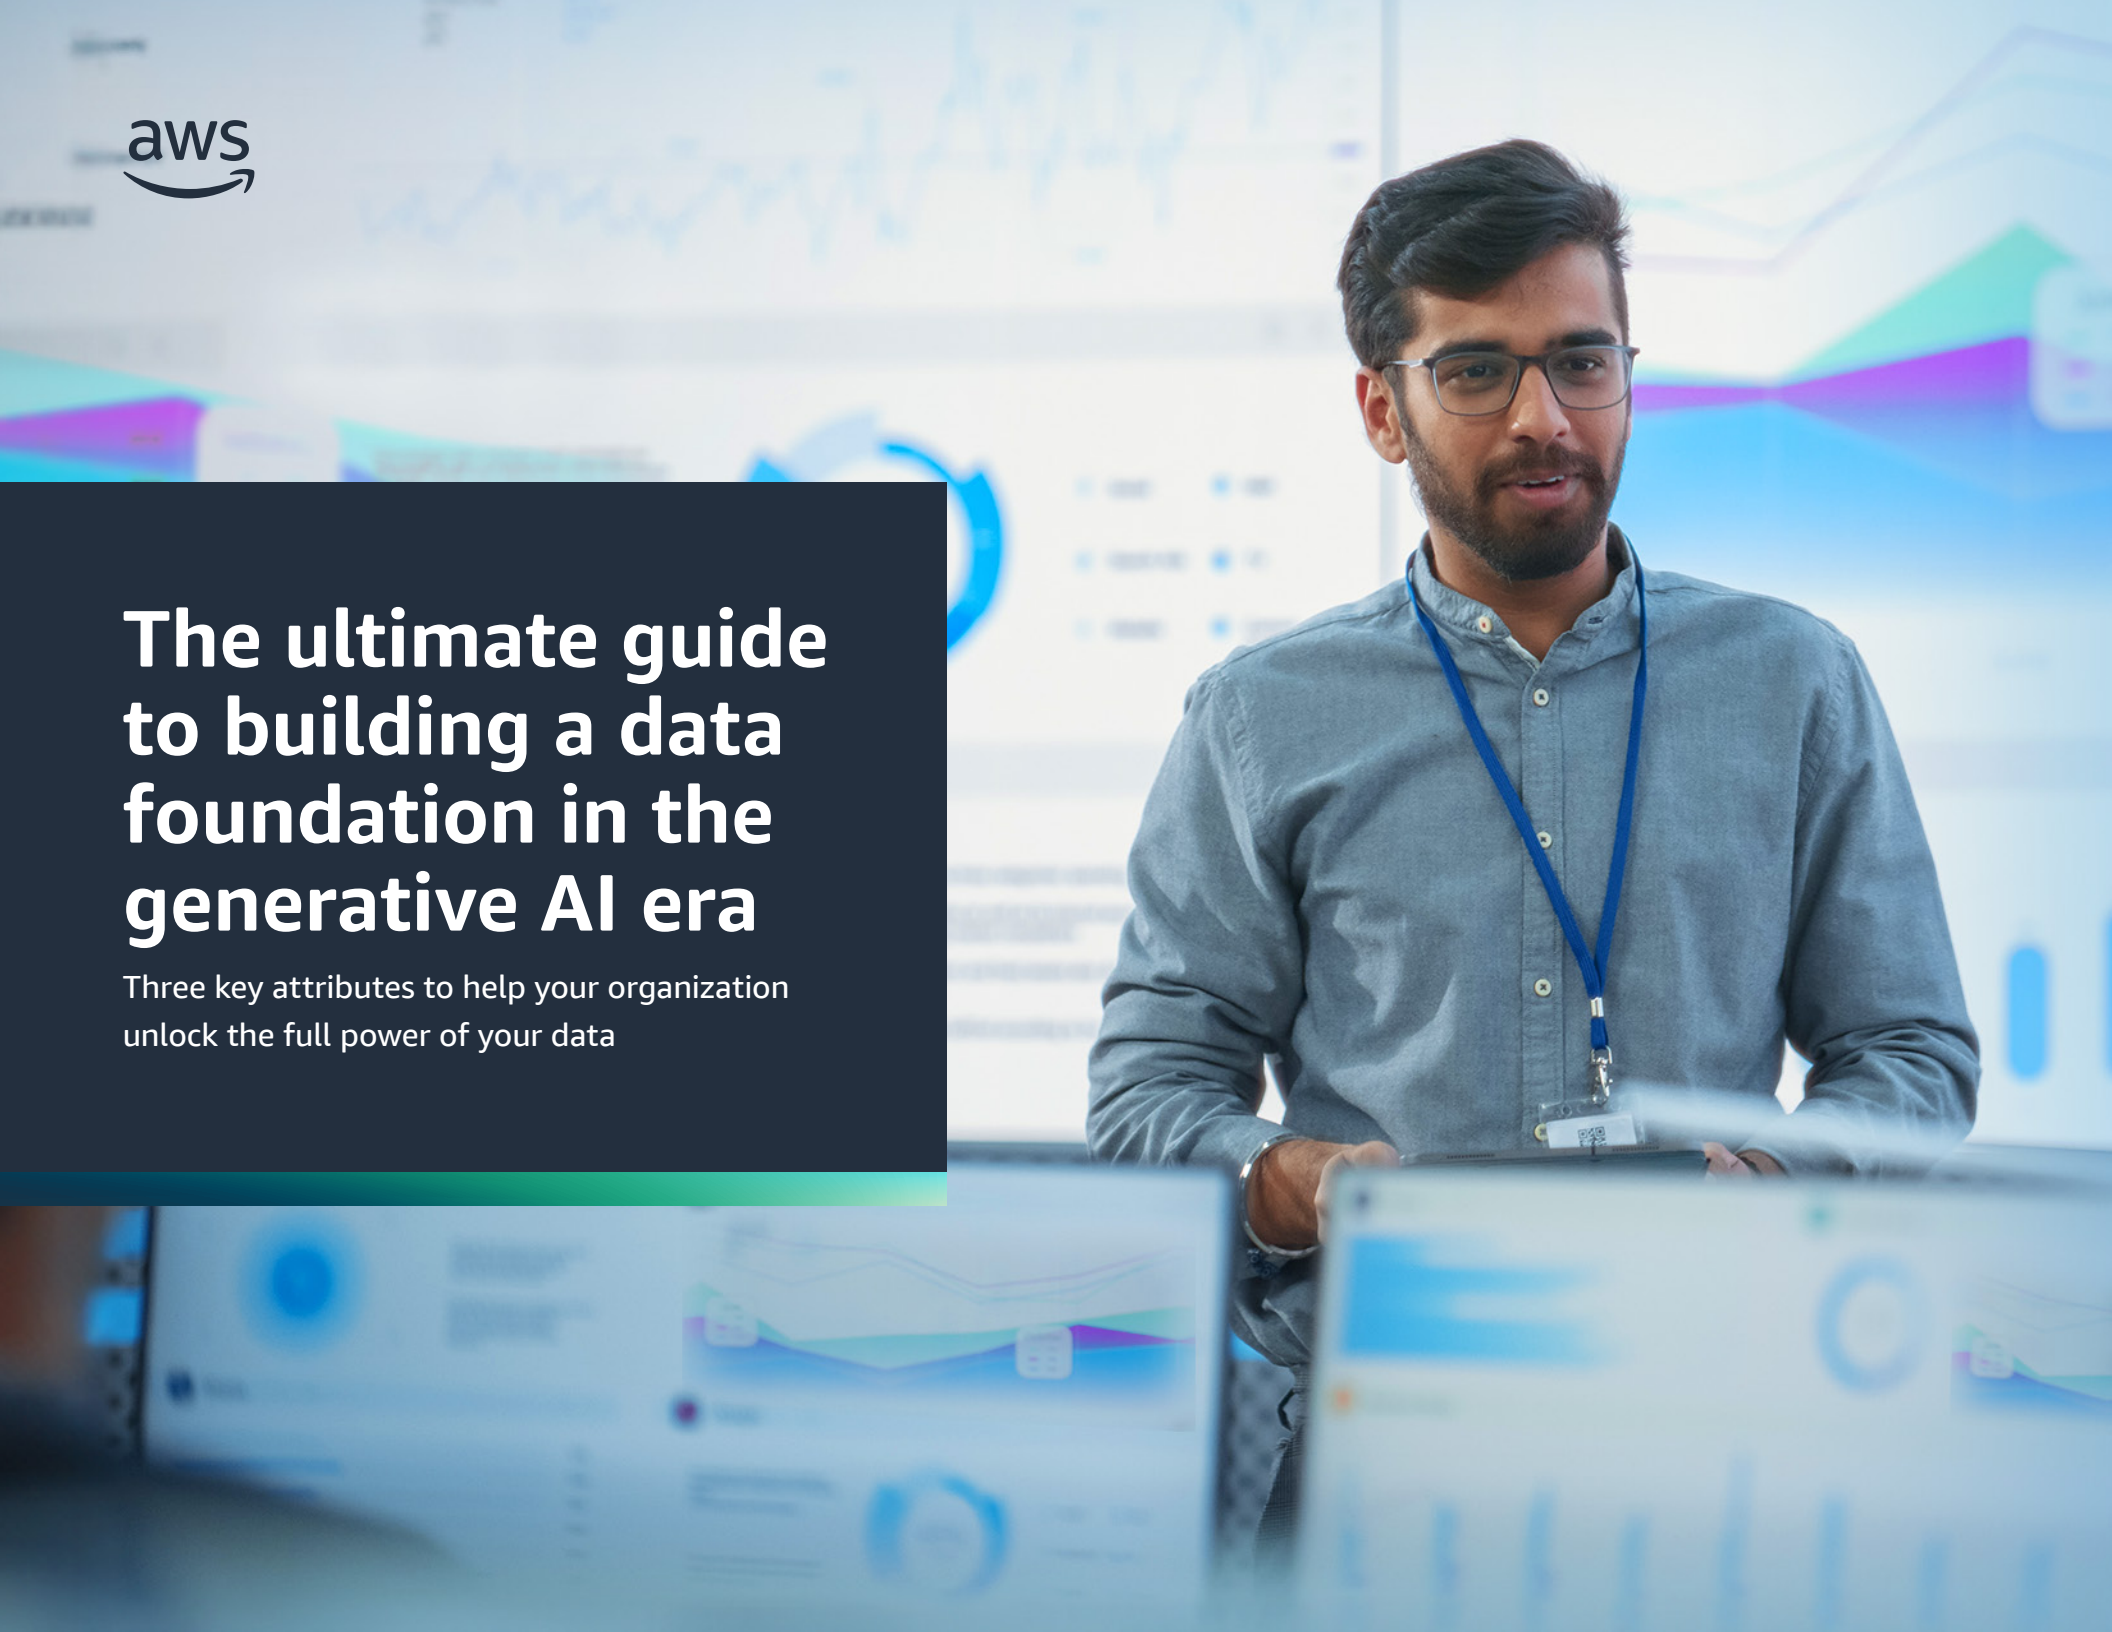

The ultimate guide to building a data foundation in the generative AI era

Three key attributes to help your organization unlock the full power of your data

Table of contents

INTRODUCTION

Data is the spark that leads to meaningful innovation

Now more than ever, data is at the center of every application, process, and business decision. It's the genesis for modern invention, and in today's fast-changing and complicated landscape, how you put your organization's data to work can be the key to accelerating innovation and accomplishing your organizational goals. The stakes are high. According to Forrester Research,"firms [that] have advanced insights-driven business capabilities are 8.5 times more likely than... firms at the beginner stage to report that their firm's annual revenue grew by 20% or more."[1]

The advent of generative AI also puts renewed emphasis on the importance of data. When you want to build generative AI applications that are unique to your business needs, data is the differentiator. Data is the key to moving from generic applications to generative AI applications that create real value for your customers and your business.

1 "Is Your B2B Organization Insights-Driven?," Forrester.com, August 17, 2023

There's no need to reinvent the wheel

Becoming a data-driven organization begins with the right data foundation. The good news is that a proven, data foundation already exists—and organizations are already capturing its benefits using Amazon Web Services (AWS). For example, AstraZeneca is integrating and scaling its data and artificial intelligence (AI) capabilities across the business to innovate faster and improve patient outcomes. They now run more than 51 billion statistical tests in less than 30 hours, facilitating the delivery of genomic insights to drug discovery projects. BMW Group is using data to optimize its supply chain and improve production capacity. And LG AI Research is harnessing its data to develop generative AI applications to transform business processes—broadening access to AI in various industries such as fashion, manufacturing, research, education, and finance.

Building the right data foundation that will transform your organization is attainable. Read on to explore the fundamentals.

Key challenges and considerations

More data than ever is being generated and stored

On-premises tools and legacy data stores can't meet today's demands. Organizations need data stores that can scale to keep up with petabytes to exabytes of data. And they need to be able to store this data in a cost-efficient way without sacrificing performance.

Data siloed across multiple sources creates productivity and cost inefficiencies

Organizations must deal with diverse data types—including log files, clickstreams, voice, and video—which are typically stored in silos across multiple data stores and departments. This makes it difficult to harness the data and extract actionable insights. To transform the infrastructure from a source of complexity and expense to an engine of value creation, organizations must break down these silos to unify all their data.

The current state of decision making is unsustainable

Gartner reports that 65 percent of decisions made today are more complex (involving more stakeholders or choices) than they were five years ago.[2] To make better and faster decisions, organizations must be able to perform analytics and machine learning (ML) operations in an agile, cost-effective way—using optimal tools and performance to scale for each use case. They can no longer waste time constantly redeploying and reconfiguring infrastructure to scale performance and capacity.

2 "How to make better business decisions," Gartner, October 20, 2021

Trying to maintain data governance is a full-time job

Traditional data architectures require risky, complicated management procedures because data is accessed from so many places. Granting, tracking, auditing, and removing employee access—while simultaneously remaining in compliance with a growing number of regulations—is a full-time job. Automating these mandatory data governance tasks frees modern teams to shift their focus back to innovation.

Analytics and machine learning adoption is still impeded by a lack of skills and inertia

Many businesses are struggling to make progress with scaling analytics and ML tools. Gartner finds that organizations investing in AI moved just 54 percent of their AI proof-of-concept pilots into production.[3] A continued lack of data and ML skills and quantity or quality of data to train on are some of the issues slowing progress in this important area. Still, the need to help business users leverage data-driven decision making (DDDM) is growing.

Data is increasingly difficult to secure

There was a time when IT teams chose between making their architectures fast or making them secure. Now, they need to do both. Meanwhile, 97 percent of organizations experienced increased cybersecurity threats from 2022 to 2023, according to Accenture's State of Cybersecurity Resilience 2023 report—while the percentage of successful breaches from external networks remains high, at 61 percent.[4] How can organizations maximize privacy and security?

3 "How to make better business decisions," Gartner, October 20, 2021
4 "State of Cybersecurity Resilience 2023: How cybersecurity boosts enterprise reinvention to drive business resilience," Accenture, 2023

Three key attributes that can help your organization unlock more value from data

Highly data-driven organizations are three times more likely to report significant improvements in decision making compared to those that rely less on data.[5]

Implementing a data foundation that makes data management easier at every step of the journey—from ingesting, storing, and querying data to analyzing, visualizing, and running ML models—requires three fundamental attributes. Regardless of your business challenges, your data foundation should be:

Comprehensive: Provide the right tools with the optimal price performance for any user, type of data, and use case

Integrated: Break down silos to connect all your organization's data, so it can be put to work effectively

Governed: Free your teams to move faster with governed data access when and where your users need it to speed innovation

Successfully becoming a data-driven organization may also require a broader mindset shift—in which both goals and decisions are supported by a data foundation that encompasses people, processes, tools, and education.

5 "Data and Intuition: Good Decisions Need Both," Harvard Business Publishing, January 6, 2023

COMPREHENSIVE

Provide the right tools, with optimal price and performance, for any user and data type

Businesses need to build a sustainable data foundation that can meet their needs now and in the future. It takes more than just a single data lake, data warehouse, or business intelligence (BI) tools to harness data effectively. It requires a data foundation with a comprehensive set of tools that accounts for the scale and variety of data and the many purposes for which you want to use it.

Building with a cloud provider that innovates to continuously bring you all the data tools you'll need with the right price performance for your use case ensures you have a data foundation that grows with you. AWS has the broadest and deepest set of data capabilities to support any data workload or use case. From databases for applications to storage for data lakes to analytics to ML and end-user tools, AWS provides the right capability in each area, so you don't have to compromise on performance, cost, or results. AWS is continually accelerating its pace of innovation, so you will never outgrow AWS for your data needs.

Scaling data-driven applications

Build applications on a modern data foundation for the best price and performance for your use case at scale using AWS databases. More than 100,000 organizations, for example, achieve the performance and availability of commercial-grade databases at one-tenth the cost using Amazon Aurora. For use cases such as graphs, streaming, and documents, AWS offers eight purpose-built database engines, each uniquely designed to provide optimal performance for your applications, transforming the economics of database ownership.

Data Foundation Tip

Achieve a cost-effective data foundation without sacrificing performance. Enable your organization to maximize its current capabilities by optimizing cost:

Samsung saved 44 percent on monthly operational costs and an additional 22 percent on maintenance fees when migrating to Amazon Aurora PostgreSQL.

Carrier connected its cold chain logistics network to help its customers optimize cold chain operations, decrease their energy use, and enhance their outcomes with a reduction in costs, delays, cargo loss, and spoilage in transit.

United Airlines created an intelligent airport with more than 20,000 sensors producing data to drive real-time insights, optimizing ground equipment capabilities and resulting in $120 million in savings for equipment that was no longer required.

AWS works with more than

2 million customers

to solve some of the most complex data problems in the world.

Powering data lakes and data warehouses

Organizations use data lakes and data warehouses to unify datasets and apply analytics and ML to aid in decision making and create new data-driven customer experiences.

For years, AWS has helped hundreds of thousands of customers build data lakes to store data reliably and securely store structured and unstructured data with services like Amazon Simple Storage Service (Amazon S3), AWS Glue, and AWS Lake Formation. There are hundreds of thousands of customers with AWS-powered data lakes. For data warehousing with super fast query results across many different data sources.

Optimize storage performance and speed up data access

For your most frequently accessed data, storage classes such as S3 Express One Zone can improve data access speeds by 10 times and reduce request costs by 50 percent compared to S3 Standard.

Providing analytics for all use cases

True agility helps organizations adapt quickly to changing business needs. Empower your organization's teams to ingest, combine, and run historical, real-time, and predictive analytics on your data with AWS analytics services. This includes services for SQL querying, log analytics, streaming, and Apache Spark. For data warehousing with super fast query results across many different data sources, Amazon Redshift, a petabyte-scale data warehouse, delivers up to six times better price performance than other cloud data warehouses. With generative AI inside of Amazon Redshift, you can generate SQL queries using natural language. And AI-driven scaling and optimizations for Amazon Redshift Serverless helps you optimize between cost and performance by learning from your patterns.

For big-data querying, you can support more big-data frameworks than any other provider using Amazon EMR, with up to two times faster time-to-insights. Our customers achieve more than three times performance with Apache Spark when they run our fully supported and AWS-optimized runtimes for EMR.

Innovating faster with services that make ML and AI more accessible

Organizations have been using ML to add intelligence to existing processes, automate time-intensive manual tasks, and accelerate innovation using data. Now, with generative AI, they have the opportunity to reinvent customer experiences and applications.

With AWS, you have access to the most comprehensive set of AI and ML services. Amazon Bedrock is the easiest way to build and scale generative AI applications with foundation models (FMs) to create new content and ideas, including conversations, stories, and images. With Bedrock, you can use your own data to easily and securely customize FMs from AI21 Labs, Anthropic, and Stability AI, as well as Amazon Titan models via an API. AWS also offers a wide range of services that allow you to add AI capabilities like image recognition, forecasting, and intelligent search to applications with a simple API call. To build your own ML models, Amazon SageMaker provides all the tools necessary to easily build, train, and deploy ML models at scale.

Enabling data insights throughout the organization

It's no longer just data-savvy individuals who can rapidly extract valuable, relevant insights from data to help inform decision making. ML-powered BI solutions, such as Amazon QuickSight, enable easy connectivity to data sources. Business analysts can use this data to showcase fresh trends and predictive insights on interactive BI visualizations and dashboards.

Amazon Q, our generative AI-powered assistant, helps you in QuickSight Q to author dashboards and create compelling visual stories from your dashboard data using natural language. Business users can self-serve meaningful insights with ease. Even if they ask vague questions in natural language, they will receive comprehensive and contextual answers that explain the data completely using visuals and narratives.

Boosting data proficiency

When your employees know how to use data effectively, they can help your organization achieve its data objectives. Invest in educating and upskilling your workforce in data, analytics, and ML with AWS Training and Certification.

Scale data-driven decision making throughout your organization

Empower business users to:

- Meet varying analytic needs from the same source of truth with unified business intelligence on Amazon QuickSight
- Extract key insights faster using Amazon Q, a generative AI powered assistant
- Build dashboards and create compelling visual stories using natural language with Amazon Q in QuickSight
- Generate accurate ML predictions without prior experience using Amazon SageMaker Canvas—with no coding required
- Simplify governed access to data for business users with Amazon DataZone
- Access more than 150 professional development courses related to data, analytics, and ML via AWS Training and Certification

ADP makes 312 trillion decisions a month with analytics processes

ADP helps more than 900,000 businesses manage 70 million employees through its people and payroll process. That management generates a massive amount of data. In fact, ADP processes more than 2.5 petabytes of data with more than 25 billion individual data points represented. To perform aspects of its overall data processing, ADP uses Amazon Redshift and Amazon Neptune. These data services help companies measure, compare, predict, and apply insights about their workforces. ADP also enables organizations to create Pay Equity Dashboards using AWS services—helping more than two-thirds of companies show improvement in pay equity.

Read the full story >

"Now is the time to use data to help people to understand what actions we can take to create a more diverse, more equitable, and a more inclusive work environment and to build the future we all want to create."

Jack Berkowitz,
Chief Data Officer, ADP

BMW Group democratizes data usage at scale

BMW Group made anonymized data from vehicle sensors and other sources across the enterprise easily accessible for internal teams who create customer-facing and internal applications with AWS. The company moved to an AWS-based centralized data lake for its agility—and its ability to process terabytes of telemetry data from millions of vehicles daily. Building up a human-readable data catalog and clearly displaying data resources proved essential, boosting the productivity of data analysts, data scientists, and engineers.

"We are just starting our journey with AWS, and we look forward to helping our business fulfill its strategy of driving innovation into the future."

Kai Demtröder,
Vice President of Data Transformation, Artificial Intelligence, Data and DevOps Platforms, BMW Group

Read the full story >

INTEGRATED

Break down silos to connect all your organization's data, so it can be put to work effectively

Opportunities to transform your business with data exist all along the value chain. But making such a transformation requires you to see the full picture of your customer and business. With data spread across multiple departments, services, on-premises databases, and third-party applications, you need to be able to easily integrate data across silos to get the best insights. Companies have various approaches to how they are unifying data—data mesh, lake house, data fabric, and so on—but typically, it involves a data lake as a foundational element. Data lakes allow you to collect, store, organize, and process valuable data from your data silos and make it available to analytics, visualization, and AI tools in a governed way.

Embracing a zero-ETL future

Many organizations have multiple data lakes in addition to data warehouses, databases, analytics tools, ML tools, and SaaS applications. Integrating data across silos requires complex extract, transform, and load (ETL) pipelines, which can take hours (or even days) to complete. That's just not fast enough for modern decision making. Organizations should adopt technologies that automate or eliminate ETL where possible.

AWS is investing in a zero-ETL future so you can quickly and easily connect and act on all your data, no matter where it lives. This includes:

- Enabling near real-time analytics on your transactional data with zero-ETL integrations between Amazon Redshift and several AWS databases, such as Amazon Aurora MySQL, Amazon Aurora PostgreSQL, Amazon RDS for MySQL, and Amazon DynamoDB
- Bringing ML to the data source with no data movement or ML exporting using SageMaker integration into Amazon Redshift, Aurora, Athena, and Neptune

- Using generative AI to make data integration and processing easier and faster with Amazon Glue
- Providing zero-ETL integration between Amazon OpenSearch Service and Amazon DynamoDB so you can run near real-time search queries and the zero-ETL integration between Amazon OpenSearch Service and S3 lets you easily perform complex queries and visualizations and identify anomalies without moving data
- Making it possible to run queries across data stored in operational databases, data warehouses, and data lakes to provide insights across multiple data sources with no data movement using Amazon Athena and Amazon Redshift

When ETL is necessary for use cases where you need to combine multiple datasets or add value through transformations, AWS continues to invest in AWS Glue—a serverless, scalable ETL and data integration service that makes it easier to discover, prepare, move, and integrate data from multiple sources for analytics and machine learning.

Analyzing all your data and third-party data

To break down data silos, you can't have connections to only some of your data sources—you need to be able to seamlessly connect to all of them, whether they live in AWS or external third-party applications, on premises, or even in another cloud environment. AWS is investing in a future where you can automatically integrate hundreds of data sources—no matter where they live.

Increasingly, organizations are also harnessing third-party data to deepen insights by joining this third-party data with their own data. AWS Data Exchange enables AWS customers to access third-party data through files, tables, and APIs from more than 300 data providers and more than 3,500 data products, all from one place. Third-party data from partners and customers is also being used, which increases the need for comprehensive governance policies to protect the data. Data clean rooms—protected environments where multiple parties can analyze combined data without ever exposing the raw datasets—have emerged as a solution. AWS Clean Rooms helps companies and their business partners securely analyze and collaborate on their datasets—without sharing or revealing the underlying data.

KINTO

About KINTO Technologies Corporation

KINTO Technologies Corporation is a leading player of the mobility platform industry and is the technology company responsible for the development of the KINTO service as Toyota's financial services company.

Challenge

KINTO faced persistent performance challenges and impacts on production workload due to a custom-built solution for streaming changes from core databases.

"Using the Aurora MySQL zero-ETL integration with Amazon Redshift, we are able to always have near real-time data available in Amazon Redshift, eliminating developer hours spent manually managing data pipelines for ETL operations or dealing with performance impacts to our workloads, which helps reduce our operational burden."

Hitoshi Kageyama,
Executive Vice President for KINTO Technologies Corporation

The AWS solution

KINTO used Amazon Aurora MySQL Zero-ETL integration with Amazon Redshift to achieve a more resilient data pipeline. KINTO can now apply Amazon Redshift's advanced analytics features to its operational data in near real-time.

Result

Using Aurora MySQL Zero-ETL integration with Amazon Redshift, KINTO is equipped with near real-time data, eliminating developer hours spent manually managing data pipelines for ETL operations. As well as significantly decreasing the hours spent dealing with performance impacts to their workloads, reducing their operational burden.

AWS products used

Amazon Aurora MySQL

Amazon Redshift

How data delivery enables Goldman Sachs to work smarter

Goldman Sachs

Goldman Sachs has seen demand for financial market data increase exponentially, typically doubling year over year. To improve the overall efficiency of consuming third-party data, the company migrated its existing market data feeds to AWS Data Exchange, a managed service that allows data processes to be centralized, streamlined, and automated wherever possible. The efficiency improvements allow teams to focus more time on the value-add analytics of the data—and less time wrangling the data to get it ready for use. Frictionless data delivery is critical to the organization's continued success.

"AWS Data Exchange is a key component of Goldman Sachs's financial cloud strategy because it reduces friction for sourcing financial data from new and existing third-party providers, and allows us to focus on delivering our core services and differentiated data analytics to better serve our clients."

Marco Argenti,
Co-Chief Information Officer, Goldman Sachs

Read the full story >

aws

ENGIE accelerates its zero-carbon transition

ENGIE, a global utility company in the process of a zero-carbon transition, built its Common Data Hub data lake on AWS. ENGIE worked with AWS Professional Services during the solution design and implementation and formed an internal service team to oversee the platform. With more than one thousand projects worldwide currently on the Common Data Hub, ENGIE eliminated its data silos, empowering every department with equal access to a common data framework.

"We were convinced that AWS was a good solution for many reasons, including the cost model—and especially in terms of data storage."

Gregory Wolowiec,
Technology Team Leader at ENGIE Data Programs

Read the full story >

GOVERNED

Free your teams to move faster with governed data access

Beyond being comprehensive and integrated, it's equally important to ensure that all consumers of your data—whether human users, applications, engines, or AI/ML models—can access data where and when it is needed with the right level of control. With the right data governance strategy in place, you can move faster with the right data access—right when it's needed.

As more data migrates to the cloud and new AI/ML models consume vast amounts of data, enterprise data governance models must evolve in lockstep. IT and business leaders need up-to-date policies to protect data as it moves back and forth among different repositories and to accommodate changing privacy and data security regulations about where data can be stored.

Simplifying data access permissions

Implementing a successful data governance strategy presents a unique set of challenges. Giving internal or external consumers their data—with the right level of access to specific datasets—is complex and time-consuming. Teams often engage in heavy lifting, such as manual scripts or investigating individual data clusters, to figure out which consumers have access to what data.

Manual work can also lead to costly data quality issues across different teams and departments. Without centralized governance tools, data gets locked down in siloes, which means you won't be able to access and analyze all the data you may need to solve problems or identify large areas of opportunity.

"The key to good [data] governance is figuring out how to define access, then getting out of the way. By that I mean creating exception processes, rather than taking an approach that any time you need data, you have to ask someone in a central organization for it."[6]

Rahul Pathak
VP of Relational Database Engines at AWS

6 Wexler, J., "A unified approach to data governance," CIO Magazine, August 2021

AWS governance portfolio

- Catalog, discover, share, and govern data to allow everyone across your organization to act on it using Amazon DataZone
- Govern and audit the complete AI/ML development cycle using features in Amazon SageMaker
- Catalog data across data lakes, data warehouses, and databases using AWS Glue
- Govern and audit actions taken with data in your data lake on Amazon S3—or govern data sharing in Amazon Redshift—using AWS Lake Formation

Developing a data governance strategy

An AWS/MIT insights report shows that data governance is the top priority of chief data officers (CDOs), with more than 65 percent noting "establishing clear and effective data governance" as their leading responsibility. CDOs also spend a lot of their time on governance—with more than 63 percent saying data governance initiatives are a top focus area.[7]

Without a data governance approach that supports innovation, organizations will find it hard to be data-driven and, ultimately, to remain competitive. After all, the more time workers spend grappling with data, the less time they spend innovating with it. For success, a data governance strategy must align with funded business initiatives. Because a standalone data governance strategy without integration into funded business initiatives risks failure.

AWS is investing across the data journey to enable end-to-end data governance with less effort. With data governance on AWS, organizations have control over where their data sits, who or what has access to it, and what can be done with it at every step of the data workflow. AWS offers a portfolio of services that help organizations understand, curate, and protect the data.

For example, with Amazon DataZone, administrators and data stewards can manage and govern access to data—while data engineers, data scientists, product managers, analysts, and other business users can discover, use, and collaborate with that data to drive insights for your business. With the new generative AI capability in Amazon DataZone—AI recommendations for descriptions—you can automatically generate business descriptions and context for your data sources, including highlighting essential columns, suggesting relevant analytical use cases, and even highlighting potential risks such as PII data. For your AI and ML models, Amazon SageMaker can help you govern and audit the end-to-end ML development cycle. Additionally, data engineers and developers can catalog data across data lakes, data warehouses, and databases using AWS Glue—which comes with data quality rules that check for data freshness, accuracy, and integrity. And with AWS Lake Formation, you can govern and audit the actions taken with data in your data lake on Amazon S3 and data sharing in Amazon Redshift.

7 "2024 CDO Insights: Data & Generative AI" AWS/MIT Insight Report, 2024

Using AI and ML responsibly

AI and ML are some of the most transformative technologies we will encounter in our generation—to tackle business and societal problems, improve customer experiences, and spur innovation. With the widespread use and growing scale of AI, it's critical that the technology is built in a responsible way.

While most companies have begun their journey to responsible AI, the majority (95 percent) are struggling to operationalize across key elements of responsible AI, including fairness, explainability, robustness, security and privacy, transparency, and governance. Ensure your AI and ML models are built in a responsible way—with data practices that mitigate bias and protect data privacy—using AWS resources.

- Leverage purpose-built capabilities like Amazon SageMaker Clarify to identify and address potential biases in their machine learning models
- Enable transparency with tools like AWS AI Service Cards and Amazon SageMaker Model Cards
- Train your teams on fairness and bias with a course from Machine Learning University
- Give data scientists governance controls to gain end-to-end visibility into ML models (including training, version history, and model performance) with SageMaker
- Build generative AI applications responsibly with Amazon Titan foundation models, which can detect and remove harmful content in the data, reject inappropriate content in the user input, and filter the models' outputs that contain inappropriate content such as hate speech, profanity, and violence. Titan Image Generator has built-in safeguards against biases and includes invisible watermarks on all images to identify them as AI-generated
- Use Guardrails for Amazon Bedrock to implement safeguards to your use cases and responsible AI policies

Data Foundation Tip

Through 2025, Gartner estimates 80 percent of organizations seeking to scale digital business will fail because they don't take a modern approach to data and analytics governance.[8] For success, you need to make sure your data is in the right condition to succeed—with the right data available to the right people and applications. Download the Data Governance Master Class guide for step-by-step guidance on modern data governance—from planning and organizing to building and deploying.

8 "12 Data and Analytics Trends to Keep on Your Radar," Gartner, April 5, 2022

By simplifying data governance, OneFootball saw a 40 percent increase in the utilization of its analytics platform

ONEFOOTBALL

OneFootball has grown rapidly to become one of the world's most popular digital media platforms for soccer ("fútbol") enthusiasts. To better use data for the benefit of the company and 70 million fans of "the beautiful game," OneFootball built a nimbler solution on AWS in just a few days. Since integrating data from its inefficient backend databases into its cloud-based data lake, OneFootball has radically simplified data ingestion and eliminated legacy ETL workloads altogether. Beautiful game, indeed.

OneFootball simplified security management and data governance at scale with AWS Lake Formation—making data more accessible across the company. This ultimately helped the team see substantial growth in weekly active analytics users and other internal key performance indicators, increasing usage of the analytics platform by 40 percent.

"Now we can spend more time understanding business problems rather than maintaining different types of database extractions. Seeing more and more people across the organization make use of analytics on a daily basis is a great achievement. Having all relevant data sources reliably integrated was a prerequisite."

Stephan Durry,
Head of Data & Insights at OneFootball

Read the full story >

Pinterest puts customers first with data governance

Pinterest

To ensure its growing data won't outgrow its existing controls, Pinterest built a scalable, automated, fine-grained access control (FGAC) system using Amazon S3. FGAC controls access to data and is based on multiple criteria, offering options such as role-based access control plus security for petabyte-scale datasets. The company also enabled creators and businesses on the platform to self-identify as members of an underrepresented group—while ensuring sensitive data wouldn't be used for any other purpose, such as advertising.

Read the full story >

"Customer-facing impacts of Pinterest's [data] governance efforts include using self-identifying data in a "very controlled way" to support Black-owned businesses for Juneteenth. Creators can also add badges to their profiles—which allows creator content to appear in themed spaces on Pinterest—to show that businesses are owned by someone who identifies with an underrepresented group."

David Chaiken,
Chief Architect, Pinterest

aws

Making security more strategic

AWS has prioritized security since day one—with continuously protected, high-performing, resilient, and efficient infrastructure for your workloads and applications. World-class security experts who monitor the AWS infrastructure also build and maintain our broad selection of innovative security services—which can help simplify the complexities of your own security and regulatory requirements.

AWS Security services and solutions can enable a mix of important advantages:

- **Getting to insights faster** – Provide the right level of access to your resources at all times while maintaining confidence that your data is protected. AWS Security is built with performance in mind, so you get maximum protection and data governance that doesn't slow you down.
- **Reducing downtime** – Tougher, more modern cloud security helps keep your enterprise moving, so you don't have to stop analyzing data to perform a discrete security process—it can be integrated every step of the way.
- **Staying within your budget** – AWS keeps security cost-effective and scales with the evolving needs of your security risks and requirements, protecting your organization's investments and its commitment to data initiatives.
- **Keeping your focus** – From infrastructure to services, AWS is secure by taking security into account at every step, so you can spend more time transforming data into better decisions that drive business results and less time worrying about security and governance.

A history of unmatched reliability and security

- **Amazon S3**:
  Store and retrieve any amount of data with the best security
- **AWS Lake Formation**:
  Build a secure data lake in days with fine-grained access control
- **Multi-AZ Regions**:
  Ensure seamless failovers if an Availability Zone (AZ) is disrupted

CONCLUSION

The next wave of innovation will be driven by data and AI

Leaders have been on a quest to make data a strategic asset to fuel innovation with data and AI. Innovation is accelerating everywhere. But whether it's ML, AI, or generative AI, success depends on relevant, high-quality data, which is why leaders must be tenacious about building a solid data foundation.

Building the right data foundation for your organization is possible—no matter its size, location, or business needs. AWS provides the most comprehensive set of services for any workload, type of data, and desired outcome.

Learn more about why AWS is the best place to unlock value from your data and turn real-time insights into meaningful innovation. And explore how we can help your teams with infrastructure, tooling, and implementation support via the world's leading professional services and partner network. When it comes to data and AI, AWS customers know how to do it better.

Learn more about reinventing your organization to be data-driven >

Discover the top data use cases to maximize business value

How can you take advantage of your data to improve customer experiences, optimize and reinvent supply chains, elevate decision making, build modern applications, and more? Explore prime improvement opportunities for your organization in the **8 essential, data-driven solution areas for leaders: Maximizing business value with data** eBook.

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.