API Reference

Amazon Redshift Data API

API Version 2019-12-20

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon Redshift Data API: API Reference

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Services or capabilities described in AWS documentation might vary by Region. To see the differences applicable to the AWS European Sovereign Cloud Region, see the AWS European Sovereign Cloud User Guide.

Table of Contents

Welcome **1**
Actions **2**
BatchExecuteStatement 3
Request Syntax 3
Request Parameters 4
Response Syntax 8
Response Elements 8
Errors 11
See Also 12
CancelStatement 13
Request Syntax 13
Request Parameters 13
Response Syntax 13
Response Elements 14
Errors 14
See Also 15
DescribeStatement 16
Request Syntax 16
Request Parameters 16
Response Syntax 17
Response Elements 18
Errors 22
See Also 23
DescribeTable 24
Request Syntax 24
Request Parameters 25
Response Syntax 27
Response Elements 28
Errors 29
See Also 30
ExecuteStatement 31
Request Syntax 31
Request Parameters 32
Response Syntax 35

Response Elements ... 36
Errors ... 38
See Also ... 39
GetStatementResult ... 41
Request Syntax ... 41
Request Parameters ... 41
Response Syntax ... 42
Response Elements ... 43
Errors ... 44
See Also ... 45
GetStatementResultV2 ... 46
Request Syntax ... 46
Request Parameters ... 46
Response Syntax ... 47
Response Elements ... 48
Errors ... 49
See Also ... 50
ListDatabases ... 51
Request Syntax ... 51
Request Parameters ... 52
Response Syntax ... 54
Response Elements ... 54
Errors ... 54
See Also ... 55
ListSchemas ... 57
Request Syntax ... 57
Request Parameters ... 58
Response Syntax ... 60
Response Elements ... 60
Errors ... 61
See Also ... 62
ListSessions ... 63
Request Syntax ... 63
Request Parameters ... 63
Response Syntax ... 66
Response Elements ... 66

Errors ... 67
See Also ... 67
ListStatements ... 69
Request Syntax ... 69
Request Parameters ... 69
Response Syntax ... 72
Response Elements ... 72
Errors ... 73
See Also ... 74
ListTables ... 75
Request Syntax ... 75
Request Parameters ... 76
Response Syntax ... 78
Response Elements ... 79
Errors ... 79
See Also ... 80
Data Types ... 82
ColumnMetadata ... 83
Contents ... 83
See Also ... 85
Field ... 86
Contents ... 86
See Also ... 87
QueryRecords ... 88
Contents ... 88
See Also ... 88
SessionData ... 89
Contents ... 89
See Also ... 91
SqlParameter ... 92
Contents ... 92
See Also ... 92
StatementData ... 94
Contents ... 94
See Also ... 96
SubStatementData ... 97

Contents 97
See Also 99
TableMember 100
Contents 100
See Also 100
Common Parameters 102
Common Error Types 105

Welcome

You can use the Amazon Redshift Data API to run queries on Amazon Redshift tables. You can run SQL statements, which are committed if the statement succeeds.

For more information about the Amazon Redshift Data API and AWS CLI usage examples, see Using the Amazon Redshift Data API in the *Amazon Redshift Management Guide*.

This document was last published on October 15, 2026.

Actions

The following actions are supported:

- BatchExecuteStatement
- CancelStatement
- DescribeStatement
- DescribeTable
- ExecuteStatement
- GetStatementResult
- GetStatementResultV2
- ListDatabases
- ListSchemas
- ListSessions
- ListStatements
- ListTables

BatchExecuteStatement

Runs one or more SQL statements, which can be data manipulation language (DML) or data definition language (DDL). Depending on the authorization method, use one of the following combinations of request parameters:

- AWS Secrets Manager - when connecting to a cluster, provide the `secret-arn` of a secret stored in AWS Secrets Manager which has `username` and `password`. The specified secret contains credentials to connect to the `database` you specify. When you are connecting to a cluster, you also supply the database name, If you provide a cluster identifier (`dbClusterIdentifier`), it must match the cluster identifier stored in the secret. When you are connecting to a serverless workgroup, you also supply the database name.
- Temporary credentials - when connecting to your data warehouse, choose one of the following options:
  - When connecting to a serverless workgroup, specify the workgroup name and database name. The database user name is derived from the IAM identity. For example, `arn:iam::123456789012:user:foo` has the database user name `IAM:foo`. Also, permission to call the `redshift-serverless:GetCredentials` operation is required.
  - When connecting to a cluster as an IAM identity, specify the cluster identifier and the database name. The database user name is derived from the IAM identity. For example, `arn:iam::123456789012:user:foo` has the database user name `IAM:foo`. Also, permission to call the `redshift:GetClusterCredentialsWithIAM` operation is required.
  - When connecting to a cluster as a database user, specify the cluster identifier, the database name, and the database user name. Also, permission to call the `redshift:GetClusterCredentials` operation is required.

For more information about the Amazon Redshift Data API and AWS CLI usage examples, see Using the Amazon Redshift Data API in the *Amazon Redshift Management Guide*.

Request Syntax

```
{
   "ClientToken": "string",
   "ClusterIdentifier": "string",
   "Database": "string",
   "DbUser": "string",
   "ExecutionMode": "string",
```

```
    "Parameters": [
       {
          "name": "string",
          "value": "string"
       }
    ],
    "ResultFormat": "string",
    "SecretArn": "string",
    "SessionId": "string",
    "SessionKeepAliveSeconds": number,
    "Sqls": [ "string" ],
    "StatementName": "string",
    "WaitTimeSeconds": number,
    "WithEvent": boolean,
    "WorkgroupName": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

The request accepts the following data in JSON format.

> **Note**
>
> In the following list, the required parameters are described first.

Sqls

One or more SQL statements to run. The SQL statements run serially in the order of the array. Subsequent SQL statements don't start until the previous statement in the array completes. By default, the SQL statements are run as a single transaction. If any SQL statement fails, all work is rolled back. To change this behavior, see the `ExecutionMode` parameter.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 40 items.

Required: Yes

ClientToken

A unique, case-sensitive identifier that you provide to ensure the idempotency of the request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Required: No

ClusterIdentifier

The cluster identifier. This parameter is required when connecting to a cluster and authenticating using either AWS Secrets Manager or temporary credentials.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[a-z][a-z0-9]*(-[a-z0-9]+)*`

Required: No

Database

The name of the database. This parameter is required when authenticating using either AWS Secrets Manager or temporary credentials.

Type: String

Required: No

DbUser

The database user name. This parameter is required when connecting to a cluster as a database user and authenticating using temporary credentials.

Type: String

Required: No

ExecutionMode

Determines how the SQL statements in the batch are run. If set to `TRANSACTION` (the default), all SQL statements are run as a single transaction and they are committed or rolled back

together. If set to `AUTO_COMMIT`, each SQL statement is committed individually, and a failure of one statement does not affect the others.

Type: String

Valid Values: `TRANSACTION | AUTO_COMMIT`

Required: No

Parameters

The parameters for the SQL statements. The parameters are available to all SQL statements in the batch. Each statement can reference any subset of the provided parameters. Each provided parameter must be referenced by at least one SQL statement in the batch.

Type: Array of SqlParameter objects

Array Members: Minimum number of 1 item.

Required: No

ResultFormat

The data format of the result of the SQL statement. If no format is specified, the default is JSON.

Type: String

Valid Values: `JSON | CSV`

Required: No

SecretArn

The name or ARN of the secret that enables access to the database. This parameter is required when authenticating using AWS Secrets Manager.

Type: String

Required: No

SessionId

The session identifier of the query.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

Required: No

SessionKeepAliveSeconds

The number of seconds to keep the session alive after the query finishes. The maximum time a session can keep alive is 24 hours. After 24 hours, the session is forced closed and the query is terminated.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 86400.

Required: No

StatementName

The name of the SQL statements. You can name the SQL statements when you create them to identify the query.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Required: No

WaitTimeSeconds

The number of seconds to wait for all SQL statements in the batch to complete execution before returning the response. If the SQL statements do not complete within the specified time, the response returns the current status. The maximum value is 30 seconds.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 30.

Required: No

WithEvent

A value that indicates whether to send an event to the Amazon EventBridge event bus after the SQL statements run.

Type: Boolean

Required: No

WorkgroupName

The serverless workgroup name or Amazon Resource Name (ARN). This parameter is required when connecting to a serverless workgroup and authenticating using either AWS Secrets Manager or temporary credentials.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `([a-z0-9-]{3,63}|arn:(aws(-[a-z]+)*):redshift-serverless:([a-z]{2}(-gov|(-iso[a-z]?))?|eusc-[a-z]+)-[a-z]+-\d{1}:\d{12}:workgroup/[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12})`

Required: No

Response Syntax

```
{
   "ClusterIdentifier": "string",
   "CreatedAt": number,
   "Database": "string",
   "DbGroups": [ "string" ],
   "DbUser": "string",
   "HasResultSet": boolean,
   "Id": "string",
   "RedshiftPid": number,
   "SecretArn": "string",
   "SessionId": "string",
   "Status": "string",
   "WorkgroupName": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

ClusterIdentifier

The cluster identifier. This element is not returned when connecting to a serverless workgroup.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[a-z][a-z0-9]*(-[a-z0-9]+)*`

CreatedAt

The date and time (UTC) the statement was created.

Type: Timestamp

Database

The name of the database.

Type: String

DbGroups

A list of colon (:) separated names of database groups.

Type: Array of strings

DbUser

The database user name.

Type: String

HasResultSet

A value that indicates whether the statement has a result set. The result set can be empty. The value is true for an empty result set. The value is true if any substatement returns a result set.

Type: Boolean

Id

The identifier of the SQL statement whose results are to be fetched. This value is a universally unique identifier (UUID) generated by Amazon Redshift Data API. This identifier is returned by `BatchExecuteStatment`.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

RedshiftPid

The process identifier from Amazon Redshift.

Type: Long

SecretArn

The name or ARN of the secret that enables access to the database.

Type: String

SessionId

The session identifier of the query.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

Status

The status of the SQL statement. Status values are defined as follows:

- ABORTED - The query run was stopped by the user.
- FAILED - The query run failed.
- FINISHED - The query has finished running.
- PICKED - The query has been chosen to be run.
- STARTED - The query run has started.
- SUBMITTED - The query was submitted, but not yet processed.

Type: String

Valid Values: `SUBMITTED | PICKED | STARTED | FINISHED | ABORTED | FAILED`

WorkgroupName

The serverless workgroup name or Amazon Resource Name (ARN). This element is not returned when connecting to a provisioned cluster.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `([a-z0-9-]{3,63}|arn:(aws(-[a-z]+)*):redshift-serverless:([a-z]{2}(-gov|(-iso[a-z]?))?|eusc-[a-z]+)-[a-z]+-\d{1}:\d{12}:workgroup/[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12})`

Errors

For information about the errors that are common to all actions, see Common Error Types.

ActiveSessionsExceededException

The Amazon Redshift Data API operation failed because the maximum number of active sessions exceeded.

HTTP Status Code: 400

ActiveStatementsExceededException

The number of active statements exceeds the limit.

HTTP Status Code: 400

BatchExecuteStatementException

An SQL statement encountered an environmental error while running.

StatementId

Statement identifier of the exception.

HTTP Status Code: 500

InternalServerException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 500

ResourceNotFoundException

The Amazon Redshift Data API operation failed due to a missing resource.

Message

The exception message.

ResourceId

Resource identifier associated with the exception.

HTTP Status Code: 400

ValidationException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS Command Line Interface V2
- AWS SDK for .NET V4
- AWS SDK for C++
- AWS SDK for Go v2
- AWS SDK for Java V2
- AWS SDK for JavaScript V3
- AWS SDK for Kotlin
- AWS SDK for PHP V3
- AWS SDK for Python (Boto3)
- AWS SDK for Ruby V3

CancelStatement

Cancels a running query. To be canceled, a query must be running.

For more information about the Amazon Redshift Data API and AWS CLI usage examples, see Using the Amazon Redshift Data API in the *Amazon Redshift Management Guide*.

Request Syntax

```
{
   "Id": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

The request accepts the following data in JSON format.

> **Note**
>
> In the following list, the required parameters are described first.

Id

The identifier of the SQL statement to cancel. This value is a universally unique identifier (UUID) generated by Amazon Redshift Data API. This identifier is returned by `BatchExecuteStatment`, `ExecuteStatment`, and `ListStatements`.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

Required: Yes

Response Syntax

```
{
```

```
    "Status": boolean
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Status

A value that indicates whether the cancel statement succeeded (true).

Type: Boolean

Errors

For information about the errors that are common to all actions, see Common Error Types.

DatabaseConnectionException

Connection to a database failed.

HTTP Status Code: 500

InternalServerException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 500

QueryTimeoutException

The Amazon Redshift Data API operation failed due to timeout.

HTTP Status Code: 400

ResourceNotFoundException

The Amazon Redshift Data API operation failed due to a missing resource.

Message

The exception message.

ResourceId

Resource identifier associated with the exception.

HTTP Status Code: 400

ValidationException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS Command Line Interface V2
- AWS SDK for .NET V4
- AWS SDK for C++
- AWS SDK for Go v2
- AWS SDK for Java V2
- AWS SDK for JavaScript V3
- AWS SDK for Kotlin
- AWS SDK for PHP V3
- AWS SDK for Python (Boto3)
- AWS SDK for Ruby V3

DescribeStatement

Describes the details about a specific instance when a query was run by the Amazon Redshift Data API. The information includes when the query started, when it finished, the query status, the number of rows returned, and the SQL statement.

For more information about the Amazon Redshift Data API and AWS CLI usage examples, see Using the Amazon Redshift Data API in the *Amazon Redshift Management Guide*.

Request Syntax

```
{
   "Id": "string",
   "WaitTimeSeconds": number
}
```

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

The request accepts the following data in JSON format.

> **Note**
>
> In the following list, the required parameters are described first.

Id

The identifier of the SQL statement to describe. This value is a universally unique identifier (UUID) generated by Amazon Redshift Data API. A suffix indicates the number of the SQL statement. For example, `d9b6c0c9-0747-4bf4-b142-e8883122f766:2` has a suffix of `:2` that indicates the second SQL statement of a batch query. This identifier is returned by `BatchExecuteStatment`, `ExecuteStatement`, and `ListStatements`.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

Required: Yes

WaitTimeSeconds

The number of seconds to wait for the SQL statement to complete execution before returning the description. The maximum value is 30 seconds.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 30.

Required: No

Response Syntax

```
{
   "ClusterIdentifier": "string",
   "CreatedAt": number,
   "Database": "string",
   "DbUser": "string",
   "Duration": number,
   "Error": "string",
   "ExecutionMode": "string",
   "HasResultSet": boolean,
   "Id": "string",
   "QueryParameters": [
      {
         "name": "string",
         "value": "string"
      }
   ],
   "QueryString": "string",
   "RedshiftPid": number,
   "RedshiftQueryId": number,
   "ResultFormat": "string",
   "ResultRows": number,
   "ResultSize": number,
   "SecretArn": "string",
   "SessionId": "string",
   "Status": "string",
   "SubStatements": [
      {
         "CreatedAt": number,
         "Duration": number,
         "Error": "string",
```

```
            "HasResultSet": boolean,
            "Id": "string",
            "QueryString": "string",
            "RedshiftQueryId": number,
            "ResultRows": number,
            "ResultSize": number,
            "Status": "string",
            "UpdatedAt": number
        }
    ],
    "UpdatedAt": number,
    "WorkgroupName": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Id

The identifier of the SQL statement described. This value is a universally unique identifier (UUID) generated by Amazon Redshift Data API.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

ClusterIdentifier

The cluster identifier.

Type: String

CreatedAt

The date and time (UTC) when the SQL statement was submitted to run.

Type: Timestamp

Database

The name of the database.

Type: String

DbUser

The database user name.

Type: String

Duration

The amount of time in nanoseconds that the statement ran.

Type: Long

Error

The error message from the cluster if the SQL statement encountered an error while running.

Type: String

ExecutionMode

The execution mode of the batch request. `TRANSACTION` indicates all SQL statements are run as a single transaction. `AUTO_COMMIT` indicates each SQL statement is committed individually.

Type: String

Valid Values: `TRANSACTION | AUTO_COMMIT`

HasResultSet

A value that indicates whether the statement has a result set. The result set can be empty. The value is true for an empty result set. The value is true if any substatement returns a result set.

Type: Boolean

QueryParameters

The parameters for the SQL statement.

Type: Array of SqlParameter objects

Array Members: Minimum number of 1 item.

QueryString

The SQL statement text.

Type: String

RedshiftPid

The process identifier from Amazon Redshift.

Type: Long

RedshiftQueryId

The identifier of the query generated by Amazon Redshift. These identifiers are also available in the `query` column of the `STL_QUERY` system view.

Type: Long

ResultFormat

The data format of the result of the SQL statement.

Type: String

Valid Values: `JSON | CSV`

ResultRows

Either the number of rows returned from the SQL statement or the number of rows affected. If result size is greater than zero, the result rows can be the number of rows affected by SQL statements such as INSERT, UPDATE, DELETE, COPY, and others. A `-1` indicates the value is null.

Type: Long

ResultSize

The size in bytes of the returned results. A `-1` indicates the value is null.

Type: Long

SecretArn

The name or Amazon Resource Name (ARN) of the secret that enables access to the database.

Type: String

SessionId

The session identifier of the query.

Type: String

Status

The status of the SQL statement being described. Status values are defined as follows:

- ABORTED - The query run was stopped by the user.
- ALL - A status value that includes all query statuses. This value can be used to filter results.
- FAILED - The query run failed.
- FINISHED - The query has finished running.
- PICKED - The query has been chosen to be run.
- STARTED - The query run has started.
- SUBMITTED - The query was submitted, but not yet processed.

Type: String

Valid Values: `SUBMITTED | PICKED | STARTED | FINISHED | ABORTED | FAILED | ALL`

SubStatements

The SQL statements from a multiple statement run.

Type: Array of SubStatementData objects

UpdatedAt

The date and time (UTC) that the metadata for the SQL statement was last updated. An example is the time the status last changed.

Type: Timestamp

WorkgroupName

The serverless workgroup name or Amazon Resource Name (ARN).

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `([a-z0-9-]{3,63}|arn:(aws(-[a-z]+)*):redshift-serverless:([a-z]{2}(-gov|(-iso[a-z]?))?|eusc-[a-z]+)-[a-z]+-\d{1}:\d{12}:workgroup/[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12})`

Errors

For information about the errors that are common to all actions, see Common Error Types.

ActiveWaitingRequestsExceededException

The number of active requests with `WaitTimeSeconds` for the same SQL statement exceeds the limit.

HTTP Status Code: 400

InternalServerException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 500

ResourceNotFoundException

The Amazon Redshift Data API operation failed due to a missing resource.

Message

The exception message.

ResourceId

Resource identifier associated with the exception.

HTTP Status Code: 400

ValidationException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS Command Line Interface V2
- AWS SDK for .NET V4
- AWS SDK for C++
- AWS SDK for Go v2
- AWS SDK for Java V2
- AWS SDK for JavaScript V3
- AWS SDK for Kotlin
- AWS SDK for PHP V3
- AWS SDK for Python (Boto3)
- AWS SDK for Ruby V3

DescribeTable

Describes the detailed information about a table from metadata in the cluster. The information includes its columns. A token is returned to page through the column list. Depending on the authorization method, use one of the following combinations of request parameters:

- AWS Secrets Manager - when connecting to a cluster, provide the `secret-arn` of a secret stored in AWS Secrets Manager which has `username` and `password`. The specified secret contains credentials to connect to the `database` you specify. When you are connecting to a cluster, you also supply the database name, If you provide a cluster identifier (`dbClusterIdentifier`), it must match the cluster identifier stored in the secret. When you are connecting to a serverless workgroup, you also supply the database name.
- Temporary credentials - when connecting to your data warehouse, choose one of the following options:
  - When connecting to a serverless workgroup, specify the workgroup name and database name. The database user name is derived from the IAM identity. For example, `arn:iam::123456789012:user:foo` has the database user name `IAM:foo`. Also, permission to call the `redshift-serverless:GetCredentials` operation is required.
  - When connecting to a cluster as an IAM identity, specify the cluster identifier and the database name. The database user name is derived from the IAM identity. For example, `arn:iam::123456789012:user:foo` has the database user name `IAM:foo`. Also, permission to call the `redshift:GetClusterCredentialsWithIAM` operation is required.
  - When connecting to a cluster as a database user, specify the cluster identifier, the database name, and the database user name. Also, permission to call the `redshift:GetClusterCredentials` operation is required.

For more information about the Amazon Redshift Data API and AWS CLI usage examples, see Using the Amazon Redshift Data API in the *Amazon Redshift Management Guide*.

Request Syntax

```
{
   "ClusterIdentifier": "string",
   "ConnectedDatabase": "string",
   "Database": "string",
   "DbUser": "string",
   "MaxResults": number,
```

```
    "NextToken": "string",
    "Schema": "string",
    "SecretArn": "string",
    "Table": "string",
    "WorkgroupName": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

The request accepts the following data in JSON format.

> **Note**
>
> In the following list, the required parameters are described first.

Database

The name of the database that contains the tables to be described. If `ConnectedDatabase` is not specified, this is also the database to connect to with your authentication credentials.

Type: String

Required: Yes

ClusterIdentifier

The cluster identifier. This parameter is required when connecting to a cluster and authenticating using either AWS Secrets Manager or temporary credentials.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[a-z][a-z0-9]*(-[a-z0-9]+)*`

Required: No

ConnectedDatabase

A database name. The connected database is specified when you connect with your authentication credentials.

Type: String

Required: No

DbUser

The database user name. This parameter is required when connecting to a cluster as a database user and authenticating using temporary credentials.

Type: String

Required: No

MaxResults

The maximum number of tables to return in the response. If more tables exist than fit in one response, then `NextToken` is returned to page through the results.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

Required: No

NextToken

A value that indicates the starting point for the next set of response records in a subsequent request. If a value is returned in a response, you can retrieve the next set of records by providing this returned NextToken value in the next NextToken parameter and retrying the command. If the NextToken field is empty, all response records have been retrieved for the request.

Type: String

Required: No

Schema

The schema that contains the table. If no schema is specified, then matching tables for all schemas are returned.

Type: String

Required: No

SecretArn

The name or ARN of the secret that enables access to the database. This parameter is required when authenticating using AWS Secrets Manager.

Type: String

Required: No

Table

The table name. If no table is specified, then all tables for all matching schemas are returned. If no table and no schema is specified, then all tables for all schemas in the database are returned

Type: String

Required: No

WorkgroupName

The serverless workgroup name or Amazon Resource Name (ARN). This parameter is required when connecting to a serverless workgroup and authenticating using either AWS Secrets Manager or temporary credentials.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `([a-z0-9-]{3,63}|arn:(aws(-[a-z]+)*):redshift-serverless:([a-z]{2}(-gov|(-iso[a-z]?))?|eusc-[a-z]+)-[a-z]+-\d{1}:\d{12}:workgroup/[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12})`

Required: No

Response Syntax

```
{
   "ColumnList": [
      {
         "columnDefault": "string",
         "isCaseSensitive": boolean,
```

```
            "isCurrency": boolean,
            "isSigned": boolean,
            "label": "string",
            "length": number,
            "name": "string",
            "nullable": number,
            "precision": number,
            "scale": number,
            "schemaName": "string",
            "tableName": "string",
            "typeName": "string"
        }
    ],
    "NextToken": "string",
    "TableName": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

ColumnList

A list of columns in the table.

Type: Array of ColumnMetadata objects

NextToken

A value that indicates the starting point for the next set of response records in a subsequent request. If a value is returned in a response, you can retrieve the next set of records by providing this returned NextToken value in the next NextToken parameter and retrying the command. If the NextToken field is empty, all response records have been retrieved for the request.

Type: String

TableName

The table name.

Type: String

Errors

For information about the errors that are common to all actions, see Common Error Types.

DatabaseConnectionException

Connection to a database failed.

HTTP Status Code: 500

InternalServerException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 500

QueryTimeoutException

The Amazon Redshift Data API operation failed due to timeout.

HTTP Status Code: 400

ResourceNotFoundException

The Amazon Redshift Data API operation failed due to a missing resource.

Message

The exception message.

ResourceId

Resource identifier associated with the exception.

HTTP Status Code: 400

ValidationException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS Command Line Interface V2
- AWS SDK for .NET V4
- AWS SDK for C++
- AWS SDK for Go v2
- AWS SDK for Java V2
- AWS SDK for JavaScript V3
- AWS SDK for Kotlin
- AWS SDK for PHP V3
- AWS SDK for Python (Boto3)
- AWS SDK for Ruby V3

ExecuteStatement

Runs an SQL statement, which can be data manipulation language (DML) or data definition language (DDL). This statement must be a single SQL statement. Depending on the authorization method, use one of the following combinations of request parameters:

- AWS Secrets Manager - when connecting to a cluster, provide the `secret-arn` of a secret stored in AWS Secrets Manager which has `username` and `password`. The specified secret contains credentials to connect to the `database` you specify. When you are connecting to a cluster, you also supply the database name, If you provide a cluster identifier (`dbClusterIdentifier`), it must match the cluster identifier stored in the secret. When you are connecting to a serverless workgroup, you also supply the database name.
- Temporary credentials - when connecting to your data warehouse, choose one of the following options:
  - When connecting to a serverless workgroup, specify the workgroup name and database name. The database user name is derived from the IAM identity. For example, `arn:iam::123456789012:user:foo` has the database user name `IAM:foo`. Also, permission to call the `redshift-serverless:GetCredentials` operation is required.
  - When connecting to a cluster as an IAM identity, specify the cluster identifier and the database name. The database user name is derived from the IAM identity. For example, `arn:iam::123456789012:user:foo` has the database user name `IAM:foo`. Also, permission to call the `redshift:GetClusterCredentialsWithIAM` operation is required.
  - When connecting to a cluster as a database user, specify the cluster identifier, the database name, and the database user name. Also, permission to call the `redshift:GetClusterCredentials` operation is required.

For more information about the Amazon Redshift Data API and AWS CLI usage examples, see Using the Amazon Redshift Data API in the *Amazon Redshift Management Guide*.

Request Syntax

```
{
   "ClientToken": "string",
   "ClusterIdentifier": "string",
   "Database": "string",
   "DbUser": "string",
   "Parameters": [
```

```
        {
            "name": "string",
            "value": "string"
        }
    ],
    "ResultFormat": "string",
    "SecretArn": "string",
    "SessionId": "string",
    "SessionKeepAliveSeconds": number,
    "Sql": "string",
    "StatementName": "string",
    "WaitTimeSeconds": number,
    "WithEvent": boolean,
    "WorkgroupName": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

The request accepts the following data in JSON format.

> **Note**
>
> In the following list, the required parameters are described first.

Sql

The SQL statement text to run.

Type: String

Required: Yes

ClientToken

A unique, case-sensitive identifier that you provide to ensure the idempotency of the request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Required: No

ClusterIdentifier

The cluster identifier. This parameter is required when connecting to a cluster and authenticating using either AWS Secrets Manager or temporary credentials.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[a-z][a-z0-9]*(-[a-z0-9]+)*`

Required: No

Database

The name of the database. This parameter is required when authenticating using either AWS Secrets Manager or temporary credentials.

Type: String

Required: No

DbUser

The database user name. This parameter is required when connecting to a cluster as a database user and authenticating using temporary credentials.

Type: String

Required: No

Parameters

The parameters for the SQL statement.

Type: Array of SqlParameter objects

Array Members: Minimum number of 1 item.

Required: No

ResultFormat

The data format of the result of the SQL statement. If no format is specified, the default is JSON.

Type: String

Valid Values: `JSON | CSV`

Required: No

SecretArn

The name or ARN of the secret that enables access to the database. This parameter is required when authenticating using AWS Secrets Manager.

Type: String

Required: No

SessionId

The session identifier of the query.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

Required: No

SessionKeepAliveSeconds

The number of seconds to keep the session alive after the query finishes. The maximum time a session can keep alive is 24 hours. After 24 hours, the session is forced closed and the query is terminated.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 86400.

Required: No

StatementName

The name of the SQL statement. You can name the SQL statement when you create it to identify the query.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Required: No

WaitTimeSeconds

The number of seconds to wait for the SQL statement to complete execution before returning the response. If the SQL statement does not complete within the specified time, the response returns the current status. The maximum value is 30 seconds.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 30.

Required: No

WithEvent

A value that indicates whether to send an event to the Amazon EventBridge event bus after the SQL statement runs.

Type: Boolean

Required: No

WorkgroupName

The serverless workgroup name or Amazon Resource Name (ARN). This parameter is required when connecting to a serverless workgroup and authenticating using either AWS Secrets Manager or temporary credentials.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `([a-z0-9-]{3,63}|arn:(aws(-[a-z]+)*):redshift-serverless:([a-z]{2}(-gov|(-iso[a-z]?))?|eusc-[a-z]+)-[a-z]+-\d{1}:\d{12}:workgroup/[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12})`

Required: No

Response Syntax

```
{
   "ClusterIdentifier": "string",
```

```
    "CreatedAt": number,
    "Database": "string",
    "DbGroups": [ "string" ],
    "DbUser": "string",
    "HasResultSet": boolean,
    "Id": "string",
    "RedshiftPid": number,
    "SecretArn": "string",
    "SessionId": "string",
    "Status": "string",
    "WorkgroupName": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

ClusterIdentifier

The cluster identifier. This element is not returned when connecting to a serverless workgroup.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[a-z][a-z0-9]*(-[a-z0-9]+)*`

CreatedAt

The date and time (UTC) the statement was created.

Type: Timestamp

Database

The name of the database.

Type: String

DbGroups

A list of colon (:) separated names of database groups.

Type: Array of strings

DbUser

The database user name.

Type: String

HasResultSet

A value that indicates whether the statement has a result set. The result set can be empty. The value is true for an empty result set.

Type: Boolean

Id

The identifier of the SQL statement whose results are to be fetched. This value is a universally unique identifier (UUID) generated by Amazon Redshift Data API.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

RedshiftPid

The process identifier from Amazon Redshift.

Type: Long

SecretArn

The name or ARN of the secret that enables access to the database.

Type: String

SessionId

The session identifier of the query.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

Status

The status of the SQL statement. Status values are defined as follows:

- ABORTED - The query run was stopped by the user.
- FAILED - The query run failed.

- FINISHED - The query has finished running.
- PICKED - The query has been chosen to be run.
- STARTED - The query run has started.
- SUBMITTED - The query was submitted, but not yet processed.

Type: String

Valid Values: `SUBMITTED | PICKED | STARTED | FINISHED | ABORTED | FAILED`

WorkgroupName

The serverless workgroup name or Amazon Resource Name (ARN). This element is not returned when connecting to a provisioned cluster.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `([a-z0-9-]{3,63}|arn:(aws(-[a-z]+)*):redshift-serverless:([a-z]{2}(-gov|(-iso[a-z]?))?|eusc-[a-z]+)-[a-z]+-\d{1}:\d{12}:workgroup/[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12})`

Errors

For information about the errors that are common to all actions, see Common Error Types.

ActiveSessionsExceededException

The Amazon Redshift Data API operation failed because the maximum number of active sessions exceeded.

HTTP Status Code: 400

ActiveStatementsExceededException

The number of active statements exceeds the limit.

HTTP Status Code: 400

ExecuteStatementException

The SQL statement encountered an environmental error while running.

Message

The exception message.

StatementId

Statement identifier of the exception.

HTTP Status Code: 500

InternalServerException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 500

ResourceNotFoundException

The Amazon Redshift Data API operation failed due to a missing resource.

Message

The exception message.

ResourceId

Resource identifier associated with the exception.

HTTP Status Code: 400

ValidationException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS Command Line Interface V2
- AWS SDK for .NET V4
- AWS SDK for C++
- AWS SDK for Go v2
- AWS SDK for Java V2
- AWS SDK for JavaScript V3
- AWS SDK for Kotlin
- AWS SDK for PHP V3
- AWS SDK for Python (Boto3)
- AWS SDK for Ruby V3

GetStatementResult

Fetches the temporarily cached result of an SQL statement in JSON format. The `ExecuteStatement` or `BatchExecuteStatement` operation that ran the SQL statement must have specified `ResultFormat` as `JSON` , or let the format default to JSON. A token is returned to page through the statement results.

For more information about the Amazon Redshift Data API and AWS CLI usage examples, see Using the Amazon Redshift Data API in the *Amazon Redshift Management Guide*.

Request Syntax

```
{
   "Id": "string",
   "NextToken": "string",
   "WaitTimeSeconds": number
}
```

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

The request accepts the following data in JSON format.

> **Note**
>
> In the following list, the required parameters are described first.

Id

The identifier of the SQL statement whose results are to be fetched. This value is a universally unique identifier (UUID) generated by Amazon Redshift Data API. A suffix indicates then number of the SQL statement. For example, `d9b6c0c9-0747-4bf4-b142-e8883122f766:2` has a suffix of `:2` that indicates the second SQL statement of a batch query. This identifier is returned by `BatchExecuteStatment`, `ExecuteStatment`, and `ListStatements`.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

Required: Yes

NextToken

A value that indicates the starting point for the next set of response records in a subsequent request. If a value is returned in a response, you can retrieve the next set of records by providing this returned NextToken value in the next NextToken parameter and retrying the command. If the NextToken field is empty, all response records have been retrieved for the request.

Type: String

Required: No

WaitTimeSeconds

The number of seconds to wait for the SQL statement to complete execution before returning the result. The maximum value is 30 seconds.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 30.

Required: No

Response Syntax

```
{
   "ColumnMetadata": [
      {
         "columnDefault": "string",
         "isCaseSensitive": boolean,
         "isCurrency": boolean,
         "isSigned": boolean,
         "label": "string",
         "length": number,
         "name": "string",
         "nullable": number,
         "precision": number,
         "scale": number,
         "schemaName": "string",
         "tableName": "string",
         "typeName": "string"
```

```
      }
   ],
   "NextToken": "string",
   "Records": [
      [
         { ... }
      ]
   ],
   "TotalNumRows": number
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Records

The results of the SQL statement in JSON format.

Type: Array of arrays of Field objects

ColumnMetadata

The properties (metadata) of a column.

Type: Array of ColumnMetadata objects

NextToken

A value that indicates the starting point for the next set of response records in a subsequent request. If a value is returned in a response, you can retrieve the next set of records by providing this returned NextToken value in the next NextToken parameter and retrying the command. If the NextToken field is empty, all response records have been retrieved for the request.

Type: String

TotalNumRows

The total number of rows in the result set returned from a query. You can use this number to estimate the number of calls to the `GetStatementResult` operation needed to page through the results.

Type: Long

Errors

For information about the errors that are common to all actions, see Common Error Types.

ActiveWaitingRequestsExceededException

The number of active requests with `WaitTimeSeconds` for the same SQL statement exceeds the limit.

HTTP Status Code: 400

InternalServerException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 500

ResourceNotFoundException

The Amazon Redshift Data API operation failed due to a missing resource.

Message

The exception message.

ResourceId

Resource identifier associated with the exception.

HTTP Status Code: 400

ValidationException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS Command Line Interface V2
- AWS SDK for .NET V4
- AWS SDK for C++
- AWS SDK for Go v2
- AWS SDK for Java V2
- AWS SDK for JavaScript V3
- AWS SDK for Kotlin
- AWS SDK for PHP V3
- AWS SDK for Python (Boto3)
- AWS SDK for Ruby V3

GetStatementResultV2

Fetches the temporarily cached result of an SQL statement in CSV format. The `ExecuteStatement` or `BatchExecuteStatement` operation that ran the SQL statement must have specified `ResultFormat` as `CSV`. A token is returned to page through the statement results.

For more information about the Amazon Redshift Data API and AWS CLI usage examples, see Using the Amazon Redshift Data API in the *Amazon Redshift Management Guide*.

Request Syntax

```
{
   "Id": "string",
   "NextToken": "string",
   "WaitTimeSeconds": number
}
```

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

The request accepts the following data in JSON format.

> **Note**
>
> In the following list, the required parameters are described first.

Id

The identifier of the SQL statement whose results are to be fetched. This value is a universally unique identifier (UUID) generated by Amazon Redshift Data API. A suffix indicates then number of the SQL statement. For example, `d9b6c0c9-0747-4bf4-b142-e8883122f766:2` has a suffix of `:2` that indicates the second SQL statement of a batch query. This identifier is returned by `BatchExecuteStatment`, `ExecuteStatment`, and `ListStatements`.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

Required: Yes

NextToken

A value that indicates the starting point for the next set of response records in a subsequent request. If a value is returned in a response, you can retrieve the next set of records by providing this returned NextToken value in the next NextToken parameter and retrying the command. If the NextToken field is empty, all response records have been retrieved for the request.

Type: String

Required: No

WaitTimeSeconds

The number of seconds to wait for the SQL statement to complete execution before returning the result. The maximum value is 30 seconds.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 30.

Required: No

Response Syntax

```
{
   "ColumnMetadata": [
      {
         "columnDefault": "string",
         "isCaseSensitive": boolean,
         "isCurrency": boolean,
         "isSigned": boolean,
         "label": "string",
         "length": number,
         "name": "string",
         "nullable": number,
         "precision": number,
         "scale": number,
         "schemaName": "string",
         "tableName": "string",
         "typeName": "string"
      }
```

```
    ],
    "NextToken": "string",
    "Records": [
       { ... }
    ],
    "ResultFormat": "string",
    "TotalNumRows": number
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Records

The results of the SQL statement in CSV format.

Type: Array of QueryRecords objects

ColumnMetadata

The properties (metadata) of a column.

Type: Array of ColumnMetadata objects

NextToken

A value that indicates the starting point for the next set of response records in a subsequent request. If a value is returned in a response, you can retrieve the next set of records by providing this returned NextToken value in the next NextToken parameter and retrying the command. If the NextToken field is empty, all response records have been retrieved for the request.

Type: String

ResultFormat

The data format of the result of the SQL statement.

Type: String

Valid Values: `JSON` | `CSV`

TotalNumRows

The total number of rows in the result set returned from a query. You can use this number to estimate the number of calls to the `GetStatementResultV2` operation needed to page through the results.

Type: Long

Errors

For information about the errors that are common to all actions, see Common Error Types.

ActiveWaitingRequestsExceededException

The number of active requests with `WaitTimeSeconds` for the same SQL statement exceeds the limit.

HTTP Status Code: 400

InternalServerException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 500

ResourceNotFoundException

The Amazon Redshift Data API operation failed due to a missing resource.

Message

The exception message.

ResourceId

Resource identifier associated with the exception.

HTTP Status Code: 400

ValidationException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS Command Line Interface V2
- AWS SDK for .NET V4
- AWS SDK for C++
- AWS SDK for Go v2
- AWS SDK for Java V2
- AWS SDK for JavaScript V3
- AWS SDK for Kotlin
- AWS SDK for PHP V3
- AWS SDK for Python (Boto3)
- AWS SDK for Ruby V3

ListDatabases

List the databases in a cluster. A token is returned to page through the database list. Depending on the authorization method, use one of the following combinations of request parameters:

- AWS Secrets Manager - when connecting to a cluster, provide the `secret-arn` of a secret stored in AWS Secrets Manager which has `username` and `password`. The specified secret contains credentials to connect to the `database` you specify. When you are connecting to a cluster, you also supply the database name, If you provide a cluster identifier (`dbClusterIdentifier`), it must match the cluster identifier stored in the secret. When you are connecting to a serverless workgroup, you also supply the database name.
- Temporary credentials - when connecting to your data warehouse, choose one of the following options:
  - When connecting to a serverless workgroup, specify the workgroup name and database name. The database user name is derived from the IAM identity. For example, `arn:iam::123456789012:user:foo` has the database user name `IAM:foo`. Also, permission to call the `redshift-serverless:GetCredentials` operation is required.
  - When connecting to a cluster as an IAM identity, specify the cluster identifier and the database name. The database user name is derived from the IAM identity. For example, `arn:iam::123456789012:user:foo` has the database user name `IAM:foo`. Also, permission to call the `redshift:GetClusterCredentialsWithIAM` operation is required.
  - When connecting to a cluster as a database user, specify the cluster identifier, the database name, and the database user name. Also, permission to call the `redshift:GetClusterCredentials` operation is required.

For more information about the Amazon Redshift Data API and AWS CLI usage examples, see Using the Amazon Redshift Data API in the *Amazon Redshift Management Guide*.

Request Syntax

```
{
   "ClusterIdentifier": "string",
   "Database": "string",
   "DbUser": "string",
   "MaxResults": number,
   "NextToken": "string",
   "SecretArn": "string",
```

```
    "WorkgroupName": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

The request accepts the following data in JSON format.

> **Note**
>
> In the following list, the required parameters are described first.

Database

The name of the database. This parameter is required when authenticating using either AWS Secrets Manager or temporary credentials.

Type: String

Required: Yes

ClusterIdentifier

The cluster identifier. This parameter is required when connecting to a cluster and authenticating using either AWS Secrets Manager or temporary credentials.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[a-z][a-z0-9]*(-[a-z0-9]+)*`

Required: No

DbUser

The database user name. This parameter is required when connecting to a cluster as a database user and authenticating using temporary credentials.

Type: String

Required: No

MaxResults

The maximum number of databases to return in the response. If more databases exist than fit in one response, then `NextToken` is returned to page through the results.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

Required: No

NextToken

A value that indicates the starting point for the next set of response records in a subsequent request. If a value is returned in a response, you can retrieve the next set of records by providing this returned NextToken value in the next NextToken parameter and retrying the command. If the NextToken field is empty, all response records have been retrieved for the request.

Type: String

Required: No

SecretArn

The name or ARN of the secret that enables access to the database. This parameter is required when authenticating using AWS Secrets Manager.

Type: String

Required: No

WorkgroupName

The serverless workgroup name or Amazon Resource Name (ARN). This parameter is required when connecting to a serverless workgroup and authenticating using either AWS Secrets Manager or temporary credentials.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `([a-z0-9-]{3,63}|arn:(aws(-[a-z]+)*):redshift-serverless:([a-z]{2}(-gov|(-iso[a-z]?))?|eusc-[a-z]+)-[a-z]+-\d{1}:\d{12}:workgroup/[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12})`

Required: No

Response Syntax

```
{
   "Databases": [ "string" ],
   "NextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Databases

The names of databases.

Type: Array of strings

NextToken

A value that indicates the starting point for the next set of response records in a subsequent request. If a value is returned in a response, you can retrieve the next set of records by providing this returned NextToken value in the next NextToken parameter and retrying the command. If the NextToken field is empty, all response records have been retrieved for the request.

Type: String

Errors

For information about the errors that are common to all actions, see Common Error Types.

DatabaseConnectionException

Connection to a database failed.

HTTP Status Code: 500

InternalServerException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 500

QueryTimeoutException

The Amazon Redshift Data API operation failed due to timeout.

HTTP Status Code: 400

ResourceNotFoundException

The Amazon Redshift Data API operation failed due to a missing resource.

Message

The exception message.

ResourceId

Resource identifier associated with the exception.

HTTP Status Code: 400

ValidationException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS Command Line Interface V2

- AWS SDK for .NET V4
- AWS SDK for C++
- AWS SDK for Go v2
- AWS SDK for Java V2
- AWS SDK for JavaScript V3
- AWS SDK for Kotlin
- AWS SDK for PHP V3
- AWS SDK for Python (Boto3)
- AWS SDK for Ruby V3

ListSchemas

Lists the schemas in a database. A token is returned to page through the schema list. Depending on the authorization method, use one of the following combinations of request parameters:

- AWS Secrets Manager - when connecting to a cluster, provide the `secret-arn` of a secret stored in AWS Secrets Manager which has `username` and `password`. The specified secret contains credentials to connect to the `database` you specify. When you are connecting to a cluster, you also supply the database name, If you provide a cluster identifier (`dbClusterIdentifier`), it must match the cluster identifier stored in the secret. When you are connecting to a serverless workgroup, you also supply the database name.
- Temporary credentials - when connecting to your data warehouse, choose one of the following options:
  - When connecting to a serverless workgroup, specify the workgroup name and database name. The database user name is derived from the IAM identity. For example, `arn:iam::123456789012:user:foo` has the database user name `IAM:foo`. Also, permission to call the `redshift-serverless:GetCredentials` operation is required.
  - When connecting to a cluster as an IAM identity, specify the cluster identifier and the database name. The database user name is derived from the IAM identity. For example, `arn:iam::123456789012:user:foo` has the database user name `IAM:foo`. Also, permission to call the `redshift:GetClusterCredentialsWithIAM` operation is required.
  - When connecting to a cluster as a database user, specify the cluster identifier, the database name, and the database user name. Also, permission to call the `redshift:GetClusterCredentials` operation is required.

For more information about the Amazon Redshift Data API and AWS CLI usage examples, see Using the Amazon Redshift Data API in the *Amazon Redshift Management Guide*.

Request Syntax

```
{
   "ClusterIdentifier": "string",
   "ConnectedDatabase": "string",
   "Database": "string",
   "DbUser": "string",
   "MaxResults": number,
   "NextToken": "string",
```

```
    "SchemaPattern": "string",
    "SecretArn": "string",
    "WorkgroupName": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

The request accepts the following data in JSON format.

> **Note**
>
> In the following list, the required parameters are described first.

Database

The name of the database that contains the schemas to list. If `ConnectedDatabase` is not specified, this is also the database to connect to with your authentication credentials.

Type: String

Required: Yes

ClusterIdentifier

The cluster identifier. This parameter is required when connecting to a cluster and authenticating using either AWS Secrets Manager or temporary credentials.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[a-z][a-z0-9]*(-[a-z0-9]+)*`

Required: No

ConnectedDatabase

A database name. The connected database is specified when you connect with your authentication credentials.

Type: String

Required: No

DbUser

The database user name. This parameter is required when connecting to a cluster as a database user and authenticating using temporary credentials.

Type: String

Required: No

MaxResults

The maximum number of schemas to return in the response. If more schemas exist than fit in one response, then `NextToken` is returned to page through the results.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

Required: No

NextToken

A value that indicates the starting point for the next set of response records in a subsequent request. If a value is returned in a response, you can retrieve the next set of records by providing this returned NextToken value in the next NextToken parameter and retrying the command. If the NextToken field is empty, all response records have been retrieved for the request.

Type: String

Required: No

SchemaPattern

A pattern to filter results by schema name. Within a schema pattern, "%" means match any substring of 0 or more characters and "_" means match any one character. Only schema name entries matching the search pattern are returned.

Type: String

Required: No

SecretArn

The name or ARN of the secret that enables access to the database. This parameter is required when authenticating using AWS Secrets Manager.

Type: String

Required: No

WorkgroupName

The serverless workgroup name or Amazon Resource Name (ARN). This parameter is required when connecting to a serverless workgroup and authenticating using either AWS Secrets Manager or temporary credentials.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `([a-z0-9-]{3,63}|arn:(aws(-[a-z]+)*):redshift-serverless:([a-z]{2}(-gov|(-iso[a-z]?))?|eusc-[a-z]+)-[a-z]+-\d{1}:\d{12}:workgroup/[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12})`

Required: No

Response Syntax

```
{
   "NextToken": "string",
   "Schemas": [ "string" ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

NextToken

A value that indicates the starting point for the next set of response records in a subsequent request. If a value is returned in a response, you can retrieve the next set of records by providing

this returned NextToken value in the next NextToken parameter and retrying the command. If the NextToken field is empty, all response records have been retrieved for the request.

Type: String

Schemas

The schemas that match the request pattern.

Type: Array of strings

Errors

For information about the errors that are common to all actions, see Common Error Types.

DatabaseConnectionException

Connection to a database failed.

HTTP Status Code: 500

InternalServerException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 500

QueryTimeoutException

The Amazon Redshift Data API operation failed due to timeout.

HTTP Status Code: 400

ResourceNotFoundException

The Amazon Redshift Data API operation failed due to a missing resource.

Message

The exception message.

ResourceId

Resource identifier associated with the exception.

HTTP Status Code: 400

ValidationException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS Command Line Interface V2
- AWS SDK for .NET V4
- AWS SDK for C++
- AWS SDK for Go v2
- AWS SDK for Java V2
- AWS SDK for JavaScript V3
- AWS SDK for Kotlin
- AWS SDK for PHP V3
- AWS SDK for Python (Boto3)
- AWS SDK for Ruby V3

ListSessions

Lists the sessions that the caller created in the last 24 hours. By default, only sessions with a status of `AVAILABLE` or `BUSY` are returned. You can filter the results by session status, compute target (cluster or serverless workgroup), or database. To retrieve the metadata for a single session, provide the `SessionId` parameter. Use `NextToken` to page through the session list.

Returns only the sessions that the caller created. When identity-enhanced role sessions are used, you must provide either the `ClusterIdentifier` or `WorkgroupName` parameter to ensure that the IAM Identity Center user can only access the Amazon Redshift IAM Identity Center applications they are assigned. For more information, see Trusted identity propagation overview.

Request Syntax

```
{
   "ClusterIdentifier": "string",
   "Database": "string",
   "MaxResults": number,
   "NextToken": "string",
   "RoleLevel": boolean,
   "SessionId": "string",
   "Status": "string",
   "WorkgroupName": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

The request accepts the following data in JSON format.

> **Note**
>
> In the following list, the required parameters are described first.

ClusterIdentifier

The cluster identifier. Only sessions on this cluster are returned. When providing `ClusterIdentifier`, then `WorkgroupName` can't be specified.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[a-z][a-z0-9]*(-[a-z0-9]+)*`

Required: No

Database

The name of the database. Only sessions connected to this database are returned.

Type: String

Required: No

MaxResults

The maximum number of sessions to return in the response. If more sessions exist than fit in one response, the operation returns `NextToken` to paginate the results.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 100.

Required: No

NextToken

A value that indicates the starting point for the next set of response records in a subsequent request. If a value is returned in a response, you can retrieve the next set of records by providing this returned NextToken value in the next NextToken parameter and retrying the command. If the NextToken field is empty, all response records have been retrieved for the request.

Type: String

Required: No

RoleLevel

Specifies whether to return all sessions created by the caller's IAM role, including sessions from previous IAM sessions. If false, only sessions created in the current IAM session are returned. The default is true.

Type: Boolean

Required: No

SessionId

The identifier of a specific session to return metadata for. This value is a universally unique identifier (UUID) generated by Amazon Redshift Data API. When you provide `SessionId`, you can't specify `Status`, `ClusterIdentifier`, `WorkgroupName`, or `Database`.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

Required: No

Status

The status of the sessions to list. If no status is specified, sessions with a status of `AVAILABLE` or BUSY are returned. Status values are defined as follows:

- AVAILABLE – The session is open and ready to run a SQL statement.
- BUSY – The session is currently running a SQL statement.
- CLOSED – The session is closed and can no longer run SQL statements.

Type: String

Valid Values: `AVAILABLE | BUSY | CLOSED`

Required: No

WorkgroupName

The serverless workgroup name or Amazon Resource Name (ARN). Only sessions on this workgroup are returned. When providing `WorkgroupName`, then `ClusterIdentifier` can't be specified.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `([a-z0-9-]{3,63}|arn:(aws(-[a-z]+)*):redshift-serverless:([a-z]{2}(-gov|(-iso[a-z]?))?|eusc-[a-z]+)-[a-z]+-\d{1}:\d{12}:workgroup/[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12})`

Required: No

Response Syntax

```
{
   "NextToken": "string",
   "Sessions": [
      {
         "ClusterIdentifier": "string",
         "CreatedAt": number,
         "CurrentStatementId": "string",
         "Database": "string",
         "DbUser": "string",
         "SessionAliveSeconds": number,
         "SessionId": "string",
         "SessionTtl": number,
         "Status": "string",
         "UpdatedAt": number,
         "WorkgroupName": "string"
      }
   ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Sessions

The sessions that match the request.

Type: Array of SessionData objects

NextToken

A value that indicates the starting point for the next set of response records in a subsequent request. If a value is returned in a response, you can retrieve the next set of records by providing this returned NextToken value in the next NextToken parameter and retrying the command. If the NextToken field is empty, all response records have been retrieved for the request.

Type: String

Errors

For information about the errors that are common to all actions, see Common Error Types.

InternalServerException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 500

ResourceNotFoundException

The Amazon Redshift Data API operation failed due to a missing resource.

Message

The exception message.

ResourceId

Resource identifier associated with the exception.

HTTP Status Code: 400

ValidationException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS Command Line Interface V2
- AWS SDK for .NET V4
- AWS SDK for C++
- AWS SDK for Go v2
- AWS SDK for Java V2
- AWS SDK for JavaScript V3
- AWS SDK for Kotlin
- AWS SDK for PHP V3
- AWS SDK for Python (Boto3)
- AWS SDK for Ruby V3

ListStatements

List of SQL statements. By default, only finished statements are shown. A token is returned to page through the statement list.

When you use identity-enhanced role sessions to list statements, you must provide either the `cluster-identifier` or `workgroup-name` parameter. This ensures that the IdC user can only access the Amazon Redshift IdC applications they are assigned. For more information, see Trusted identity propagation overview.

For more information about the Amazon Redshift Data API and AWS CLI usage examples, see Using the Amazon Redshift Data API in the *Amazon Redshift Management Guide*.

Request Syntax

```
{
   "ClusterIdentifier": "string",
   "Database": "string",
   "MaxResults": number,
   "NextToken": "string",
   "RoleLevel": boolean,
   "StatementName": "string",
   "Status": "string",
   "WorkgroupName": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

The request accepts the following data in JSON format.

> **Note**
>
> In the following list, the required parameters are described first.

ClusterIdentifier

The cluster identifier. Only statements that ran on this cluster are returned. When providing `ClusterIdentifier`, then `WorkgroupName` can't be specified.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[a-z][a-z0-9]*(-[a-z0-9]+)*`

Required: No

Database

The name of the database when listing statements run against a `ClusterIdentifier` or `WorkgroupName`.

Type: String

Required: No

MaxResults

The maximum number of SQL statements to return in the response. If more SQL statements exist than fit in one response, then `NextToken` is returned to page through the results.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 100.

Required: No

NextToken

A value that indicates the starting point for the next set of response records in a subsequent request. If a value is returned in a response, you can retrieve the next set of records by providing this returned NextToken value in the next NextToken parameter and retrying the command. If the NextToken field is empty, all response records have been retrieved for the request.

Type: String

Required: No

RoleLevel

A value that filters which statements to return in the response. If true, all statements run by the caller's IAM role are returned. If false, only statements run by the caller's IAM role in the current IAM session are returned. The default is true.

Type: Boolean

Required: No

StatementName

The name of the SQL statement specified as input to `BatchExecuteStatement` or `ExecuteStatement` to identify the query. You can list multiple statements by providing a prefix that matches the beginning of the statement name. For example, to list myStatement1, myStatement2, myStatement3, and so on, then provide the a value of `myStatement`. Data API does a case-sensitive match of SQL statement names to the prefix value you provide.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Required: No

Status

The status of the SQL statement to list. Status values are defined as follows:

- ABORTED - The query run was stopped by the user.
- ALL - A status value that includes all query statuses. This value can be used to filter results.
- FAILED - The query run failed.
- FINISHED - The query has finished running.
- PICKED - The query has been chosen to be run.
- STARTED - The query run has started.
- SUBMITTED - The query was submitted, but not yet processed.

Type: String

Valid Values: `SUBMITTED | PICKED | STARTED | FINISHED | ABORTED | FAILED | ALL`

Required: No

WorkgroupName

The serverless workgroup name or Amazon Resource Name (ARN). Only statements that ran on this workgroup are returned. When providing `WorkgroupName`, then `ClusterIdentifier` can't be specified.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `([a-z0-9-]{3,63}|arn:(aws(-[a-z]+)*):redshift-serverless:([a-z]{2}(-gov|(-iso[a-z]?))?|eusc-[a-z]+)-[a-z]+-\d{1}:\d{12}:workgroup/[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12})`

Required: No

Response Syntax

```
{
   "NextToken": "string",
   "Statements": [
      {
         "CreatedAt": number,
         "Id": "string",
         "IsBatchStatement": boolean,
         "QueryParameters": [
            {
               "name": "string",
               "value": "string"
            }
         ],
         "QueryString": "string",
         "QueryStrings": [ "string" ],
         "ResultFormat": "string",
         "SecretArn": "string",
         "SessionId": "string",
         "StatementName": "string",
         "Status": "string",
         "UpdatedAt": number
      }
   ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

Statements

The SQL statements.

Type: Array of StatementData objects

NextToken

A value that indicates the starting point for the next set of response records in a subsequent request. If a value is returned in a response, you can retrieve the next set of records by providing this returned NextToken value in the next NextToken parameter and retrying the command. If the NextToken field is empty, all response records have been retrieved for the request.

Type: String

Errors

For information about the errors that are common to all actions, see Common Error Types.

InternalServerException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 500

ResourceNotFoundException

The Amazon Redshift Data API operation failed due to a missing resource.

Message

The exception message.

ResourceId

Resource identifier associated with the exception.

HTTP Status Code: 400

ValidationException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS Command Line Interface V2
- AWS SDK for .NET V4
- AWS SDK for C++
- AWS SDK for Go v2
- AWS SDK for Java V2
- AWS SDK for JavaScript V3
- AWS SDK for Kotlin
- AWS SDK for PHP V3
- AWS SDK for Python (Boto3)
- AWS SDK for Ruby V3

ListTables

List the tables in a database. If neither `SchemaPattern` nor `TablePattern` are specified, then all tables in the database are returned. A token is returned to page through the table list. Depending on the authorization method, use one of the following combinations of request parameters:

- AWS Secrets Manager - when connecting to a cluster, provide the `secret-arn` of a secret stored in AWS Secrets Manager which has `username` and `password`. The specified secret contains credentials to connect to the `database` you specify. When you are connecting to a cluster, you also supply the database name, If you provide a cluster identifier (`dbClusterIdentifier`), it must match the cluster identifier stored in the secret. When you are connecting to a serverless workgroup, you also supply the database name.
- Temporary credentials - when connecting to your data warehouse, choose one of the following options:
  - When connecting to a serverless workgroup, specify the workgroup name and database name. The database user name is derived from the IAM identity. For example, `arn:iam::123456789012:user:foo` has the database user name `IAM:foo`. Also, permission to call the `redshift-serverless:GetCredentials` operation is required.
  - When connecting to a cluster as an IAM identity, specify the cluster identifier and the database name. The database user name is derived from the IAM identity. For example, `arn:iam::123456789012:user:foo` has the database user name `IAM:foo`. Also, permission to call the `redshift:GetClusterCredentialsWithIAM` operation is required.
  - When connecting to a cluster as a database user, specify the cluster identifier, the database name, and the database user name. Also, permission to call the `redshift:GetClusterCredentials` operation is required.

For more information about the Amazon Redshift Data API and AWS CLI usage examples, see Using the Amazon Redshift Data API in the *Amazon Redshift Management Guide*.

Request Syntax

```
{
   "ClusterIdentifier": "string",
   "ConnectedDatabase": "string",
   "Database": "string",
   "DbUser": "string",
   "MaxResults": number,
```

```
    "NextToken": "string",
    "SchemaPattern": "string",
    "SecretArn": "string",
    "TablePattern": "string",
    "WorkgroupName": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters.

The request accepts the following data in JSON format.

> **Note**
>
> In the following list, the required parameters are described first.

Database

The name of the database that contains the tables to list. If `ConnectedDatabase` is not specified, this is also the database to connect to with your authentication credentials.

Type: String

Required: Yes

ClusterIdentifier

The cluster identifier. This parameter is required when connecting to a cluster and authenticating using either AWS Secrets Manager or temporary credentials.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 63.

Pattern: `[a-z][a-z0-9]*(-[a-z0-9]+)*`

Required: No

ConnectedDatabase

A database name. The connected database is specified when you connect with your authentication credentials.

Type: String

Required: No

DbUser

The database user name. This parameter is required when connecting to a cluster as a database user and authenticating using temporary credentials.

Type: String

Required: No

MaxResults

The maximum number of tables to return in the response. If more tables exist than fit in one response, then `NextToken` is returned to page through the results.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 1000.

Required: No

NextToken

A value that indicates the starting point for the next set of response records in a subsequent request. If a value is returned in a response, you can retrieve the next set of records by providing this returned NextToken value in the next NextToken parameter and retrying the command. If the NextToken field is empty, all response records have been retrieved for the request.

Type: String

Required: No

SchemaPattern

A pattern to filter results by schema name. Within a schema pattern, "%" means match any substring of 0 or more characters and "_" means match any one character. Only schema name entries matching the search pattern are returned. If `SchemaPattern` is not specified, then all tables that match `TablePattern` are returned. If neither `SchemaPattern` or `TablePattern` are specified, then all tables are returned.

Type: String

Required: No

SecretArn

The name or ARN of the secret that enables access to the database. This parameter is required when authenticating using AWS Secrets Manager.

Type: String

Required: No

TablePattern

A pattern to filter results by table name. Within a table pattern, "%" means match any substring of 0 or more characters and "_" means match any one character. Only table name entries matching the search pattern are returned. If `TablePattern` is not specified, then all tables that match `SchemaPattern`are returned. If neither `SchemaPattern` or `TablePattern` are specified, then all tables are returned.

Type: String

Required: No

WorkgroupName

The serverless workgroup name or Amazon Resource Name (ARN). This parameter is required when connecting to a serverless workgroup and authenticating using either AWS Secrets Manager or temporary credentials.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `([a-z0-9-]{3,63}|arn:(aws(-[a-z]+)*):redshift-serverless:([a-z]{2}(-gov|(-iso[a-z]?))?|eusc-[a-z]+)-[a-z]+-\d{1}:\d{12}:workgroup/[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12})`

Required: No

Response Syntax

```
{
   "NextToken": "string",
```

```
    "Tables": [
       {
          "name": "string",
          "schema": "string",
          "type": "string"
       }
    ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

NextToken

A value that indicates the starting point for the next set of response records in a subsequent request. If a value is returned in a response, you can retrieve the next set of records by providing this returned NextToken value in the next NextToken parameter and retrying the command. If the NextToken field is empty, all response records have been retrieved for the request.

Type: String

Tables

The tables that match the request pattern.

Type: Array of TableMember objects

Errors

For information about the errors that are common to all actions, see Common Error Types.

DatabaseConnectionException

Connection to a database failed.

HTTP Status Code: 500

InternalServerException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 500

QueryTimeoutException

The Amazon Redshift Data API operation failed due to timeout.

HTTP Status Code: 400

ResourceNotFoundException

The Amazon Redshift Data API operation failed due to a missing resource.

Message

The exception message.

ResourceId

Resource identifier associated with the exception.

HTTP Status Code: 400

ValidationException

The Amazon Redshift Data API operation failed due to invalid input.

Message

The exception message.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS Command Line Interface V2
- AWS SDK for .NET V4
- AWS SDK for C++
- AWS SDK for Go v2

- AWS SDK for Java V2
- AWS SDK for JavaScript V3
- AWS SDK for Kotlin
- AWS SDK for PHP V3
- AWS SDK for Python (Boto3)
- AWS SDK for Ruby V3

Data Types

The Redshift Data API Service API contains several data types that various actions use. This section describes each data type in detail.

> **Note**
>
> The order of each element in a data type structure is not guaranteed. Applications should not assume a particular order.

The following data types are supported:

- ColumnMetadata
- Field
- QueryRecords
- SessionData
- SqlParameter
- StatementData
- SubStatementData
- TableMember

ColumnMetadata

The properties (metadata) of a column.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

columnDefault

The default value of the column.

Type: String

Required: No

isCaseSensitive

A value that indicates whether the column is case-sensitive.

Type: Boolean

Required: No

isCurrency

A value that indicates whether the column contains currency values.

Type: Boolean

Required: No

isSigned

A value that indicates whether an integer column is signed.

Type: Boolean

Required: No

label

The label for the column.

Type: String

Required: No

length

The length of the column.

Type: Integer

Required: No

name

The name of the column.

Type: String

Required: No

nullable

A value that indicates whether the column is nullable.

Type: Integer

Required: No

precision

The precision value of a decimal number column, or the column length for a non-numeric column.

Type: Integer

Required: No

scale

The scale value of a decimal number column.

Type: Integer

Required: No

schemaName

The name of the schema that contains the table that includes the column.

Type: String

Required: No

tableName

The name of the table that includes the column.

Type: String

Required: No

typeName

The database-specific data type of the column.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

Field

A data value in a column.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

> **Important**
>
> This data type is a UNION, so only one of the following members can be specified when used or returned.

blobValue

A value of the BLOB data type.

Type: Base64-encoded binary data object

Required: No

booleanValue

A value of the Boolean data type.

Type: Boolean

Required: No

doubleValue

A value of the double data type.

Type: Double

Required: No

isNull

A value that indicates whether the data is NULL.

Type: Boolean

Required: No

longValue

A value of the long data type.

Type: Long

Required: No

stringValue

A value of the string data type.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

QueryRecords

The results of the SQL statement.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

> **Important**
>
> This data type is a UNION, so only one of the following members can be specified when used or returned.

CSVRecords

The results of the SQL statement in CSV format.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

SessionData

Contains the metadata for a session returned by `ListSessions`, including its status, compute target, database connection, and lifecycle timestamps.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

CreatedAt

The date and time (UTC) when the session was created.

Type: Timestamp

Required: Yes

SessionId

The session identifier. This value is a universally unique identifier (UUID) generated by Amazon Redshift Data API.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

Required: Yes

Status

The status of the session. Status values are defined as follows:

- AVAILABLE – The session is open and ready to run a SQL statement.
- BUSY – The session is currently running a SQL statement.
- CLOSED – The session is closed and can no longer run SQL statements.

Type: String

Valid Values: `AVAILABLE | BUSY | CLOSED`

Required: Yes

ClusterIdentifier

The cluster identifier. This element is not returned when connecting to a serverless workgroup.

Type: String

Required: No

CurrentStatementId

The identifier of the SQL statement currently running in the session. This value is a universally unique identifier (UUID) generated by Amazon Redshift Data API. This element is returned only when the session status is BUSY.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

Required: No

Database

The name of the database that the session is connected to.

Type: String

Required: No

DbUser

The database user name.

Type: String

Required: No

SessionAliveSeconds

The number of seconds that the session is kept alive after a query finishes.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 86400.

Required: No

SessionTtl

The date and time (UTC) when the session is set to expire and be closed.

Type: Timestamp

Required: No

UpdatedAt

The date and time (UTC) that the session metadata was last updated. An example is the time the status last changed.

Type: Timestamp

Required: No

WorkgroupName

The serverless workgroup name or Amazon Resource Name (ARN). This element is not returned when connecting to a provisioned cluster.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 128.

Pattern: `([a-z0-9-]{3,63}|arn:(aws(-[a-z]+)*):redshift-serverless:([a-z]{2}(-gov|(-iso[a-z]?))?|eusc-[a-z]+)-[a-z]+-\d{1}:\d{12}:workgroup/[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12})`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

SqlParameter

A parameter used in a SQL statement.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

name

The name of the parameter.

Type: String

Pattern: `[0-9a-zA-Z_]+`

Required: Yes

value

The value of the parameter. Amazon Redshift implicitly converts to the proper data type. For more information, see Data types in the *Amazon Redshift Database Developer Guide*.

Type: String

Length Constraints: Minimum length of 1.

Required: Yes

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

StatementData

The SQL statement to run.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

Id

The SQL statement identifier. This value is a universally unique identifier (UUID) generated by Amazon Redshift Data API.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

Required: Yes

CreatedAt

The date and time (UTC) the statement was created.

Type: Timestamp

Required: No

IsBatchStatement

A value that indicates whether the statement is a batch query request.

Type: Boolean

Required: No

QueryParameters

The parameters used in a SQL statement.

Type: Array of SqlParameter objects

Array Members: Minimum number of 1 item.

Required: No

QueryString

The SQL statement.

Type: String

Required: No

QueryStrings

One or more SQL statements. Each query string in the array corresponds to one of the queries in a batch query request.

Type: Array of strings

Required: No

ResultFormat

The data format of the result of the SQL statement.

Type: String

Valid Values: `JSON | CSV`

Required: No

SecretArn

The name or Amazon Resource Name (ARN) of the secret that enables access to the database.

Type: String

Required: No

SessionId

The session identifier of the query.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

Required: No

StatementName

The name of the SQL statement.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Required: No

Status

The status of the SQL statement. An example is the that the SQL statement finished.

Type: String

Valid Values: `SUBMITTED | PICKED | STARTED | FINISHED | ABORTED | FAILED | ALL`

Required: No

UpdatedAt

The date and time (UTC) that the statement metadata was last updated.

Type: Timestamp

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

SubStatementData

Information about an SQL statement.

Contents

> **Note**
>
> In the following list, the required parameters are described first.

Id

The identifier of the SQL statement. This value is a universally unique identifier (UUID) generated by Amazon Redshift Data API. A suffix indicates the number of the SQL statement. For example, `d9b6c0c9-0747-4bf4-b142-e8883122f766:2` has a suffix of `:2` that indicates the second SQL statement of a batch query.

Type: String

Pattern: `[a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d{0,2})?`

Required: Yes

CreatedAt

The date and time (UTC) the statement was created.

Type: Timestamp

Required: No

Duration

The amount of time in nanoseconds that the statement ran.

Type: Long

Required: No

Error

The error message from the cluster if the SQL statement encountered an error while running.

Type: String

Required: No

HasResultSet

A value that indicates whether the statement has a result set. The result set can be empty. The value is true for an empty result set.

Type: Boolean

Required: No

QueryString

The SQL statement text.

Type: String

Required: No

RedshiftQueryId

The SQL statement identifier. This value is a universally unique identifier (UUID) generated by Amazon Redshift Data API.

Type: Long

Required: No

ResultRows

Either the number of rows returned from the SQL statement or the number of rows affected. If result size is greater than zero, the result rows can be the number of rows affected by SQL statements such as INSERT, UPDATE, DELETE, COPY, and others. A `-1` indicates the value is null.

Type: Long

Required: No

ResultSize

The size in bytes of the returned results. A `-1` indicates the value is null.

Type: Long

Required: No

Status

The status of the SQL statement. An example is the that the SQL statement finished.

Type: String

Valid Values: `SUBMITTED | PICKED | STARTED | FINISHED | ABORTED | FAILED`

Required: No

UpdatedAt

The date and time (UTC) that the statement metadata was last updated.

Type: Timestamp

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

TableMember

The properties of a table.

Contents

> **Note**
> In the following list, the required parameters are described first.

name

The name of the table.

Type: String

Required: No

schema

The schema containing the table.

Type: String

Required: No

type

The type of the table. Possible values include TABLE, VIEW, SYSTEM TABLE, GLOBAL TEMPORARY, LOCAL TEMPORARY, ALIAS, and SYNONYM.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- AWS SDK for C++

- AWS SDK for Java V2
- AWS SDK for Ruby V3

Common Parameters

The following list contains the parameters that all actions use for signing Signature Version 4 requests with a query string. Any action-specific parameters are listed in the topic for that action. For more information about Signature Version 4, see Signing AWS API requests in the *IAM User Guide*.

X-Amz-Algorithm

The hash algorithm that you used to create the request signature.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Valid Values: `AWS4-HMAC-SHA256`

Required: Conditional

X-Amz-Credential

The credential scope value, which is a string that includes your access key, the date, the region you are targeting, the service you are requesting, and a termination string ("aws4_request"). The value is expressed in the following format: *access_key*/*YYYYMMDD*/*region*/*service*/aws4_request.

For more information, see Create a signed AWS API request in the *IAM User Guide*.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Required: Conditional

X-Amz-Date

The date that is used to create the signature. The format must be ISO 8601 basic format (YYYYMMDD'T'HHMMSS'Z'). For example, the following date time is a valid X-Amz-Date value: `20120325T120000Z`.

Condition: X-Amz-Date is optional for all requests; it can be used to override the date used for signing requests. If the Date header is specified in the ISO 8601 basic format, X-Amz-Date is not required. When X-Amz-Date is used, it always overrides the value of the Date header. For more information, see Elements of an AWS API request signature in the *IAM User Guide*.

Type: string

Required: Conditional

X-Amz-Security-Token

The temporary security token that was obtained through a call to AWS Security Token Service (AWS STS). For a list of services that support temporary security credentials from AWS STS, see AWS services that work with IAM in the *IAM User Guide*.

Condition: If you're using temporary security credentials from AWS STS, you must include the security token.

Type: string

Required: Conditional

X-Amz-Signature

Specifies the hex-encoded signature that was calculated from the string to sign and the derived signing key.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Required: Conditional

X-Amz-SignedHeaders

Specifies all the HTTP headers that were included as part of the canonical request. For more information about specifying signed headers, see Create a signed AWS API request in the *IAM User Guide*.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Required: Conditional

Common Error Types

This section lists common error types that this AWS service may return. Not all services return all error types listed here. For errors specific to an API action for this service, see the topic for that API action.

AccessDeniedException

You don't have permission to perform this action. Verify that your IAM policy includes the required permissions.

HTTP Status Code: 403

ExpiredTokenException

The security token included in the request has expired. Request a new security token and try again.

HTTP Status Code: 403

IncompleteSignature

The request signature doesn't conform to AWS standards. Verify that you're using valid AWS credentials and that your request is properly formatted. If you're using an SDK, ensure it's up to date.

HTTP Status Code: 403

InternalFailure

The request can't be processed right now because of an internal server issue. Try again later. If the problem persists, contact AWS Support.

HTTP Status Code: 500

MalformedHttpRequestException

The request body can't be processed. This typically happens when the request body can't be decompressed using the specified content encoding algorithm. Verify that the content encoding header matches the compression format used.

HTTP Status Code: 400

NotAuthorized

You don't have permissions to perform this action. Verify that your IAM policy includes the required permissions.

HTTP Status Code: 401

OptInRequired

Your AWS account needs a subscription for this service. Verify that you've enabled the service in your account.

HTTP Status Code: 403

RequestAbortedException

The request was aborted before a response could be returned. This typically happens when the client closes the connection.

HTTP Status Code: 400

RequestEntityTooLargeException

The request entity is too large. Reduce the size of the request body and try again.

HTTP Status Code: 413

RequestTimeoutException

The request timed out. The server didn't receive the complete request within the expected time frame. Try again.

HTTP Status Code: 408

ServiceUnavailable

The service is temporarily unavailable. Try again later.

HTTP Status Code: 503

ThrottlingException

Your request rate is too high. The AWS SDKs automatically retry requests that receive this exception. Reduce the frequency of requests.

HTTP Status Code: 400

UnknownOperationException

The action or operation isn't recognized. Verify that the action name is spelled correctly and that it's supported by the API version you're using.

HTTP Status Code: 404

UnrecognizedClientException

The X.509 certificate or AWS access key ID you provided doesn't exist in our records. Verify that you're using valid credentials and that they haven't expired.

HTTP Status Code: 403

ValidationError

The input doesn't meet the required format or constraints. Check that all required parameters are included and that values are valid.

HTTP Status Code: 400