³.

1 Parsons Corporation managing on-premises and AWS data protection with Commvault
 2 <https://aws.amazon.com/blogs/storage/accelerating-backups-with-amazon-ebs-direct-apis-and-commvault/>
 3 <https://www.commvault.com/itleaders>

Commvault extends this to AWS, as well, supporting an unparalleled number of AWS services. From compute to database to storage, chances are your AWS workloads can be protected by Commvault. Here's a list of current AWS services supported by Commvault, with additional support added regularly:

Compute	Database	Storage
Amazon Elastic Compute Cloud (Amazon EC2)	Amazon Relational Database Service (Amazon RDS)	Amazon Simple Storage Service (Amazon S3)
Amazon Elastic Kubernetes Service (Amazon EKS)	Amazon DocumentDB	Amazon S3 Glacier; S3 Glacier Deep Archive
VMware Cloud on AWS	Amazon Redshift	Amazon Elastic Block Store (Amazon EBS)
AWS Outposts	Amazon DynamoDB	Amazon Elastic File System (Amazon EFS)
	Amazon RDS on VMware	Amazon FSx for Windows File Server
	Amazon Aurora	AWS Storage Gateway

See the architecture diagram here illustrating how Commvault can backup a cloud-native environment in AWS and recover in the cloud or to an on-premises location.



Protection from threats like ransomware

With ransomware attacks on the rise globally, you need a solution that ensures protection and recoverability of your data at all times. Commvault provides a multi-layered approach to ransomware protection, including:

- **Identity security** – Commvault leverages IAM roles, policies, and STS:AssumeRole to effectively eliminate the challenge of leaking credentials.
- **Smart detection** – AI-based anomaly detection and honeypots to detect abnormal activity quickly.
- **Advanced alerting and reporting** – Report across your entire environment, including your cloud instances. Automated alerts instantly notify your SecOps team as soon as a threat is detected.
- **Air gapped backups** – Store backup copies in isolated Amazon S3 infrastructure not typically online or accessible for ultimate protection.
- **Parallel recovery** – Commvault recovers multiple systems in parallel for **fast** organization – wide recovery.

Summary

With native support across a wide array of AWS services, Commvault enables efficient management and protection of data used and stored in AWS services with the same SLAs and processes used to manage data residing on-premises. Commvault provides a single solution that manages data across an organization’s entire data estate – all from a single pane of glass, with global policies, alerting and reporting, AI and ML, and automation – to simplify operations and reduce complexity and cost. Easily deployed from the AWS Marketplace, Commvault is the ideal solution for you to modernize your backup & recovery so your organization can do great things with your data.

Visit the [AWS Marketplace page](#) for Commvault solutions to start your modernization today.