



Amazon Timestream for LiveAnalytics

Amazon Timestream Write



API Version 2018-11-01

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

Amazon Timestream Write: Amazon Timestream for LiveAnalytics

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.

Table of Contents

Welcome	1
Amazon Timestream Write	1
Amazon Timestream Query	1
Actions	2
Amazon Timestream Write	3
CreateBatchLoadTask	5
CreateDatabase	11
CreateTable	15
DeleteDatabase	21
DeleteTable	24
DescribeBatchLoadTask	27
DescribeDatabase	32
DescribeEndpoints	35
DescribeTable	37
ListBatchLoadTasks	41
ListDatabases	45
ListTables	48
ListTagsForResource	52
ResumeBatchLoadTask	55
TagResource	58
UntagResource	61
UpdateDatabase	64
UpdateTable	68
WriteRecords	73
Amazon Timestream Query	78
CancelQuery	80
CreateScheduledQuery	83
DeleteScheduledQuery	90
DescribeAccountSettings	93
DescribeEndpoints	96
DescribeScheduledQuery	98
ExecuteScheduledQuery	104
ListScheduledQueries	111
ListTagsForResource	115

PrepareQuery	118
Query	122
TagResource	130
UntagResource	133
UpdateAccountSettings	136
UpdateScheduledQuery	141
Data Types	144
Amazon Timestream Write	146
BatchLoadProgressReport	148
BatchLoadTask	150
BatchLoadTaskDescription	152
CsvConfiguration	155
Database	157
DataModel	159
DataModelConfiguration	161
DataModelS3Configuration	162
DataSourceConfiguration	163
DataSourceS3Configuration	165
Dimension	166
DimensionMapping	168
Endpoint	169
MagneticStoreRejectedDataLocation	170
MagneticStoreWriteProperties	171
MeasureValue	172
MixedMeasureMapping	174
MultiMeasureAttributeMapping	176
MultiMeasureMappings	177
PartitionKey	178
Record	180
RecordsIngested	183
RejectedRecord	184
ReportConfiguration	186
ReportS3Configuration	187
RetentionProperties	189
S3Configuration	190
Schema	192

Table	193
Tag	196
Amazon Timestream Query	196
AccountSettingsNotificationConfiguration	199
ColumnInfo	200
Datum	201
DimensionMapping	203
Endpoint	204
ErrorReportConfiguration	205
ErrorReportLocation	206
ExecutionStats	207
LastUpdate	209
MixedMeasureMapping	211
MultiMeasureAttributeMapping	213
MultiMeasureMappings	215
NotificationConfiguration	216
ParameterMapping	217
ProvisionedCapacityRequest	218
ProvisionedCapacityResponse	219
QueryComputeRequest	221
QueryComputeResponse	222
QueryInsights	223
QueryInsightsResponse	225
QuerySpatialCoverage	228
QuerySpatialCoverageMax	229
QueryStatus	231
QueryTemporalRange	233
QueryTemporalRangeMax	234
Row	235
S3Configuration	236
S3ReportLocation	238
ScheduleConfiguration	239
ScheduledQuery	240
ScheduledQueryDescription	243
ScheduledQueryInsights	247
ScheduledQueryInsightsResponse	248

ScheduledQueryRunSummary	250
SelectColumn	252
SnsConfiguration	254
Tag	255
TargetConfiguration	256
TargetDestination	257
TimeSeriesDataPoint	258
TimestreamConfiguration	259
TimestreamDestination	261
Type	262
Common Parameters	264
Common Error Types	267

Welcome

Amazon Timestream Write

Amazon Timestream is a fast, scalable, fully managed time-series database service that makes it easy to store and analyze trillions of time-series data points per day. With Timestream, you can easily store and analyze IoT sensor data to derive insights from your IoT applications. You can analyze industrial telemetry to streamline equipment management and maintenance. You can also store and analyze log data and metrics to improve the performance and availability of your applications.

Timestream is built from the ground up to effectively ingest, process, and store time-series data. It organizes data to optimize query processing. It automatically scales based on the volume of data ingested and on the query volume to ensure you receive optimal performance while inserting and querying data. As your data grows over time, Timestream's adaptive query processing engine spans across storage tiers to provide fast analysis while reducing costs.

Amazon Timestream Query

Actions

The following actions are supported by Amazon Timestream Write:

- [CreateBatchLoadTask](#)
- [CreateDatabase](#)
- [CreateTable](#)
- [DeleteDatabase](#)
- [DeleteTable](#)
- [DescribeBatchLoadTask](#)
- [DescribeDatabase](#)
- [DescribeEndpoints](#)
- [DescribeTable](#)
- [ListBatchLoadTasks](#)
- [ListDatabases](#)
- [ListTables](#)
- [ListTagsForResource](#)
- [ResumeBatchLoadTask](#)
- [TagResource](#)
- [UntagResource](#)
- [UpdateDatabase](#)
- [UpdateTable](#)
- [WriteRecords](#)

The following actions are supported by Amazon Timestream Query:

- [CancelQuery](#)
- [CreateScheduledQuery](#)
- [DeleteScheduledQuery](#)
- [DescribeAccountSettings](#)
- [DescribeEndpoints](#)

- [DescribeScheduledQuery](#)
- [ExecuteScheduledQuery](#)
- [ListScheduledQueries](#)
- [ListTagsForResource](#)
- [PrepareQuery](#)
- [Query](#)
- [TagResource](#)
- [UntagResource](#)
- [UpdateAccountSettings](#)
- [UpdateScheduledQuery](#)

Amazon Timestream Write

The following actions are supported by Amazon Timestream Write:

- [CreateBatchLoadTask](#)
- [CreateDatabase](#)
- [CreateTable](#)
- [DeleteDatabase](#)
- [DeleteTable](#)
- [DescribeBatchLoadTask](#)
- [DescribeDatabase](#)
- [DescribeEndpoints](#)
- [DescribeTable](#)
- [ListBatchLoadTasks](#)
- [ListDatabases](#)
- [ListTables](#)
- [ListTagsForResource](#)
- [ResumeBatchLoadTask](#)
- [TagResource](#)
- [UntagResource](#)

- [UpdateDatabase](#)
- [UpdateTable](#)
- [WriteRecords](#)

CreateBatchLoadTask

Service: Amazon Timestream Write

Creates a new Timestream batch load task. A batch load task processes data from a CSV source in an S3 location and writes to a Timestream table. A mapping from source to target is defined in a batch load task. Errors and events are written to a report at an S3 location. For the report, if the AWS KMS key is not specified, the report will be encrypted with an S3 managed key when `SSE_S3` is the option. Otherwise an error is thrown. For more information, see [AWS managed keys](#). [Service quotas apply](#). For details, see [code sample](#).

Request Syntax

```
{
  "ClientToken": "string",
  "DataModelConfiguration": {
    "DataModel": {
      "DimensionMappings": [
        {
          "DestinationColumn": "string",
          "SourceColumn": "string"
        }
      ],
      "MeasureNameColumn": "string",
      "MixedMeasureMappings": [
        {
          "MeasureName": "string",
          "MeasureValueType": "string",
          "MultiMeasureAttributeMappings": [
            {
              "MeasureValueType": "string",
              "SourceColumn": "string",
              "TargetMultiMeasureAttributeName": "string"
            }
          ],
          "SourceColumn": "string",
          "TargetMeasureName": "string"
        }
      ],
      "MultiMeasureMappings": {
        "MultiMeasureAttributeMappings": [
          {
            "MeasureValueType": "string",
```

```

        "SourceColumn": "string",
        "TargetMultiMeasureAttributeName": "string"
    }
],
    "TargetMultiMeasureName": "string"
},
    "TimeColumn": "string",
    "TimeUnit": "string"
},
    "DataModelS3Configuration": {
        "BucketName": "string",
        "ObjectKey": "string"
    }
},
    "DataSourceConfiguration": {
        "CsvConfiguration": {
            "ColumnSeparator": "string",
            "EscapeChar": "string",
            "NullValue": "string",
            "QuoteChar": "string",
            "TrimWhiteSpace": boolean
        },
        "DataFormat": "string",
        "DataSourceS3Configuration": {
            "BucketName": "string",
            "ObjectKeyPrefix": "string"
        }
    },
    "RecordVersion": number,
    "ReportConfiguration": {
        "ReportS3Configuration": {
            "BucketName": "string",
            "EncryptionOption": "string",
            "KmsKeyId": "string",
            "ObjectKeyPrefix": "string"
        }
    },
    "TargetDatabaseName": "string",
    "TargetTableName": "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.

ClientToken

Type: String

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

Required: No

DataModelConfiguration

Type: [DataModelConfiguration](#) object

Required: No

DataSourceConfiguration

Defines configuration details about the data source for a batch load task.

Type: [DataSourceConfiguration](#) object

Required: Yes

RecordVersion

Type: Long

Required: No

ReportConfiguration

Report configuration for a batch load task. This contains details about where error reports are stored.

Type: [ReportConfiguration](#) object

Required: Yes

TargetDatabaseName

Target Timestream database for a batch load task.

Type: String

Pattern: [a-zA-Z0-9_.-]+

Required: Yes

TargetTableName

Target Timestream table for a batch load task.

Type: String

Pattern: [a-zA-Z0-9_.-]+

Required: Yes

Response Syntax

```
{  
  "TaskId": "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.

TaskId

The ID of the batch load task.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 32.

Pattern: [A-Z0-9]+

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You are not authorized to perform this action.

HTTP Status Code: 400

ConflictException

Timestream was unable to process this request because it contains resource that already exists.

HTTP Status Code: 400

InternalServerErrorException

Timestream was unable to fully process this request because of an internal server error.

HTTP Status Code: 500

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ResourceNotFoundException

The operation tried to access a nonexistent resource. The resource might not be specified correctly, or its status might not be ACTIVE.

HTTP Status Code: 400

ServiceQuotaExceededException

The instance quota of resource exceeded for this account.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

CreateDatabase

Service: Amazon Timestream Write

Creates a new Timestream database. If the AWS KMS key is not specified, the database will be encrypted with a Timestream managed AWS KMS key located in your account. For more information, see [AWS managed keys](#). [Service quotas apply](#). For details, see [code sample](#).

Request Syntax

```
{
  "DatabaseName": "string",
  "KmsKeyId": "string",
  "Tags": [
    {
      "Key": "string",
      "Value": "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.

DatabaseName

The name of the Timestream database.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Pattern: [a-zA-Z0-9_.-]+

Required: Yes

KmsKeyId

The AWS KMS key for the database. If the AWS KMS key is not specified, the database will be encrypted with a Timestream managed AWS KMS key located in your account. For more information, see [AWS managed keys](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

Tags

A list of key-value pairs to label the table.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Required: No

Response Syntax

```
{
  "Database": {
    "Arn": "string",
    "CreationTime": number,
    "DatabaseName": "string",
    "KmsKeyId": "string",
    "LastUpdatedTime": number,
    "TableCount": 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.

Database

The newly created Timestream database.

Type: [Database](#) object

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You are not authorized to perform this action.

HTTP Status Code: 400

ConflictException

Timestream was unable to process this request because it contains resource that already exists.

HTTP Status Code: 400

InternalServerErrorException

Timestream was unable to fully process this request because of an internal server error.

HTTP Status Code: 500

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ServiceQuotaExceededException

The instance quota of resource exceeded for this account.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

CreateTable

Service: Amazon Timestream Write

Adds a new table to an existing database in your account. In an AWS account, table names must be at least unique within each Region if they are in the same database. You might have identical table names in the same Region if the tables are in separate databases. While creating the table, you must specify the table name, database name, and the retention properties. [Service quotas apply](#). See [code sample](#) for details.

Request Syntax

```
{
  "DatabaseName": "string",
  "MagneticStoreWriteProperties": {
    "EnableMagneticStoreWrites": boolean,
    "MagneticStoreRejectedDataLocation": {
      "S3Configuration": {
        "BucketName": "string",
        "EncryptionOption": "string",
        "KmsKeyId": "string",
        "ObjectKeyPrefix": "string"
      }
    }
  },
  "RetentionProperties": {
    "MagneticStoreRetentionPeriodInDays": number,
    "MemoryStoreRetentionPeriodInHours": number
  },
  "Schema": {
    "CompositePartitionKey": [
      {
        "EnforcementInRecord": "string",
        "Name": "string",
        "Type": "string"
      }
    ]
  },
  "TableName": "string",
  "Tags": [
    {
      "Key": "string",
      "Value": "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.

DatabaseName

The name of the Timestream database.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Pattern: [a-zA-Z0-9_.-]+

Required: Yes

MagneticStoreWriteProperties

Contains properties to set on the table when enabling magnetic store writes.

Type: [MagneticStoreWriteProperties](#) object

Required: No

RetentionProperties

The duration for which your time-series data must be stored in the memory store and the magnetic store.

Type: [RetentionProperties](#) object

Required: No

Schema

The schema of the table.

Type: [Schema](#) object

Required: No

TableName

The name of the Timestream table.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Pattern: [a-zA-Z0-9_.-]+

Required: Yes

Tags

A list of key-value pairs to label the table.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Required: No

Response Syntax

```
{
  "Table": {
    "Arn": "string",
    "CreationTime": number,
    "DatabaseName": "string",
    "LastUpdatedTime": number,
    "MagneticStoreWriteProperties": {
      "EnableMagneticStoreWrites": boolean,
      "MagneticStoreRejectedDataLocation": {
        "S3Configuration": {
          "BucketName": "string",
          "EncryptionOption": "string",
          "KmsKeyId": "string",
          "ObjectKeyPrefix": "string"
        }
      }
    }
  },
}
```

```
"RetentionProperties": {
  "MagneticStoreRetentionPeriodInDays": number,
  "MemoryStoreRetentionPeriodInHours": number
},
"Schema": {
  "CompositePartitionKey": [
    {
      "EnforcementInRecord": "string",
      "Name": "string",
      "Type": "string"
    }
  ]
},
"TableName": "string",
"TableStatus": "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.

Table

The newly created Timestream table.

Type: [Table](#) object

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You are not authorized to perform this action.

HTTP Status Code: 400

ConflictException

Timestream was unable to process this request because it contains resource that already exists.

HTTP Status Code: 400

InternalServerErrorException

Timestream was unable to fully process this request because of an internal server error.

HTTP Status Code: 500

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ResourceNotFoundException

The operation tried to access a nonexistent resource. The resource might not be specified correctly, or its status might not be ACTIVE.

HTTP Status Code: 400

ServiceQuotaExceededException

The instance quota of resource exceeded for this account.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

DeleteDatabase

Service: Amazon Timestream Write

Deletes a given Timestream database. *This is an irreversible operation. After a database is deleted, the time-series data from its tables cannot be recovered.*

Note

All tables in the database must be deleted first, or a `ValidationException` error will be thrown.

Due to the nature of distributed retries, the operation can return either success or a `ResourceNotFoundException`. Clients should consider them equivalent.

See [code sample](#) for details.

Request Syntax

```
{
  "DatabaseName": "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.

DatabaseName

The name of the Timestream database to be deleted.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Required: Yes

Response Elements

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

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You are not authorized to perform this action.

HTTP Status Code: 400

InternalServerErrorException

Timestream was unable to fully process this request because of an internal server error.

HTTP Status Code: 500

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ResourceNotFoundException

The operation tried to access a nonexistent resource. The resource might not be specified correctly, or its status might not be ACTIVE.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

DeleteTable

Service: Amazon Timestream Write

Deletes a given Timestream table. This is an irreversible operation. After a Timestream database table is deleted, the time-series data stored in the table cannot be recovered.

Note

Due to the nature of distributed retries, the operation can return either success or a `ResourceNotFoundException`. Clients should consider them equivalent.

See [code sample](#) for details.

Request Syntax

```
{
  "DatabaseName": "string",
  "TableName": "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.

DatabaseName

The name of the database where the Timestream database is to be deleted.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Required: Yes

TableName

The name of the Timestream table to be deleted.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Required: Yes

Response Elements

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

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You are not authorized to perform this action.

HTTP Status Code: 400

InternalServerError

Timestream was unable to fully process this request because of an internal server error.

HTTP Status Code: 500

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ResourceNotFoundException

The operation tried to access a nonexistent resource. The resource might not be specified correctly, or its status might not be ACTIVE.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

DescribeBatchLoadTask

Service: Amazon Timestream Write

Returns information about the batch load task, including configurations, mappings, progress, and other details. [Service quotas apply](#). See [code sample](#) for details.

Request Syntax

```
{
  "TaskId": "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.

TaskId

The ID of the batch load task.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 32.

Pattern: [A-Z0-9]+

Required: Yes

Response Syntax

```
{
  "BatchLoadTaskDescription": {
    "CreationTime": number,
    "DataModelConfiguration": {
      "DataModel": {
        "DimensionMappings": [
          {
            "DestinationColumn": "string",
            "SourceColumn": "string"
          }
        ]
      }
    }
  }
}
```

```

    ],
    "MeasureNameColumn": "string",
    "MixedMeasureMappings": [
      {
        "MeasureName": "string",
        "MeasureValueType": "string",
        "MultiMeasureAttributeMappings": [
          {
            "MeasureValueType": "string",
            "SourceColumn": "string",
            "TargetMultiMeasureAttributeName": "string"
          }
        ],
        "SourceColumn": "string",
        "TargetMeasureName": "string"
      }
    ],
    "MultiMeasureMappings": {
      "MultiMeasureAttributeMappings": [
        {
          "MeasureValueType": "string",
          "SourceColumn": "string",
          "TargetMultiMeasureAttributeName": "string"
        }
      ],
      "TargetMultiMeasureName": "string"
    },
    "TimeColumn": "string",
    "TimeUnit": "string"
  },
  "DataModelS3Configuration": {
    "BucketName": "string",
    "ObjectKey": "string"
  }
},
"DataSourceConfiguration": {
  "CsvConfiguration": {
    "ColumnSeparator": "string",
    "EscapeChar": "string",
    "NullValue": "string",
    "QuoteChar": "string",
    "TrimWhiteSpace": boolean
  },
  "DataFormat": "string",

```

```
    "DataSourceS3Configuration": {
      "BucketName": "string",
      "ObjectKeyPrefix": "string"
    },
    "ErrorMessage": "string",
    "LastUpdateTime": number,
    "ProgressReport": {
      "BytesMetered": number,
      "FileFailures": number,
      "ParseFailures": number,
      "RecordIngestionFailures": number,
      "RecordsIngested": number,
      "RecordsProcessed": number
    },
    "RecordVersion": number,
    "ReportConfiguration": {
      "ReportS3Configuration": {
        "BucketName": "string",
        "EncryptionOption": "string",
        "KmsKeyId": "string",
        "ObjectKeyPrefix": "string"
      }
    },
    "ResumableUntil": number,
    "TargetDatabaseName": "string",
    "TargetTableName": "string",
    "TaskId": "string",
    "TaskStatus": "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.

BatchLoadTaskDescription

Description of the batch load task.

Type: [BatchLoadTaskDescription](#) object

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You are not authorized to perform this action.

HTTP Status Code: 400

InternalServerErrorException

Timestream was unable to fully process this request because of an internal server error.

HTTP Status Code: 500

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ResourceNotFoundException

The operation tried to access a nonexistent resource. The resource might not be specified correctly, or its status might not be ACTIVE.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

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](#)
- [AWS SDK for Ruby V3](#)

DescribeDatabase

Service: Amazon Timestream Write

Returns information about the database, including the database name, time that the database was created, and the total number of tables found within the database. [Service quotas apply](#). See [code sample](#) for details.

Request Syntax

```
{  
  "DatabaseName": "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.

[DatabaseName](#)

The name of the Timestream database.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Required: Yes

Response Syntax

```
{  
  "Database": {  
    "Arn": "string",  
    "CreationTime": number,  
    "DatabaseName": "string",  
    "KmsKeyId": "string",  
    "LastUpdatedTime": number,  
    "TableCount": 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.

Database

The name of the Timestream table.

Type: [Database](#) object

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You are not authorized to perform this action.

HTTP Status Code: 400

InternalServerErrorException

Timestream was unable to fully process this request because of an internal server error.

HTTP Status Code: 500

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ResourceNotFoundException

The operation tried to access a nonexistent resource. The resource might not be specified correctly, or its status might not be ACTIVE.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

DescribeEndpoints

Service: Amazon Timestream Write

Returns a list of available endpoints to make Timestream API calls against. This API operation is available through both the Write and Query APIs.

Because the Timestream SDKs are designed to transparently work with the service's architecture, including the management and mapping of the service endpoints, *we don't recommend that you use this API operation unless:*

- You are using [VPC endpoints \(AWS PrivateLink\) with Timestream](#)
- Your application uses a programming language that does not yet have SDK support
- You require better control over the client-side implementation

For detailed information on how and when to use and implement DescribeEndpoints, see [The Endpoint Discovery Pattern](#).

Response Syntax

```
{
  "Endpoints": [
    {
      "Address": "string",
      "CachePeriodInMinutes": 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.

[Endpoints](#)

An Endpoints object is returned when a DescribeEndpoints request is made.

Type: Array of [Endpoint](#) objects

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerErrorException

Timestream was unable to fully process this request because of an internal server error.

HTTP Status Code: 500

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

DescribeTable

Service: Amazon Timestream Write

Returns information about the table, including the table name, database name, retention duration of the memory store and the magnetic store. [Service quotas apply](#). See [code sample](#) for details.

Request Syntax

```
{  
  "DatabaseName": "string",  
  "TableName": "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.

DatabaseName

The name of the Timestream database.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Required: Yes

TableName

The name of the Timestream table.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Required: Yes

Response Syntax

```
{  
  "Table": {
```

```

    "Arn": "string",
    "CreationTime": number,
    "DatabaseName": "string",
    "LastUpdatedTime": number,
    "MagneticStoreWriteProperties": {
      "EnableMagneticStoreWrites": boolean,
      "MagneticStoreRejectedDataLocation": {
        "S3Configuration": {
          "BucketName": "string",
          "EncryptionOption": "string",
          "KmsKeyId": "string",
          "ObjectKeyPrefix": "string"
        }
      }
    },
    "RetentionProperties": {
      "MagneticStoreRetentionPeriodInDays": number,
      "MemoryStoreRetentionPeriodInHours": number
    },
    "Schema": {
      "CompositePartitionKey": [
        {
          "EnforcementInRecord": "string",
          "Name": "string",
          "Type": "string"
        }
      ]
    },
    "TableName": "string",
    "TableStatus": "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.

Table

The Timestream table.

Type: [Table](#) object

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You are not authorized to perform this action.

HTTP Status Code: 400

InternalServerErrorException

Timestream was unable to fully process this request because of an internal server error.

HTTP Status Code: 500

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ResourceNotFoundException

The operation tried to access a nonexistent resource. The resource might not be specified correctly, or its status might not be ACTIVE.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

ListBatchLoadTasks

Service: Amazon Timestream Write

Provides a list of batch load tasks, along with the name, status, when the task is resumable until, and other details. See [code sample](#) for details.

Request Syntax

```
{  
  "MaxResults": number,  
  "NextToken": "string",  
  "TaskStatus": "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.

MaxResults

The total number of items to return in the output. If the total number of items available is more than the value specified, a NextToken is provided in the output. To resume pagination, provide the NextToken value as argument of a subsequent API invocation.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 100.

Required: No

NextToken

A token to specify where to start paginating. This is the NextToken from a previously truncated response.

Type: String

Required: No

TaskStatus

Status of the batch load task.

Type: String

Valid Values: CREATED | IN_PROGRESS | FAILED | SUCCEEDED | PROGRESS_STOPPED
| PENDING_RESUME

Required: No

Response Syntax

```
{
  "BatchLoadTasks": [
    {
      "CreationTime": number,
      "DatabaseName": "string",
      "LastUpdatedTime": number,
      "ResumableUntil": number,
      "TableName": "string",
      "TaskId": "string",
      "TaskStatus": "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.

BatchLoadTasks

A list of batch load task details.

Type: Array of [BatchLoadTask](#) objects

NextToken

A token to specify where to start paginating. Provide the next ListBatchLoadTasksRequest.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You are not authorized to perform this action.

HTTP Status Code: 400

InternalServerErrorException

Timestream was unable to fully process this request because of an internal server error.

HTTP Status Code: 500

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

ListDatabases

Service: Amazon Timestream Write

Returns a list of your Timestream databases. [Service quotas apply](#). See [code sample](#) for details.

Request Syntax

```
{  
  "MaxResults": number,  
  "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.

[MaxResults](#)

The total number of items to return in the output. If the total number of items available is more than the value specified, a NextToken is provided in the output. To resume pagination, provide the NextToken value as argument of a subsequent API invocation.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 20.

Required: No

[NextToken](#)

The pagination token. To resume pagination, provide the NextToken value as argument of a subsequent API invocation.

Type: String

Required: No

Response Syntax

```
{
```

```
"Databases": [  
  {  
    "Arn": "string",  
    "CreationTime": number,  
    "DatabaseName": "string",  
    "KmsKeyId": "string",  
    "LastUpdatedTime": number,  
    "TableCount": number  
  }  
],  
"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

A list of database names.

Type: Array of [Database](#) objects

NextToken

The pagination token. This parameter is returned when the response is truncated.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You are not authorized to perform this action.

HTTP Status Code: 400

InternalServerError

Timestream was unable to fully process this request because of an internal server error.

HTTP Status Code: 500

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

ListTables

Service: Amazon Timestream Write

Provides a list of tables, along with the name, status, and retention properties of each table. See [code sample](#) for details.

Request Syntax

```
{  
  "DatabaseName": "string",  
  "MaxResults": number,  
  "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.

DatabaseName

The name of the Timestream database.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Required: No

MaxResults

The total number of items to return in the output. If the total number of items available is more than the value specified, a NextToken is provided in the output. To resume pagination, provide the NextToken value as argument of a subsequent API invocation.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 20.

Required: No

NextToken

The pagination token. To resume pagination, provide the NextToken value as argument of a subsequent API invocation.

Type: String

Required: No

Response Syntax

```
{
  "NextToken": "string",
  "Tables": [
    {
      "Arn": "string",
      "CreationTime": number,
      "DatabaseName": "string",
      "LastUpdatedTime": number,
      "MagneticStoreWriteProperties": {
        "EnableMagneticStoreWrites": boolean,
        "MagneticStoreRejectedDataLocation": {
          "S3Configuration": {
            "BucketName": "string",
            "EncryptionOption": "string",
            "KmsKeyId": "string",
            "ObjectKeyPrefix": "string"
          }
        }
      },
      "RetentionProperties": {
        "MagneticStoreRetentionPeriodInDays": number,
        "MemoryStoreRetentionPeriodInHours": number
      },
      "Schema": {
        "CompositePartitionKey": [
          {
            "EnforcementInRecord": "string",
            "Name": "string",
            "Type": "string"
          }
        ]
      }
    }
  ],
}
```

```
    "TableName": "string",  
    "TableStatus": "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 token to specify where to start paginating. This is the NextToken from a previously truncated response.

Type: String

Tables

A list of tables.

Type: Array of [Table](#) objects

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You are not authorized to perform this action.

HTTP Status Code: 400

InternalServerError

Timestream was unable to fully process this request because of an internal server error.

HTTP Status Code: 500

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ResourceNotFoundException

The operation tried to access a nonexistent resource. The resource might not be specified correctly, or its status might not be ACTIVE.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

ListTagsForResource

Service: Amazon Timestream Write

Lists all tags on a Timestream resource.

Request Syntax

```
{
  "ResourceARN": "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.

ResourceARN

The Timestream resource with tags to be listed. This value is an Amazon Resource Name (ARN).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1011.

Required: Yes

Response Syntax

```
{
  "Tags": [
    {
      "Key": "string",
      "Value": "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.

Tags

The tags currently associated with the Timestream resource.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ResourceNotFoundException

The operation tried to access a nonexistent resource. The resource might not be specified correctly, or its status might not be ACTIVE.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

ResumeBatchLoadTask

Service: Amazon Timestream Write

Request Syntax

```
{  
  "TaskId": "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.

TaskId

The ID of the batch load task to resume.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 32.

Pattern: [A-Z0-9]+

Required: Yes

Response Elements

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

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You are not authorized to perform this action.

HTTP Status Code: 400

InternalServerErrorException

Timestream was unable to fully process this request because of an internal server error.

HTTP Status Code: 500

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ResourceNotFoundException

The operation tried to access a nonexistent resource. The resource might not be specified correctly, or its status might not be ACTIVE.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

TagResource

Service: Amazon Timestream Write

Associates a set of tags with a Timestream resource. You can then activate these user-defined tags so that they appear on the Billing and Cost Management console for cost allocation tracking.

Request Syntax

```
{
  "ResourceARN": "string",
  "Tags": [
    {
      "Key": "string",
      "Value": "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.

[ResourceARN](#)

Identifies the Timestream resource to which tags should be added. This value is an Amazon Resource Name (ARN).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1011.

Required: Yes

[Tags](#)

The tags to be assigned to the Timestream resource.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Required: Yes

Response Elements

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

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ResourceNotFoundException

The operation tried to access a nonexistent resource. The resource might not be specified correctly, or its status might not be ACTIVE.

HTTP Status Code: 400

ServiceQuotaExceededException

The instance quota of resource exceeded for this account.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

UntagResource

Service: Amazon Timestream Write

Removes the association of tags from a Timestream resource.

Request Syntax

```
{
  "ResourceARN": "string",
  "TagKeys": [ "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.

ResourceARN

The Timestream resource that the tags will be removed from. This value is an Amazon Resource Name (ARN).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1011.

Required: Yes

TagKeys

A list of tags keys. Existing tags of the resource whose keys are members of this list will be removed from the Timestream resource.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Length Constraints: Minimum length of 1. Maximum length of 128.

Required: Yes

Response Elements

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

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ResourceNotFoundException

The operation tried to access a nonexistent resource. The resource might not be specified correctly, or its status might not be ACTIVE.

HTTP Status Code: 400

ServiceQuotaExceededException

The instance quota of resource exceeded for this account.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

UpdateDatabase

Service: Amazon Timestream Write

Modifies the AWS KMS key for an existing database. While updating the database, you must specify the database name and the identifier of the new AWS KMS key to be used (`KmsKeyId`). If there are any concurrent `UpdateDatabase` requests, first writer wins.

See [code sample](#) for details.

Request Syntax

```
{
  "DatabaseName": "string",
  "KmsKeyId": "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.

DatabaseName

The name of the database.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Required: Yes

KmsKeyId

The identifier of the new AWS KMS key (`KmsKeyId`) to be used to encrypt the data stored in the database. If the `KmsKeyId` currently registered with the database is the same as the `KmsKeyId` in the request, there will not be any update.

You can specify the `KmsKeyId` using any of the following:

- Key ID: 1234abcd-12ab-34cd-56ef-1234567890ab
- Key ARN: arn:aws:kms:us-east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab

- Alias name: `alias/ExampleAlias`
- Alias ARN: `arn:aws:kms:us-east-1:111122223333:alias/ExampleAlias`

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: Yes

Response Syntax

```
{
  "Database": {
    "Arn": "string",
    "CreationTime": number,
    "DatabaseName": "string",
    "KmsKeyId": "string",
    "LastUpdatedTime": number,
    "TableCount": 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.

Database

A top-level container for a table. Databases and tables are the fundamental management concepts in Amazon Timestream. All tables in a database are encrypted with the same AWS KMS key.

Type: [Database](#) object

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You are not authorized to perform this action.

HTTP Status Code: 400

InternalServerErrorException

Timestream was unable to fully process this request because of an internal server error.

HTTP Status Code: 500

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ResourceNotFoundException

The operation tried to access a nonexistent resource. The resource might not be specified correctly, or its status might not be ACTIVE.

HTTP Status Code: 400

ServiceQuotaExceededException

The instance quota of resource exceeded for this account.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

UpdateTable

Service: Amazon Timestream Write

Modifies the retention duration of the memory store and magnetic store for your Timestream table. Note that the change in retention duration takes effect immediately. For example, if the retention period of the memory store was initially set to 2 hours and then changed to 24 hours, the memory store will be capable of holding 24 hours of data, but will be populated with 24 hours of data 22 hours after this change was made. Timestream does not retrieve data from the magnetic store to populate the memory store.

See [code sample](#) for details.

Request Syntax

```
{
  "DatabaseName": "string",
  "MagneticStoreWriteProperties": {
    "EnableMagneticStoreWrites": boolean,
    "MagneticStoreRejectedDataLocation": {
      "S3Configuration": {
        "BucketName": "string",
        "EncryptionOption": "string",
        "KmsKeyId": "string",
        "ObjectKeyPrefix": "string"
      }
    }
  },
  "RetentionProperties": {
    "MagneticStoreRetentionPeriodInDays": number,
    "MemoryStoreRetentionPeriodInHours": number
  },
  "Schema": {
    "CompositePartitionKey": [
      {
        "EnforcementInRecord": "string",
        "Name": "string",
        "Type": "string"
      }
    ]
  },
  "TableName": "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.

DatabaseName

The name of the Timestream database.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Required: Yes

MagneticStoreWriteProperties

Contains properties to set on the table when enabling magnetic store writes.

Type: [MagneticStoreWriteProperties](#) object

Required: No

RetentionProperties

The retention duration of the memory store and the magnetic store.

Type: [RetentionProperties](#) object

Required: No

Schema

The schema of the table.

Type: [Schema](#) object

Required: No

TableName

The name of the Timestream table.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Required: Yes

Response Syntax

```
{
  "Table": {
    "Arn": "string",
    "CreationTime": number,
    "DatabaseName": "string",
    "LastUpdatedTime": number,
    "MagneticStoreWriteProperties": {
      "EnableMagneticStoreWrites": boolean,
      "MagneticStoreRejectedDataLocation": {
        "S3Configuration": {
          "BucketName": "string",
          "EncryptionOption": "string",
          "KmsKeyId": "string",
          "ObjectKeyPrefix": "string"
        }
      }
    },
    "RetentionProperties": {
      "MagneticStoreRetentionPeriodInDays": number,
      "MemoryStoreRetentionPeriodInHours": number
    },
    "Schema": {
      "CompositePartitionKey": [
        {
          "EnforcementInRecord": "string",
          "Name": "string",
          "Type": "string"
        }
      ]
    },
    "TableName": "string",
    "TableStatus": "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.

Table

The updated Timestream table.

Type: [Table](#) object

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You are not authorized to perform this action.

HTTP Status Code: 400

InternalServerError

Timestream was unable to fully process this request because of an internal server error.

HTTP Status Code: 500

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

ResourceNotFoundException

The operation tried to access a nonexistent resource. The resource might not be specified correctly, or its status might not be ACTIVE.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

WriteRecords

Service: Amazon Timestream Write

Enables you to write your time-series data into Timestream. You can specify a single data point or a batch of data points to be inserted into the system. Timestream offers you a flexible schema that auto detects the column names and data types for your Timestream tables based on the dimension names and data types of the data points you specify when invoking writes into the database.

Timestream supports eventual consistency read semantics. This means that when you query data immediately after writing a batch of data into Timestream, the query results might not reflect the results of a recently completed write operation. The results may also include some stale data. If you repeat the query request after a short time, the results should return the latest data. [Service quotas apply](#).

See [code sample](#) for details.

Upserts

You can use the `Version` parameter in a `WriteRecords` request to update data points. Timestream tracks a version number with each record. `Version` defaults to 1 when it's not specified for the record in the request. Timestream updates an existing record's measure value along with its `Version` when it receives a write request with a higher `Version` number for that record. When it receives an update request where the measure value is the same as that of the existing record, Timestream still updates `Version`, if it is greater than the existing value of `Version`. You can update a data point as many times as desired, as long as the value of `Version` continuously increases.

For example, suppose you write a new record without indicating `Version` in the request. Timestream stores this record, and set `Version` to 1. Now, suppose you try to update this record with a `WriteRecords` request of the same record with a different measure value but, like before, do not provide `Version`. In this case, Timestream will reject this update with a `RejectedRecordsException` since the updated record's version is not greater than the existing value of `Version`.

However, if you were to resend the update request with `Version` set to 2, Timestream would then succeed in updating the record's value, and the `Version` would be set to 2. Next, suppose you sent a `WriteRecords` request with this same record and an identical measure value, but with `Version` set to 3. In this case, Timestream would only update `Version` to 3. Any further

updates would need to send a version number greater than 3, or the update requests would receive a `RejectedRecordsException`.

Request Syntax

```
{
  "CommonAttributes": {
    "Dimensions": [
      {
        "DimensionValueType": "string",
        "Name": "string",
        "Value": "string"
      }
    ],
    "MeasureName": "string",
    "MeasureValue": "string",
    "MeasureValues": [
      {
        "Name": "string",
        "Type": "string",
        "Value": "string"
      }
    ],
    "MeasureValueType": "string",
    "Time": "string",
    "TimeUnit": "string",
    "Version": number
  },
  "DatabaseName": "string",
  "Records": [
    {
      "Dimensions": [
        {
          "DimensionValueType": "string",
          "Name": "string",
          "Value": "string"
        }
      ],
      "MeasureName": "string",
      "MeasureValue": "string",
      "MeasureValues": [
        {
          "Name": "string",
```

```
        "Type": "string",
        "Value": "string"
    }
],
"MeasureValueType": "string",
"Time": "string",
"TimeUnit": "string",
"Version": number
}
],
"TableName": "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.

CommonAttributes

A record that contains the common measure, dimension, time, and version attributes shared across all the records in the request. The measure and dimension attributes specified will be merged with the measure and dimension attributes in the records object when the data is written into Timestream. Dimensions may not overlap, or a `ValidationException` will be thrown. In other words, a record must contain dimensions with unique names.

Type: [Record](#) object

Required: No

DatabaseName

The name of the Timestream database.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Required: Yes

Records

An array of records that contain the unique measure, dimension, time, and version attributes for each time-series data point.

Type: Array of [Record](#) objects

Array Members: Minimum number of 1 item. Maximum number of 100 items.

Required: Yes

TableName

The name of the Timestream table.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Required: Yes

Response Syntax

```
{
  "RecordsIngested": {
    "MagneticStore": number,
    "MemoryStore": number,
    "Total": 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.

RecordsIngested

Information on the records ingested by this request.

Type: [RecordsIngested](#) object

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You are not authorized to perform this action.

HTTP Status Code: 400

InternalServerErrorException

Timestream was unable to fully process this request because of an internal server error.

HTTP Status Code: 500

InvalidEndpointException

The requested endpoint was not valid.

HTTP Status Code: 400

RejectedRecordsException

WriteRecords would throw this exception in the following cases:

- Records with duplicate data where there are multiple records with the same dimensions, timestamps, and measure names but:
 - Measure values are different
 - Version is not present in the request *or* the value of version in the new record is equal to or lower than the existing value

In this case, if Timestream rejects data, the `ExistingVersion` field in the `RejectedRecords` response will indicate the current record's version. To force an update, you can resend the request with a version for the record set to a value greater than the `ExistingVersion`.

- Records with timestamps that lie outside the retention duration of the memory store.
- Records with dimensions or measures that exceed the Timestream defined limits.

For more information, see [Quotas](#) in the Amazon Timestream Developer Guide.

RejectedRecords

HTTP Status Code: 400

ResourceNotFoundException

The operation tried to access a nonexistent resource. The resource might not be specified correctly, or its status might not be `ACTIVE`.

HTTP Status Code: 400

ThrottlingException

Too many requests were made by a user and they exceeded the service quotas. The request was throttled.

HTTP Status Code: 400

ValidationException

An invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

Amazon Timestream Query

The following actions are supported by Amazon Timestream Query:

- [CancelQuery](#)
- [CreateScheduledQuery](#)

- [DeleteScheduledQuery](#)
- [DescribeAccountSettings](#)
- [DescribeEndpoints](#)
- [DescribeScheduledQuery](#)
- [ExecuteScheduledQuery](#)
- [ListScheduledQueries](#)
- [ListTagsForResource](#)
- [PrepareQuery](#)
- [Query](#)
- [TagResource](#)
- [UntagResource](#)
- [UpdateAccountSettings](#)
- [UpdateScheduledQuery](#)

CancelQuery

Service: Amazon Timestream Query

Cancels a query that has been issued. Cancellation is provided only if the query has not completed running before the cancellation request was issued. Because cancellation is an idempotent operation, subsequent cancellation requests will return a `CancellationMessage`, indicating that the query has already been canceled. See [code sample](#) for details.

Request Syntax

```
{  
  "QueryId": "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.

QueryId

The ID of the query that needs to be cancelled. QueryID is returned as part of the query result.

Type: String

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

Pattern: `[a-zA-Z0-9]+`

Required: Yes

Response Syntax

```
{  
  "CancellationMessage": "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.

CancellationMessage

A CancellationMessage is returned when a CancelQuery request for the query specified by QueryId has already been issued.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You do not have the necessary permissions to access the account settings.

HTTP Status Code: 400

InternalServerError

An internal server error occurred while processing the request.

HTTP Status Code: 400

InvalidEndpointException

The requested endpoint is invalid.

HTTP Status Code: 400

ThrottlingException

The request was throttled due to excessive requests.

HTTP Status Code: 400

ValidationException

Invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

CreateScheduledQuery

Service: Amazon Timestream Query

Create a scheduled query that will be run on your behalf at the configured schedule. Timestream assumes the execution role provided as part of the `ScheduledQueryExecutionRoleArn` parameter to run the query. You can use the `NotificationConfiguration` parameter to configure notification for your scheduled query operations.

Request Syntax

```
{
  "ClientToken": "string",
  "ErrorReportConfiguration": {
    "S3Configuration": {
      "BucketName": "string",
      "EncryptionOption": "string",
      "ObjectKeyPrefix": "string"
    }
  },
  "KmsKeyId": "string",
  "Name": "string",
  "NotificationConfiguration": {
    "SnsConfiguration": {
      "TopicArn": "string"
    }
  },
  "QueryString": "string",
  "ScheduleConfiguration": {
    "ScheduleExpression": "string"
  },
  "ScheduledQueryExecutionRoleArn": "string",
  "Tags": [
    {
      "Key": "string",
      "Value": "string"
    }
  ],
  "TargetConfiguration": {
    "TimestreamConfiguration": {
      "DatabaseName": "string",
      "DimensionMappings": [
        {
```

```

        "DimensionValueType": "string",
        "Name": "string"
    }
],
"MeasureNameColumn": "string",
"MixedMeasureMappings": [
    {
        "MeasureName": "string",
        "MeasureValueType": "string",
        "MultiMeasureAttributeMappings": [
            {
                "MeasureValueType": "string",
                "SourceColumn": "string",
                "TargetMultiMeasureAttributeName": "string"
            }
        ],
        "SourceColumn": "string",
        "TargetMeasureName": "string"
    }
],
"MultiMeasureMappings": {
    "MultiMeasureAttributeMappings": [
        {
            "MeasureValueType": "string",
            "SourceColumn": "string",
            "TargetMultiMeasureAttributeName": "string"
        }
    ],
    "TargetMultiMeasureName": "string"
},
"TableName": "string",
"TimeColumn": "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.

ClientToken

Using a `ClientToken` makes the call to `CreateScheduledQuery` idempotent, in other words, making the same request repeatedly will produce the same result. Making multiple identical `CreateScheduledQuery` requests has the same effect as making a single request.

- If `CreateScheduledQuery` is called without a `ClientToken`, the Query SDK generates a `ClientToken` on your behalf.
- After 8 hours, any request with the same `ClientToken` is treated as a new request.

Type: String

Length Constraints: Minimum length of 32. Maximum length of 128.

Required: No

ErrorReportConfiguration

Configuration for error reporting. Error reports will be generated when a problem is encountered when writing the query results.

Type: [ErrorReportConfiguration](#) object

Required: Yes

KmsKeyId

The Amazon KMS key used to encrypt the scheduled query resource, at-rest. If the Amazon KMS key is not specified, the scheduled query resource will be encrypted with a Timestream owned Amazon KMS key. To specify a KMS key, use the key ID, key ARN, alias name, or alias ARN. When using an alias name, prefix the name with *alias/*

If `ErrorReportConfiguration` uses `SSE_KMS` as encryption type, the same `KmsKeyId` is used to encrypt the error report at rest.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

Name

Name of the scheduled query.

Type: String

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

Pattern: `[a-zA-Z0-9|!_ * '\\(\\)]([a-zA-Z0-9|!_ * '\\(\\)\\/\\.])+`

Required: Yes

NotificationConfiguration

Notification configuration for the scheduled query. A notification is sent by Timestream when a query run finishes, when the state is updated or when you delete it.

Type: [NotificationConfiguration](#) object

Required: Yes

QueryString

The query string to run. Parameter names can be specified in the query string @ character followed by an identifier. The named Parameter @scheduled_runtime is reserved and can be used in the query to get the time at which the query is scheduled to run.

The timestamp calculated according to the ScheduleConfiguration parameter, will be the value of @scheduled_runtime parameter for each query run. For example, consider an instance of a scheduled query executing on 2021-12-01 00:00:00. For this instance, the @scheduled_runtime parameter is initialized to the timestamp 2021-12-01 00:00:00 when invoking the query.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 262144.

Required: Yes

ScheduleConfiguration

The schedule configuration for the query.

Type: [ScheduleConfiguration](#) object

Required: Yes

ScheduledQueryExecutionRoleArn

The ARN for the IAM role that Timestream will assume when running the scheduled query.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: Yes

Tags

A list of key-value pairs to label the scheduled query.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Required: No

TargetConfiguration

Configuration used for writing the result of a query.

Type: [TargetConfiguration](#) object

Required: No

Response Syntax

```
{  
  "Arn": "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.

Arn

ARN for the created scheduled query.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You do not have the necessary permissions to access the account settings.

HTTP Status Code: 400

ConflictException

Unable to poll results for a cancelled query.

HTTP Status Code: 400

InternalServerErrorException

An internal server error occurred while processing the request.

HTTP Status Code: 400

InvalidEndpointException

The requested endpoint is invalid.

HTTP Status Code: 400

ServiceQuotaExceededException

You have exceeded the service quota.

HTTP Status Code: 400

ThrottlingException

The request was throttled due to excessive requests.

HTTP Status Code: 400

ValidationException

Invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

DeleteScheduledQuery

Service: Amazon Timestream Query

Deletes a given scheduled query. This is an irreversible operation.

Request Syntax

```
{  
  "ScheduledQueryArn": "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.

ScheduledQueryArn

The ARN of the scheduled query.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: Yes

Response Elements

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

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You do not have the necessary permissions to access the account settings.

HTTP Status Code: 400

InternalServerErrorException

An internal server error occurred while processing the request.

HTTP Status Code: 400

InvalidEndpointException

The requested endpoint is invalid.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found.

ScheduledQueryArn

The ARN of the scheduled query.

HTTP Status Code: 400

ThrottlingException

The request was throttled due to excessive requests.

HTTP Status Code: 400

ValidationException

Invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

DescribeAccountSettings

Service: Amazon Timestream Query

Describes the settings for your account that include the query pricing model and the configured maximum TCUs the service can use for your query workload.

You're charged only for the duration of compute units used for your workloads.

Response Syntax

```
{
  "MaxQueryTCU": number,
  "QueryCompute": {
    "ComputeMode": "string",
    "ProvisionedCapacity": {
      "ActiveQueryTCU": number,
      "LastUpdate": {
        "Status": "string",
        "StatusMessage": "string",
        "TargetQueryTCU": number
      },
      "NotificationConfiguration": {
        "RoleArn": "string",
        "SnsConfiguration": {
          "TopicArn": "string"
        }
      }
    }
  },
  "QueryPricingModel": "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.

MaxQueryTCU

The maximum number of [Timestream compute units](#) (TCUs) the service will use at any point in time to serve your queries. To run queries, you must set a minimum capacity of 4 TCU. You

can set the maximum number of TCU in multiples of 4, for example, 4, 8, 16, 32, and so on. This configuration is applicable only for on-demand usage of (TCUs).

Type: Integer

QueryCompute

An object that contains the usage settings for Timestream Compute Units (TCUs) in your account for the query workload. QueryCompute is available only in the Asia Pacific (Mumbai) region.

Type: [QueryComputeResponse](#) object

QueryPricingModel

The pricing model for queries in your account.

Note

The `QueryPricingModel` parameter is used by several Timestream operations; however, the `UpdateAccountSettings` API operation doesn't recognize any values other than `COMPUTE_UNITS`.

Type: String

Valid Values: `BYTES_SCANNED` | `COMPUTE_UNITS`

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You do not have the necessary permissions to access the account settings.

HTTP Status Code: 400

InternalServerError

An internal server error occurred while processing the request.

HTTP Status Code: 400

InvalidEndpointException

The requested endpoint is invalid.

HTTP Status Code: 400

ThrottlingException

The request was throttled due to excessive requests.

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](#)
- [AWS SDK for Ruby V3](#)

DescribeEndpoints

Service: Amazon Timestream Query

DescribeEndpoints returns a list of available endpoints to make Timestream API calls against. This API is available through both Write and Query.

Because the Timestream SDKs are designed to transparently work with the service's architecture, including the management and mapping of the service endpoints, *it is not recommended that you use this API unless:*

- You are using [VPC endpoints \(AWS PrivateLink\) with Timestream](#)
- Your application uses a programming language that does not yet have SDK support
- You require better control over the client-side implementation

For detailed information on how and when to use and implement DescribeEndpoints, see [The Endpoint Discovery Pattern](#).

Response Syntax

```
{
  "Endpoints": [
    {
      "Address": "string",
      "CachePeriodInMinutes": 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.

[Endpoints](#)

An Endpoints object is returned when a DescribeEndpoints request is made.

Type: Array of [Endpoint](#) objects

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerError

An internal server error occurred while processing the request.

HTTP Status Code: 400

ThrottlingException

The request was throttled due to excessive requests.

HTTP Status Code: 400

ValidationException

Invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

DescribeScheduledQuery

Service: Amazon Timestream Query

Provides detailed information about a scheduled query.

Request Syntax

```
{
  "ScheduledQueryArn": "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.

ScheduledQueryArn

The ARN of the scheduled query.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: Yes

Response Syntax

```
{
  "ScheduledQuery": {
    "Arn": "string",
    "CreationTime": number,
    "ErrorReportConfiguration": {
      "S3Configuration": {
        "BucketName": "string",
        "EncryptionOption": "string",
        "ObjectKeyPrefix": "string"
      }
    },
    "KmsKeyId": "string",
    "LastRunSummary": {
```

```

    "ErrorReportLocation": {
      "S3ReportLocation": {
        "BucketName": "string",
        "ObjectKey": "string"
      }
    },
    "ExecutionStats": {
      "BytesMetered": number,
      "CumulativeBytesScanned": number,
      "DataWrites": number,
      "ExecutionTimeInMillis": number,
      "QueryResultRows": number,
      "RecordsIngested": number
    },
    "FailureReason": "string",
    "InvocationTime": number,
    "QueryInsightsResponse": {
      "OutputBytes": number,
      "OutputRows": number,
      "QuerySpatialCoverage": {
        "Max": {
          "PartitionKey": [ "string" ],
          "TableArn": "string",
          "Value": number
        }
      }
    },
    "QueryTableCount": number,
    "QueryTemporalRange": {
      "Max": {
        "TableArn": "string",
        "Value": number
      }
    }
  },
  "RunStatus": "string",
  "TriggerTime": number
},
"Name": "string",
"NextInvocationTime": number,
"NotificationConfiguration": {
  "SnsConfiguration": {
    "TopicArn": "string"
  }
}
},

```

```

    "PreviousInvocationTime": number,
    "QueryString": "string",
    "RecentlyFailedRuns": [
      {
        "ErrorReportLocation": {
          "S3ReportLocation": {
            "BucketName": "string",
            "ObjectKey": "string"
          }
        },
        "ExecutionStats": {
          "BytesMetered": number,
          "CumulativeBytesScanned": number,
          "DataWrites": number,
          "ExecutionTimeInMillis": number,
          "QueryResultRows": number,
          "RecordsIngested": number
        },
        "FailureReason": "string",
        "InvocationTime": number,
        "QueryInsightsResponse": {
          "OutputBytes": number,
          "OutputRows": number,
          "QuerySpatialCoverage": {
            "Max": {
              "PartitionKey": [ "string" ],
              "TableArn": "string",
              "Value": number
            }
          }
        },
        "QueryTableCount": number,
        "QueryTemporalRange": {
          "Max": {
            "TableArn": "string",
            "Value": number
          }
        }
      },
      {
        "RunStatus": "string",
        "TriggerTime": number
      }
    ],
    "ScheduleConfiguration": {
      "ScheduleExpression": "string"
    }
  }

```

```

    },
    "ScheduledQueryExecutionRoleArn": "string",
    "State": "string",
    "TargetConfiguration": {
      "TimestreamConfiguration": {
        "DatabaseName": "string",
        "DimensionMappings": [
          {
            "DimensionValueType": "string",
            "Name": "string"
          }
        ],
        "MeasureNameColumn": "string",
        "MixedMeasureMappings": [
          {
            "MeasureName": "string",
            "MeasureValueType": "string",
            "MultiMeasureAttributeMappings": [
              {
                "MeasureValueType": "string",
                "SourceColumn": "string",
                "TargetMultiMeasureAttributeName": "string"
              }
            ],
            "SourceColumn": "string",
            "TargetMeasureName": "string"
          }
        ],
        "MultiMeasureMappings": {
          "MultiMeasureAttributeMappings": [
            {
              "MeasureValueType": "string",
              "SourceColumn": "string",
              "TargetMultiMeasureAttributeName": "string"
            }
          ],
          "TargetMultiMeasureName": "string"
        }
      },
      "TableName": "string",
      "TimeColumn": "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.

ScheduledQuery

The scheduled query.

Type: [ScheduledQueryDescription](#) object

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You do not have the necessary permissions to access the account settings.

HTTP Status Code: 400

InternalServerError

An internal server error occurred while processing the request.

HTTP Status Code: 400

InvalidEndpointException

The requested endpoint is invalid.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found.

ScheduledQueryArn

The ARN of the scheduled query.

HTTP Status Code: 400

ThrottlingException

The request was throttled due to excessive requests.

HTTP Status Code: 400

ValidationException

Invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

ExecuteScheduledQuery

Service: Amazon Timestream Query

You can use this API to run a scheduled query manually.

If you enabled QueryInsights, this API also returns insights and metrics related to the query that you executed as part of an Amazon SNS notification. QueryInsights helps with performance tuning of your query. For more information about QueryInsights, see [Using query insights to optimize queries in Amazon Timestream](#).

Request Syntax

```
{
  "ClientToken": "string",
  "InvocationTime": number,
  "QueryInsights": {
    "Mode": "string"
  },
  "ScheduledQueryArn": "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.

ClientToken

Not used.

Type: String

Length Constraints: Minimum length of 32. Maximum length of 128.

Required: No

InvocationTime

The timestamp in UTC. Query will be run as if it was invoked at this timestamp.

Type: Timestamp

Required: Yes

QueryInsights

Encapsulates settings for enabling QueryInsights.

Enabling QueryInsights returns insights and metrics as a part of the Amazon SNS notification for the query that you executed. You can use QueryInsights to tune your query performance and cost.

Type: [ScheduledQueryInsights](#) object

Required: No

ScheduledQueryArn

ARN of the scheduled query.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: Yes

Response Elements

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

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You do not have the necessary permissions to access the account settings.

HTTP Status Code: 400

InternalServerError

An internal server error occurred while processing the request.

HTTP Status Code: 400

InvalidEndpointException

The requested endpoint is invalid.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found.

ScheduledQueryArn

The ARN of the scheduled query.

HTTP Status Code: 400

ThrottlingException

The request was throttled due to excessive requests.

HTTP Status Code: 400

ValidationException

Invalid or malformed request.

HTTP Status Code: 400

Examples

Scheduled query notification message for the ENABLED_WITH_RATE_CONTROL mode

The following example shows a successful scheduled query notification message for the ENABLED_WITH_RATE_CONTROL mode of the QueryInsights parameter.

```
"SuccessNotificationMessage": {
  "type": "MANUAL_TRIGGER_SUCCESS",
  "arn": "arn:aws:timestream:<Region>:<Account>:scheduled-query/sq-test-49c6ed55-
c2e7-4cc2-9956-4a0ecea13420-80e05b035236a4c3",
  "scheduledQueryRunSummary": {
    "invocationEpochSecond": 1723710546,
    "triggerTimeMillis": 1723710547490,
    "runStatus": "MANUAL_TRIGGER_SUCCESS",
    "executionStats": {
      "executionTimeInMillis": 17343,
```

```

        "dataWrites": 1024,
        "bytesMetered": 0,
        "cumulativeBytesScanned": 600,
        "recordsIngested": 1,
        "queryResultRows": 1
    },
    "queryInsightsResponse": {
        "querySpatialCoverage": {
            "max": {
                "value": 1.0,
                "tableArn": "arn:aws:timestream:<Region>:<Account>:database/BaseDb/
table/BaseTable",
                "partitionKey": [
                    "measure_name"
                ]
            }
        },
        "queryTemporalRange": {
            "max": {
                "value": 2399999999999,
                "tableArn": "arn:aws:timestream:<Region>:<Account>:database/BaseDb/
table/BaseTable"
            }
        },
        "queryTableCount": 1,
        "outputRows": 1,
        "outputBytes": 59
    }
}

```

Scheduled query notification message for the DISABLED mode

The following example shows a successful scheduled query notification message for the DISABLED mode of the QueryInsights parameter.

```

"SuccessNotificationMessage": {
    "type": "MANUAL_TRIGGER_SUCCESS",
    "arn": "arn:aws:timestream:<Region>:<Account>:scheduled-query/sq-test-
fa109d9e-6528-4a0d-ac40-482fa05e657f-140faaeecdc5b2a7",
    "scheduledQueryRunSummary": {
        "invocationEpochSecond": 1723711401,
        "triggerTimeMillis": 1723711402144,
    }
}

```

```

    "runStatus": "MANUAL_TRIGGER_SUCCESS",
    "executionStats": {
      "executionTimeInMillis": 17992,
      "dataWrites": 1024,
      "bytesMetered": 0,
      "cumulativeBytesScanned": 600,
      "recordsIngested": 1,
      "queryResultRows": 1
    }
  }
}

```

Failure notification message for the ENABLED_WITH_RATE_CONTROL mode

The following example shows a failed scheduled query notification message for the ENABLED_WITH_RATE_CONTROL mode of the QueryInsights parameter.

```

"FailureNotificationMessage": {
  "type": "MANUAL_TRIGGER_FAILURE",
  "arn": "arn:aws:timestream:<Region>:<Account>:scheduled-query/sq-test-
b261670d-790c-4116-9db5-0798071b18b1-b7e27a1d79be226d",
  "scheduledQueryRunSummary": {
    "invocationEpochSecond": 1727915513,
    "triggerTimeInMillis": 1727915513894,
    "runStatus": "MANUAL_TRIGGER_FAILURE",
    "executionStats": {
      "executionTimeInMillis": 10777,
      "dataWrites": 0,
      "bytesMetered": 0,
      "cumulativeBytesScanned": 0,
      "recordsIngested": 0,
      "queryResultRows": 4
    },
    "errorReportLocation": {
      "s3ReportLocation": {
        "bucketName": "amzn-s3-demo-bucket",
        "objectKey": "4my-organization-f7a3c5d065a1a95e/1727915513/
MANUAL/1727915513894/5e14b3df-b147-49f4-9331-784f749b68ae"
      }
    },
    "failureReason": "Schedule encountered some errors and is incomplete. Please
take a look at error report for further details"
  }
}

```

```
}
```

Failure notification message for the DISABLED mode

The following example shows a failed scheduled query notification message for the DISABLED mode of the QueryInsights parameter.

```
"FailureNotificationMessage": {
  "type": "MANUAL_TRIGGER_FAILURE",
  "arn": "arn:aws:timestream:<Region>:<Account>:scheduled-query/sq-test-
b261670d-790c-4116-9db5-0798071b18b1-b7e27a1d79be226d",
  "scheduledQueryRunSummary": {
    "invocationEpochSecond": 1727915194,
    "triggerTimeMillis": 1727915195119,
    "runStatus": "MANUAL_TRIGGER_FAILURE",
    "executionStats": {
      "executionTimeInMillis": 10777,
      "dataWrites": 0,
      "bytesMetered": 0,
      "cumulativeBytesScanned": 0,
      "recordsIngested": 0,
      "queryResultRows": 4
    },
    "errorReportLocation": {
      "s3ReportLocation": {
        "bucketName": "amzn-s3-demo-bucket",
        "objectKey": "4my-organization-b7e27a1d79be226d/1727915194/
MANUAL/1727915195119/08dea9f5-9a0a-4e63-a5f7-ded23247bb98"
      }
    },
    "failureReason": "Schedule encountered some errors and is incomplete. Please
take a look at error report for further details"
  }
}
```

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](#)
- [AWS SDK for Ruby V3](#)

ListScheduledQueries

Service: Amazon Timestream Query

Gets a list of all scheduled queries in the caller's Amazon account and Region.

ListScheduledQueries is eventually consistent.

Request Syntax

```
{  
  "MaxResults": number,  
  "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.

MaxResults

The maximum number of items to return in the output. If the total number of items available is more than the value specified, a NextToken is provided in the output. To resume pagination, provide the NextToken value as the argument to the subsequent call to ListScheduledQueriesRequest.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 1000.

Required: No

NextToken

A pagination token to resume pagination.

Type: String

Required: No

Response Syntax

```
{
```

```

"NextToken": "string",
"ScheduledQueries": [
  {
    "Arn": "string",
    "CreationTime": number,
    "ErrorReportConfiguration": {
      "S3Configuration": {
        "BucketName": "string",
        "EncryptionOption": "string",
        "ObjectKeyPrefix": "string"
      }
    },
    "LastRunStatus": "string",
    "Name": "string",
    "NextInvocationTime": number,
    "PreviousInvocationTime": number,
    "State": "string",
    "TargetDestination": {
      "TimestreamDestination": {
        "DatabaseName": "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.

NextToken

A token to specify where to start paginating. This is the NextToken from a previously truncated response.

Type: String

ScheduledQueries

A list of scheduled queries.

Type: Array of [ScheduledQuery](#) objects

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You do not have the necessary permissions to access the account settings.

HTTP Status Code: 400

InternalServerError

An internal server error occurred while processing the request.

HTTP Status Code: 400

InvalidEndpointException

The requested endpoint is invalid.

HTTP Status Code: 400

ThrottlingException

The request was throttled due to excessive requests.

HTTP Status Code: 400

ValidationException

Invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

ListTagsForResource

Service: Amazon Timestream Query

List all tags on a Timestream query resource.

Request Syntax

```
{  
  "MaxResults": number,  
  "NextToken": "string",  
  "ResourceARN": "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.

[MaxResults](#)

The maximum number of tags to return.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 200.

Required: No

[NextToken](#)

A pagination token to resume pagination.

Type: String

Required: No

[ResourceARN](#)

The Timestream resource with tags to be listed. This value is an Amazon Resource Name (ARN).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: Yes

Response Syntax

```
{
  "NextToken": "string",
  "Tags": [
    {
      "Key": "string",
      "Value": "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 pagination token to resume pagination with a subsequent call to `ListTagsForResourceResponse`.

Type: String

Tags

The tags currently associated with the Timestream resource.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InvalidEndpointException

The requested endpoint is invalid.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found.

ScheduledQueryArn

The ARN of the scheduled query.

HTTP Status Code: 400

ThrottlingException

The request was throttled due to excessive requests.

HTTP Status Code: 400

ValidationException

Invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

PrepareQuery

Service: Amazon Timestream Query

A synchronous operation that allows you to submit a query with parameters to be stored by Timestream for later running. Timestream only supports using this operation with `ValidateOnly` set to `true`.

Request Syntax

```
{
  "QueryString": "string",
  "ValidateOnly": boolean
}
```

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.

QueryString

The Timestream query string that you want to use as a prepared statement. Parameter names can be specified in the query string @ character followed by an identifier.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 262144.

Required: Yes

ValidateOnly

By setting this value to `true`, Timestream will only validate that the query string is a valid Timestream query, and not store the prepared query for later use.

Type: Boolean

Required: No

Response Syntax

```
{
```

```
"Columns": [  
  {  
    "Aliased": boolean,  
    "DatabaseName": "string",  
    "Name": "string",  
    "TableName": "string",  
    "Type": {  
      "ArrayColumnInfo": {  
        "Name": "string",  
        "Type": "Type"  
      },  
      "RowColumnInfo": [  
        {  
          "Name": "string",  
          "Type": "Type"  
        }  
      ],  
      "ScalarType": "string",  
      "TimeSeriesMeasureValueColumnInfo": {  
        "Name": "string",  
        "Type": "Type"  
      }  
    }  
  }  
],  
"Parameters": [  
  {  
    "Name": "string",  
    "Type": {  
      "ArrayColumnInfo": {  
        "Name": "string",  
        "Type": "Type"  
      },  
      "RowColumnInfo": [  
        {  
          "Name": "string",  
          "Type": "Type"  
        }  
      ],  
      "ScalarType": "string",  
      "TimeSeriesMeasureValueColumnInfo": {  
        "Name": "string",  
        "Type": "Type"  
      }  
    }  
  }  
]
```

```
    }  
  }  
],  
"QueryString": "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.

Columns

A list of SELECT clause columns of the submitted query string.

Type: Array of [SelectColumn](#) objects

Parameters

A list of parameters used in the submitted query string.

Type: Array of [ParameterMapping](#) objects

QueryString

The query string that you want prepare.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 262144.

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You do not have the necessary permissions to access the account settings.

HTTP Status Code: 400

InternalServerError

An internal server error occurred while processing the request.

HTTP Status Code: 400

InvalidEndpointException

The requested endpoint is invalid.

HTTP Status Code: 400

ThrottlingException

The request was throttled due to excessive requests.

HTTP Status Code: 400

ValidationException

Invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

Query

Service: Amazon Timestream Query

Query is a synchronous operation that enables you to run a query against your Amazon Timestream data.

If you enabled `QueryInsights`, this API also returns insights and metrics related to the query that you executed. `QueryInsights` helps with performance tuning of your query. For more information about `QueryInsights`, see [Using query insights to optimize queries in Amazon Timestream](#).

Note

The maximum number of Query API requests you're allowed to make with `QueryInsights` enabled is 1 query per second (QPS). If you exceed this query rate, it might result in throttling.

Query will time out after 60 seconds. You must update the default timeout in the SDK to support a timeout of 60 seconds. See the [code sample](#) for details.

Your query request will fail in the following cases:

- If you submit a Query request with the same client token outside of the 5-minute idempotency window.
- If you submit a Query request with the same client token, but change other parameters, within the 5-minute idempotency window.
- If the size of the row (including the query metadata) exceeds 1 MB, then the query will fail with the following error message:

```
Query aborted as max page response size has been exceeded by the output
result row
```

- If the IAM principal of the query initiator and the result reader are not the same and/or the query initiator and the result reader do not have the same query string in the query requests, the query will fail with an `Invalid pagination token` error.

Request Syntax

```
{
  "ClientToken": "string",
  "MaxRows": number,
  "NextToken": "string",
  "QueryInsights": {
    "Mode": "string"
  },
  "QueryString": "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.

ClientToken

Unique, case-sensitive string of up to 64 ASCII characters specified when a Query request is made. Providing a ClientToken makes the call to Query *idempotent*. This means that running the same query repeatedly will produce the same result. In other words, making multiple identical Query requests has the same effect as making a single request. When using ClientToken in a query, note the following:

- If the Query API is instantiated without a ClientToken, the Query SDK generates a ClientToken on your behalf.
- If the Query invocation only contains the ClientToken but does not include a NextToken, that invocation of Query is assumed to be a new query run.
- If the invocation contains NextToken, that particular invocation is assumed to be a subsequent invocation of a prior call to the Query API, and a result set is returned.
- After 4 hours, any request with the same ClientToken is treated as a new request.

Type: String

Length Constraints: Minimum length of 32. Maximum length of 128.

Required: No

MaxRows

The total number of rows to be returned in the Query output. The initial run of Query with a MaxRows value specified will return the result set of the query in two cases:

- The size of the result is less than 1MB.
- The number of rows in the result set is less than the value of maxRows.

Otherwise, the initial invocation of Query only returns a NextToken, which can then be used in subsequent calls to fetch the result set. To resume pagination, provide the NextToken value in the subsequent command.

If the row size is large (e.g. a row has many columns), Timestream may return fewer rows to keep the response size from exceeding the 1 MB limit. If MaxRows is not provided, Timestream will send the necessary number of rows to meet the 1 MB limit.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 1000.

Required: No

NextToken

A pagination token used to return a set of results. When the Query API is invoked using NextToken, that particular invocation is assumed to be a subsequent invocation of a prior call to Query, and a result set is returned. However, if the Query invocation only contains the ClientToken, that invocation of Query is assumed to be a new query run.

Note the following when using NextToken in a query:

- A pagination token can be used for up to five Query invocations, OR for a duration of up to 1 hour – whichever comes first.
- Using the same NextToken will return the same set of records. To keep paginating through the result set, you must to use the most recent nextToken.
- Suppose a Query invocation returns two NextToken values, TokenA and TokenB. If TokenB is used in a subsequent Query invocation, then TokenA is invalidated and cannot be reused.
- To request a previous result set from a query after pagination has begun, you must re-invoke the Query API.
- The latest NextToken should be used to paginate until null is returned, at which point a new NextToken should be used.

- If the IAM principal of the query initiator and the result reader are not the same and/or the query initiator and the result reader do not have the same query string in the query requests, the query will fail with an `Invalid pagination token` error.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

QueryInsights

Encapsulates settings for enabling QueryInsights.

Enabling QueryInsights returns insights and metrics in addition to query results for the query that you executed. You can use QueryInsights to tune your query performance.

Type: [QueryInsights](#) object

Required: No

QueryString

The query to be run by Timestream.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 262144.

Required: Yes

Response Syntax

```
{
  "ColumnInfo": [
    {
      "Name": "string",
      "Type": {
        "ArrayColumnInfo": "ColumnInfo",
        "RowColumnInfo": [
          "ColumnInfo"
        ],
        "ScalarType": "string",
        "TimeSeriesMeasureValueColumnInfo": "ColumnInfo"
      }
    }
  ]
}
```

```

    }
  ],
  "NextToken": "string",
  "QueryId": "string",
  "QueryInsightsResponse": {
    "OutputBytes": number,
    "OutputRows": number,
    "QuerySpatialCoverage": {
      "Max": {
        "PartitionKey": [ "string" ],
        "TableArn": "string",
        "Value": number
      }
    }
  },
  "QueryTableCount": number,
  "QueryTemporalRange": {
    "Max": {
      "TableArn": "string",
      "Value": number
    }
  },
  "UnloadPartitionCount": number,
  "UnloadWrittenBytes": number,
  "UnloadWrittenRows": number
},
"QueryStatus": {
  "CumulativeBytesMetered": number,
  "CumulativeBytesScanned": number,
  "ProgressPercentage": number
},
"Rows": [
  {
    "Data": [
      {
        "ArrayValue": [
          "Datum"
        ],
        "NullValue": boolean,
        "RowValue": "Row",
        "ScalarValue": "string",
        "TimeSeriesValue": [
          {
            "Time": "string",
            "Value": "Datum"
          }
        ]
      }
    ]
  }
]

```

```
}
  ]
  }
  ]
  }
  ]
  }
```

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.

ColumnInfo

The column data types of the returned result set.

Type: Array of [ColumnInfo](#) objects

NextToken

A pagination token that can be used again on a Query call to get the next set of results.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

QueryId

A unique ID for the given query.

Type: String

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

Pattern: [a-zA-Z0-9]+

QueryInsightsResponse

Encapsulates QueryInsights containing insights and metrics related to the query that you executed.

Type: [QueryInsightsResponse](#) object

QueryStatus

Information about the status of the query, including progress and bytes scanned.

Type: [QueryStatus](#) object

Rows

The result set rows returned by the query.

Type: Array of [Row](#) objects

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You do not have the necessary permissions to access the account settings.

HTTP Status Code: 400

ConflictException

Unable to poll results for a cancelled query.

HTTP Status Code: 400

InternalServerErrorException

An internal server error occurred while processing the request.

HTTP Status Code: 400

InvalidEndpointException

The requested endpoint is invalid.

HTTP Status Code: 400

QueryExecutionException

Timestream was unable to run the query successfully.

HTTP Status Code: 400

ThrottlingException

The request was throttled due to excessive requests.

HTTP Status Code: 400

ValidationException

Invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

TagResource

Service: Amazon Timestream Query

Associate a set of tags with a Timestream resource. You can then activate these user-defined tags so that they appear on the Billing and Cost Management console for cost allocation tracking.

Request Syntax

```
{
  "ResourceARN": "string",
  "Tags": [
    {
      "Key": "string",
      "Value": "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.

ResourceARN

Identifies the Timestream resource to which tags should be added. This value is an Amazon Resource Name (ARN).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: Yes

Tags

The tags to be assigned to the Timestream resource.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Required: Yes

Response Elements

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

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InvalidEndpointException

The requested endpoint is invalid.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found.

ScheduledQueryArn

The ARN of the scheduled query.

HTTP Status Code: 400

ServiceQuotaExceededException

You have exceeded the service quota.

HTTP Status Code: 400

ThrottlingException

The request was throttled due to excessive requests.

HTTP Status Code: 400

ValidationException

Invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

UntagResource

Service: Amazon Timestream Query

Removes the association of tags from a Timestream query resource.

Request Syntax

```
{
  "ResourceARN": "string",
  "TagKeys": [ "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.

ResourceARN

The Timestream resource that the tags will be removed from. This value is an Amazon Resource Name (ARN).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: Yes

TagKeys

A list of tags keys. Existing tags of the resource whose keys are members of this list will be removed from the Timestream resource.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Length Constraints: Minimum length of 1. Maximum length of 128.

Required: Yes

Response Elements

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

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InvalidEndpointException

The requested endpoint is invalid.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found.

ScheduledQueryArn

The ARN of the scheduled query.

HTTP Status Code: 400

ThrottlingException

The request was throttled due to excessive requests.

HTTP Status Code: 400

ValidationException

Invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

UpdateAccountSettings

Service: Amazon Timestream Query

Transitions your account to use TCUs for query pricing and modifies the maximum query compute units that you've configured. If you reduce the value of MaxQueryTCU to a desired configuration, the new value can take up to 24 hours to be effective.

Note

After you've transitioned your account to use TCUs for query pricing, you can't transition to using bytes scanned for query pricing.

Request Syntax

```
{
  "MaxQueryTCU": number,
  "QueryCompute": {
    "ComputeMode": "string",
    "ProvisionedCapacity": {
      "NotificationConfiguration": {
        "RoleArn": "string",
        "SnsConfiguration": {
          "TopicArn": "string"
        }
      },
      "TargetQueryTCU": number
    }
  },
  "QueryPricingModel": "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.

MaxQueryTCU

The maximum number of compute units the service will use at any point in time to serve your queries. To run queries, you must set a minimum capacity of 4 TCU. You can set the maximum number of TCU in multiples of 4, for example, 4, 8, 16, 32, and so on. The maximum value supported for MaxQueryTCU is 1000. To request an increase to this soft limit, contact AWS Support. For information about the default quota for maxQueryTCU, see [Default quotas](#). This configuration is applicable only for on-demand usage of Timestream Compute Units (TCUs).

The maximum value supported for MaxQueryTCU is 1000. To request an increase to this soft limit, contact AWS Support. For information about the default quota for maxQueryTCU, see [Default quotas](#).

Type: Integer

Required: No

QueryCompute

Modifies the query compute settings configured in your account, including the query pricing model and provisioned Timestream Compute Units (TCUs) in your account. QueryCompute is available only in the Asia Pacific (Mumbai) region.

Note

This API is idempotent, meaning that making the same request multiple times will have the same effect as making the request once.

Type: [QueryComputeRequest](#) object

Required: No

QueryPricingModel

The pricing model for queries in an account.

Note

The QueryPricingModel parameter is used by several Timestream operations; however, the UpdateAccountSettings API operation doesn't recognize any values other than COMPUTE_UNITS.

Type: String

Valid Values: BYTES_SCANNED | COMPUTE_UNITS

Required: No

Response Syntax

```
{
  "MaxQueryTCU": number,
  "QueryCompute": {
    "ComputeMode": "string",
    "ProvisionedCapacity": {
      "ActiveQueryTCU": number,
      "LastUpdate": {
        "Status": "string",
        "StatusMessage": "string",
        "TargetQueryTCU": number
      },
      "NotificationConfiguration": {
        "RoleArn": "string",
        "SnsConfiguration": {
          "TopicArn": "string"
        }
      }
    }
  },
  "QueryPricingModel": "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.

MaxQueryTCU

The configured maximum number of compute units the service will use at any point in time to serve your queries.

Type: Integer

QueryCompute

Confirms the updated account settings for querying data in your account. QueryCompute is available only in the Asia Pacific (Mumbai) region.

Type: [QueryComputeResponse](#) object

QueryPricingModel

The pricing model for an account.

Type: String

Valid Values: BYTES_SCANNED | COMPUTE_UNITS

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You do not have the necessary permissions to access the account settings.

HTTP Status Code: 400

InternalServerError

An internal server error occurred while processing the request.

HTTP Status Code: 400

InvalidEndpointException

The requested endpoint is invalid.

HTTP Status Code: 400

ThrottlingException

The request was throttled due to excessive requests.

HTTP Status Code: 400

ValidationException

Invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

UpdateScheduledQuery

Service: Amazon Timestream Query

Update a scheduled query.

Request Syntax

```
{  
  "ScheduledQueryArn": "string",  
  "State": "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.

ScheduledQueryArn

ARN of the scheduled query.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: Yes

State

State of the scheduled query.

Type: String

Valid Values: ENABLED | DISABLED

Required: Yes

Response Elements

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

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

AccessDeniedException

You do not have the necessary permissions to access the account settings.

HTTP Status Code: 400

InternalServerErrorException

An internal server error occurred while processing the request.

HTTP Status Code: 400

InvalidEndpointException

The requested endpoint is invalid.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found.

ScheduledQueryArn

The ARN of the scheduled query.

HTTP Status Code: 400

ThrottlingException

The request was throttled due to excessive requests.

HTTP Status Code: 400

ValidationException

Invalid or malformed request.

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](#)
- [AWS SDK for Ruby V3](#)

Data Types

The following data types are supported by Amazon Timestream Write:

- [BatchLoadProgressReport](#)
- [BatchLoadTask](#)
- [BatchLoadTaskDescription](#)
- [CsvConfiguration](#)
- [Database](#)
- [DataModel](#)
- [DataModelConfiguration](#)
- [DataModelS3Configuration](#)
- [DataSourceConfiguration](#)
- [DataSourceS3Configuration](#)
- [Dimension](#)
- [DimensionMapping](#)
- [Endpoint](#)
- [MagneticStoreRejectedDataLocation](#)
- [MagneticStoreWriteProperties](#)
- [MeasureValue](#)
- [MixedMeasureMapping](#)
- [MultiMeasureAttributeMapping](#)
- [MultiMeasureMappings](#)
- [PartitionKey](#)
- [Record](#)
- [RecordsIngested](#)
- [RejectedRecord](#)
- [ReportConfiguration](#)
- [ReportS3Configuration](#)
- [RetentionProperties](#)
- [S3Configuration](#)

- [Schema](#)
- [Table](#)
- [Tag](#)

The following data types are supported by Amazon Timestream Query:

- [AccountSettingsNotificationConfiguration](#)
- [ColumnInfo](#)
- [Datum](#)
- [DimensionMapping](#)
- [Endpoint](#)
- [ErrorReportConfiguration](#)
- [ErrorReportLocation](#)
- [ExecutionStats](#)
- [LastUpdate](#)
- [MixedMeasureMapping](#)
- [MultiMeasureAttributeMapping](#)
- [MultiMeasureMappings](#)
- [NotificationConfiguration](#)
- [ParameterMapping](#)
- [ProvisionedCapacityRequest](#)
- [ProvisionedCapacityResponse](#)
- [QueryComputeRequest](#)
- [QueryComputeResponse](#)
- [QueryInsights](#)
- [QueryInsightsResponse](#)
- [QuerySpatialCoverage](#)
- [QuerySpatialCoverageMax](#)
- [QueryStatus](#)
- [QueryTemporalRange](#)
- [QueryTemporalRangeMax](#)

- [Row](#)
- [S3Configuration](#)
- [S3ReportLocation](#)
- [ScheduleConfiguration](#)
- [ScheduledQuery](#)
- [ScheduledQueryDescription](#)
- [ScheduledQueryInsights](#)
- [ScheduledQueryInsightsResponse](#)
- [ScheduledQueryRunSummary](#)
- [SelectColumn](#)
- [SnsConfiguration](#)
- [Tag](#)
- [TargetConfiguration](#)
- [TargetDestination](#)
- [TimeSeriesDataPoint](#)
- [TimestreamConfiguration](#)
- [TimestreamDestination](#)
- [Type](#)

Amazon Timestream Write

The following data types are supported by Amazon Timestream Write:

- [BatchLoadProgressReport](#)
- [BatchLoadTask](#)
- [BatchLoadTaskDescription](#)
- [CsvConfiguration](#)
- [Database](#)
- [DataModel](#)
- [DataModelConfiguration](#)
- [DataModelS3Configuration](#)

- [DataSourceConfiguration](#)
- [DataSourceS3Configuration](#)
- [Dimension](#)
- [DimensionMapping](#)
- [Endpoint](#)
- [MagneticStoreRejectedDataLocation](#)
- [MagneticStoreWriteProperties](#)
- [MeasureValue](#)
- [MixedMeasureMapping](#)
- [MultiMeasureAttributeMapping](#)
- [MultiMeasureMappings](#)
- [PartitionKey](#)
- [Record](#)
- [RecordsIngested](#)
- [RejectedRecord](#)
- [ReportConfiguration](#)
- [ReportS3Configuration](#)
- [RetentionProperties](#)
- [S3Configuration](#)
- [Schema](#)
- [Table](#)
- [Tag](#)

BatchLoadProgressReport

Service: Amazon Timestream Write

Details about the progress of a batch load task.

Contents

BytesMetered

Type: Long

Required: No

FileFailures

Type: Long

Required: No

ParseFailures

Type: Long

Required: No

RecordIngestionFailures

Type: Long

Required: No

RecordsIngested

Type: Long

Required: No

RecordsProcessed

Type: Long

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](#)

BatchLoadTask

Service: Amazon Timestream Write

Details about a batch load task.

Contents

CreationTime

The time when the Timestream batch load task was created.

Type: Timestamp

Required: No

DatabaseName

Database name for the database into which a batch load task loads data.

Type: String

Required: No

LastUpdatedTime

The time when the Timestream batch load task was last updated.

Type: Timestamp

Required: No

ResumableUntil

Type: Timestamp

Required: No

TableName

Table name for the table into which a batch load task loads data.

Type: String

Required: No

TaskId

The ID of the batch load task.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 32.

Pattern: [A-Z0-9]+

Required: No

TaskStatus

Status of the batch load task.

Type: String

Valid Values: CREATED | IN_PROGRESS | FAILED | SUCCEEDED | PROGRESS_STOPPED
| PENDING_RESUME

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](#)

BatchLoadTaskDescription

Service: Amazon Timestream Write

Details about a batch load task.

Contents

CreationTime

The time when the Timestream batch load task was created.

Type: Timestamp

Required: No

DataModelConfiguration

Data model configuration for a batch load task. This contains details about where a data model for a batch load task is stored.

Type: [DataModelConfiguration](#) object

Required: No

DataSourceConfiguration

Configuration details about the data source for a batch load task.

Type: [DataSourceConfiguration](#) object

Required: No

ErrorMessage

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

LastUpdatedTime

The time when the Timestream batch load task was last updated.

Type: Timestamp

Required: No

ProgressReport

Type: [BatchLoadProgressReport](#) object

Required: No

RecordVersion

Type: Long

Required: No

ReportConfiguration

Report configuration for a batch load task. This contains details about where error reports are stored.

Type: [ReportConfiguration](#) object

Required: No

ResumableUntil

Type: Timestamp

Required: No

TargetDatabaseName

Type: String

Required: No

TargetTableName

Type: String

Required: No

TaskId

The ID of the batch load task.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 32.

Pattern: [A-Z0-9]+

Required: No

TaskStatus

Status of the batch load task.

Type: String

Valid Values: CREATED | IN_PROGRESS | FAILED | SUCCEEDED | PROGRESS_STOPPED
| PENDING_RESUME

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](#)

CsvConfiguration

Service: Amazon Timestream Write

A delimited data format where the column separator can be a comma and the record separator is a newline character.

Contents

ColumnSeparator

Column separator can be one of comma (','), pipe ('|'), semicolon (';'), tab('/t'), or blank space (' ').

Type: String

Length Constraints: Fixed length of 1.

Required: No

EscapeChar

Escape character can be one of

Type: String

Length Constraints: Fixed length of 1.

Required: No

NullValue

Can be blank space (' ').

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Required: No

QuoteChar

Can be single quote (') or double quote (").

Type: String

Length Constraints: Fixed length of 1.

Required: No

TrimWhiteSpace

Specifies to trim leading and trailing white space.

Type: Boolean

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](#)

Database

Service: Amazon Timestream Write

A top-level container for a table. Databases and tables are the fundamental management concepts in Amazon Timestream. All tables in a database are encrypted with the same AWS KMS key.

Contents

Arn

The Amazon Resource Name that uniquely identifies this database.

Type: String

Required: No

CreationTime

The time when the database was created, calculated from the Unix epoch time.

Type: Timestamp

Required: No

DatabaseName

The name of the Timestream database.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Required: No

KmsKeyId

The identifier of the AWS KMS key used to encrypt the data stored in the database.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

LastUpdatedTime

The last time that this database was updated.

Type: Timestamp

Required: No

TableCount

The total number of tables found within a Timestream database.

Type: Long

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](#)

DataModel

Service: Amazon Timestream Write

Data model for a batch load task.

Contents

DimensionMappings

Source to target mappings for dimensions.

Type: Array of [DimensionMapping](#) objects

Array Members: Minimum number of 1 item.

Required: Yes

MeasureNameColumn

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Required: No

MixedMeasureMappings

Source to target mappings for measures.

Type: Array of [MixedMeasureMapping](#) objects

Array Members: Minimum number of 1 item.

Required: No

MultiMeasureMappings

Source to target mappings for multi-measure records.

Type: [MultiMeasureMappings](#) object

Required: No

TimeColumn

Source column to be mapped to time.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Required: No

TimeUnit

The granularity of the timestamp unit. It indicates if the time value is in seconds, milliseconds, nanoseconds, or other supported values. Default is `MILLISECONDS`.

Type: String

Valid Values: `MILLISECONDS` | `SECONDS` | `MICROSECONDS` | `NANOSECONDS`

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](#)

DataModelConfiguration

Service: Amazon Timestream Write

Contents

DataModel

Type: [DataModel](#) object

Required: No

DataModelS3Configuration

Type: [DataModelS3Configuration](#) object

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](#)

DataModelS3Configuration

Service: Amazon Timestream Write

Contents

BucketName

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Pattern: `[a-z0-9][\.\-a-z0-9]{1,61}[a-z0-9]`

Required: No

ObjectKey

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `[a-zA-Z0-9|!\\-_*'\\(\\)]([a-zA-Z0-9]|[!\\-_*'\\(\\)\\/.])+`

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](#)

DataSourceConfiguration

Service: Amazon Timestream Write

Defines configuration details about the data source.

Contents

DataFormat

This is currently CSV.

Type: String

Valid Values: CSV

Required: Yes

DataSourceS3Configuration

Configuration of an S3 location for a file which contains data to load.

Type: [DataSourceS3Configuration](#) object

Required: Yes

CsvConfiguration

A delimited data format where the column separator can be a comma and the record separator is a newline character.

Type: [CsvConfiguration](#) object

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](#)

DataSourceS3Configuration

Service: Amazon Timestream Write

Contents

BucketName

The bucket name of the customer S3 bucket.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Pattern: `[a-z0-9][\.\-a-z0-9]{1,61}[a-z0-9]`

Required: Yes

ObjectKeyPrefix

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `[a-zA-Z0-9|!\\-_*'\\(\\)]([a-zA-Z0-9]|[!\\-_*'\\(\\)\\/\\.])+`

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](#)

Dimension

Service: Amazon Timestream Write

Represents the metadata attributes of the time series. For example, the name and Availability Zone of an EC2 instance or the name of the manufacturer of a wind turbine are dimensions.

Contents

Name

Dimension represents the metadata attributes of the time series. For example, the name and Availability Zone of an EC2 instance or the name of the manufacturer of a wind turbine are dimensions.

For constraints on dimension names, see [Naming Constraints](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 60.

Required: Yes

Value

The value of the dimension.

Type: String

Required: Yes

DimensionValueType

The data type of the dimension for the time-series data point.

Type: String

Valid Values: VARCHAR

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](#)

DimensionMapping

Service: Amazon Timestream Write

Contents

DestinationColumn

Type: String

Length Constraints: Minimum length of 1.

Required: No

SourceColumn

Type: String

Length Constraints: Minimum length of 1.

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](#)

Endpoint

Service: Amazon Timestream Write

Represents an available endpoint against which to make API calls against, as well as the TTL for that endpoint.

Contents

Address

An endpoint address.

Type: String

Required: Yes

CachePeriodInMinutes

The TTL for the endpoint, in minutes.

Type: Long

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](#)

MagneticStoreRejectedDataLocation

Service: Amazon Timestream Write

The location to write error reports for records rejected, asynchronously, during magnetic store writes.

Contents

S3Configuration

Configuration of an S3 location to write error reports for records rejected, asynchronously, during magnetic store writes.

Type: [S3Configuration](#) object

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](#)

MagneticStoreWriteProperties

Service: Amazon Timestream Write

The set of properties on a table for configuring magnetic store writes.

Contents

EnableMagneticStoreWrites

A flag to enable magnetic store writes.

Type: Boolean

Required: Yes

MagneticStoreRejectedDataLocation

The location to write error reports for records rejected asynchronously during magnetic store writes.

Type: [MagneticStoreRejectedDataLocation](#) object

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](#)

MeasureValue

Service: Amazon Timestream Write

Represents the data attribute of the time series. For example, the CPU utilization of an EC2 instance or the RPM of a wind turbine are measures. MeasureValue has both name and value.

MeasureValue is only allowed for type MULTI. Using MULTI type, you can pass multiple data attributes associated with the same time series in a single record

Contents

Name

The name of the MeasureValue.

For constraints on MeasureValue names, see [Naming Constraints](#) in the Amazon Timestream Developer Guide.

Type: String

Length Constraints: Minimum length of 1.

Required: Yes

Type

Contains the data type of the MeasureValue for the time-series data point.

Type: String

Valid Values: DOUBLE | BIGINT | VARCHAR | BOOLEAN | TIMESTAMP | MULTI

Required: Yes

Value

The value for the MeasureValue. For information, see [Data types](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

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](#)

MixedMeasureMapping

Service: Amazon Timestream Write

Contents

MeasureValueType

Type: String

Valid Values: DOUBLE | BIGINT | VARCHAR | BOOLEAN | TIMESTAMP | MULTI

Required: Yes

MeasureName

Type: String

Length Constraints: Minimum length of 1.

Required: No

MultiMeasureAttributeMappings

Type: Array of [MultiMeasureAttributeMapping](#) objects

Array Members: Minimum number of 1 item.

Required: No

SourceColumn

Type: String

Length Constraints: Minimum length of 1.

Required: No

TargetMeasureName

Type: String

Length Constraints: Minimum length of 1.

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](#)

MultiMeasureAttributeMapping

Service: Amazon Timestream Write

Contents

SourceColumn

Type: String

Length Constraints: Minimum length of 1.

Required: Yes

MeasureValueType

Type: String

Valid Values: DOUBLE | BIGINT | BOOLEAN | VARCHAR | TIMESTAMP

Required: No

TargetMultiMeasureAttributeName

Type: String

Length Constraints: Minimum length of 1.

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](#)

MultiMeasureMappings

Service: Amazon Timestream Write

Contents

MultiMeasureAttributeMappings

Type: Array of [MultiMeasureAttributeMapping](#) objects

Array Members: Minimum number of 1 item.

Required: Yes

TargetMultiMeasureName

Type: String

Length Constraints: Minimum length of 1.

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](#)

PartitionKey

Service: Amazon Timestream Write

An attribute used in partitioning data in a table. A dimension key partitions data using the values of the dimension specified by the dimension-name as partition key, while a measure key partitions data using measure names (values of the 'measure_name' column).

Contents

Type

The type of the partition key. Options are DIMENSION (dimension key) and MEASURE (measure key).

Type: String

Valid Values: DIMENSION | MEASURE

Required: Yes

EnforcementInRecord

The level of enforcement for the specification of a dimension key in ingested records. Options are REQUIRED (dimension key must be specified) and OPTIONAL (dimension key does not have to be specified).

Type: String

Valid Values: REQUIRED | OPTIONAL

Required: No

Name

The name of the attribute used for a dimension key.

Type: String

Length Constraints: Minimum length of 1.

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](#)

Record

Service: Amazon Timestream Write

Represents a time-series data point being written into Timestream. Each record contains an array of dimensions. Dimensions represent the metadata attributes of a time-series data point, such as the instance name or Availability Zone of an EC2 instance. A record also contains the measure name, which is the name of the measure being collected (for example, the CPU utilization of an EC2 instance). Additionally, a record contains the measure value and the value type, which is the data type of the measure value. Also, the record contains the timestamp of when the measure was collected and the timestamp unit, which represents the granularity of the timestamp.

Records have a `Version` field, which is a 64-bit long that you can use for updating data points. Writes of a duplicate record with the same dimension, timestamp, and measure name but different measure value will only succeed if the `Version` attribute of the record in the write request is higher than that of the existing record. Timestream defaults to a `Version` of 1 for records without the `Version` field.

Contents

Dimensions

Contains the list of dimensions for time-series data points.

Type: Array of [Dimension](#) objects

Array Members: Maximum number of 128 items.

Required: No

MeasureName

Measure represents the data attribute of the time series. For example, the CPU utilization of an EC2 instance or the RPM of a wind turbine are measures.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Required: No

MeasureValue

Contains the measure value for the time-series data point.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

MeasureValues

Contains the list of MeasureValue for time-series data points.

This is only allowed for type MULTI. For scalar values, use MeasureValue attribute of the record directly.

Type: Array of [MeasureValue](#) objects

Required: No

MeasureValueType

Contains the data type of the measure value for the time-series data point. Default type is DOUBLE. For more information, see [Data types](#).

Type: String

Valid Values: DOUBLE | BIGINT | VARCHAR | BOOLEAN | TIMESTAMP | MULTI

Required: No

Time

Contains the time at which the measure value for the data point was collected. The time value plus the unit provides the time elapsed since the epoch. For example, if the time value is 12345 and the unit is ms, then 12345 ms have elapsed since the epoch.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Required: No

TimeUnit

The granularity of the timestamp unit. It indicates if the time value is in seconds, milliseconds, nanoseconds, or other supported values. Default is MILLISECONDS.

Type: String

Valid Values: MILLISECONDS | SECONDS | MICROSECONDS | NANoseconds

Required: No

Version

64-bit attribute used for record updates. Write requests for duplicate data with a higher version number will update the existing measure value and version. In cases where the measure value is the same, Version will still be updated. Default value is 1.

Note

Version must be 1 or greater, or you will receive a ValidationException error.

Type: Long

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](#)

RecordsIngested

Service: Amazon Timestream Write

Information on the records ingested by this request.

Contents

MagneticStore

Count of records ingested into the magnetic store.

Type: Integer

Required: No

MemoryStore

Count of records ingested into the memory store.

Type: Integer

Required: No

Total

Total count of successfully ingested records.

Type: Integer

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](#)

RejectedRecord

Service: Amazon Timestream Write

Represents records that were not successfully inserted into Timestream due to data validation issues that must be resolved before reinserting time-series data into the system.

Contents

ExistingVersion

The existing version of the record. This value is populated in scenarios where an identical record exists with a higher version than the version in the write request.

Type: Long

Required: No

Reason

The reason why a record was not successfully inserted into Timestream. Possible causes of failure include:

- Records with duplicate data where there are multiple records with the same dimensions, timestamps, and measure names but:
 - Measure values are different
 - Version is not present in the request, *or* the value of version in the new record is equal to or lower than the existing value

If Timestream rejects data for this case, the `ExistingVersion` field in the `RejectedRecords` response will indicate the current record's version. To force an update, you can resend the request with a version for the record set to a value greater than the `ExistingVersion`.

- Records with timestamps that lie outside the retention duration of the memory store.

Note

When the retention window is updated, you will receive a `RejectedRecords` exception if you immediately try to ingest data within the new window. To avoid a `RejectedRecords` exception, wait until the duration of the new window to ingest

new data. For further information, see [Best Practices for Configuring Timestream](#) and [the explanation of how storage works in Timestream](#).

- Records with dimensions or measures that exceed the Timestream defined limits.

For more information, see [Access Management](#) in the Timestream Developer Guide.

Type: String

Required: No

RecordIndex

The index of the record in the input request for WriteRecords. Indexes begin with 0.

Type: Integer

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](#)

ReportConfiguration

Service: Amazon Timestream Write

Report configuration for a batch load task. This contains details about where error reports are stored.

Contents

ReportS3Configuration

Configuration of an S3 location to write error reports and events for a batch load.

Type: [ReportS3Configuration](#) object

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](#)

ReportS3Configuration

Service: Amazon Timestream Write

Contents

BucketName

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Pattern: `[a-z0-9][\.\-a-z0-9]{1,61}[a-z0-9]`

Required: Yes

EncryptionOption

Type: String

Valid Values: `SSE_S3` | `SSE_KMS`

Required: No

KmsKeyId

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

ObjectKeyPrefix

Type: String

Length Constraints: Minimum length of 1. Maximum length of 928.

Pattern: `[a-zA-Z0-9|!\\-_*'\\(\\)]([a-zA-Z0-9]|[!\\-_*'\\(\\)\\/.])+`

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](#)

RetentionProperties

Service: Amazon Timestream Write

Retention properties contain the duration for which your time-series data must be stored in the magnetic store and the memory store.

Contents

MagneticStoreRetentionPeriodInDays

The duration for which data must be stored in the magnetic store.

Type: Long

Valid Range: Minimum value of 1. Maximum value of 73000.

Required: Yes

MemoryStoreRetentionPeriodInHours

The duration for which data must be stored in the memory store.

Type: Long

Valid Range: Minimum value of 1. Maximum value of 8766.

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](#)

S3Configuration

Service: Amazon Timestream Write

The configuration that specifies an S3 location.

Contents

BucketName

The bucket name of the customer S3 bucket.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Pattern: `[a-z0-9][\.\-a-z0-9]{1,61}[a-z0-9]`

Required: No

EncryptionOption

The encryption option for the customer S3 location. Options are S3 server-side encryption with an S3 managed key or AWS managed key.

Type: String

Valid Values: `SSE_S3` | `SSE_KMS`

Required: No

KmsKeyId

The AWS KMS key ID for the customer S3 location when encrypting with an AWS managed key.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

ObjectKeyPrefix

The object key preview for the customer S3 location.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 928.

Pattern: `[a-zA-Z0-9|!_*'\\(\\)]([a-zA-Z0-9|!_*'\\(\\)\\/\\.])+`

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](#)

Schema

Service: Amazon Timestream Write

A Schema specifies the expected data model of the table.

Contents

CompositePartitionKey

A non-empty list of partition keys defining the attributes used to partition the table data. The order of the list determines the partition hierarchy. The name and type of each partition key as well as the partition key order cannot be changed after the table is created. However, the enforcement level of each partition key can be changed.

Type: Array of [PartitionKey](#) objects

Array Members: Minimum number of 1 item.

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](#)

Table

Service: Amazon Timestream Write

Represents a database table in Timestream. Tables contain one or more related time series. You can modify the retention duration of the memory store and the magnetic store for a table.

Contents

Arn

The Amazon Resource Name that uniquely identifies this table.

Type: String

Required: No

CreationTime

The time when the Timestream table was created.

Type: Timestamp

Required: No

DatabaseName

The name of the Timestream database that contains this table.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Required: No

LastUpdatedTime

The time when the Timestream table was last updated.

Type: Timestamp

Required: No

MagneticStoreWriteProperties

Contains properties to set on the table when enabling magnetic store writes.

Type: [MagneticStoreWriteProperties](#) object

Required: No

RetentionProperties

The retention duration for the memory store and magnetic store.

Type: [RetentionProperties](#) object

Required: No

Schema

The schema of the table.

Type: [Schema](#) object

Required: No

TableName

The name of the Timestream table.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 256.

Required: No

TableStatus

The current state of the table:

- DELETING - The table is being deleted.
- ACTIVE - The table is ready for use.

Type: String

Valid Values: ACTIVE | DELETING | RESTORING

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](#)

Tag

Service: Amazon Timestream Write

A tag is a label that you assign to a Timestream database and/or table. Each tag consists of a key and an optional value, both of which you define. With tags, you can categorize databases and/or tables, for example, by purpose, owner, or environment.

Contents

Key

The key of the tag. Tag keys are case sensitive.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 128.

Required: Yes

Value

The value of the tag. Tag values are case-sensitive and can be null.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

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](#)

Amazon Timestream Query

The following data types are supported by Amazon Timestream Query:

- [AccountSettingsNotificationConfiguration](#)
- [ColumnInfo](#)
- [Datum](#)
- [DimensionMapping](#)
- [Endpoint](#)
- [ErrorReportConfiguration](#)
- [ErrorReportLocation](#)
- [ExecutionStats](#)
- [LastUpdate](#)
- [MixedMeasureMapping](#)
- [MultiMeasureAttributeMapping](#)
- [MultiMeasureMappings](#)
- [NotificationConfiguration](#)
- [ParameterMapping](#)
- [ProvisionedCapacityRequest](#)
- [ProvisionedCapacityResponse](#)
- [QueryComputeRequest](#)
- [QueryComputeResponse](#)
- [QueryInsights](#)
- [QueryInsightsResponse](#)
- [QuerySpatialCoverage](#)
- [QuerySpatialCoverageMax](#)
- [QueryStatus](#)
- [QueryTemporalRange](#)
- [QueryTemporalRangeMax](#)
- [Row](#)
- [S3Configuration](#)
- [S3ReportLocation](#)
- [ScheduleConfiguration](#)
- [ScheduledQuery](#)

- [ScheduledQueryDescription](#)
- [ScheduledQueryInsights](#)
- [ScheduledQueryInsightsResponse](#)
- [ScheduledQueryRunSummary](#)
- [SelectColumn](#)
- [SnsConfiguration](#)
- [Tag](#)
- [TargetConfiguration](#)
- [TargetDestination](#)
- [TimeSeriesDataPoint](#)
- [TimestreamConfiguration](#)
- [TimestreamDestination](#)
- [Type](#)

AccountSettingsNotificationConfiguration

Service: Amazon Timestream Query

Configuration settings for notifications related to account settings.

Contents

RoleArn

An Amazon Resource Name (ARN) that grants Timestream permission to publish notifications. This field is only visible if SNS Topic is provided when updating the account settings.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: Yes

SnsConfiguration

Details on SNS that are required to send the notification.

Type: [SnsConfiguration](#) object

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](#)

ColumnInfo

Service: Amazon Timestream Query

Contains the metadata for query results such as the column names, data types, and other attributes.

Contents

Type

The data type of the result set column. The data type can be a scalar or complex. Scalar data types are integers, strings, doubles, Booleans, and others. Complex data types are types such as arrays, rows, and others.

Type: [Type](#) object

Required: Yes

Name

The name of the result set column. The name of the result set is available for columns of all data types except for arrays.

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](#)

Datum

Service: Amazon Timestream Query

Datum represents a single data point in a query result.

Contents

ArrayValue

Indicates if the data point is an array.

Type: Array of [Datum](#) objects

Required: No

NullValue

Indicates if the data point is null.

Type: Boolean

Required: No

RowValue

Indicates if the data point is a row.

Type: [Row](#) object

Required: No

ScalarValue

Indicates if the data point is a scalar value such as integer, string, double, or Boolean.

Type: String

Required: No

TimeSeriesValue

Indicates if the data point is a timeseries data type.

Type: Array of [TimeSeriesDataPoint](#) objects

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](#)

DimensionMapping

Service: Amazon Timestream Query

This type is used to map column(s) from the query result to a dimension in the destination table.

Contents

DimensionValueType

Type for the dimension.

Type: String

Valid Values: VARCHAR

Required: Yes

Name

Column name from query result.

Type: String

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](#)

Endpoint

Service: Amazon Timestream Query

Represents an available endpoint against which to make API calls against, as well as the TTL for that endpoint.

Contents

Address

An endpoint address.

Type: String

Required: Yes

CachePeriodInMinutes

The TTL for the endpoint, in minutes.

Type: Long

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](#)

ErrorReportConfiguration

Service: Amazon Timestream Query

Configuration required for error reporting.

Contents

S3Configuration

The S3 configuration for the error reports.

Type: [S3Configuration](#) object

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](#)

ErrorReportLocation

Service: Amazon Timestream Query

This contains the location of the error report for a single scheduled query call.

Contents

S3ReportLocation

The S3 location where error reports are written.

Type: [S3ReportLocation](#) object

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](#)

ExecutionStats

Service: Amazon Timestream Query

Statistics for a single scheduled query run.

Contents

BytesMetered

Bytes metered for a single scheduled query run.

Type: Long

Required: No

CumulativeBytesScanned

Bytes scanned for a single scheduled query run.

Type: Long

Required: No

DataWrites

Data writes metered for records ingested in a single scheduled query run.

Type: Long

Required: No

ExecutionTimeInMillis

Total time, measured in milliseconds, that was needed for the scheduled query run to complete.

Type: Long

Required: No

QueryResultRows

Number of rows present in the output from running a query before ingestion to destination data source.

Type: Long

Required: No

RecordsIngested

The number of records ingested for a single scheduled query run.

Type: Long

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](#)

LastUpdate

Service: Amazon Timestream Query

Configuration object that contains the most recent account settings update, visible only if settings have been updated previously.

Contents

Status

The status of the last update. Can be either PENDING, FAILED, or SUCCEEDED.

Type: String

Valid Values: PENDING | FAILED | SUCCEEDED

Required: No

StatusMessage

Error message describing the last account settings update status, visible only if an error occurred.

Type: String

Required: No

TargetQueryTCU

The number of TimeStream Compute Units (TCUs) requested in the last account settings update.

Type: Integer

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](#)

MixedMeasureMapping

Service: Amazon Timestream Query

MixedMeasureMappings are mappings that can be used to ingest data into a mixture of narrow and multi measures in the derived table.

Contents

MeasureValueType

Type of the value that is to be read from sourceColumn. If the mapping is for MULTI, use MeasureValueType.MULTI.

Type: String

Valid Values: BIGINT | BOOLEAN | DOUBLE | VARCHAR | MULTI

Required: Yes

MeasureName

Refers to the value of measure_name in a result row. This field is required if MeasureNameColumn is provided.

Type: String

Required: No

MultiMeasureAttributeMappings

Required when measureValueType is MULTI. Attribute mappings for MULTI value measures.

Type: Array of [MultiMeasureAttributeMapping](#) objects

Array Members: Minimum number of 1 item.

Required: No

SourceColumn

This field refers to the source column from which measure-value is to be read for result materialization.

Type: String

Required: No

TargetMeasureName

Target measure name to be used. If not provided, the target measure name by default would be measure-name if provided, or sourceColumn otherwise.

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](#)

MultiMeasureAttributeMapping

Service: Amazon Timestream Query

Attribute mapping for MULTI value measures.

Contents

MeasureValueType

Type of the attribute to be read from the source column.

Type: String

Valid Values: BIGINT | BOOLEAN | DOUBLE | VARCHAR | TIMESTAMP

Required: Yes

SourceColumn

Source column from where the attribute value is to be read.

Type: String

Required: Yes

TargetMultiMeasureAttributeName

Custom name to be used for attribute name in derived table. If not provided, source column name would be used.

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](#)

MultiMeasureMappings

Service: Amazon Timestream Query

Only one of MixedMeasureMappings or MultiMeasureMappings is to be provided. MultiMeasureMappings can be used to ingest data as multi measures in the derived table.

Contents

MultiMeasureAttributeMappings

Required. Attribute mappings to be used for mapping query results to ingest data for multi-measure attributes.

Type: Array of [MultiMeasureAttributeMapping](#) objects

Array Members: Minimum number of 1 item.

Required: Yes

TargetMultiMeasureName

The name of the target multi-measure name in the derived table. This input is required when measureNameColumn is not provided. If MeasureNameColumn is provided, then value from that column will be used as multi-measure name.

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](#)

NotificationConfiguration

Service: Amazon Timestream Query

Notification configuration for a scheduled query. A notification is sent by Timestream when a scheduled query is created, its state is updated or when it is deleted.

Contents

SnsConfiguration

Details about the Amazon Simple Notification Service (SNS) configuration. This field is visible only when SNS Topic is provided when updating the account settings.

Type: [SnsConfiguration](#) object

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](#)

ParameterMapping

Service: Amazon Timestream Query

Mapping for named parameters.

Contents

Name

Parameter name.

Type: String

Required: Yes

Type

Contains the data type of a column in a query result set. The data type can be scalar or complex. The supported scalar data types are integers, Boolean, string, double, timestamp, date, time, and intervals. The supported complex data types are arrays, rows, and timeseries.

Type: [Type](#) object

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](#)

ProvisionedCapacityRequest

Service: Amazon Timestream Query

A request to update the provisioned capacity settings for querying data.

Contents

TargetQueryTCU

The target compute capacity for querying data, specified in Timestream Compute Units (TCUs).

Type: Integer

Required: Yes

NotificationConfiguration

Configuration settings for notifications related to the provisioned capacity update.

Type: [AccountSettingsNotificationConfiguration](#) object

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](#)

ProvisionedCapacityResponse

Service: Amazon Timestream Query

The response to a request to update the provisioned capacity settings for querying data.

Contents

ActiveQueryTCU

The number of Timestream Compute Units (TCUs) provisioned in the account. This field is only visible when the compute mode is PROVISIONED.

Type: Integer

Required: No

LastUpdate

Information about the last update to the provisioned capacity settings.

Type: [LastUpdate](#) object

Required: No

NotificationConfiguration

An object that contains settings for notifications that are sent whenever the provisioned capacity settings are modified. This field is only visible when the compute mode is PROVISIONED.

Type: [AccountSettingsNotificationConfiguration](#) object

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](#)

QueryComputeRequest

Service: Amazon Timestream Query

A request to retrieve or update the compute capacity settings for querying data. QueryCompute is available only in the Asia Pacific (Mumbai) region.

Contents

ComputeMode

The mode in which Timestream Compute Units (TCUs) are allocated and utilized within an account. Note that in the Asia Pacific (Mumbai) region, the API operation only recognizes the value PROVISIONED. QueryCompute is available only in the Asia Pacific (Mumbai) region.

Type: String

Valid Values: ON_DEMAND | PROVISIONED

Required: No

ProvisionedCapacity

Configuration object that contains settings for provisioned Timestream Compute Units (TCUs) in your account. QueryCompute is available only in the Asia Pacific (Mumbai) region.

Type: [ProvisionedCapacityRequest](#) object

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](#)

QueryComputeResponse

Service: Amazon Timestream Query

The response to a request to retrieve or update the compute capacity settings for querying data. QueryCompute is available only in the Asia Pacific (Mumbai) region.

Contents

ComputeMode

The mode in which Timestream Compute Units (TCUs) are allocated and utilized within an account. Note that in the Asia Pacific (Mumbai) region, the API operation only recognizes the value PROVISIONED. QueryCompute is available only in the Asia Pacific (Mumbai) region.

Type: String

Valid Values: ON_DEMAND | PROVISIONED

Required: No

ProvisionedCapacity

Configuration object that contains settings for provisioned Timestream Compute Units (TCUs) in your account. QueryCompute is available only in the Asia Pacific (Mumbai) region.

Type: [ProvisionedCapacityResponse](#) object

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](#)

QueryInsights

Service: Amazon Timestream Query

QueryInsights is a performance tuning feature that helps you optimize your queries, reducing costs and improving performance. With QueryInsights, you can assess the pruning efficiency of your queries and identify areas for improvement to enhance query performance. With QueryInsights, you can also analyze the effectiveness of your queries in terms of temporal and spatial pruning, and identify opportunities to improve performance. Specifically, you can evaluate how well your queries use time-based and partition key-based indexing strategies to optimize data retrieval. To optimize query performance, it's essential that you fine-tune both the temporal and spatial parameters that govern query execution.

The key metrics provided by QueryInsights are QuerySpatialCoverage and QueryTemporalRange. QuerySpatialCoverage indicates how much of the spatial axis the query scans, with lower values being more efficient. QueryTemporalRange shows the time range scanned, with narrower ranges being more performant.

Benefits of QueryInsights

The following are the key benefits of using QueryInsights:

- **Identifying inefficient queries** – QueryInsights provides information on the time-based and attribute-based pruning of the tables accessed by the query. This information helps you identify the tables that are sub-optimally accessed.
- **Optimizing your data model and partitioning** – You can use the QueryInsights information to access and fine-tune your data model and partitioning strategy.
- **Tuning queries** – QueryInsights highlights opportunities to use indexes more effectively.

Note

The maximum number of Query API requests you're allowed to make with QueryInsights enabled is 1 query per second (QPS). If you exceed this query rate, it might result in throttling.

Contents

Mode

Provides the following modes to enable QueryInsights:

- `ENABLED_WITH_RATE_CONTROL` – Enables QueryInsights for the queries being processed. This mode also includes a rate control mechanism, which limits the QueryInsights feature to 1 query per second (QPS).
- `DISABLED` – Disables QueryInsights.

Type: String

Valid Values: `ENABLED_WITH_RATE_CONTROL` | `DISABLED`

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](#)

QueryInsightsResponse

Service: Amazon Timestream Query

Provides various insights and metrics related to the query that you executed.

Contents

OutputBytes

Indicates the size of query result set in bytes. You can use this data to validate if the result set has changed as part of the query tuning exercise.

Type: Long

Required: No

OutputRows

Indicates the total number of rows returned as part of the query result set. You can use this data to validate if the number of rows in the result set have changed as part of the query tuning exercise.

Type: Long

Required: No

QuerySpatialCoverage

Provides insights into the spatial coverage of the query, including the table with sub-optimal (max) spatial pruning. This information can help you identify areas for improvement in your partitioning strategy to enhance spatial pruning.

Type: [QuerySpatialCoverage](#) object

Required: No

QueryTableCount

Indicates the number of tables in the query.

Type: Long

Required: No

QueryTemporalRange

Provides insights into the temporal range of the query, including the table with the largest (max) time range. Following are some of the potential options for optimizing time-based pruning:

- Add missing time-predicates.
- Remove functions around the time predicates.
- Add time predicates to all the sub-queries.

Type: [QueryTemporalRange](#) object

Required: No

UnloadPartitionCount

Indicates the partitions created by the Unload operation.

Type: Long

Required: No

UnloadWrittenBytes

Indicates the size, in bytes, written by the Unload operation.

Type: Long

Required: No

UnloadWrittenRows

Indicates the rows written by the Unload query.

Type: Long

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](#)

QuerySpatialCoverage

Service: Amazon Timestream Query

Provides insights into the spatial coverage of the query, including the table with sub-optimal (max) spatial pruning. This information can help you identify areas for improvement in your partitioning strategy to enhance spatial pruning

For example, you can do the following with the QuerySpatialCoverage information:

- Add `measure_name` or use [customer-defined partition key](#) (CDPK) predicates.
- If you've already done the preceding action, remove functions around them or clauses, such as `LIKE`.

Contents

Max

Provides insights into the spatial coverage of the executed query and the table with the most inefficient spatial pruning.

- `Value` – The maximum ratio of spatial coverage.
- `TableArn` – The Amazon Resource Name (ARN) of the table with sub-optimal spatial pruning.
- `PartitionKey` – The partition key used for partitioning, which can be a default `measure_name` or a CDPK.

Type: [QuerySpatialCoverageMax](#) object

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](#)

QuerySpatialCoverageMax

Service: Amazon Timestream Query

Provides insights into the table with the most sub-optimal spatial range scanned by your query.

Contents

PartitionKey

The partition key used for partitioning, which can be a default `measure_name` or a [customer defined partition key](#).

Type: Array of strings

Required: No

TableArn

The Amazon Resource Name (ARN) of the table with the most sub-optimal spatial pruning.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

Value

The maximum ratio of spatial coverage.

Type: Double

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](#)

QueryStatus

Service: Amazon Timestream Query

Information about the status of the query, including progress and bytes scanned.

Contents

CumulativeBytesMetered

The amount of data scanned by the query in bytes that you will be charged for. This is a cumulative sum and represents the total amount of data that you will be charged for since the query was started. The charge is applied only once and is either applied when the query completes running or when the query is cancelled.

Type: Long

Required: No

CumulativeBytesScanned

The amount of data scanned by the query in bytes. This is a cumulative sum and represents the total amount of bytes scanned since the query was started.

Type: Long

Required: No

ProgressPercentage

The progress of the query, expressed as a percentage.

Type: Double

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](#)

QueryTemporalRange

Service: Amazon Timestream Query

Provides insights into the temporal range of the query, including the table with the largest (max) time range.

Contents

Max

Encapsulates the following properties that provide insights into the most sub-optimal performing table on the temporal axis:

- `Value` – The maximum duration in nanoseconds between the start and end of the query.
- `TableArn` – The Amazon Resource Name (ARN) of the table which is queried with the largest time range.

Type: [QueryTemporalRangeMax](#) object

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](#)

QueryTemporalRangeMax

Service: Amazon Timestream Query

Provides insights into the table with the most sub-optimal temporal pruning scanned by your query.

Contents

TableArn

The Amazon Resource Name (ARN) of the table which is queried with the largest time range.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

Value

The maximum duration in nanoseconds between the start and end of the query.

Type: Long

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](#)

Row

Service: Amazon Timestream Query

Represents a single row in the query results.

Contents

Data

List of data points in a single row of the result set.

Type: Array of [Datum](#) objects

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](#)

S3Configuration

Service: Amazon Timestream Query

Details on S3 location for error reports that result from running a query.

Contents

BucketName

Name of the S3 bucket under which error reports will be created.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Pattern: `[a-z0-9][\.\-a-z0-9]{1,61}[a-z0-9]`

Required: Yes

EncryptionOption

Encryption at rest options for the error reports. If no encryption option is specified, Timestream will choose SSE_S3 as default.

Type: String

Valid Values: SSE_S3 | SSE_KMS

Required: No

ObjectKeyPrefix

Prefix for the error report key. Timestream by default adds the following prefix to the error report path.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 896.

Pattern: `[a-zA-Z0-9|!\\-_*'\\(\\)]([a-zA-Z0-9]|[!\\-_*'\\(\\)\\./])+`

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](#)

S3ReportLocation

Service: Amazon Timestream Query

S3 report location for the scheduled query run.

Contents

BucketName

S3 bucket name.

Type: String

Length Constraints: Minimum length of 3. Maximum length of 63.

Pattern: `[a-z0-9][\.\-a-z0-9]{1,61}[a-z0-9]`

Required: No

ObjectKey

S3 key.

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](#)

ScheduleConfiguration

Service: Amazon Timestream Query

Configuration of the schedule of the query.

Contents

ScheduleExpression

An expression that denotes when to trigger the scheduled query run. This can be a cron expression or a rate expression.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

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](#)

ScheduledQuery

Service: Amazon Timestream Query

Scheduled Query

Contents

Arn

The Amazon Resource Name.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: Yes

Name

The name of the scheduled query.

Type: String

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

Pattern: `[a-zA-Z0-9|!_-*'\\(\\)]([a-zA-Z0-9|!_-*'\\(\\)\\./])+`

Required: Yes

State

State of scheduled query.

Type: String

Valid Values: ENABLED | DISABLED

Required: Yes

CreationTime

The creation time of the scheduled query.

Type: Timestamp

Required: No

ErrorReportConfiguration

Configuration for scheduled query error reporting.

Type: [ErrorReportConfiguration](#) object

Required: No

LastRunStatus

Status of the last scheduled query run.

Type: String

Valid Values: AUTO_TRIGGER_SUCCESS | AUTO_TRIGGER_FAILURE |
MANUAL_TRIGGER_SUCCESS | MANUAL_TRIGGER_FAILURE

Required: No

NextInvocationTime

The next time the scheduled query is to be run.

Type: Timestamp

Required: No

PreviousInvocationTime

The last time the scheduled query was run.

Type: Timestamp

Required: No

TargetDestination

Target data source where final scheduled query result will be written.

Type: [TargetDestination](#) object

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](#)

ScheduledQueryDescription

Service: Amazon Timestream Query

Structure that describes scheduled query.

Contents

Arn

Scheduled query ARN.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: Yes

Name

Name of the scheduled query.

Type: String

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

Pattern: `[a-zA-Z0-9|!\\-_*'\\(\\)]([a-zA-Z0-9|!\\-_*'\\(\\)\\./])+`

Required: Yes

NotificationConfiguration

Notification configuration.

Type: [NotificationConfiguration](#) object

Required: Yes

QueryString

The query to be run.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 262144.

Required: Yes

ScheduleConfiguration

Schedule configuration.

Type: [ScheduleConfiguration](#) object

Required: Yes

State

State of the scheduled query.

Type: String

Valid Values: ENABLED | DISABLED

Required: Yes

CreationTime

Creation time of the scheduled query.

Type: Timestamp

Required: No

ErrorReportConfiguration

Error-reporting configuration for the scheduled query.

Type: [ErrorReportConfiguration](#) object

Required: No

KmsKeyId

A customer provided KMS key used to encrypt the scheduled query resource.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

LastRunSummary

Runtime summary for the last scheduled query run.

Type: [ScheduledQueryRunSummary](#) object

Required: No

NextInvocationTime

The next time the scheduled query is scheduled to run.

Type: Timestamp

Required: No

PreviousInvocationTime

Last time the query was run.

Type: Timestamp

Required: No

RecentlyFailedRuns

Runtime summary for the last five failed scheduled query runs.

Type: Array of [ScheduledQueryRunSummary](#) objects

Required: No

ScheduledQueryExecutionRoleArn

IAM role that Timestream uses to run the schedule query.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Required: No

TargetConfiguration

Scheduled query target store configuration.

Type: [TargetConfiguration](#) object

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](#)

ScheduledQueryInsights

Service: Amazon Timestream Query

Encapsulates settings for enabling QueryInsights on an ExecuteScheduledQueryRequest.

Contents

Mode

Provides the following modes to enable ScheduledQueryInsights:

- `ENABLED_WITH_RATE_CONTROL` – Enables ScheduledQueryInsights for the queries being processed. This mode also includes a rate control mechanism, which limits the QueryInsights feature to 1 query per second (QPS).
- `DISABLED` – Disables ScheduledQueryInsights.

Type: String

Valid Values: `ENABLED_WITH_RATE_CONTROL` | `DISABLED`

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](#)

ScheduledQueryInsightsResponse

Service: Amazon Timestream Query

Provides various insights and metrics related to the `ExecuteScheduledQueryRequest` that was executed.

Contents

OutputBytes

Indicates the size of query result set in bytes. You can use this data to validate if the result set has changed as part of the query tuning exercise.

Type: Long

Required: No

OutputRows

Indicates the total number of rows returned as part of the query result set. You can use this data to validate if the number of rows in the result set have changed as part of the query tuning exercise.

Type: Long

Required: No

QuerySpatialCoverage

Provides insights into the spatial coverage of the query, including the table with sub-optimal (max) spatial pruning. This information can help you identify areas for improvement in your partitioning strategy to enhance spatial pruning.

Type: [QuerySpatialCoverage](#) object

Required: No

QueryTableCount

Indicates the number of tables in the query.

Type: Long

Required: No

QueryTemporalRange

Provides insights into the temporal range of the query, including the table with the largest (max) time range. Following are some of the potential options for optimizing time-based pruning:

- Add missing time-predicates.
- Remove functions around the time predicates.
- Add time predicates to all the sub-queries.

Type: [QueryTemporalRange](#) object

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](#)

ScheduledQueryRunSummary

Service: Amazon Timestream Query

Run summary for the scheduled query

Contents

ErrorReportLocation

S3 location for error report.

Type: [ErrorReportLocation](#) object

Required: No

ExecutionStats

Runtime statistics for a scheduled run.

Type: [ExecutionStats](#) object

Required: No

FailureReason

Error message for the scheduled query in case of failure. You might have to look at the error report to get more detailed error reasons.

Type: String

Required: No

InvocationTime

InvocationTime for this run. This is the time at which the query is scheduled to run. Parameter `@scheduled_runtime` can be used in the query to get the value.

Type: Timestamp

Required: No

QueryInsightsResponse

Provides various insights and metrics related to the run summary of the scheduled query.

Type: [ScheduledQueryInsightsResponse](#) object

Required: No

RunStatus

The status of a scheduled query run.

Type: String

Valid Values: AUTO_TRIGGER_SUCCESS | AUTO_TRIGGER_FAILURE |
MANUAL_TRIGGER_SUCCESS | MANUAL_TRIGGER_FAILURE

Required: No

TriggerTime

The actual time when the query was run.

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](#)

SelectColumn

Service: Amazon Timestream Query

Details of the column that is returned by the query.

Contents

Aliased

True, if the column name was aliased by the query. False otherwise.

Type: Boolean

Required: No

DatabaseName

Database that has this column.

Type: String

Required: No

Name

Name of the column.

Type: String

Required: No

TableName

Table within the database that has this column.

Type: String

Required: No

Type

Contains the data type of a column in a query result set. The data type can be scalar or complex. The supported scalar data types are integers, Boolean, string, double, timestamp, date, time, and intervals. The supported complex data types are arrays, rows, and timeseries.

Type: [Type](#) object

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](#)

SnsConfiguration

Service: Amazon Timestream Query

Details on SNS that are required to send the notification.

Contents

TopicArn

SNS topic ARN that the scheduled query status notifications will be sent to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

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](#)

Tag

Service: Amazon Timestream Query

A tag is a label that you assign to a Timestream database and/or table. Each tag consists of a key and an optional value, both of which you define. Tags enable you to categorize databases and/or tables, for example, by purpose, owner, or environment.

Contents

Key

The key of the tag. Tag keys are case sensitive.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 128.

Required: Yes

Value

The value of the tag. Tag values are case sensitive and can be null.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 256.

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](#)

TargetConfiguration

Service: Amazon Timestream Query

Configuration used for writing the output of a query.

Contents

TimestreamConfiguration

Configuration needed to write data into the Timestream database and table.

Type: [TimestreamConfiguration](#) object

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](#)

TargetDestination

Service: Amazon Timestream Query

Destination details to write data for a target data source. Current supported data source is Timestream.

Contents

TimestreamDestination

Query result destination details for Timestream data source.

Type: [TimestreamDestination](#) object

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](#)

TimeSeriesDataPoint

Service: Amazon Timestream Query

The timeseries data type represents the values of a measure over time. A time series is an array of rows of timestamps and measure values, with rows sorted in ascending order of time. A TimeSeriesDataPoint is a single data point in the time series. It represents a tuple of (time, measure value) in a time series.

Contents

Time

The timestamp when the measure value was collected.

Type: String

Required: Yes

Value

The measure value for the data point.

Type: [Datum](#) object

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](#)

TimestreamConfiguration

Service: Amazon Timestream Query

Configuration to write data into Timestream database and table. This configuration allows the user to map the query result select columns into the destination table columns.

Contents

DatabaseName

Name of Timestream database to which the query result will be written.

Type: String

Required: Yes

DimensionMappings

This is to allow mapping column(s) from the query result to the dimension in the destination table.

Type: Array of [DimensionMapping](#) objects

Required: Yes

TableName

Name of Timestream table that the query result will be written to. The table should be within the same database that is provided in Timestream configuration.

Type: String

Required: Yes

TimeColumn

Column from query result that should be used as the time column in destination table. Column type for this should be `TIMESTAMP`.

Type: String

Required: Yes

MeasureNameColumn

Name of the measure column.

Type: String

Required: No

MixedMeasureMappings

Specifies how to map measures to multi-measure records.

Type: Array of [MixedMeasureMapping](#) objects

Array Members: Minimum number of 1 item.

Required: No

MultiMeasureMappings

Multi-measure mappings.

Type: [MultiMeasureMappings](#) object

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](#)

TimestreamDestination

Service: Amazon Timestream Query

Destination for scheduled query.

Contents

DatabaseName

Timestream database name.

Type: String

Required: No

TableName

Timestream table name.

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](#)

Type

Service: Amazon Timestream Query

Contains the data type of a column in a query result set. The data type can be scalar or complex. The supported scalar data types are integers, Boolean, string, double, timestamp, date, time, and intervals. The supported complex data types are arrays, rows, and timeseries.

Contents

ArrayColumnInfo

Indicates if the column is an array.

Type: [ColumnInfo](#) object

Required: No

RowColumnInfo

Indicates if the column is a row.

Type: Array of [ColumnInfo](#) objects

Required: No

ScalarType

Indicates if the column is of type string, integer, Boolean, double, timestamp, date, time. For more information, see [Supported data types](#).

Type: String

Valid Values: VARCHAR | BOOLEAN | BIGINT | DOUBLE | TIMESTAMP | DATE | TIME
| INTERVAL_DAY_TO_SECOND | INTERVAL_YEAR_TO_MONTH | UNKNOWN | INTEGER

Required: No

TimeSeriesMeasureValueColumnInfo

Indicates if the column is a timeseries data type.

Type: [ColumnInfo](#) object

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