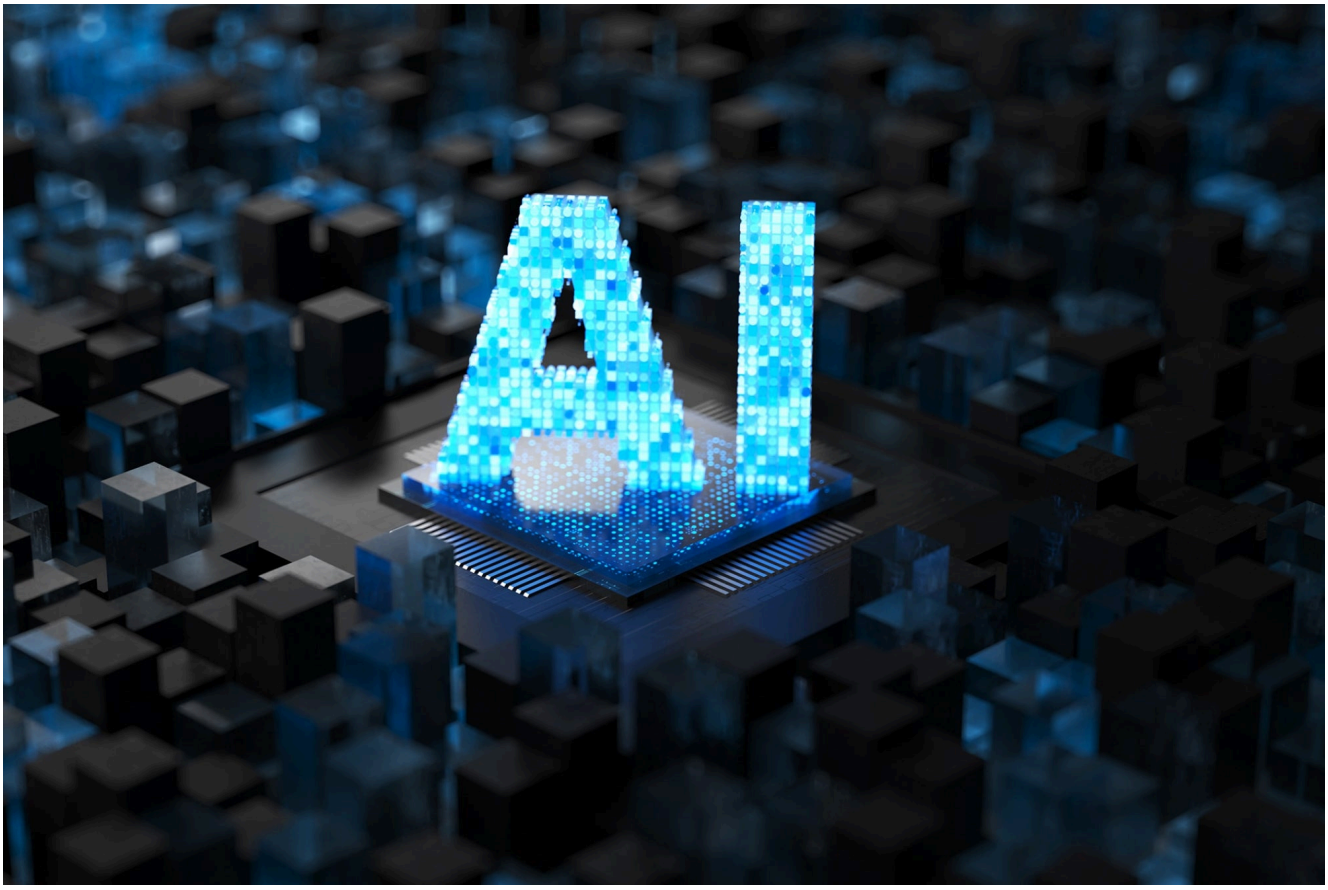
5 Key Data and AI Innovations to Keep an Eye on in 2025



Opinions expressed by Entrepreneur contributors are their own.

At the end of the first quarter in 2025, now is a good time to reflect upon the recent updates from Amazon Web Services (AWS) to their services that provide data and AI capabilities to end customers. At the end of 2024, AWS hosted 60,000+ practitioners at their annual conference, re:Invent, in Las Vegas.

Hundreds of features and services were announced during the week; I've combined these with the announcements that have come since and curated five key data and AI innovations that you should take notice of. Let's dive in.

The next generation of Amazon SageMaker

Amazon SageMaker has historically been seen as the center for everything AI in AWS. Services like Amazon Glue or Elastic MapReduce have taken care of data processing tasks, with Amazon Redshift picking up the task of SQL analytics. With an increasing number of organizations focusing efforts on data and AI, all-in-one platforms such as Databricks have understandably caught the eyes of those starting their journey.

The next generation of Amazon SageMaker is AWS's answer to these services. SageMaker Unified Studio brings together SQL analytics, data processing, AI model development and generative AI application development under one roof. This is all built on top of the foundations of another new service – SageMaker Lakehouse – with data and AI governance integrated through what previously existed standalone as Amazon DataZone.

The promise of an AWS first-party solution for customers looking to get started with, increase the capability of, or gain better control of their data and AI workloads is exciting indeed.

Amazon Bedrock Marketplace

Sticking with the theme of AI workloads, I want to highlight Amazon Bedrock Marketplace. The world of generative AI is fast-moving, and new models are being developed all the time. Through Bedrock, customers can access the most popular models on a serverless basis – only paying for the input/output tokens that they use. To do this for every specialized industry model that customers may want to access is not scalable, however.

Amazon Bedrock Marketplace is the answer to this. Previously,

customers could use Amazon SageMaker JumpStart to deploy LLMs to your AWS account in a managed way; this excluded them from the Bedrock features that were being actively developed (Agents, Flows, Knowledge Bases etc.), though. With Bedrock Marketplace, customers can select from 100+ (and growing) specialized models, including those from HuggingFace and DeepSeek, deploy them to a managed endpoint and access them through the standard Bedrock APIs.

This results in a more seamless experience and makes experimenting with different models significantly easier (including customers' own fine-tuned models).

Amazon Bedrock Data Automation

Extracting insights from unstructured data (documents, audio, images, video) is something that LLMs have proven themselves to excel at. While the potential value borne from this is enormous, setting up performant, scalable, cost-effective and secure pipelines to extract this is something that can be complicated, and customers have historically struggled with it.

In recent days – at time of writing – Amazon Bedrock Data Automation reached General Availability (GA). This service sets out to solve the exact problem I've just described. Let's focus on the document use case.

Intelligent Document Processing (IDP) isn't a new use case for AI – it existed long before GenAI was all the rage. IDP can unlock huge efficiencies for organizations that deal in paper-based forms when augmenting or replacing the manual processes that are performed by humans.

With Bedrock Data Automation, the heavy-lifting of building IDP pipelines is abstracted away from customers and provided as a managed service that's easy to consume and subsequently integrate into legacy processes and systems.

Amazon Aurora DSQL

Databases are an example of a tool where the level of complexity exposed to those leveraging it is not necessarily correlated with how complex it is behind the scenes. Often, it's an inverse relationship where the simpler and more "magic" a database is to use, the more complex it is in the areas that are unseen.

Amazon Aurora DSQL is a great example of such a tool where it's as straightforward to use as AWS's other managed database services, but the level of engineering complexity to make its feature set possible is huge. Speaking of its feature set, let's look at that.

Aurora DSQL sets out to be *the* service of choice for workloads that need durable, strongly consistent, active-active databases across multiple regions or availability zones. Multi-region, or multi-AZ databases, are already well established in active-passive configurations (i.e., one writer and many read-replicas); active-active is a problem that's much harder to solve while still being performant and retaining strong consistency.

If you're interested in reading the deep technical details of challenges that were overcome in the building of this service, I'd recommend reading Marc Brooker's (Distinguished Engineer at AWS) series of blog posts on the topic.

When announcing the service, AWS described it as providing "virtually unlimited horizontal scaling with the flexibility to independently scale reads, writes, compute, and storage. It automatically scales to meet any workload demand without database sharding or instance upgrades. Its active-active distributed architecture is designed for 99.99% single-Region and 99.999% multi-Region availability with no single point of failure, and automated failure recovery."

For organizations where global scale is an aspiration or requirement, building on top of a foundation of Aurora DSQL sets them up very nicely.

Expansion of zero-ETL features

AWS has been pushing the “zero-ETL” vision for a couple of years now, with the aspiration being to make moving data between purpose-built services as easy as possible. An example would be moving transactional data from a PostgreSQL database running on Amazon Aurora to a database designed for large-scale analytics like Amazon Redshift.

While there has been a relatively continuous flow of new announcements in this area, the end of 2024 and start of 2025 saw a flurry that accompanied the new AWS services released at re:Invent.

There are far too many to talk about here in any level of detail that’d provide value; to find out more about all of the available zero-ETL integrations between AWS services, please visit AWS’s dedicated zero-ETL page.

Wrapping this up, we’ve covered five areas relating to data and AI that AWS is innovating in to make building, growing and streamlining organizations easier. All of these areas are relevant to small and growing startups, as well as billion-dollar enterprises. AWS and other cloud service providers are there to abstract away the complexity and heavy lifting, leaving you to focus on building your business logic.

https://www.entrepreneur.com/science-technology/5-key-data-and-ai-innovations-to-keep-an-eye-on-in-2025/489228