Amazon Redshift pricing

Get started with Amazon Redshift | Contact us

Pricing overview

Amazon Redshift offers two deployment options: Provisioned and Serverless. Provisioned clusters are billed per-second, with no per-query charges, and you can choose from on-demand or 1/3 year Reserved Instance options to suit your workload profile. Serverless automatically provisions and intelligently scales data warehouse capacity to deliver fast performance for even the most demanding and unpredictable workloads, and you pay only for the compute that you use. You can select On-Demand Instances with hourly billing and no long-term commitments or get significantly lower effective hourly rates with Reserved Instances. Reserved Instances are available in three flexible options: 1 year No Upfront, 1 year All Upfront, and 3 year All Upfront. Both Provisioned and Serverless options include managed storage, and you can scale compute and storage separately. Amazon Redshift Provisioned, you can select On-Demand Instances with hourly billing and no long-term commitments or get significantly lower effective hourly rates with Reserved Instances. Reserved Instances are available in three flexible options: 1 year No Upfront, 1 year All Upfront, and 3 year All Upfront. Both Provisioned and Serverless options include managed storage, and you can scale compute and storage separately. Amazon Redshift has two billing components: instance charges and storage charges. For Amazon Redshift Serverless workloads, you pay for the compute capacity you use by the second, and you can use Amazon Redshift Serverless without having to make any capacity decisions. Amazon Redshift Serverless measures compute capacity in Redshift Processing Units (RPUs). You pay for the RPUs you use in 60-second increments. For more information about Amazon Redshift Serverless pricing, see Amazon Redshift Serverless pricing.

If this is your first time using Amazon Redshift, you may be eligible for the Amazon Redshift free trial. For more information, see the Amazon Redshift Free Trial section below.

You can optionally use Basic, Plus, Pro and Ultra plans for Amazon Redshift to get the compute capacity your data warehouse scenarios within it's runtime. Your data warehouse capacity automatically scales to provide consistent performance and unlimited concurrency. For example, you can choose the Basic plan for your data warehouse with 32 RPUs and the Pro plan for your data warehouse with 128 RPUs. Amazon Redshift Serverless then automatically scales to provide consistent performance and cost-efficiency. For more information about Amazon Redshift Serverless pricing, see Amazon Redshift Serverless pricing.

Amazon Redshift Serverless is billed for the RPUs you use in 60-second increments. You can also set a limit for the maximum number of RPUs that Amazon Redshift Serverless can use. Amazon Redshift Serverless then automatically scales to provide consistent performance and cost-efficiency. For more information about Amazon Redshift Serverless pricing, see Amazon Redshift Serverless pricing.

Amazon Redshift Serverless is billed for Amazon Redshift Managed Storage (RMS) and storage used for user snapshots is billed at the standard Amazon Redshift Managed Storage rate. For more information about Amazon Redshift Managed Storage pricing, see Amazon Redshift Managed Storage pricing.

Amazon Redshift Serverless is billed for Amazon Redshift Managed Storage (RMS) and storage used for user snapshots is billed at the standard Amazon Redshift Managed Storage rate. For more information about Amazon Redshift Managed Storage pricing, see Amazon Redshift Managed Storage pricing.

Amazon Redshift Serverless is billed for Amazon Redshift Managed Storage (RMS) and storage used for user snapshots is billed at the standard Amazon Redshift Managed Storage rate. For more information about Amazon Redshift Managed Storage pricing, see Amazon Redshift Managed Storage pricing.

Region: US East (N. Virginia)

Amazon Redshift Serverless is billed for Amazon Redshift Managed Storage (RMS) and storage used for user snapshots is billed at the standard Amazon Redshift Managed Storage rate. For more information about Amazon Redshift Managed Storage pricing, see Amazon Redshift Managed Storage pricing.

Region: US East (N. Virginia)

Serverless Reservations for Amazon Redshift Serverless

Please note the following reference to view Serverless Reservation rates for Amazon Redshift Serverless:

Select a Number type and region: US East (N. Virginia)

Select a term length and payment option for your Redshift Serverless Reservation: 1 year

Payment options: No Upfront

Viewing 1 available Serverless Reservations (US)

All upfront fee	All monthly fees	All effective hourly rate	Savings over Serverless On-Demand	On-demand rate
$1120	$0.208	0.208	20%	$0.260

*Serverless reservations can be purchased at a usage amount in Redshift Processing Units (RPUs) and are billed whether or not you use them.

Amazon Redshift managed storage pricing

Amazon Redshift Managed Storage (RMS) provides you the flexibility to scale and pay for storage and compute separately. Managed storage scales automatically with RMS node types, and you pay only the costs for the managed storage that you use. For example, if you have 2 TB of data stored in Amazon Redshift Managed Storage, your monthly charges would be the same as if you stored 2 TB of data in Amazon S3. You can calculate the amount of data stored in Amazon Redshift Clusters or the AWS Management Console. You can also log into the Amazon Redshift Console and view the Cluster Summary page. For more information about Amazon Redshift Managed Storage, see the Amazon Redshift Managed Storage page.

Region: US East (N. Virginia)

Pricing example for managed storage pricing

Let's assume that your total 100 GB data in managed storage with RMS nodes type for the 13 hours in April, and 100 TB data for the first 10 days in April.

In the first 10 days in April, you used 100 TB-Hours (100 TB x 240 hours) in managed storage. In the last 3 days in April, you used 100 GB-Hours (100 GB x 72 hours) in managed storage. Your total usage for the billing period is 100 TB-Hours + 100 GB-Hours.

At the end of April, Amazon Redshift charges you $0.024 per GB-Hour for 100 GB-Hours + $0.024 per TB-Hour for 100 TB-Hours = $2.4 + $2,400 = $2,402.40 for Amazon Redshift Managed Storage in the month of April.

If this data was stored in the US (N. Virginia) Region, managed storage will be charged at $0.024/GB-Month. Monthly storage charges for 10 TB (10,000 GB-Month) will be 10,000 GB-Month x $0.024 per GB-Month = $240.

Total RMS charges will = $2,640.

Amazon Redshift Spectrum pricing

Amazon Redshift Spectrum allows you to directly run SQL queries against exabytes of data in Amazon S3. You are charged for the number of bytes scanned by Redshift Spectrum, rounded up to the next megabyte, with a 10 MB minimum per query. There are no charges for DDL statements. Compressed data can be scanned by Redshift Spectrum, but the charges are based on the uncompressed data. For example, if you have 1 TB of data stored in Amazon S3 in compressed format and Redshift Spectrum scans 1 TB of data after decompression, you will be charged $5 for 1 TB of data scanned. For more information about Amazon Redshift Spectrum pricing, see Amazon Redshift Spectrum pricing.

Region: US East (N. Virginia)

Pricing example for managed storage pricing

Let's assume that your total 100 GB data in managed storage with RMS nodes type for the 13 hours in April, and 100 TB data for the first 10 days in April.

In the first 10 days in April, you used 100 TB-Hours (100 TB x 240 hours) in managed storage. In the last 3 days in April, you used 100 GB-Hours (100 GB x 72 hours) in managed storage. Your total usage for the billing period is 100 TB-Hours + 100 GB-Hours.

At the end of April, Amazon Redshift charges you $0.024 per GB-Hour for 100 GB-Hours + $0.024 per TB-Hour for 100 TB-Hours = $2.4 + $2,400 = $2,402.40 for Amazon Redshift Managed Storage in the month of April.

If this data was stored in the US (N. Virginia) Region, managed storage will be charged at $0.024/GB-Month. Monthly storage charges for 10 TB (10,000 GB-Month) will be 10,000 GB-Month x $0.024 per GB-Month = $240.

Total RMS charges will = $2,640.

Additional costs

You can be charged for the Amazon Redshift data copied to your data with Redshift Spectrum. Redshift Spectrum queries data directly in Amazon S3. You are charged standard S3 rates for storing data in your S3 buckets, and the standard S3 rates for requests made against your S3 buckets. See Amazon S3 rates.

If you use the AWS Glue Data Catalog with Amazon Redshift Spectrum, you are charged a standard AWS Glue Data Catalog rate. See the AWS Glue pricing.

When using Amazon Redshift Spectrum to query AWS Key Management Service (KMS) encrypted data in Amazon S3, you are charged standard AWS KMS rates. For details, refer to AWS KMS pricing.

Redshift Spectrum pricing examples for Ad-hoc Querying

Consider a table with 100 equally sized columns stored in Amazon S3 as an uncompressed text file that is 10 TB in size. You have a query that scans only 5 of the 100 columns. Redshift Spectrum recognizes that only 5 of the 100 columns are needed for the query and so it only scans those 5 columns for processing.

If you compress your file using GZIP and you see a 4:1 compression ratio, in this case, you would have a compressed file size of 1 TB. Redshift Spectrum has to scan the entire file, but it will see 4 times more data because the data is compressed at a 4:1 ratio.

Now, let's assume you store the data in Apache Parquet format and you compress it using GZIP. You see a 4:1 compression ratio. Parquet is a columnar format, so Redshift Spectrum needs to scan only the 5 columns it needs for the query, which is 5 GB of data. Redshift Spectrum needs to scan only one column in Parquet format, which is 1 GB of data after compression.

Region: US East (N. Virginia)

Note: The above pricing examples are for illustration purposes only. The compression ratio of different files and columns may vary.

Concurrency Scaling pricing

Amazon Redshift automatically and elastically scales query concurrency to provide consistently fast performance, even with thousands of concurrent users and queries. There are no resources to manage, no upfront costs, and you pay only for the concurrency scaling resources that you use. Amazon Redshift customers earn one hour of free Concurrency Scaling credits for each day they run their cluster. You can choose to pay for any concurrency usage above the free credits at on-demand rates. Concurrency Scaling credits are earned on a per-cluster basis—only when a cluster is running—and they can be shared across the account in which the cluster is running. Credits earned in one region cannot be used in another region. For more information about Concurrency Scaling, see Concurrency Scaling.

Region: US East (N. Virginia)

Amazon Redshift Serverless automatically scales resources up and down as needed to meet workload needs by default and there are no separate charges for Concurrency Scaling.

Region: US East (N. Virginia)

Concurrency Scaling credits

Redshift customers earn one hour of free Concurrency Scaling credits per day. Credits are earned on an hourly basis for each active cluster in your AWS account, and can be accumulated for the same cluster only when credits are not earned. For example, a 10-minute burst of free Concurrency Scaling credits will be earned for every 24 hours that a cluster is running.

Pricing example for Concurrency Scaling

A 10-DC8 node-type Redshift cluster in US West 1 costs $36.00 per hour. Consider a concurrency-scaling scenario in which queries are routed to the remote cluster for the first 60 minutes. The remote cluster is billed for 60 minutes at $36.00 per hour. Assume that this hour credit was used to offset the first hour credit which you earned in the past 24 hours. Therefore, you will be charged for 60 minutes - 60 minutes = 0 minutes at $36.00 per hour = $0.00.

Redshift ML pricing

When you use Redshift ML to create a prediction ML, you qualify for the Amazon SageMaker free tier if you haven't previously used Amazon SageMaker. This includes two free SAGEMAKER.M5.XLARGE instance-hours per month for the first two months. Any usage that exceeds the free tier credits will be charged at the standard SageMaker rates.

Amazon EC2 charges

Redshift ML uses Amazon SageMaker to train your model. Training a model can require the creation of new Amazon EC2 instances that are provisioned by SAGEMAKER in the name of the customer. After training is complete, these instances are terminated. Amazon S3 is used for storing the training data and storing the final ML artifact. Any Amazon S3 usage for ML will be billed at the standard rates of Amazon S3. For more information, see Amazon S3 pricing.

Cost control options

You can control the training cost by setting the MAX_RUNTIME_IN_SECONDS parameter in the CREATE MODEL command. The default value is 5,400 seconds (1.5 hours). If training doesn't complete within the maximum runtime, Amazon EC2 is used for the training time actually consumed by the EC2 instances of the CREATE MODEL. There is also no extra charge. Redshift charges a flat fee for training each model. The upfront pricing ensures that you know the cost before training begins.

	Price
Number of cells	$ per million cells
First 100K cells	$ per million cells
Next 900K cells	$ per million cells
Over 1MM cells	$ per million cells

Note: Pricing will reflect any fees that the open-launch shows above.

Examples of Redshift Models cost:

- 100,000 cells x $20 = 1 cell
- 1,000,000 cells x $20 = 1 cell
- 10,000,000 cells x $20 = 10 cells
- 50,000,000 cells x $20 = 50 cells = $1 cell
- 100,000,000 cells x $20 = 100 cells = $2 cell

If the training data exceeded by 20% MAX_RUNTIME_IN_SECONDS, Redshift ML stops training and uses the artifact. Additionally, if you did not provide one, the CREATE MODEL will, behind the scenes, create approximately 10%, 20 or 5 models of variation" results from the training dataset and will train using these similarly shown inputs. The process is called hyperparameter tuning (HPO). You can control the number of models that HPO tries by setting the MAX_NUM_JOBS parameter.

Reserved Instance pricing

Reserved Instances are available for Amazon Redshift clusters running in multiple availability zones and are available in all supported AWS regions. For more information about Reserved Instances, see Reserved Instances.

You can benefit from significant savings over on-demand rates by committing to use Amazon Redshift for 1 year or three year terms. Reserved Instance pricing is specific to the node type, number of nodes, and region in which you run your cluster. Reserved Instances are available in three flexible options: No Upfront, All Upfront, and Partial Upfront. For more information about Reserved Instances, see Reserved Instances.

Reserved Instances are available in three flexible options: No Upfront, All Upfront, and Partial Upfront. For more information about Reserved Instances, see Reserved Instances.

Reserved Instances are available in three flexible options: No Upfront, All Upfront, and Partial Upfront. For more information about Reserved Instances, see Reserved Instances.

Region: US East (N. Virginia)

This is the concept monthly payment rate for the duration of the Reserved Instance term. Your monthly payment will usually be the product of the concept rate and the number of hours in the month. Because the number of hours in a month varies, your monthly payment will vary slightly. For example, the concept rate is based on the average monthly payments over the term of the Reserved Instance divided by the total number of hours (based on a 365 day year) over the term of the Reserved Instance.

Effective hourly pricing takes into account the Reserved Instance upfront fee and the concept rate. For example, if you purchase a Reserved Instance with a 1 year term, you pay a fixed price for the Reserved Instance and then pay the concept rate for every hour during the entire Reserved Instance term. The effective hourly price is the total of the upfront fee amortized over the hours in the term and the concept rate. The following table shows the effective hourly rate for each Reserved Instance option.

If the Reserved Instance offering you purchase has a partial upfront payment option, you will be charged the associated upfront price and monthly fee. Each month you will be billed for the monthly fee until the end of the term. Partial upfront Reserved Instances are the same as All Upfront Reserved Instances, but with a smaller upfront payment.

We may terminate the Reserved Instance program at any time. In addition to being subject to Reserved Instance pricing, Reserved Instances are subject to all other terms and other fees applicable under the AWS Customer Agreement or other agreement with us governing your use of our services.

Zero-ETL integration costs

Amazon Redshift has zero-ETL integrations with online transactional processing (OLTP) databases and streaming applications, reducing the need to build ETL pipelines. Amazon Aurora, and Amazon DynamoDB. Supported Applications include Amazon Aurora, Amazon DynamoDB, and Amazon Relational Database Service (Amazon RDS), and Oracle DatabaseAWS.

Zero-ETL integrations are billed per-second, with no per-query charges. You pay for compute capacity in Redshift Processing Units (RPUs) that you configure for your Amazon Redshift data warehouse. Amazon Aurora, and Amazon DynamoDB. Supported Applications include Amazon Aurora, Amazon DynamoDB, and Amazon Relational Database Service (Amazon RDS), and Oracle DatabaseAWS. You pay for the RPUs you use in 60-second increments. For more information about Amazon Redshift Serverless, see Amazon Redshift Serverless pricing. For more information about Amazon Redshift Serverless, see Amazon Redshift Serverless pricing.

Amazon Redshift Serverless is billed for the RPUs you use in 60-second increments. You can also set a limit for the maximum number of RPUs that Amazon Redshift Serverless can use. Amazon Redshift Serverless then automatically scales to provide consistent performance and cost-efficiency. For more information about Amazon Redshift Serverless, see Amazon Redshift Serverless pricing.

Amazon Redshift Serverless is billed for the RPUs you use in 60-second increments. You can also set a limit for the maximum number of RPUs that Amazon Redshift Serverless can use. Amazon Redshift Serverless then automatically scales to provide consistent performance and cost-efficiency. For more information about Amazon Redshift Serverless, see Amazon Redshift Serverless pricing.

Region: US East (N. Virginia)

Oracle DatabaseAWS zero-ETL integration pricing

The Cross DatabaseAWS zero-ETL integration is charged for data ingested and billed per storage and billed per EC2 per GB.

	Data ingested (zero-ETL)	CDC per GB
Full Zero-ETL Ongoing	$0.04	$0.08
Full zero-ETL Ongoing	$0.04	$0.08

Backup storage

Backup storage is the storage associated with the snapshots taken for your data warehouse, including your backup retention period or taking additional snapshots to increase the backup storage. Backup storage is billed per GB-month, and you are billed for backup storage beyond the storage allocated for free for your provisioned cluster. Backup storage is billed for the AWS Free Tier. Amazon Redshift Serverless, which is only billed at usage rates for Amazon Redshift Managed Storage (RMS) and storage used for user snapshots is billed at the standard Amazon Redshift Managed Storage rate. For more information about Amazon Redshift Managed Storage pricing, see Amazon Redshift Managed Storage pricing.

Data transfer

There is no charge for data transferred between Amazon Redshift and Amazon S3 within the same AWS Region for backup, restore, load, and unload operations. For all other data transfers into and out of Amazon Redshift, you will be billed at standard AWS data transfer rates. For more information about data transfer pricing, see Amazon EC2 pricing.

Amazon Redshift charges for data transferred between Amazon Redshift and Amazon S3 within the same AWS Region for backup, restore, load, and unload operations. For all other data transfers into and out of Amazon Redshift, you will be billed at standard AWS data transfer rates. For more information about data transfer pricing, see Amazon EC2 pricing.

Amazon Redshift charges for data transferred between Amazon Redshift and Amazon S3 within the same AWS Region for backup, restore, load, and unload operations. For all other data transfers into and out of Amazon Redshift, you will be billed at standard AWS data transfer rates. For more information about data transfer pricing, see Amazon EC2 pricing.

Amazon Redshift charges for data transferred between Amazon Redshift and Amazon S3 within the same AWS Region for backup, restore, load, and unload operations. For all other data transfers into and out of Amazon Redshift, you will be billed at standard AWS data transfer rates. For more information about data transfer pricing, see Amazon EC2 pricing.

Region: US East (N. Virginia)

Pricing Examples

Example for RPU

You can find out what price rates and 50 RPUs for a month. During the month, you also use 20 of that data using Redshift Spectrum and use 20 TB of data. You can be observed pricing.

Your charges would be calculated as follows:

- Redshift Serverless cost = 50 RPUs * 744 hours in a month * $0.375 per RPU-hour = $6,578.00 USD
- RPU cost = 20 TB * 1,000 GB per TB * $5.00 per GB = $100.00 USD
- Redshift Spectrum cost = 20 TB * $5.00 per TB = $100.00 USD

Total monthly cost: $6,682.00 USD

Example for Pro 4 x2

You use a Redshift ML Create Model is deployed in two AZ simultaneously. You cluster the four x2 nodes storage and you can use 50 TB of RMS for a month. You can store observed pricing.

Your charges would be calculated as follows:

- Redshift ML cost = 744 hours in a month * 2 instances (AZ) * $1.00 per hour = $1,488.00 USD
- Redshift RMS storage cost for 50 TB = 50 instances (50 GB each) per 50 hours = $1,200.00 USD
- RMS cost = 50 TB * 1,024 GB per TB * $0.024 USD = $1,228.80 USD

Total monthly cost: $2,621.44 USD

Example for Amazon Redshift Serverless

This example demonstrates how Amazon Redshift charges you when you execute capacity inside for the serverless option.

Basic example

Let's assume that a data warehouse processing job that needs to be executed on a serverless data warehouse in the US East (N. Virginia) Region. For this example, it takes 1 hour and 30 minutes to complete the job when the data warehouse is set to 32 RPUs. It runs once each 24 months, and you have Amazon Redshift Serverless 32 RPUs capacity as its Region.

The following table summarizes your total usage for the day:

Total price	Charges are 1.5 times because 1 hour and 1 PM and 1 are taking. 32 minutes and 30 seconds = 1.5 hours at $1.50 per hour
Capacity used	1.50 RPU
Daily charges	$1.50 (1.50 RPU * 1.50 RPU = $2.25 per RPU-hour) ($450)

Example with time configuration and automatic scaling

Let's assume that a data warehouse processing job that needs to be executed on a serverless data warehouse in the US East (N. Virginia) Region. For this example, it takes 1 hour and 30 minutes to complete the job when the data warehouse is set to 32 RPUs. It runs once each 24 months, and you have Amazon Redshift Serverless 32 RPUs capacity as its Region. Automatic scaling is turned on and the maximum capacity is set to 64 RPUs. For the first 1.5 hours, the job runs on 32 RPUs as the base capacity. For the next 30 minutes, the job runs on 64 RPUs as the maximum capacity. For the last 30 minutes, the job runs on 32 RPUs as the base capacity.

The following table summarizes your total usage for the day:

Total price	Total usage calculation explained
0 am – 8 am	48 RPUs for 8 hours = 384 = 48 * 128 RPU-hours
8 am – 11 am	192 RPUs for 3 hours = 384=192 hours
11 am – 2 pm	Excluding the 1pm 11 min interval of process, the activity time is 2 hours. 64 RPUs for 2 hours = 64 = 128 RPU-hours
2 pm – 5 pm	128 RPU for 1 hour = 128 RPU-hours
5 pm – 10 pm	48 RPUs for 7 hours = 64 = 112 = 448 RPU-hours
10 pm – 1 am	16 x 3 RPU
Daily charges	$344.32 ($344 RPU * $1.50 = 1,152 RPU-hours)

Example for Amazon Redshift Concurrency Scaling

A 10-DC8 node-type Redshift cluster in US West 1 costs $36.00 per hour. Consider a concurrency-scaling scenario in which queries are routed to the remote cluster for the first 60 minutes. The remote cluster is billed for 60 minutes at $36.00 per hour. Assume that this hour credit was used to offset the first hour credit which you earned in the past 24 hours. Therefore, you will be charged for 60 minutes - 60 minutes = 0 minutes at $36.00 per hour = $0.00.

Example for Amazon Redshift Spectrum

If you compress your file and convert it to a columnar format such as Apache Parquet, you may see a 4:1 compression ratio and now a compressed file size of 1 TB. Using the same query, Redshift Spectrum needs to scan only one column in Parquet format, which is 1 GB of data after compression. You are charged $5 for 1 TB of data scanned.

If you compress your file using GZIP and you see a 4:1 compression ratio, in this case, you would have a compressed file size of 1 TB. Redshift Spectrum has to scan the entire file, but it will see 4 times more data because the data is compressed at a 4:1 ratio.

Now, let's assume you store the data in Apache Parquet format and you compress it using GZIP. You see a 4:1 compression ratio. Parquet is a columnar format, so Redshift Spectrum needs to scan only the 5 columns it needs for the query, which is 5 GB of data. Redshift Spectrum needs to scan only one column in Parquet format, which is 1 GB of data after compression.

Region: US East (N. Virginia)

Note: The above pricing examples are for illustration purposes only. The compression ratio of different files and columns may vary.

Additional pricing resources

AWS Pricing Calculator
Easily estimate your monthly costs with AWS.

Get Pricing Assistance
Connect with specialists who can recommend options.

Get started with Amazon Redshift

Get Started with AWS
Learn how to start using AWS in minutes.

Start for free
Get hands-on experience with AWS for 12 months.

Contact Sales
Get help from AWS sales experts.

Find the resources you need to build, run, and scale your applications on the most secure cloud ecosystem leveraged by millions.