

Amazon S3 - Cloud Object Storage - AWS

Archived: 2026-04-05 14:03:46 UTC

- [Products](#)
- [Storage](#)
- Amazon S3

Amazon S3: 20 years of building what's next

From storage for the internet to the data foundation for AI and analytics, Amazon S3 scales and evolves with your ambition

What is Amazon S3?

Amazon Simple Storage Service (Amazon S3) is an object storage service offering industry-leading scalability, data availability, security, and performance. Millions of customers of all sizes and industries store, manage, analyze, and protect any amount of data for virtually any use case, such as data lakes, cloud-native applications, and mobile apps. With cost-effective storage classes and easy-to-use management features, you can optimize costs, organize and analyze data, and configure fine-tuned access controls to meet specific business and compliance requirements.

An error occurred.

Unable to execute JavaScript.



Video Player is loading.

Current Time 0:00

Duration 1:31

Remaining Time 1:31

Benefits

Amazon S3: 20 launches that defined 20 years of storage innovation

Explore key S3 launches across architecture, storage classes, security, and management that transformed how customers store, analyze, and build with data.[Learn more](#)

Durable storage at scale

Amazon S3 launched on Pi Day, March 14, 2006, as the first generally available AWS service. Amazon S3 gave developers something new: a simple, durable, and secure-by-default way to store virtually unlimited data, from anywhere, at any time. Twenty years later, S3 stores hundreds of exabytes of data and more than 500 trillion objects and serves millions of customers worldwide.

As data accumulated faster than enterprises could afford to keep it, S3 Glacier became the cloud on-ramp for financial services, healthcare, and media. S3 Glacier Deep Archive pushed costs lower still, and S3 Glacier Instant Retrieval modernized archival storage by combining archive pricing with instant access.

Data foundation for AI

S3 became the foundation of modern data lakes, enabling any organization to centralize structured and unstructured data, run analytics, and build applications without managing complex infrastructure. S3 Intelligent-Tiering rewrote the economics of storage, saving S3 customers more than \$6 billion. S3 didn't just scale with industry shifts—it helped accelerate them.

As AI training and generative AI raised the bar for what storage needed to deliver, S3 evolved. S3 Express One Zone is the fastest cloud object storage class, S3 Tables deliver native Apache Iceberg support, and S3 Vectors provides native vector storage and query for AI applications.

What began as storage for the internet is now the data foundation the world runs on.

Amazon S3 today

500T+

objects and hundreds of exabytes of data

200M+

requests per second on average

300B+

event notifications sent to serverless applications daily

\$6B+

saved in storage costs with S3 Intelligent-Tiering

Use cases


Build an open data architecture for analytics and AI

More than 1,000,000 data lakes run on AWS. Whether you're building a data lake or lakehouse architecture, Amazon S3 provides the ideal foundation for petabyte-scale analytics and AI workloads because of its unmatched durability, availability, scalability, security, and price performance. S3 supports open table formats like Apache Iceberg, allowing you to run queries directly against a single, reliable copy of data using your preferred AWS or third-party analytics engine. For a fully managed experience, S3 Tables automatically handles compaction, snapshot management, and file cleanup of Apache Iceberg tables, reducing the operational burden of managing a data lakehouse and scaling performance as your data volumes grow.



Accelerate performance-critical applications


Building on the scale and durability of Amazon S3, the S3 Express One Zone storage class accelerates performance-intensive workloads with consistent single-digit millisecond latency and up to 10x faster data access than the S3 Standard storage class. It's ideal for high-throughput workloads such as AI training and inference, real-time analytics, media processing, and interactive applications. As the fastest cloud object storage, S3 Express One Zone delivers faster access to the most frequently accessed dataset, improves compute efficiency, and lowers API request costs, reducing the total cost of ownership for the most demanding workloads.



Scale and differentiate generative AI and agentic applications


Build generative AI and agentic applications on Amazon S3, the durable and scalable foundation for your data. Access diverse data types at scale—including unstructured, structured, streaming, and vector data—to train, fine-tune, and customize models or improve contextual understanding through RAG. With direct integrations across AWS analytics and AI/ML services, S3 accelerates the path from data to intelligence without managing complex infrastructure.

S3 simplifies data pipelines, protects critical information with security by default, and delivers cost-efficient scalability so you can support evolving AI workloads. Whether you're developing domain-specific assistants, intelligent agents, or personalized generative AI experiences, S3 provides the trusted data foundation to build, customize, and deploy quickly and efficiently.




Optimize vector storage and query costs for AI and semantic search

Semantic search enables AI applications to understand the meaning and context of data by using vector embeddings to represent relationships across content such as documents, images, and videos. Amazon S3 Vectors brings native vector support to S3, allowing you to store and query vectors alongside your source data in a fully serverless architecture. By reducing the cost of storing and querying vectors by up to 90% while maintaining subsecond query performance, S3 Vectors makes it cost-effective to create and use large vector datasets to improve the memory and context of AI agents as well as conduct semantic search results of your S3 data. With S3 Vectors, developers can get started quickly, reduce operational complexity, and scale vector-driven semantic search and AI workloads efficiently on a durable, trusted foundation.



Back up, restore, and archive data while meeting compliance requirements

Meet your recovery time objective (RTO), recovery point objective (RPO), and compliance requirements with the robust replication functionality of S3, data protection with AWS Backup, and various AWS Partner Network solutions. For cost-optimized, long-term data storage, archive data using Amazon S3 Glacier storage classes to lower costs, eliminate operational complexities, and gain new insights.



Customers

Get started

Source: <https://aws.amazon.com/s3/>