



API Reference

Amazon Redshift Data API



API Version 2019-12-20

Amazon Redshift Data API: API Reference

Table of Contents

Welcome	1
Actions	2
BatchExecuteStatement	3
Request Syntax	3
Request Parameters	4
Response Syntax	7
Response Elements	7
Errors	9
See Also	10
CancelStatement	11
Request Syntax	11
Request Parameters	11
Response Syntax	11
Response Elements	12
Errors	12
See Also	13
DescribeStatement	14
Request Syntax	14
Request Parameters	14
Response Syntax	15
Response Elements	16
Errors	19
See Also	20
DescribeTable	21
Request Syntax	21
Request Parameters	22
Response Syntax	24
Response Elements	25
Errors	25
See Also	26
ExecuteStatement	27
Request Syntax	27
Request Parameters	28
Response Syntax	31

Response Elements	31
Errors	33
See Also	34
GetStatementResult	35
Request Syntax	35
Request Parameters	35
Response Syntax	36
Response Elements	37
Errors	37
See Also	38
GetStatementResultV2	39
Request Syntax	39
Request Parameters	39
Response Syntax	40
Response Elements	40
Errors	41
See Also	42
ListDatabases	43
Request Syntax	43
Request Parameters	44
Response Syntax	46
Response Elements	46
Errors	46
See Also	47
ListSchemas	48
Request Syntax	48
Request Parameters	49
Response Syntax	51
Response Elements	51
Errors	52
See Also	52
ListStatements	54
Request Syntax	54
Request Parameters	54
Response Syntax	57
Response Elements	57

Errors	58
See Also	58
ListTables	60
Request Syntax	60
Request Parameters	61
Response Syntax	63
Response Elements	64
Errors	64
See Also	65
Data Types	66
ColumnMetadata	67
Contents	67
See Also	69
Field	70
Contents	70
See Also	71
QueryRecords	72
Contents	72
See Also	72
SqlParameter	73
Contents	73
See Also	73
StatementData	75
Contents	75
See Also	77
SubStatementData	78
Contents	78
See Also	80
TableMember	81
Contents	81
See Also	81
Common Parameters	83
Common Errors	86

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 Amazon CLI usage examples, see [Using the Amazon Redshift Data API](#) in the *Amazon Redshift Management Guide*.

This document was last published on July 3, 2025.

Actions

The following actions are supported:

- [BatchExecuteStatement](#)
- [CancelStatement](#)
- [DescribeStatement](#)
- [DescribeTable](#)
- [ExecuteStatement](#)
- [GetStatementResult](#)
- [GetStatementResultV2](#)
- [ListDatabases](#)
- [ListSchemas](#)
- [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:

- Amazon Secrets Manager - when connecting to a cluster, provide the `secret-arn` of a secret stored in Amazon 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 Amazon 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",  
  "ResultFormat": "string",  
}
```

```
"SecretArn": "string",
"SessionId": "string",
"SessionKeepAliveSeconds": number,
"Sqls": [ "string" ],
"StatementName": "string",
"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 are run as a single transaction. They run serially in the order of the array. Subsequent SQL statements don't start until the previous statement in the array completes. If any SQL statement fails, then because they are run as one transaction, all work is rolled back.

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 Amazon Secrets Manager or temporary credentials.

Type: String

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

Required: No

Database

The name of the database. This parameter is required when authenticating using either Amazon 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

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 Amazon Secrets Manager.

Type: String

Required: No

SessionId

The session identifier of the query.

Type: String

Pattern: [a-zA-Z0-9]{8}(-[a-zA-Z0-9]{4}){3}-[a-zA-Z0-9]{12}(:\d+)?

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 500.

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 Amazon Secrets Manager or temporary credentials.

Type: String

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

Pattern: (([a-z0-9-]+)|(arn:(aws(-[a-z]+)*:redshift-serverless:[a-z]{2}(-gov|(-iso[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",  
    "Id": "string",  
    "SecretArn": "string",  
    "SessionId": "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.

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

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 BatchExecuteStatement.

Type: String

Pattern: [a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d+)?

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+)?

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-]+)|(arn:(aws(-[a-z]+)*:redshift-serverless:[a-z]{2}(-gov|(-iso[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 Errors](#).

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.

HTTP Status Code: 500

InternalServerError

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

HTTP Status Code: 500

ValidationException

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

HTTP Status Code: 400

See Also

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

- [Amazon Command Line Interface](#)
- [Amazon SDK for .NET](#)
- [Amazon SDK for C++](#)
- [Amazon SDK for Go v2](#)
- [Amazon SDK for Java V2](#)
- [Amazon SDK for JavaScript V3](#)
- [Amazon SDK for Kotlin](#)
- [Amazon SDK for PHP V3](#)
- [Amazon SDK for Python](#)
- [Amazon 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 Amazon 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 `BatchExecuteStatement`, `ExecuteStatement`, and `ListStatements`.

Type: String

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

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 Errors](#).

DatabaseConnectionException

Connection to a database failed.

HTTP Status Code: 500

InternalServerException

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

HTTP Status Code: 500

ResourceNotFoundException

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

HTTP Status Code: 400

ValidationException

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

HTTP Status Code: 400

See Also

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

- [Amazon Command Line Interface](#)
- [Amazon SDK for .NET](#)
- [Amazon SDK for C++](#)
- [Amazon SDK for Go v2](#)
- [Amazon SDK for Java V2](#)
- [Amazon SDK for JavaScript V3](#)
- [Amazon SDK for Kotlin](#)
- [Amazon SDK for PHP V3](#)
- [Amazon SDK for Python](#)
- [Amazon 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 Amazon 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 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 BatchExecuteStatement, ExecuteStatement, and ListStatements.

Type: String

Pattern: [a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d+)?

Required: Yes

Response Syntax

```
{  
    "ClusterIdentifier": "string",  
    "CreatedAt": number,  
    "Database": "string",  
    "DbUser": "string",  
    "Duration": number,  
    "Error": "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+)?

[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

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-]+)|(arn:(aws(-[a-z]+)*):redshift-serverless:[a-z]{2}(-gov|(-iso[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 Errors](#).

InternalServerError

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

HTTP Status Code: 500

ResourceNotFoundException

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

HTTP Status Code: 400

ValidationException

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

HTTP Status Code: 400

See Also

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

- [Amazon Command Line Interface](#)
- [Amazon SDK for .NET](#)
- [Amazon SDK for C++](#)
- [Amazon SDK for Go v2](#)
- [Amazon SDK for Java V2](#)
- [Amazon SDK for JavaScript V3](#)
- [Amazon SDK for Kotlin](#)
- [Amazon SDK for PHP V3](#)
- [Amazon SDK for Python](#)
- [Amazon 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:

- Amazon Secrets Manager - when connecting to a cluster, provide the `secret-arn` of a secret stored in Amazon 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 Amazon 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 Amazon Secrets Manager or temporary credentials.

Type: String

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

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 Amazon 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 Amazon Secrets Manager or temporary credentials.

Type: String

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

Pattern: (([a-z0-9-]+)|(arn:(aws(-[a-z]+)*:redshift-serverless:[a-z]{2}(-gov|(-iso[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  
    }  
  ]  
}
```

```
        "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 Errors](#).

DatabaseConnectionException

Connection to a database failed.

HTTP Status Code: 500

InternalServerException

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

HTTP Status Code: 500

QueryTimeoutException

The Amazon Redshift Data API operation failed due to timeout.

HTTP Status Code: 400

ValidationException

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

HTTP Status Code: 400

See Also

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

- [Amazon Command Line Interface](#)
- [Amazon SDK for .NET](#)
- [Amazon SDK for C++](#)
- [Amazon SDK for Go v2](#)
- [Amazon SDK for Java V2](#)
- [Amazon SDK for JavaScript V3](#)
- [Amazon SDK for Kotlin](#)
- [Amazon SDK for PHP V3](#)
- [Amazon SDK for Python](#)
- [Amazon 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:

- Amazon Secrets Manager - when connecting to a cluster, provide the `secret-arn` of a secret stored in Amazon 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 Amazon 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",  
"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 Amazon Secrets Manager or temporary credentials.

Type: String

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

Required: No

Database

The name of the database. This parameter is required when authenticating using either Amazon 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 Amazon Secrets Manager.

Type: String

Required: No

[SessionId](#)

The session identifier of the query.

Type: String

Pattern: [a-zA-Z0-9]{8}(-[a-zA-Z0-9]{4}){3}-[a-zA-Z0-9]{12}(:\d+)?

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 500.

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 Amazon Secrets Manager or temporary credentials.

Type: String

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

Pattern: (([a-z0-9-]+)|(arn:(aws(-[a-z]+)*):redshift-serverless:[a-z]{2}(-gov|(-iso[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",  
    "Id": "string",  
    "SecretArn": "string",  
    "SessionId": "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.

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

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+)?

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+)?

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-]+)|(arn:(aws(-[a-z]+)*):redshift-serverless:[a-z]{2}(-gov|(-iso[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 Errors](#).

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.

HTTP Status Code: 500

InternalServerException

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

HTTP Status Code: 500

ValidationException

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

HTTP Status Code: 400

See Also

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

- [Amazon Command Line Interface](#)
- [Amazon SDK for .NET](#)
- [Amazon SDK for C++](#)
- [Amazon SDK for Go v2](#)
- [Amazon SDK for Java V2](#)
- [Amazon SDK for JavaScript V3](#)
- [Amazon SDK for Kotlin](#)
- [Amazon SDK for PHP V3](#)
- [Amazon SDK for Python](#)
- [Amazon 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 Amazon CLI usage examples, see [Using the Amazon Redshift Data API](#) in the *Amazon Redshift Management Guide*.

Request Syntax

```
{  
  "Id": "string",  
  "NextToken": "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 whose results are to be fetched. 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 BatchExecuteStatement, ExecuteStatement, and ListStatements.

Type: String

Pattern: [a-zA-Z0-9]{8}(-[a-zA-Z0-9]{4}){3}-[a-zA-Z0-9]{12}(:\d+)?

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

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 Errors](#).

InternalServerError

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

HTTP Status Code: 500

ResourceNotFoundException

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

HTTP Status Code: 400

ValidationException

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

HTTP Status Code: 400

See Also

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

- [Amazon Command Line Interface](#)
- [Amazon SDK for .NET](#)
- [Amazon SDK for C++](#)
- [Amazon SDK for Go v2](#)
- [Amazon SDK for Java V2](#)
- [Amazon SDK for JavaScript V3](#)
- [Amazon SDK for Kotlin](#)
- [Amazon SDK for PHP V3](#)
- [Amazon SDK for Python](#)
- [Amazon 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 Amazon CLI usage examples, see [Using the Amazon Redshift Data API](#) in the *Amazon Redshift Management Guide*.

Request Syntax

```
{  
    "Id": "string",  
    "NextToken": "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 whose results are to be fetched. 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 BatchExecuteStatement, ExecuteStatement, and ListStatements.

Type: String

Pattern: [a-z0-9]{8}(-[a-z0-9]{4}){3}-[a-z0-9]{12}(:\d+)?

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

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 Errors](#).

InternalServerError

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

HTTP Status Code: 500

ResourceNotFoundException

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

HTTP Status Code: 400

ValidationException

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

HTTP Status Code: 400

See Also

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

- [Amazon Command Line Interface](#)
- [Amazon SDK for .NET](#)
- [Amazon SDK for C++](#)
- [Amazon SDK for Go v2](#)
- [Amazon SDK for Java V2](#)
- [Amazon SDK for JavaScript V3](#)
- [Amazon SDK for Kotlin](#)
- [Amazon SDK for PHP V3](#)
- [Amazon SDK for Python](#)
- [Amazon 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:

- Amazon Secrets Manager - when connecting to a cluster, provide the `secret-arn` of a secret stored in Amazon 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 Amazon 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 Amazon 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 Amazon Secrets Manager or temporary credentials.

Type: String

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

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 Amazon 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 Amazon Secrets Manager or temporary credentials.

Type: String

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

Pattern: (([a-z0-9-]+)|(arn:(aws(-[a-z]+)*:redshift-serverless:[a-z]{2}(-gov|(-iso[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 Errors](#).

DatabaseConnectionException

Connection to a database failed.

HTTP Status Code: 500

InternalServerException

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

HTTP Status Code: 500

QueryTimeoutException

The Amazon Redshift Data API operation failed due to timeout.

HTTP Status Code: 400

ValidationException

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

HTTP Status Code: 400

See Also

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

- [Amazon Command Line Interface](#)
- [Amazon SDK for .NET](#)
- [Amazon SDK for C++](#)
- [Amazon SDK for Go v2](#)
- [Amazon SDK for Java V2](#)
- [Amazon SDK for JavaScript V3](#)
- [Amazon SDK for Kotlin](#)
- [Amazon SDK for PHP V3](#)
- [Amazon SDK for Python](#)
- [Amazon 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:

- Amazon Secrets Manager - when connecting to a cluster, provide the `secret-arn` of a secret stored in Amazon 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 Amazon 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 Amazon Secrets Manager or temporary credentials.

Type: String

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

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 Amazon 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 Amazon Secrets Manager or temporary credentials.

Type: String

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

Pattern: (([a-z0-9-]+)|(arn:(aws(-[a-z]+)*):redshift-serverless:[a-z]{2}(-gov|(-iso[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 Errors](#).

DatabaseConnectionException

Connection to a database failed.

HTTP Status Code: 500

InternalServerException

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

HTTP Status Code: 500

QueryTimeoutException

The Amazon Redshift Data API operation failed due to timeout.

HTTP Status Code: 400

ValidationException

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

HTTP Status Code: 400

See Also

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

- [Amazon Command Line Interface](#)
- [Amazon SDK for .NET](#)
- [Amazon SDK for C++](#)

- [Amazon SDK for Go v2](#)
- [Amazon SDK for Java V2](#)
- [Amazon SDK for JavaScript V3](#)
- [Amazon SDK for Kotlin](#)
- [Amazon SDK for PHP V3](#)
- [Amazon SDK for Python](#)
- [Amazon 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 Amazon 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.

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 500.

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-]+)|(arn:(aws(-[a-z]+)*):redshift-serverless:[a-z]{2}(-gov|(-iso[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 Errors](#).

InternalServerError

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

HTTP Status Code: 500

ValidationException

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

HTTP Status Code: 400

See Also

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

- [Amazon Command Line Interface](#)
- [Amazon SDK for .NET](#)
- [Amazon SDK for C++](#)
- [Amazon SDK for Go v2](#)
- [Amazon SDK for Java V2](#)
- [Amazon SDK for JavaScript V3](#)
- [Amazon SDK for Kotlin](#)

- [Amazon SDK for PHP V3](#)
- [Amazon SDK for Python](#)
- [Amazon 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:

- Amazon Secrets Manager - when connecting to a cluster, provide the secret-arn of a secret stored in Amazon 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 Amazon 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 Amazon Secrets Manager or temporary credentials.

Type: String

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

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 Amazon 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 Amazon Secrets Manager or temporary credentials.

Type: String

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

Pattern: (([a-z0-9-]+)|(arn:(aws(-[a-z]+)*:redshift-serverless:[a-z]{2}(-gov|(-iso[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 Errors](#).

DatabaseConnectionException

Connection to a database failed.

HTTP Status Code: 500

InternalServerException

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

HTTP Status Code: 500

QueryTimeoutException

The Amazon Redshift Data API operation failed due to timeout.

HTTP Status Code: 400

ValidationException

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

HTTP Status Code: 400

See Also

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

- [Amazon Command Line Interface](#)
- [Amazon SDK for .NET](#)
- [Amazon SDK for C++](#)
- [Amazon SDK for Go v2](#)
- [Amazon SDK for Java V2](#)
- [Amazon SDK for JavaScript V3](#)
- [Amazon SDK for Kotlin](#)
- [Amazon SDK for PHP V3](#)
- [Amazon SDK for Python](#)
- [Amazon 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](#)
- [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 Amazon SDKs, see the following:

- [Amazon SDK for C++](#)
- [Amazon SDK for Java V2](#)
- [Amazon 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 Amazon SDKs, see the following:

- [Amazon SDK for C++](#)
- [Amazon SDK for Java V2](#)
- [Amazon 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 Amazon SDKs, see the following:

- [Amazon SDK for C++](#)
- [Amazon SDK for Java V2](#)
- [Amazon 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 Amazon SDKs, see the following:

- [Amazon SDK for C++](#)
- [Amazon SDK for Java V2](#)
- [Amazon 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+)?

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-zA-Z0-9]{8}(-[a-zA-Z0-9]{4}){3}-[a-zA-Z0-9]{12}(:\d+)?

Required: No

StatementName

The name of the SQL statement.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 500.

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 Amazon SDKs, see the following:

- [Amazon SDK for C++](#)
- [Amazon SDK for Java V2](#)
- [Amazon 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-zA-Z0-9]{8}(-[a-zA-Z0-9]{4}){3}-[a-zA-Z0-9]{12}(:\d+)?

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 Amazon SDKs, see the following:

- [Amazon SDK for C++](#)
- [Amazon SDK for Java V2](#)
- [Amazon 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 Amazon SDKs, see the following:

- [Amazon SDK for C++](#)

- [Amazon SDK for Java V2](#)
- [Amazon 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 Amazon API requests](#) in the *IAM User Guide*.

Action

The action to be performed.

Type: string

Required: Yes

Version

The API version that the request is written for, expressed in the format YYYY-MM-DD.

Type: string

Required: Yes

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 Amazon 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 Amazon 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 Amazon Security Token Service (Amazon STS). For a list of services that support temporary security credentials from Amazon STS, see [Amazon Web Services services that work with IAM](#) in the *IAM User Guide*.

Condition: If you're using temporary security credentials from Amazon 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 Amazon 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 Errors

This section lists the errors common to the API actions of all Amazon services. For errors specific to an API action for this service, see the topic for that API action.

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

IncompleteSignature

The request signature does not conform to Amazon standards.

HTTP Status Code: 400

InternalFailure

The request processing has failed because of an unknown error, exception or failure.

HTTP Status Code: 500

InvalidAction

The action or operation requested is invalid. Verify that the action is typed correctly.

HTTP Status Code: 400

InvalidClientTokenId

The X.509 certificate or Amazon access key ID provided does not exist in our records.

HTTP Status Code: 403

NotAuthorized

You do not have permission to perform this action.

HTTP Status Code: 400

OptInRequired

The Amazon access key ID needs a subscription for the service.

HTTP Status Code: 403

RequestExpired

The request reached the service more than 15 minutes after the date stamp on the request or more than 15 minutes after the request expiration date (such as for pre-signed URLs), or the date stamp on the request is more than 15 minutes in the future.

HTTP Status Code: 400

ServiceUnavailable

The request has failed due to a temporary failure of the server.

HTTP Status Code: 503

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

ValidationException

The input fails to satisfy the constraints specified by an Amazon service.

HTTP Status Code: 400