Hybrid Data Management with AWS

Data continues to grow exponentially worldwide as new data sources add to its rapid trajectory. Business workloads today are adopting a growing variety of data types and structures, thus efficient and secure data management practice becomes crucial to all organizations.

Today, CIOs are on a path to modernization by adjusting their existing data management models to accommodate contemporary data sources as they turn them into strategic business insights. While the cloud provides the necessary capabilities to analyze the data, some organizations within finance, healthcare and manufacturing industries, face an array of challenges when migrating their applications to the cloud. These challenges are due to legacy applications having to reside on-premises, local data residency requirements and low-latency computation needs.

To navigate through data management complexities while leveraging cloud capabilities on-premises, organizations are looking for a consistent set of infrastructure, tools, and services to make their transition to the cloud easier.

AWS offers a broad range of services that can help organizations meet their data management requirements in their own data centers, co-location spaces, or on-premises facilities. The agility and speed of AWS infrastructure and software allows data processing and analysis to occur as close as possible to the point of data orientation. AWS computing services offers a full range of services to help mitigate data management challenges organizations face today. Helping them transition their workloads effectively to a hybrid cloud environment.

Challenges of building a consistent data management solution

The journey to the cloud creates challenges for data infrastructure and development teams to design data management models that provide consistent and reliable cloud services on-premises. Challenges can vary depending on the specific industry and operational requirements, however they generally include:

1. Consistent infrastructure

Organizations require consistency of infrastructure, tools, and services to streamline their data management efforts, regardless if they manage sensitive data on-premises or run machine learning models in the cloud. With disparate infrastructure solutions, organizations have to deal with different sets of management tools and skills to maintain the infrastructure, which slows down their operations and migration to the cloud.

2. Centralize management

Controlling and easily managing access to computing resources on premises and in the cloud remains a priority. Simplifying design architecture by centralizing its management, monitoring, and points of operation is a growing requirement for companies across various industries.

3. Scale flexibility

On-premises data sharing, processing, and residency constraints can challenge organizations' infrastructure modernization journey, particularly if organizations don't have access to services, APIs, or tools for automation, deployment, and security. In this case, infrastructure scaling becomes overly complex, and can become a drain on resources.

4. Process large datasets

Within different industries, many need the ability to process and analyze large datasets on-premises, where compute and storage resides close to devices and IT equipment.

Build a truly consistent hybrid data model with AWS Outposts

Whether you need to deploy your data workloads in the cloud, or you want to build and deploy specific applications closer to your end-users with single-digit millisecond latency, AWS offers a hybrid solution in the form of AWS Outposts, to bring the same AWS infrastructure, services, APIs, management tools and support to virtually any data center, colocation space, or on-premises facility.

Benefits of hybrid data management

Organizations can deploy cloud infrastructure on-premises, determine data processing priorities, and when ready, migrate towards the cloud.

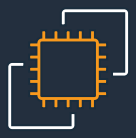

- **Cloud capabilities on-premises**
  Amazon EC2 instances featuring Intel® Xeon® Scalable processors brings the same cloud capabilities on-premises.

- **Seamless migration to the cloud**
  Build an application once and deploy it in the cloud, on-premises, or in a hybrid architecture with consistent performance.

- **Accelerated modernization**
  Companies can accelerate the adoption of cloud services on-premises across teams.

- **Reduce management overheads**
  Reduce the time, resources, operational risk, and maintenance downtime required for managing IT infrastructure, giving you the ability to focus on what differentiates your business.

How AWS Outposts works

AWS Outposts catalog includes options supporting the latest generation Intel powered EC2 instance types with or without local instance storage. Organizations can choose from a range of pre-validated Outposts configurations offering a mix of S3, EC2 and EBS capacity which are designed to meet a variety of data management needs.

Outposts deployment and set-up

Once an Outpost is deployed at your location, it connects to the designated AWS Region to extend your Amazon VPC on-premises and deploy your EC2 instances. Through VPC extension, you get EC2 instances running in subnets inside Outposts that can communicate to those in the Region via Direct Connect or public internet.

Running Outposts on-premises

AWS Outposts connects to your on-premises devices via Local Gateway (LGW). The connection to local devices is established by configuring an on-premise route and sending the collected data towards the LGW whether it be in your datacenter, colocation or on-premises facility. Outposts data traffic can also be sent elsewhere, including back to AWS via private virtual interface or Direct Connect.

Connecting Outposts to cloud services

A broad range of services can be run either from the AWS Region or locally on AWS Outposts. AWS regional services include: Amazon ECS and EKS clusters for containers, EMR for big data, Amazon Elastic Block Store (EBS) to deploy intensive workloads, and Amazon S3 for data storage. Alternatively, you can run services locally including SageMaker for AI and ML frameworks, and SageMaker Neo for local inference. By connecting to an AWS Region you can access the full range of AWS services.

How Outposts supports the AWS Cloud

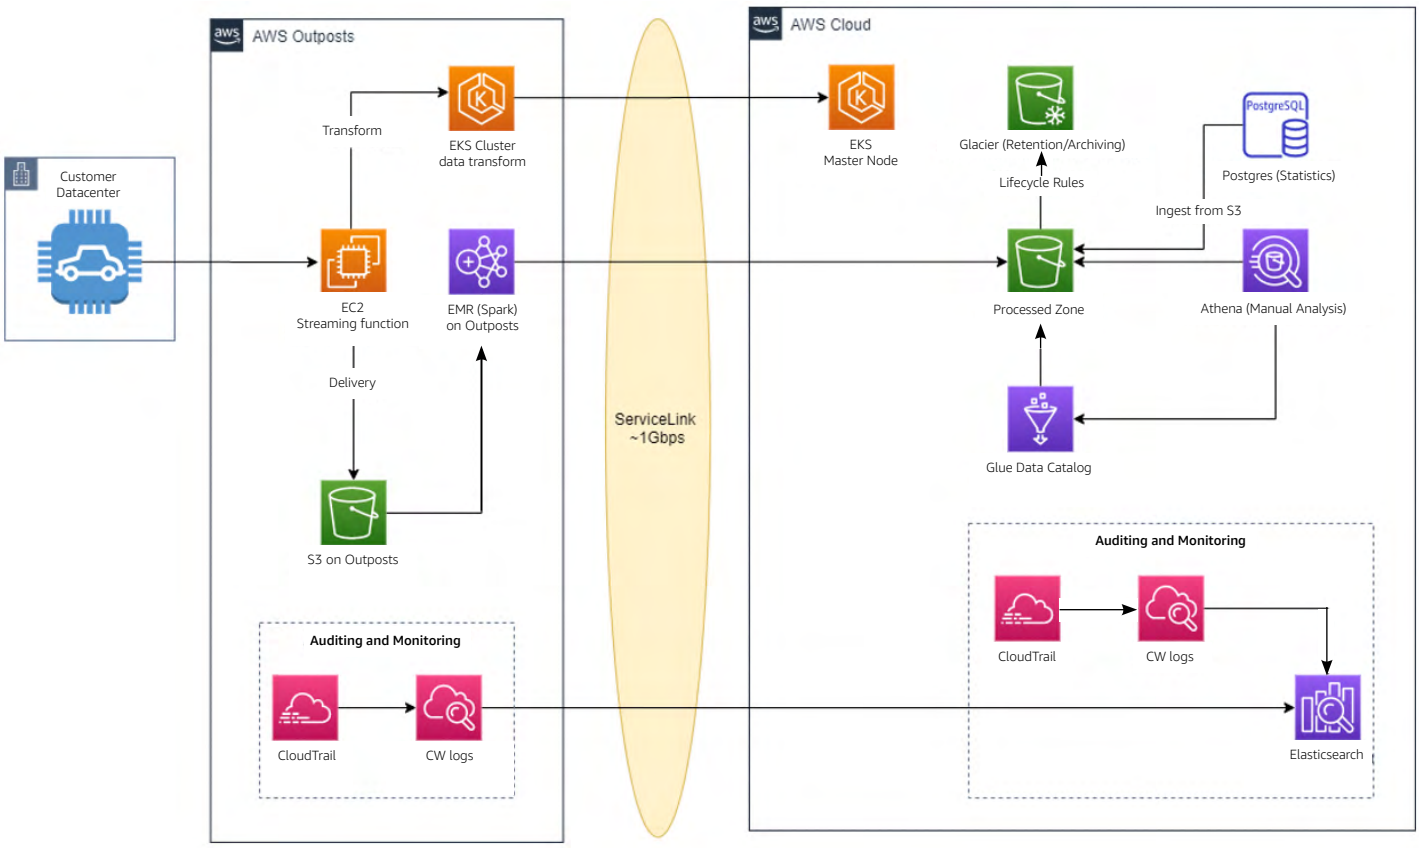

Taking the data management journey with AWS

Managing data with Outposts

Amazon S3 on Outposts delivers object storage to your on-premises AWS Outposts environment. With S3 and AWS Storage Gateway on Outposts, organizations can enable on-premises workloads to use AWS cloud storage while processing data locally for low-latency access. Moreover, organizations can build Amazon S3 data lakes to cost-effectively store all data types in their native formats. This type of data architecture allows companies to manage multiple data types and structures from a wide variety of sources, store it in a centralized repository and deploy data analytics tools or run machine learning inference for valuable business insights.

AWS regional services like Amazon EMR, AWS Glue, Amazon EBS and Amazon S3 you to scale your compute and storage, and move from an on-premises-only approach to a faster, more agile and more cost-efficient data management process, leaving you to focus resources on creating business value.

Processing data with Outposts

With Outposts, companies can set the data processing criteria based on importance and urgency, and upload to the AWS Region as per their chosen preference, to save time and increase efficiency. Services such as Amazon S3, Amazon EBS and EC2 instances (Amazon Elastic Compute Cloud) powered by Intel® Xeon® Scalable processors gives you the opportunity to launch as much or as little compute capacity as you need. AWS services seamlessly integrate with S3, so you can easily store and analyze large datasets on-premises. You can choose your data management components like EMR

for deeper analysis, Glue for data catalogue, Athena for data exploration, and Redshift Spectrum for data warehouse. Amazon S3 APIs are supported by most major third-party independent software vendors (ISVs) too, so you can deploy the tools your organization is most comfortable with.

Moving data with AWS

At each stage of an organization's journey to the cloud customers can decide whether they will lift and shift, redesign for the cloud or leverage the AWS marketplace. How quickly you migrate is completely in your control and can be dictated by business requirements.

Organizations can start from scratch, move into AWS Region or AWS Cloud incrementally, and as they take advantage of AWS tools and services, they become more capable of implementing new data management

Why AWS for hybrid data management?

Data security

AWS offers security, compliance, and governance services to the most heavily regulated industries. All AWS data-storing services provide encryption, including Outposts, where data is encrypted at rest and in transit between Outposts and the AWS Region.

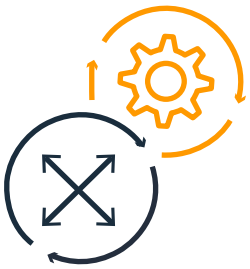

Data management enhancement

With the support of AWS hybrid cloud solutions, organizations can drive adoption of AWS services across teams, scale their data management practice, and accelerate innovation.

Data systems scalability

With AWS's range of services, organizations can avoid long hardware procurement cycles, scale on-premises infrastructure as per their business needs, and alleviate the underutilized capacity issue.

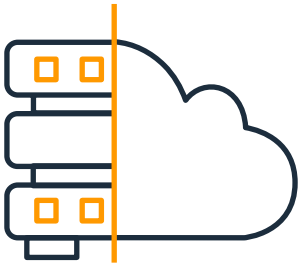

Data management consistency

AWS Outposts enable end-users to enjoy fully managed and consistent infrastructure performance across cloud, on-premises, or in a hybrid data architecture.

AWS Partner support

AWS Partners provide expertise in managing and migrating various workloads across cloud and on-premises environments. You can choose partners for strategy and technology advisory services or software-specific solutions to help you migrate and run your workload on Outposts.

Mitigate data migration challenges with Veritas Technologies LLC

Organizations with unique data migration needs such as high availability, backup and recovery, and information governance can leverage proven AWS partners like Veritas to address their needs across Outposts, VPC or AWS Region. Veritas helps customers automate and orchestrate migrations with no interruptions.

Harness machine learning with Splunk Enterprise

For companies that require specific data collection automation, data processing and artificial intelligence capabilities, Splunk offers a machine-data consumption and analytics software solution that is available on and powered by AWS.

Address your sensitive data needs with Privitar

Financial organizations that collect sensitive customer data and want to unlock its value without compromising data privacy can leverage the support from AWS partner Privitar. Optimize sensitive data management by eliminating slow, manual provisioning and enabling cloud-based analytics and machine learning.

Accelerate software delivery with CloudBees

Software-delivery companies that require an end-to-end software delivery management system to streamline and centralize their operations can leverage CloudBees list of products available on the AWS Marketplace. With CloudBees support, software teams can deliver new features and updates faster and safer, without having to worry about infrastructure restricting the growth.

Manage data intelligently with Komprise

Organizations that require greater visibility into their data points across NAS and S3 on-premises and cloud, can use Komprise analytics-driven data management software to help them intelligently tier their data and transparently archive it to AWS. Komprise offers capabilities to optimize cost and analyze, archive, migrate, replicate, and manage data at scale.

Use cases: How AWS empowers innovation

AWS success stories demonstrate that the benefits of cloud capabilities such as flexibility, scalability including geographic reach, and the access to new technologies are suited for today's evolving business needs.

Healthcare

Medical professionals manually collect structured data to store and analyze in vital fields such as cancer staging, medical/family history and patient-reported symptoms. AWS Cloud services automates data collection, where using machine learning inference models amplify data processing and extraction of valuable insights. AWS provides the tools, services and APIs to deliver real-time video analytics and pattern matching, while delivering on-premises flexibility and access to cloud capabilities when needed.

Finance

Financial or government institutions that need to comply to specific data regulations use hybrid cloud to meet their contractual obligations with their customers and demonstrate compliance with legal policies. AWS Outposts allow these organizations to maintain data visibility, process sensitive data locally, including collecting local cache and filtering, and when needed connect to Local Zones or send it to AWS Region.

IoT

Self-driving cars that ingest large amounts of sensory data can benefit from the hybrid model by deploying AWS Outposts on their depots, to process data in real time without having to send it all back to the cloud. Hybrid data management models allows customers to avoid transmitting their large volume of data (data lakes) which reduces costs while providing real-time response to on-premises entities.

Security

Companies that are interested in using Outposts to run physical security environments, such as video surveillance, badging systems or security systems, can build and run these workflows on Outposts, archiving relevant data to S3/Glacier within the AWS Region for forensic analysis.

Application Migration

Organizations with large datasets leverage the AWS Outposts consistent infrastructure to support and segment large application migrations into smaller pieces by applying latency-sensitive connectivity between parts of the application. Using the "lift and shift" methodology, organizations can build applications once, run on-premises, and seamlessly migrate to AWS Region when ready.

Modernize on-premises

Semiconductor companies create large production data lakes that must remain on-premises and therefore can't easily migrate to an AWS Region. AWS Outposts provides the solution to modernize these workloads by integrating EMR into their existing data lakes. This gives their data science/ML teams the environment to develop new applications and use cases without impacting their existing production environment.

Manufacturing

Manufacturing enterprises require edge computing so that their machinery and robotics on the manufacturing floor can detect real-time data insights. With AWS Outposts, manufacturing customers benefit from running low latency and AI/ML on 'the edge' since the reduction in latency allows them to analyze data locally, to achieve a real-time feedback loop.

NEXT STEPS

AWS offers all the advantages of hybrid cloud capabilities while still supporting the on-premises use cases of local data processing, low latency, and legacy application migration. Take the next steps to modernize your organization today with AWS Outposts.

1. Engage

Reach out to your account team or fill out our contact form. Alternatively, go into the AWS Management Console.

2. Choose

Select your size and then order the Outpost rack configuration that best suits. Custom configuration is available.

3. Install and Launch

AWS will install and deliver your configuration. Use standard AWS APIs or Management Console to launch and run AWS resources locally.