



AWS DataSync pricing



AWS DataSync is a high-speed, online data transfer service that makes it easy to securely move your data anywhere its needed: on premises, at the edge, or in the cloud. DataSync offers simple, predictable, pay-as-you-go pricing. You pay a flat, per-gigabyte fee for the amount of data transferred between your storage locations. Tasks configured for Enhanced mode are also charged a fee per task execution. These fees cover the costs to move, verify, and synchronize your data. With DataSync, there are no upfront costs and no minimum charges.

Enhanced mode data transfers

Using Enhanced mode, you can transfer virtually unlimited numbers of objects between your storage locations with higher performance than Basic mode. Enhanced mode tasks optimize and streamline the data transfer process by listing, preparing, transferring, and verifying data in parallel. Enhanced mode is currently available for transfers between Amazon S3 locations, between other clouds and Amazon S3 locations between other clouds and Amazon S3, and between on-premises NFS or SMB file servers and Amazon S3.

Basic mode data transfers

Basic mode is subject to [quotas](#) on the number of files and objects in a dataset. Basic mode sequentially prepares, transfers, and verifies files and objects in a dataset, making it slower than Enhanced mode for most workloads. Basic mode supports all DataSync location types available today.

Region:
Europe (London) ▼

Enhanced mode	Price
Data transferred	\$0.015 per gigabyte (GB)
Task executions	\$0.55 per task execution

Basic mode	Price
Data transferred	\$0.0125 per gigabyte (GB)



AWS Pricing Calculator

Calculate your DataSync and architecture cost in a single estimate.

[Create your custom estimate now »](#)

Additional charges

When using DataSync to copy data, you are charged standard request, storage, and data transfer rates to read from and write to AWS services, such as [Amazon Simple Storage Service \(Amazon S3\)](#), [Amazon Elastic File System \(Amazon EFS\)](#), [Amazon FSx for Windows File Server](#), [Amazon FSx for Lustre](#), [Amazon FSx for OpenZFS](#), [Amazon FSx for NetApp ONTAP](#), and [AWS Key Management Service \(AWS KMS\)](#).

You will be billed by [AWS PrivateLink](#) for interface Amazon Virtual Private Cloud (Amazon VPC) endpoints that you create to manage and control the traffic between your agent(s) and the DataSync service over PrivateLink. Only control traffic is subject to PrivateLink charges. This includes commands from the DataSync service to the agent, such as StartTaskExecution. The files or objects transferred by DataSync are not subject to PrivateLink charges. You are also charged standard rates for Amazon CloudWatch Logs, Amazon CloudWatch Metrics, and Amazon CloudWatch Events.

DataSync uses AWS Secrets Manager to store or access your storage credentials. You can configure DataSync to create secrets on your behalf or you can provide a custom secret. You are billed by [AWS Secrets Manager](#) only for the use of secrets that you create outside of DataSync.

Copying data to and from Amazon S3

When copying data to or from S3, you are charged standard request charges for the following S3 requests: LIST, HEAD, GET, PUT, and COPY. The exact operations performed depend upon the configuration of your DataSync task. For more information, read about [evaluating S3 request costs](#) in our documentation. Requests and storage charges can vary between S3 storage classes. You may also incur additional charges when reading and writing objects in different S3 storage classes. Read about [working with Amazon S3 storage classes](#) in our documentation.

Copying data to and from AWS Regions

When copying data from AWS to an on-premises storage system, you pay for AWS Data Transfer OUT from your source AWS Region. When copying data between Regions, you pay for Data Transfer between Regions (billed as Data Transfer OUT from your source Region to your destination Region). When copying data between AWS services using the AWS DataSync EC2 agent, you pay for Data Transfer between Regions and for Data Transfer between Availability Zones. To learn more, see [Data Transfer pricing](#).

Enhanced mode pricing examples

The pricing examples below assume that the DataSync data transfer task has been configured for Enhanced mode to achieve higher performance, scale, and enhanced monitoring and metrics.

Example 1: Copy a 50 TB dataset from one S3 bucket to another, within the same Region, using the S3 Standard storage class

DataSync will perform a one-time copy of objects in one S3 bucket to a destination bucket in the same Region. Objects in both the source and destination buckets are created with the S3 Standard storage class. When the DataSync task executes, it will scan the existing objects in the source bucket to determine which objects should be copied. The scan process involves, at a minimum, LIST and HEAD requests on objects in each S3 bucket, as explained in [the AWS DataSync documentation](#). Assume the source dataset contains 5 million objects.

(50 TB * 1024 GB copied using DataSync * \$0.015/GB) +
(5,000 S3 LIST requests * \$0.005 / 1000) +
(5 million S3 HEAD requests * \$0.0004 / 1000) +
(5 million S3 GET requests * \$0.0004 / 1000) +
(5 million S3 PUT requests * \$0.005 / 1000) +
(1 task execution * \$0.55)
= \$768 + \$0.025 + \$2 + \$2 + \$25 + \$0.55
= \$797.58

Example 2: Perform a daily replication of a 5 TB dataset from one S3 bucket to another, within the same Region, using the S3 Glacier Instant Retrieval storage class

DataSync will replicate objects in one S3 bucket to a destination bucket in the same Region, using a schedule configured to run once per day. Objects in both the source and destination buckets were created with the Glacier Instant Retrieval storage class. For each daily task execution, DataSync will scan the existing objects in the source and destination buckets to determine which objects should be copied. The scan process involves, at a minimum, LIST and HEAD requests on objects in each S3 bucket, as explained in [the AWS DataSync documentation](#). In this example, we assume a 5% change rate per day.

Cost to use DataSync for each daily replication of 5 TB of 16 MB files into Amazon S3 in the US East (Ohio) Region, assuming 4000 LIST requests in each S3 bucket, one GET per object copied, and one HEAD request per object in each S3 bucket:

Amount of data copied incrementally is 5% of 5TB = 256 GB

(256 GB copied using DataSync * \$0.015/GB) +
(256 GB retrieved from S3 Glacier Instant Retrieval * \$0.03) +
(4000 S3 LIST requests * 2 * \$0.02 / 1000) +
(5 TB * 1024 GB * 1024 MB / 16 MB S3 HEAD requests * 2 * \$0.01 / 1000) +
(256 GB * 1024 MB / 16 MB S3 GET requests * \$0.01 / 1000) +
(256 GB * 1024 MB / 16 MB S3 PUT requests * \$0.02 / 1000) +
(1 task execution per day * \$0.55)
= \$3.84 + \$7.68 + \$0.16 + \$6.55 + \$0.16 + \$0.32 + \$0.55
= \$19.26 / day

Basic mode pricing examples

The pricing examples below assume that the DataSync data transfer task has been configured for Basic mode.

Example 1: Migrate 50 TB from on-premises storage to an empty S3 bucket using the S3 Standard storage class

Cost to use DataSync for a one-time migration of 50 TB of 16 MB files into an empty Amazon S3 bucket in the US East (Ohio) Region:

(50 TB copied into S3 * 1024 GB * \$0.0125/GB) + (1 S3 LIST request * \$0.005 / 1000) + (50 TB / 16 MB S3 PUT requests * \$0.005 / 1000)
= \$640 + \$0 + \$16.38
= \$656.38

Example 2: Incrementally copy 10 TB from on-premises storage to an S3 bucket with 50 TB of existing data using the S3 Standard storage class

In this example, DataSync will scan the existing objects in the S3 bucket to determine which data should be copied from on premises. The scan process involves, at a minimum, LIST and HEAD requests on objects in the S3 bucket, as explained in [the AWS DataSync documentation](#).

Cost to use DataSync for each incremental copy of 10 TB of 16 MB files into S3 in the US East (Ohio) Region, assuming 500 LIST requests and one HEAD request per object in S3:

(10 TB copied into S3 using DataSync * 1024 GB * \$0.0125/GB) +
(500 S3 LIST requests * \$0.005 / 1000) +
(50 * 1024 * 1024 MB / 16 MB S3 HEAD requests * \$0.0004 / 1000) +
(10 * 1024 * 1024 MB / 16 MB S3 PUT requests * \$0.005 / 1000)
= \$128 + \$0 + \$1.31 + \$3.28
= \$132.59

Example 3: Perform a daily replication of a 5 TB dataset from one S3 bucket to another, within the same Region, using the S3 Glacier Instant Retrieval storage class

In this example, DataSync will replicate objects in one S3 bucket to a destination bucket in the same Region. Objects in both the source and destination buckets were created with the Glacier Instant Retrieval storage class. For each daily task execution, DataSync will scan the existing objects in the source and destination buckets to determine which objects should be copied. The scan process involves, at a minimum, LIST and HEAD requests on objects in each S3 bucket, as explained in [the AWS DataSync documentation](#). In this example, we assume a 5% change rate per day.

Cost to use DataSync for each daily replication of 5 TB of 16 MB files into Amazon S3 in the US East (Ohio) Region, assuming 4000 LIST requests in each S3 bucket, one GET per object copied, and one HEAD request per object in each S3 bucket:

Amount of data copied incrementally is 5% of 5TB = 256 GB
(256 GB copied using DataSync * \$0.0125/GB) +
(256 GB retrieved from S3 Glacier Instant Retrieval * \$0.03) +
(4000 S3 LIST requests * 2 * \$0.02 / 1000) +
(5 TB * 1024 GB * 1024 MB / 16 MB S3 HEAD requests * 2 * \$0.01 / 1000) +
(256 GB * 1024 MB / 16 MB S3 GET requests * \$0.01 / 1000) +
(256 GB * 1024 MB / 16 MB S3 PUT requests * \$0.02 / 1000)
= \$3.20 + \$7.68 + \$0.16 + \$6.55 + \$0.16 + \$0.32
= \$18.07 / day

Example 4: Migrate 100 TB from a self-managed server to Amazon FSx for Windows File Server

Cost to migrate 100 TB from an on-premises Windows File Server to Amazon FSx, using DataSync:

(100 TB copied into Amazon FSx For Windows File Server * 1024 GB * \$0.0125/GB)
= \$1,280.00

Example 5: Ongoing data transfers for an in-cloud processing and data distribution workflow

Suppose you perform a daily 1 TB transfer into Amazon EFS. Your cost for 31 days will be:

(1 TB copied into EFS * 1024 GB * \$0.0125/GB) * 31 days
= \$396.80

Example 6: File system replication for business continuity with incremental transfers

Suppose you have a data replication job that copies a 10 TB file system into Amazon EFS in the Asia Pacific (Sydney) Region once and replicates 1 TB of changes every day. Your cost for the initial copy and the 31 following days of normal usage will be:

10 TB * 1024 GB * \$0.0125/GB = \$128 for the initial transfer
+ 1 TB * 1024 GB * \$0.0125/GB * 31 days for the incremental transfers
= \$396.80 per month on an ongoing basis

Additional pricing resources

[AWS Pricing Calculator](#)
Easily calculate your monthly costs with AWS.

[Get pricing assistance](#)
Contact AWS specialists to get a personalized quote.

Get Started with AWS

Learn how to start using AWS in minutes





AWS Free Tier

Gain free, hands-on experience with AWS for 12 months



Get started



Learn

What Is AWS?
What Is Cloud Computing?
What Is Agentic AI?
Cloud Computing Concepts Hub
AWS Cloud Security
What's New
Blogs
Press Releases

Resources

Getting Started
Training
AWS Trust Center
AWS Solutions Library
Architecture Center
Product and Technical FAQs
Analyst Reports
AWS Partners

Developers

Builder Center
SDKs & Tools
.NET on AWS
Python on AWS
Java on AWS
PHP on AWS
JavaScript on AWS

Help

Contact Us
File a Support Ticket
AWS rePost
Knowledge Center
AWS Support Overview
Get Expert Help
AWS Accessibility
Legal

[Back to top](#) ↑