
Application Auto Scaling

API Reference

API Version 2016-02-06



Application Auto Scaling: API Reference

Copyright © 2019 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
Actions	2
DeleteScalingPolicy	3
Request Syntax	3
Request Parameters	3
Response Elements	5
Errors	5
Example	5
See Also	6
DeleteScheduledAction	7
Request Syntax	7
Request Parameters	7
Response Elements	9
Errors	9
See Also	9
DeregisterScalableTarget	10
Request Syntax	10
Request Parameters	10
Response Elements	11
Errors	11
Example	12
See Also	12
DescribeScalableTargets	14
Request Syntax	14
Request Parameters	14
Response Syntax	16
Response Elements	16
Errors	16
Example	17
See Also	18
DescribeScalingActivities	19
Request Syntax	19
Request Parameters	19
Response Syntax	21
Response Elements	21
Errors	22
Example	22
See Also	24
DescribeScalingPolicies	25
Request Syntax	25
Request Parameters	25
Response Syntax	27
Response Elements	28
Errors	28
Example	29
See Also	31
DescribeScheduledActions	32
Request Syntax	32
Request Parameters	32
Response Syntax	34
Response Elements	34
Errors	35
See Also	35
PutScalingPolicy	37

Request Syntax	37
Request Parameters	38
Response Syntax	40
Response Elements	40
Errors	40
Examples	41
See Also	44
PutScheduledAction	45
Request Syntax	45
Request Parameters	45
Response Elements	48
Errors	48
See Also	48
RegisterScalableTarget	50
Request Syntax	50
Request Parameters	50
Response Elements	52
Errors	52
Example	53
See Also	53
Data Types	54
Alarm	55
Contents	55
See Also	55
CustomizedMetricSpecification	56
Contents	56
See Also	57
MetricDimension	58
Contents	58
See Also	58
PredefinedMetricSpecification	59
Contents	59
See Also	59
ScalableTarget	60
Contents	60
See Also	62
ScalableTargetAction	63
Contents	63
See Also	63
ScalingActivity	64
Contents	64
See Also	66
ScalingPolicy	67
Contents	67
See Also	69
ScheduledAction	70
Contents	70
See Also	72
StepAdjustment	74
Contents	74
See Also	75
StepScalingPolicyConfiguration	76
Contents	76
See Also	77
TargetTrackingScalingPolicyConfiguration	78
Contents	78
See Also	79

Common Parameters	80
Common Errors	82
Logging API Calls with CloudTrail	84
Application Auto Scaling Information in CloudTrail	84
Understanding Application Auto Scaling Log File Entries	85
.....	85

Welcome

With Application Auto Scaling, you can configure automatic scaling for your scalable resources. You can use Application Auto Scaling to accomplish the following tasks:

- Define scaling policies to automatically scale your AWS or custom resources
- Scale your resources in response to CloudWatch alarms
- Schedule one-time or recurring scaling actions
- View the history of your scaling events

Application Auto Scaling can scale the following resources:

- Amazon ECS services. For more information, see [Service Auto Scaling](#) in the *Amazon Elastic Container Service Developer Guide*.
- Amazon EC2 Spot fleets. For more information, see [Automatic Scaling for Spot Fleet](#) in the *Amazon EC2 User Guide*.
- Amazon EMR clusters. For more information, see [Using Automatic Scaling in Amazon EMR](#) in the *Amazon EMR Management Guide*.
- AppStream 2.0 fleets. For more information, see [Fleet Auto Scaling for Amazon AppStream 2.0](#) in the *Amazon AppStream 2.0 Developer Guide*.
- Provisioned read and write capacity for Amazon DynamoDB tables and global secondary indexes. For more information, see [Managing Throughput Capacity Automatically with DynamoDB Auto Scaling](#) in the *Amazon DynamoDB Developer Guide*.
- Amazon Aurora Replicas. For more information, see [Using Amazon Aurora Auto Scaling with Aurora Replicas](#).
- Amazon SageMaker endpoint variants. For more information, see [Automatically Scaling Amazon SageMaker Models](#).
- Custom resources provided by your own applications or services. More information is available in our [GitHub repository](#).

To learn more about Application Auto Scaling, including information about granting IAM users required permissions for Application Auto Scaling actions, see the [Application Auto Scaling User Guide](#).

This document was last published on April 18, 2019.

Actions

The following actions are supported:

- [DeleteScalingPolicy](#) (p. 3)
- [DeleteScheduledAction](#) (p. 7)
- [DeregisterScalableTarget](#) (p. 10)
- [DescribeScalableTargets](#) (p. 14)
- [DescribeScalingActivities](#) (p. 19)
- [DescribeScalingPolicies](#) (p. 25)
- [DescribeScheduledActions](#) (p. 32)
- [PutScalingPolicy](#) (p. 37)
- [PutScheduledAction](#) (p. 45)
- [RegisterScalableTarget](#) (p. 50)

- Amazon SageMaker endpoint variants - The resource type is `variant` and the unique identifier is the resource ID. Example: `endpoint/my-end-point/variant/KMeansClustering`.
- Custom resources are not supported with a resource type. This parameter must specify the `OutputValue` from the CloudFormation template stack used to access the resources. The unique identifier is defined by the service provider. More information is available in our [GitHub repository](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1600.

Pattern: `[\u0020-\u007F\u00E0-\u00FF\u0100-\u017F\u0180-\u01FF\u0200-\u02FF\u0300-\u03FF\u0400-\u04FF\u0500-\u05FF\u0600-\u06FF\u0700-\u07FF\u0800-\u08FF\u0900-\u09FF\u0A00-\u0AFF\u0B00-\u0BFF\u0C00-\u0CFF\u0D00-\u0DFF\u0E00-\u0EFF\u0F00-\u0FFF\u1000-\u10FF\u1100-\u11FF\u1200-\u12FF\u1300-\u13FF\u1400-\u14FF\u1500-\u15FF\u1600-\u16FF\u1700-\u17FF\u1800-\u18FF\u1900-\u19FF\u1A00-\u1AFF\u1B00-\u1BFF\u1C00-\u1CFF\u1D00-\u1DFF\u1E00-\u1EFF\u1F00-\u1FFF\u2000-\u206F\u2070-\u209F\u2100-\u214F\u2150-\u218F\u2190-\u21FF\u2200-\u223F\u2240-\u227F\u2280-\u229F\u22A0-\u22FF\u2300-\u232F\u2330-\u236F\u2370-\u239F\u23A0-\u23FF\u2400-\u243F\u2440-\u247F\u2480-\u249F\u24A0-\u24FF\u2500-\u254F\u2550-\u258F\u2590-\u25FF\u2600-\u264F\u2650-\u268F\u2690-\u26FF\u2700-\u273F\u2740-\u277F\u2780-\u279F\u27A0-\u27FF\u2800-\u283F\u2840-\u287F\u2880-\u289F\u28A0-\u28FF\u2900-\u293F\u2940-\u297F\u2980-\u299F\u29A0-\u29FF\u2A00-\u2A3F\u2A40-\u2A7F\u2A80-\u2A9F\u2AA0-\u2AFF\u2B00-\u2B3F\u2B40-\u2B7F\u2B80-\u2B9F\u2BA0-\u2BFF\u2C00-\u2C3F\u2C40-\u2C7F\u2C80-\u2C9F\u2CA0-\u2CFF\u2D00-\u2D3F\u2D40-\u2D7F\u2D80-\u2D9F\u2DA0-\u2DAF\u2DB0-\u2DBF\u2DC0-\u2DCF\u2DD0-\u2DDF\u2DE0-\u2DEF\u2DF0-\u2DFB\u2E00-\u2E3F\u2E40-\u2E7F\u2E80-\u2E9F\u2EA0-\u2EBF\u2EC0-\u2ECF\u2ED0-\u2EDF\u2EE0-\u2EEB\u2EF0-\u2EFB\u2F00-\u2F3F\u2F40-\u2F7F\u2F80-\u2F9F\u2FA0-\u2FBF\u2FC0-\u2FCF\u2FD0-\u2FDF\u2FE0-\u2FEF\u2FF0-\u2FFB]`*

Required: Yes

ScalableDimension (p. 3)

The scalable dimension. This string consists of the service namespace, resource type, and scaling property.

- `ecs:service:DesiredCount` - The desired task count of an ECS service.
- `ec2:spot-fleet-request:TargetCapacity` - The target capacity of a Spot fleet request.
- `elasticmapreduce:instancegroup:InstanceCount` - The instance count of an EMR Instance Group.
- `appstream:fleet:DesiredCapacity` - The desired capacity of an AppStream 2.0 fleet.
- `dynamodb:table:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB table.
- `dynamodb:table:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB table.
- `dynamodb:index:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB global secondary index.
- `dynamodb:index:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB global secondary index.
- `rds:cluster:ReadReplicaCount` - The count of Aurora Replicas in an Aurora DB cluster. Available for Aurora MySQL-compatible edition and Aurora PostgreSQL-compatible edition.
- `sagemaker:variant:DesiredInstanceCount` - The number of EC2 instances for an Amazon SageMaker model endpoint variant.
- `custom-resource:ResourceType:Property` - The scalable dimension for a custom resource provided by your own application or service.

Type: String

Valid Values: `ecs:service:DesiredCount | ec2:spot-fleet-request:TargetCapacity | elasticmapreduce:instancegroup:InstanceCount | appstream:fleet:DesiredCapacity | dynamodb:table:ReadCapacityUnits | dynamodb:table:WriteCapacityUnits | dynamodb:index:ReadCapacityUnits | dynamodb:index:WriteCapacityUnits | rds:cluster:ReadReplicaCount | sagemaker:variant:DesiredInstanceCount | custom-resource:ResourceType:Property`

Required: Yes

ServiceNamespace (p. 3)

The namespace of the AWS service that provides the resource or `custom-resource` for a resource provided by your own application or service. For more information, see [AWS Service Namespaces](#) in the *Amazon Web Services General Reference*.

Type: String

Valid Values: `ecs` | `elasticmapreduce` | `ec2` | `appstream` | `dynamodb` | `rds` | `sagemaker` | `custom-resource`

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 Errors \(p. 82\)](#).

ConcurrentUpdateException

Concurrent updates caused an exception, for example, if you request an update to an Application Auto Scaling resource that already has a pending update.

HTTP Status Code: 400

InternalServiceException

The service encountered an internal error.

HTTP Status Code: 400

ObjectNotFoundException

The specified object could not be found. For any operation that depends on the existence of a scalable target, this exception is thrown if the scalable target with the specified service namespace, resource ID, and scalable dimension does not exist. For any operation that deletes or deregisters a resource, this exception is thrown if the resource cannot be found.

HTTP Status Code: 400

ValidationException

An exception was thrown for a validation issue. Review the available parameters for the API request.

HTTP Status Code: 400

Example

If you plan to create requests manually, you must replace the Authorization header contents in the examples (`AUTHPARAMS`) with a signature. For more information, see [Signature Version 4 Signing Process](#) in the *Amazon Web Services General Reference*. If you plan to use the [AWS CLI](#) or one of the [AWS SDKs](#), these tools sign the requests for you.

Example

The following example deletes a scaling policy for the Amazon ECS service `web-app` running in the `default` cluster.

Sample Request

```
POST / HTTP/1.1
Host: autoscaling.us-west-2.amazonaws.com
```

```
Accept-Encoding: identity
Content-Length: 152
X-Amz-Target: AnyScaleFrontendService.DeleteScalingPolicy
X-Amz-Date: 20160506T205712Z
User-Agent: aws-cli/1.10.23 Python/2.7.11 Darwin/15.4.0 botocore/1.4.8
Content-Type: application/x-amz-json-1.1
Authorization: AUTHPARAMS

{
  "PolicyName": "my-scaleout-policy",
  "ServiceNamespace": "ecs",
  "ScalableDimension": "ecs:service:DesiredCount",
  "ResourceId": "service/default/web-app"
}
```

See Also

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

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DeleteScheduledAction

Deletes the specified Application Auto Scaling scheduled action.

Request Syntax

```
{  
  "ResourceId": "string",  
  "ScalableDimension": "string",  
  "ScheduledActionName": "string",  
  "ServiceNamespace": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 80).

The request accepts the following data in JSON format.

ResourceId (p. 7)

The identifier of the resource associated with the scheduled action. This string consists of the resource type and unique identifier.

- ECS service - The resource type is `service` and the unique identifier is the cluster name and service name. Example: `service/default/sample-webapp`.
- Spot fleet request - The resource type is `spot-fleet-request` and the unique identifier is the Spot fleet request ID. Example: `spot-fleet-request/sfr-73fbd2ce-aa30-494c-8788-1cee4EXAMPLE`.
- EMR cluster - The resource type is `instancegroup` and the unique identifier is the cluster ID and instance group ID. Example: `instancegroup/j-2EEZNYKUA1NTV/ig-1791Y4E1L8YI0`.
- AppStream 2.0 fleet - The resource type is `fleet` and the unique identifier is the fleet name. Example: `fleet/sample-fleet`.
- DynamoDB table - The resource type is `table` and the unique identifier is the resource ID. Example: `table/my-table`.
- DynamoDB global secondary index - The resource type is `index` and the unique identifier is the resource ID. Example: `table/my-table/index/my-table-index`.
- Aurora DB cluster - The resource type is `cluster` and the unique identifier is the cluster name. Example: `cluster:my-db-cluster`.
- Amazon SageMaker endpoint variants - The resource type is `variant` and the unique identifier is the resource ID. Example: `endpoint/my-end-point/variant/KMeansClustering`.
- Custom resources are not supported with a resource type. This parameter must specify the `OutputValue` from the CloudFormation template stack used to access the resources. The unique identifier is defined by the service provider. More information is available in our [GitHub repository](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1600.

Pattern: `[\u0020-\u007F\u00E0-\u00FF\u0100-\u017F\u0180-\u01FF\u0200-\u02FF\u0300-\u03FF\u0400-\u04FF\u0500-\u05FF\u0600-\u06FF\u0700-\u07FF\u0800-\u08FF\u0900-\u09FF\u0A00-\u0AFF\u0B00-\u0BFF\u0C00-\u0CFF\u0D00-\u0DFF\u0E00-\u0EFF\u0F00-\u0FFF\u1000-\u10FF\u1100-\u11FF\u1200-\u12FF\u1300-\u13FF\u1400-\u14FF\u1500-\u15FF\u1600-\u16FF\u1700-\u17FF\u1800-\u18FF\u1900-\u19FF\u1A00-\u1AFF\u1B00-\u1BFF\u1C00-\u1CFF\u1D00-\u1DFF\u1E00-\u1EFF\u1F00-\u1FFF\u2000-\u20FF\u2100-\u21FF\u2200-\u22FF\u2300-\u23FF\u2400-\u24FF\u2500-\u25FF\u2600-\u26FF\u2700-\u27FF\u2800-\u28FF\u2900-\u29FF\u2A00-\u2AFF\u2B00-\u2BFF\u2C00-\u2CFF\u2D00-\u2DFF\u2E00-\u2EFF\u2F00-\u2FFF\u3000-\u30FF\u3100-\u31FF\u3200-\u32FF\u3300-\u33FF\u3400-\u34FF\u3500-\u35FF\u3600-\u36FF\u3700-\u37FF\u3800-\u38FF\u3900-\u39FF\u3A00-\u3AFF\u3B00-\u3BFF\u3C00-\u3CFF\u3D00-\u3DFF\u3E00-\u3EFF\u3F00-\u3FFF\u4000-\u40FF\u4100-\u41FF\u4200-\u42FF\u4300-\u43FF\u4400-\u44FF\u4500-\u45FF\u4600-\u46FF\u4700-\u47FF\u4800-\u48FF\u4900-\u49FF\u4A00-\u4AFF\u4B00-\u4BFF\u4C00-\u4CFF\u4D00-\u4DFF\u4E00-\u4EFF\u4F00-\u4FFF\u5000-\u50FF\u5100-\u51FF\u5200-\u52FF\u5300-\u53FF\u5400-\u54FF\u5500-\u55FF\u5600-\u56FF\u5700-\u57FF\u5800-\u58FF\u5900-\u59FF\u5A00-\u5AFF\u5B00-\u5BFF\u5C00-\u5CFF\u5D00-\u5DFF\u5E00-\u5EFF\u5F00-\u5FFF\u6000-\u60FF\u6100-\u61FF\u6200-\u62FF\u6300-\u63FF\u6400-\u64FF\u6500-\u65FF\u6600-\u66FF\u6700-\u67FF\u6800-\u68FF\u6900-\u69FF\u6A00-\u6AFF\u6B00-\u6BFF\u6C00-\u6CFF\u6D00-\u6DFF\u6E00-\u6EFF\u6F00-\u6FFF\u7000-\u70FF\u7100-\u71FF\u7200-\u72FF\u7300-\u73FF\u7400-\u74FF\u7500-\u75FF\u7600-\u76FF\u7700-\u77FF\u7800-\u78FF\u7900-\u79FF\u7A00-\u7AFF\u7B00-\u7BFF\u7C00-\u7CFF\u7D00-\u7DFF\u7E00-\u7EFF\u7F00-\u7FFF\u8000-\u80FF\u8100-\u81FF\u8200-\u82FF\u8300-\u83FF\u8400-\u84FF\u8500-\u85FF\u8600-\u86FF\u8700-\u87FF\u8800-\u88FF\u8900-\u89FF\u8A00-\u8AFF\u8B00-\u8BFF\u8C00-\u8CFF\u8D00-\u8DFF\u8E00-\u8EFF\u8F00-\u8FFF\u9000-\u90FF\u9100-\u91FF\u9200-\u92FF\u9300-\u93FF\u9400-\u94FF\u9500-\u95FF\u9600-\u96FF\u9700-\u97FF\u9800-\u98FF\u9900-\u99FF\u9A00-\u9AFF\u9B00-\u9BFF\u9C00-\u9CFF\u9D00-\u9DFF\u9E00-\u9EFF\u9F00-\u9FFF\uA000-\uA0FF\uA100-\uA1FF\uA200-\uA2FF\uA300-\uA3FF\uA400-\uA4FF\uA500-\uA5FF\uA600-\uA6FF\uA700-\uA7FF\uA800-\uA8FF\uA900-\uA9FF\uAA00-\uAAFF\uAB00-\uABFF\uAC00-\uACFF\uAD00-\uADFF\uAE00-\uAEFF\uAF00-\uAFFF\uB000-\uB0FF\uB100-\uB1FF\uB200-\uB2FF\uB300-\uB3FF\uB400-\uB4FF\uB500-\uB5FF\uB600-\uB6FF\uB700-\uB7FF\uB800-\uB8FF\uB900-\uB9FF\uBA00-\uBAFF\uBB00-\uBBFF\uBC00-\uBCFF\uBD00-\uBDFF\uBE00-\uBEFF\uBF00-\uBFFF\uC000-\u00FF\uC001-\u007F\uC002-\u007F\uC003-\u007F\uC004-\u007F\uC005-\u007F\uC006-\u007F\uC007-\u007F\uC008-\u007F\uC009-\u007F\uC00A-\u007F\uC00B-\u007F\uC00C-\u007F\uC00D-\u007F\uC00E-\u007F\uC00F-\u007F\uC010-\u007F\uC011-\u007F\uC012-\u007F\uC013-\u007F\uC014-\u007F\uC015-\u007F\uC016-\u007F\uC017-\u007F\uC018-\u007F\uC019-\u007F\uC01A-\u007F\uC01B-\u007F\uC01C-\u007F\uC01D-\u007F\uC01E-\u007F\uC01F-\u007F\uC020-\u007F\uC021-\u007F\uC022-\u007F\uC023-\u007F\uC024-\u007F\uC025-\u007F\uC026-\u007F\uC027-\u007F\uC028-\u007F\uC029-\u007F\uC02A-\u007F\uC02B-\u007F\uC02C-\u007F\uC02D-\u007F\uC02E-\u007F\uC02F-\u007F\uC030-\u007F\uC031-\u007F\uC032-\u007F\uC033-\u007F\uC034-\u007F\uC035-\u007F\uC036-\u007F\uC037-\u007F\uC038-\u007F\uC039-\u007F\uC03A-\u007F\uC03B-\u007F\uC03C-\u007F\uC03D-\u007F\uC03E-\u007F\uC03F-\u007F\uC040-\u007F\uC041-\u007F\uC042-\u007F\uC043-\u007F\uC044-\u007F\uC045-\u007F\uC046-\u007F\uC047-\u007F\uC048-\u007F\uC049-\u007F\uC04A-\u007F\uC04B-\u007F\uC04C-\u007F\uC04D-\u007F\uC04E-\u007F\uC04F-\u007F\uC050-\u007F\uC051-\u007F\uC052-\u007F\uC053-\u007F\uC054-\u007F\uC055-\u007F\uC056-\u007F\uC057-\u007F\uC058-\u007F\uC059-\u007F\uC05A-\u007F\uC05B-\u007F\uC05C-\u007F\uC05D-\u007F\uC05E-\u007F\uC05F-\u007F\uC060-\u007F\uC061-\u007F\uC062-\u007F\uC063-\u007F\uC064-\u007F\uC065-\u007F\uC066-\u007F\uC067-\u007F\uC068-\u007F\uC069-\u007F\uC06A-\u007F\uC06B-\u007F\uC06C-\u007F\uC06D-\u007F\uC06E-\u007F\uC06F-\u007F\uC070-\u007F\uC071-\u007F\uC072-\u007F\uC073-\u007F\uC074-\u007F\uC075-\u007F\uC076-\u007F\uC077-\u007F\uC078-\u007F\uC079-\u007F\uC07A-\u007F\uC07B-\u007F\uC07C-\u007F\uC07D-\u007F\uC07E-\u007F\uC07F\u0080-\u00FF\u0081-\u00FF\u0082-\u00FF\u0083-\u00FF\u0084-\u00FF\u0085-\u00FF\u0086-\u00FF\u0087-\u00FF\u0088-\u00FF\u0089-\u00FF\u008A-\u00FF\u008B-\u00FF\u008C-\u00FF\u008D-\u00FF\u008E-\u00FF\u008F-\u00FF\u0090-\u00FF\u0091-\u00FF\u0092-\u00FF\u0093-\u00FF\u0094-\u00FF\u0095-\u00FF\u0096-\u00FF\u0097-\u00FF\u0098-\u00FF\u0099-\u00FF\u009A-\u00FF\u009B-\u00FF\u009C-\u00FF\u009D-\u00FF\u009E-\u00FF\u009F\u00A0-\u00FF\u00A1-\u00FF\u00A2-\u00FF\u00A3-\u00FF\u00A4-\u00FF\u00A5-\u00FF\u00A6-\u00FF\u00A7-\u00FF\u00A8-\u00FF\u00A9-\u00FF\u00AA-\u00FF\u00AB-\u00FF\u00AC-\u00FF\u00AD-\u00FF\u00AE-\u00FF\u00AF\u00B0-\u00FF\u00B1-\u00FF\u00B2-\u00FF\u00B3-\u00FF\u00B4-\u00FF\u00B5-\u00FF\u00B6-\u00FF\u00B7-\u00FF\u00B8-\u00FF\u00B9-\u00FF\u00BA-\u00FF\u00BB-\u00FF\u00BC-\u00FF\u00BD-\u00FF\u00BE-\u00FF\u00BF\u00C0-\u00FF\u00C1-\u00FF\u00C2-\u00FF\u00C3-\u00FF\u00C4-\u00FF\u00C5-\u00FF\u00C6-\u00FF\u00C7-\u00FF\u00C8-\u00FF\u00C9-\u00FF\u00CA-\u00FF\u00CB-\u00FF\u00CC-\u00FF\u00CD-\u00FF\u00CE-\u00FF\u00CF\u00D0-\u00FF\u00D1-\u00FF\u00D2-\u00FF\u00D3-\u00FF\u00D4-\u00FF\u00D5-\u00FF\u00D6-\u00FF\u00D7-\u00FF\u00D8-\u00FF\u00D9-\u00FF\u00DA-\u00FF\u00DB-\u00FF\u00DC-\u00FF\u00DD-\u00FF\u00DE-\u00FF\u00DF\u00E0-\u00FF\u00E1-\u00FF\u00E2-\u00FF\u00E3-\u00FF\u00E4-\u00FF\u00E5-\u00FF\u00E6-\u00FF\u00E7-\u00FF\u00E8-\u00FF\u00E9-\u00FF\u00EA-\u00FF\u00EB-\u00FF\u00EC-\u00FF\u00ED-\u00FF\u00EE-\u00FF\u00EF\u00F0-\u00FF\u00F1-\u00FF\u00F2-\u00FF\u00F3-\u00FF\u00F4-\u00FF\u00F5-\u00FF\u00F6-\u00FF\u00F7-\u00FF\u00F8-\u00FF\u00F9-\u00FF\u00FA-\u00FF\u00FB-\u00FF\u00FC-\u00FF\u00FD-\u00FF\u00FE-\u00FF\u00FF]`*

Required: Yes

ScalableDimension (p. 7)

The scalable dimension. This string consists of the service namespace, resource type, and scaling property.

- `ecs:service:DesiredCount` - The desired task count of an ECS service.
- `ec2:spot-fleet-request:TargetCapacity` - The target capacity of a Spot fleet request.
- `elasticmapreduce:instancegroup:InstanceCount` - The instance count of an EMR Instance Group.
- `appstream:fleet:DesiredCapacity` - The desired capacity of an AppStream 2.0 fleet.
- `dynamodb:table:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB table.
- `dynamodb:table:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB table.
- `dynamodb:index:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB global secondary index.
- `dynamodb:index:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB global secondary index.
- `rds:cluster:ReadReplicaCount` - The count of Aurora Replicas in an Aurora DB cluster. Available for Aurora MySQL-compatible edition and Aurora PostgreSQL-compatible edition.
- `sagemaker:variant:DesiredInstanceCount` - The number of EC2 instances for an Amazon SageMaker model endpoint variant.
- `custom-resource:ResourceType:Property` - The scalable dimension for a custom resource provided by your own application or service.

Type: String

Valid Values: `ecs:service:DesiredCount` | `ec2:spot-fleet-request:TargetCapacity` | `elasticmapreduce:instancegroup:InstanceCount` | `appstream:fleet:DesiredCapacity` | `dynamodb:table:ReadCapacityUnits` | `dynamodb:table:WriteCapacityUnits` | `dynamodb:index:ReadCapacityUnits` | `dynamodb:index:WriteCapacityUnits` | `rds:cluster:ReadReplicaCount` | `sagemaker:variant:DesiredInstanceCount` | `custom-resource:ResourceType:Property`

Required: Yes

ScheduledActionName (p. 7)

The name of the scheduled action.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1600.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*`

Required: Yes

ServiceNamespace (p. 7)

The namespace of the AWS service that provides the resource or `custom-resource` for a resource provided by your own application or service. For more information, see [AWS Service Namespaces](#) in the *Amazon Web Services General Reference*.

Type: String

Valid Values: `ecs` | `elasticmapreduce` | `ec2` | `appstream` | `dynamodb` | `rds` | `sagemaker` | `custom-resource`

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 Errors \(p. 82\)](#).

ConcurrentUpdateException

Concurrent updates caused an exception, for example, if you request an update to an Application Auto Scaling resource that already has a pending update.

HTTP Status Code: 400

InternalServiceException

The service encountered an internal error.

HTTP Status Code: 400

ObjectNotFoundException

The specified object could not be found. For any operation that depends on the existence of a scalable target, this exception is thrown if the scalable target with the specified service namespace, resource ID, and scalable dimension does not exist. For any operation that deletes or deregisters a resource, this exception is thrown if the resource cannot be found.

HTTP Status Code: 400

ValidationException

An exception was thrown for a validation issue. Review the available parameters for the API 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](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DeregisterScalableTarget

Deregisters a scalable target.

Deregistering a scalable target deletes the scaling policies that are associated with it.

To create a scalable target or update an existing one, see [RegisterScalableTarget \(p. 50\)](#).

Request Syntax

```
{
  "ResourceId": "string",
  "ScalableDimension": "string",
  "ServiceNamespace": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 80\)](#).

The request accepts the following data in JSON format.

ResourceId (p. 10)

The identifier of the resource associated with the scalable target. This string consists of the resource type and unique identifier.

- ECS service - The resource type is `service` and the unique identifier is the cluster name and service name. Example: `service/default/sample-webapp`.
- Spot fleet request - The resource type is `spot-fleet-request` and the unique identifier is the Spot fleet request ID. Example: `spot-fleet-request/sfr-73fbd2ce-aa30-494c-8788-1cee4EXAMPLE`.
- EMR cluster - The resource type is `instancegroup` and the unique identifier is the cluster ID and instance group ID. Example: `instancegroup/j-2EEZNYKUA1NTV/ig-1791Y4E1L8YI0`.
- AppStream 2.0 fleet - The resource type is `fleet` and the unique identifier is the fleet name. Example: `fleet/sample-fleet`.
- DynamoDB table - The resource type is `table` and the unique identifier is the resource ID. Example: `table/my-table`.
- DynamoDB global secondary index - The resource type is `index` and the unique identifier is the resource ID. Example: `table/my-table/index/my-table-index`.
- Aurora DB cluster - The resource type is `cluster` and the unique identifier is the cluster name. Example: `cluster:my-db-cluster`.
- Amazon SageMaker endpoint variants - The resource type is `variant` and the unique identifier is the resource ID. Example: `endpoint/my-end-point/variant/KMeansClustering`.
- Custom resources are not supported with a resource type. This parameter must specify the `OutputValue` from the CloudFormation template stack used to access the resources. The unique identifier is defined by the service provider. More information is available in our [GitHub repository](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1600.

Pattern: `[\u0020-\u007F\u00E0-\u00FF\u0100-\u017F\u0180-\u01FF\u0200-\u02FF\u0300-\u037F\u0380-\u03FF\u0400-\u047F\u0480-\u04FF\u0500-\u057F\u0580-\u05FF\u0600-\u06FF\u0700-\u077F\u0780-\u07FF\u0800-\u087F\u0880-\u08FF\u0900-\u097F\u0980-\u09FF\u0A00-\u0A7F\u0A80-\u0AFF\u0B00-\u0B7F\u0B80-\u0BFF\u0C00-\u0C7F\u0C80-\u0CFF\u0D00-\u0D7F\u0D80-\u0DBFF\u0E00-\u0E7F\u0E80-\u0EFF\u0F00-\u0F7F\u0F80-\u0FFF\u1000-\u107F\u1080-\u10FF\u1100-\u117F\u1180-\u11FF\u1200-\u127F\u1280-\u12FF\u1300-\u137F\u1380-\u13FF\u1400-\u147F\u1480-\u14FF\u1500-\u157F\u1580-\u15FF\u1600-\u167F\u1680-\u16FF\u1700-\u177F\u1780-\u17FF\u1800-\u187F\u1880-\u18FF\u1900-\u197F\u1980-\u19FF\u1A00-\u1A7F\u1A80-\u1AFF\u1B00-\u1B7F\u1B80-\u1BFF\u1C00-\u1C7F\u1C80-\u1CFF\u1D00-\u1D7F\u1D80-\u1DBFF\u1E00-\u1E7F\u1E80-\u1EFF\u1F00-\u1F7F\u1F80-\u1FFF\u2000-\u207F\u2080-\u20FF\u2100-\u217F\u2180-\u21FF\u2200-\u227F\u2280-\u22FF\u2300-\u237F\u2380-\u23FF\u2400-\u247F\u2480-\u24FF\u2500-\u257F\u2580-\u25FF\u2600-\u267F\u2680-\u26FF\u2700-\u277F\u2780-\u27FF\u2800-\u287F\u2880-\u28FF\u2900-\u297F\u2980-\u29FF\u2A00-\u2A7F\u2A80-\u2AFF\u2B00-\u2B7F\u2B80-\u2BFF\u2C00-\u2C7F\u2C80-\u2CFF\u2D00-\u2D7F\u2D80-\u2DBFF\u2E00-\u2E7F\u2E80-\u2EFF\u2F00-\u2F7F\u2F80-\u2FFF\u3000-\u307F\u3080-\u30FF\u3100-\u317F\u3180-\u31FF\u3200-\u327F\u3280-\u32FF\u3300-\u337F\u3380-\u33FF\u3400-\u347F\u3480-\u34FF\u3500-\u357F\u3580-\u35FF\u3600-\u367F\u3680-\u36FF\u3700-\u377F\u3780-\u37FF\u3800-\u387F\u3880-\u38FF\u3900-\u397F\u3980-\u39FF\u3A00-\u3A7F\u3A80-\u3AFF\u3B00-\u3B7F\u3B80-\u3BFF\u3C00-\u3C7F\u3C80-\u3CFF\u3D00-\u3D7F\u3D80-\u3DBFF\u3E00-\u3E7F\u3E80-\u3EFF\u3F00-\u3F7F\u3F80-\u3FFF\u4000-\u407F\u4080-\u40FF\u4100-\u417F\u4180-\u41FF\u4200-\u427F\u4280-\u42FF\u4300-\u437F\u4380-\u43FF\u4400-\u447F\u4480-\u44FF\u4500-\u457F\u4580-\u45FF\u4600-\u467F\u4680-\u46FF\u4700-\u477F\u4780-\u47FF\u4800-\u487F\u4880-\u48FF\u4900-\u497F\u4980-\u49FF\u4A00-\u4A7F\u4A80-\u4AFF\u4B00-\u4B7F\u4B80-\u4BFF\u4C00-\u4C7F\u4C80-\u4CFF\u4D00-\u4D7F\u4D80-\u4DBFF\u4E00-\u4E7F\u4E80-\u4EFF\u4F00-\u4F7F\u4F80-\u4FFF\u5000-\u507F\u5080-\u50FF\u5100-\u517F\u5180-\u51FF\u5200-\u527F\u5280-\u52FF\u5300-\u537F\u5380-\u53FF\u5400-\u547F\u5480-\u54FF\u5500-\u557F\u5580-\u55FF\u5600-\u567F\u5680-\u56FF\u5700-\u577F\u5780-\u57FF\u5800-\u587F\u5880-\u58FF\u5900-\u597F\u5980-\u59FF\u5A00-\u5A7F\u5A80-\u5AFF\u5B00-\u5B7F\u5B80-\u5BFF\u5C00-\u5C7F\u5C80-\u5CFF\u5D00-\u5D7F\u5D80-\u5DBFF\u5E00-\u5E7F\u5E80-\u5EFF\u5F00-\u5F7F\u5F80-\u5FFF\u6000-\u607F\u6080-\u60FF\u6100-\u617F\u6180-\u61FF\u6200-\u627F\u6280-\u62FF\u6300-\u637F\u6380-\u63FF\u6400-\u647F\u6480-\u64FF\u6500-\u657F\u6580-\u65FF\u6600-\u667F\u6680-\u66FF\u6700-\u677F\u6780-\u67FF\u6800-\u687F\u6880-\u68FF\u6900-\u697F\u6980-\u69FF\u6A00-\u6A7F\u6A80-\u6AFF\u6B00-\u6B7F\u6B80-\u6BFF\u6C00-\u6C7F\u6C80-\u6CFF\u6D00-\u6D7F\u6D80-\u6DBFF\u6E00-\u6E7F\u6E80-\u6EFF\u6F00-\u6F7F\u6F80-\u6FFF\u7000-\u707F\u7080-\u70FF\u7100-\u717F\u7180-\u71FF\u7200-\u727F\u7280-\u72FF\u7300-\u737F\u7380-\u73FF\u7400-\u747F\u7480-\u74FF\u7500-\u757F\u7580-\u75FF\u7600-\u767F\u7680-\u76FF\u7700-\u777F\u7780-\u77FF\u7800-\u787F\u7880-\u78FF\u7900-\u797F\u7980-\u79FF\u7A00-\u7A7F\u7A80-\u7AFF\u7B00-\u7B7F\u7B80-\u7BFF\u7C00-\u7C7F\u7C80-\u7CFF\u7D00-\u7D7F\u7D80-\u7DBFF\u7E00-\u7E7F\u7E80-\u7EFF\u7F00-\u7F7F\u7F80-\u7FFF\u8000-\u807F\u8080-\u80FF\u8100-\u817F\u8180-\u81FF\u8200-\u827F\u8280-\u82FF\u8300-\u837F\u8380-\u83FF\u8400-\u847F\u8480-\u84FF\u8500-\u857F\u8580-\u85FF\u8600-\u867F\u8680-\u86FF\u8700-\u877F\u8780-\u87FF\u8800-\u887F\u8880-\u88FF\u8900-\u897F\u8980-\u89FF\u8A00-\u8A7F\u8A80-\u8AFF\u8B00-\u8B7F\u8B80-\u8BFF\u8C00-\u8C7F\u8C80-\u8CFF\u8D00-\u8D7F\u8D80-\u8DBFF\u8E00-\u8E7F\u8E80-\u8EFF\u8F00-\u8F7F\u8F80-\u8FFF\u9000-\u907F\u9080-\u90FF\u9100-\u917F\u9180-\u91FF\u9200-\u927F\u9280-\u92FF\u9300-\u937F\u9380-\u93FF\u9400-\u947F\u9480-\u94FF\u9500-\u957F\u9580-\u95FF\u9600-\u967F\u9680-\u96FF\u9700-\u977F\u9780-\u97FF\u9800-\u987F\u9880-\u98FF\u9900-\u997F\u9980-\u99FF\u9A00-\u9A7F\u9A80-\u9AFF\u9B00-\u9B7F\u9B80-\u9BFF\u9C00-\u9C7F\u9C80-\u9CFF\u9D00-\u9D7F\u9D80-\u9DBFF\u9E00-\u9E7F\u9E80-\u9EFF\u9F00-\u9F7F\u9F80-\u9FFF\uA000-\uA07F\uA080-\uA0FF\uA100-\uA17F\uA180-\uA1FF\uA200-\uA27F\uA280-\uA2FF\uA300-\uA37F\uA380-\uA3FF\uA400-\uA47F\uA480-\uA4FF\uA500-\uA57F\uA580-\uA5FF\uA600-\uA67F\uA680-\uA6FF\uA700-\uA77F\uA780-\uA7FF\uA800-\uA87F\uA880-\uA8FF\uA900-\uA97F\uA980-\uA9FF\uAA00-\uAA7F\uAA80-\uAAFF\uAB00-\uAB7F\uAB80-\uABFF\uAC00-\uAC7F\uAC80-\uACFF\uAD00-\uAD7F\uAD80-\uADBFF\uAE00-\uAE7F\uAE80-\uAEFF\uAF00-\uAF7F\uAF80-\uAFFF\uB000-\uB07F\uB080-\uB0FF\uB100-\uB17F\uB180-\uB1FF\uB200-\uB27F\uB280-\uB2FF\uB300-\uB37F\uB380-\uB3FF\uB400-\uB47F\uB480-\uB4FF\uB500-\uB57F\uB580-\uB5FF\uB600-\uB67F\uB680-\uB6FF\uB700-\uB77F\uB780-\uB7FF\uB800-\uB87F\uB880-\uB8FF\uB900-\uB97F\uB980-\uB9FF\uBA00-\uBA7F\uBA80-\uBAFF\uBB00-\uBB7F\uBB80-\uBBFF\uBC00-\uBC7F\uBC80-\uBCFF\uBD00-\uBD7F\uBD80-\uBDBFF\uBE00-\uBE7F\uBE80-\uBEFF\uBF00-\uBF7F\uBF80-\uBFFF\uC000-\uC07F\uC080-\uC0FF\uC100-\uC17F\uC180-\uC1FF\uC200-\uC27F\uC280-\uC2FF\uC300-\uC37F\uC380-\uC3FF\uC400-\uC47F\uC480-\uC4FF\uC500-\uC57F\uC580-\uC5FF\uC600-\uC67F\uC680-\uC6FF\uC700-\uC77F\uC780-\uC7FF\uC800-\uC87F\uC880-\uC8FF\uC900-\uC97F\uC980-\uC9FF\uCA00-\uCA7F\uCA80-\uCAFF\uCB00-\uCB7F\uCB80-\uCBFF\uCC00-\uCC7F\uCC80-\uCCFF\uCD00-\uCD7F\uCD80-\uCCBFF\uCE00-\uCE7F\uCE80-\uCEFF\uCF00-\uCF7F\uCF80-\uCFFF\uD000-\uD07F\uD080-\uD0FF\uD100-\uD17F\uD180-\uD1FF\uD200-\uD27F\uD280-\uD2FF\uD300-\uD37F\uD380-\uD3FF\uD400-\uD47F\uD480-\uD4FF\uD500-\uD57F\uD580-\uD5FF\uD600-\uD67F\uD680-\uD6FF\uD700-\uD77F\uD780-\uD7FF\uD800-\uD87F\uD880-\uD8FF\uD900-\uD97F\uD980-\uD9FF\uDA00-\uDA7F\uDA80-\uDAFF\uDB00-\uDB7F\uDB80-\uDBFF\uDC00-\uDC7F\uDC80-\uDCFF\uDD00-\uDD7F\uDD80-\uDCBFF\uDE00-\uDE7F\uDE80-\uDEFF\uDF00-\uDF7F\uDF80-\uDCFFF\uE000-\uE07F\uE080-\uE0FF\uE100-\uE17F\uE180-\uE1FF\uE200-\uE27F\uE280-\uE2FF\uE300-\uE37F\uE380-\uE3FF\uE400-\uE47F\uE480-\uE4FF\uE500-\uE57F\uE580-\uE5FF\uE600-\uE67F\uE680-\uE6FF\uE700-\uE77F\uE780-\uE7FF\uE800-\uE87F\uE880-\uE8FF\uE900-\uE97F\uE980-\uE9FF\uEA00-\uEA7F\uEA80-\uEAFF\uEB00-\uEB7F\uEB80-\uEBFF\uEC00-\uEC7F\uEC80-\uECFF\uED00-\uED7F\uED80-\uEDFF\uEE00-\uEE7F\uEE80-\uEEFF\uEF00-\uEF7F\uEF80-\uEFFF\uF000-\uF07F\uF080-\uF0FF\uF100-\uF17F\uF180-\uF1FF\uF200-\uF27F\uF280-\uF2FF\uF300-\uF37F\uF380-\uF3FF\uF400-\uF47F\uF480-\uF4FF\uF500-\uF57F\uF580-\uF5FF\uF600-\uF67F\uF680-\uF6FF\uF700-\uF77F\uF780-\uF7FF\uF800-\uF87F\uF880-\uF8FF\uF900-\uF97F\uF980-\uF9FF\uFA00-\uFA7F\uFA80-\uFAFF\uFB00-\uFB7F\uFB80-\uFBFF\uFC00-\uFC7F\uFC80-\uFCFF\uFD00-\uFD7F\uFD80-\uFCBFF\uFE00-\uFE7F\uFE80-\uFEFF\uFF00-\uFF7F\uFF80-\uFFFF]`

Required: Yes

ScalableDimension (p. 10)

The scalable dimension associated with the scalable target. This string consists of the service namespace, resource type, and scaling property.

- `ecs:service:DesiredCount` - The desired task count of an ECS service.
- `ec2:spot-fleet-request:TargetCapacity` - The target capacity of a Spot fleet request.
- `elasticmapreduce:instancegroup:InstanceCount` - The instance count of an EMR Instance Group.
- `appstream:fleet:DesiredCapacity` - The desired capacity of an AppStream 2.0 fleet.
- `dynamodb:table:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB table.
- `dynamodb:table:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB table.
- `dynamodb:index:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB global secondary index.
- `dynamodb:index:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB global secondary index.
- `rds:cluster:ReadReplicaCount` - The count of Aurora Replicas in an Aurora DB cluster. Available for Aurora MySQL-compatible edition and Aurora PostgreSQL-compatible edition.
- `sagemaker:variant:DesiredInstanceCount` - The number of EC2 instances for an Amazon SageMaker model endpoint variant.
- `custom-resource:ResourceType:Property` - The scalable dimension for a custom resource provided by your own application or service.

Type: String

Valid Values: `ecs:service:DesiredCount` | `ec2:spot-fleet-request:TargetCapacity` | `elasticmapreduce:instancegroup:InstanceCount` | `appstream:fleet:DesiredCapacity` | `dynamodb:table:ReadCapacityUnits` | `dynamodb:table:WriteCapacityUnits` | `dynamodb:index:ReadCapacityUnits` | `dynamodb:index:WriteCapacityUnits` | `rds:cluster:ReadReplicaCount` | `sagemaker:variant:DesiredInstanceCount` | `custom-resource:ResourceType:Property`

Required: Yes

ServiceNamespace (p. 10)

The namespace of the AWS service that provides the resource or `custom-resource` for a resource provided by your own application or service. For more information, see [AWS Service Namespaces](#) in the *Amazon Web Services General Reference*.

Type: String

Valid Values: `ecs` | `elasticmapreduce` | `ec2` | `appstream` | `dynamodb` | `rds` | `sagemaker` | `custom-resource`

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 Errors \(p. 82\)](#).

ConcurrentUpdateException

Concurrent updates caused an exception, for example, if you request an update to an Application Auto Scaling resource that already has a pending update.

HTTP Status Code: 400

InternalServiceException

The service encountered an internal error.

HTTP Status Code: 400

ObjectNotFoundException

The specified object could not be found. For any operation that depends on the existence of a scalable target, this exception is thrown if the scalable target with the specified service namespace, resource ID, and scalable dimension does not exist. For any operation that deletes or deregisters a resource, this exception is thrown if the resource cannot be found.

HTTP Status Code: 400

ValidationException

An exception was thrown for a validation issue. Review the available parameters for the API request.

HTTP Status Code: 400

Example

If you plan to create requests manually, you must replace the Authorization header contents in the examples (AUTHPARAMS) with a signature. For more information, see [Signature Version 4 Signing Process](#) in the *Amazon Web Services General Reference*. If you plan to use the [AWS CLI](#) or one of the [AWS SDKs](#), these tools sign the requests for you.

Example

The following example deregisters a scalable target for an Amazon ECS service called `web-app` that is running in the `default` cluster.

Sample Request

```
POST / HTTP/1.1
Host: autoscaling.us-west-2.amazonaws.com
Accept-Encoding: identity
Content-Length: 117
X-Amz-Target: AnyScaleFrontendService.DeregisterScalableTarget
X-Amz-Date: 20160506T210150Z
User-Agent: aws-cli/1.10.23 Python/2.7.11 Darwin/15.4.0 botocore/1.4.8
Content-Type: application/x-amz-json-1.1
Authorization: AUTHPARAMS

{
  "ResourceId": "service/default/web-app",
  "ServiceNamespace": "ecs",
  "ScalableDimension": "ecs:service:DesiredCount"
}
```

See Also

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

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DescribeScalableTargets

Gets information about the scalable targets in the specified namespace.

You can filter the results using the `ResourceIds` and `ScalableDimension` parameters.

To create a scalable target or update an existing one, see [RegisterScalableTarget \(p. 50\)](#). If you are no longer using a scalable target, you can deregister it using [DeregisterScalableTarget \(p. 10\)](#).

Request Syntax

```
{  
  "MaxResults": number,  
  "NextToken": "string",  
  "ResourceIds": [ "string" ],  
  "ScalableDimension": "string",  
  "ServiceNamespace": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 80\)](#).

The request accepts the following data in JSON format.

MaxResults (p. 14)

The maximum number of scalable targets. This value can be between 1 and 50. The default value is 50.

If this parameter is used, the operation returns up to `MaxResults` results at a time, along with a `NextToken` value. To get the next set of results, include the `NextToken` value in a subsequent call. If this parameter is not used, the operation returns up to 50 results and a `NextToken` value, if applicable.

Type: Integer

Required: No

NextToken (p. 14)

The token for the next set of results.

Type: String

Pattern: `[\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*`

Required: No

ResourceIds (p. 14)

The identifier of the resource associated with the scalable target. This string consists of the resource type and unique identifier. If you specify a scalable dimension, you must also specify a resource ID.

- ECS service - The resource type is `service` and the unique identifier is the cluster name and service name. Example: `service/default/sample-webapp`.
- Spot fleet request - The resource type is `spot-fleet-request` and the unique identifier is the Spot fleet request ID. Example: `spot-fleet-request/sfr-73fbd2ce-aa30-494c-8788-1cee4EXAMPLE`.

| sagemaker:variant:DesiredInstanceCount | custom-
resource:ResourceType:Property

Required: No

ServiceNamespace (p. 14)

The namespace of the AWS service that provides the resource or custom-resource for a resource provided by your own application or service. For more information, see [AWS Service Namespaces](#) in the *Amazon Web Services General Reference*.

Type: String

Valid Values: ecs | elasticmapreduce | ec2 | appstream | dynamodb | rds |
sagemaker | custom-resource

Required: Yes

Response Syntax

```
{
  "NextToken": "string",
  "ScalableTargets": [
    {
      "CreationTime": number,
      "MaxCapacity": number,
      "MinCapacity": number,
      "ResourceId": "string",
      "RoleARN": "string",
      "ScalableDimension": "string",
      "ServiceNamespace": "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 (p. 16)

The token required to get the next set of results. This value is null if there are no more results to return.

Type: String

Pattern: [\u0020-\uD7FF\uE000-\uFFFD\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*

ScalableTargets (p. 16)

The scalable targets that match the request parameters.

Type: Array of [ScalableTarget \(p. 60\)](#) objects

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 82\)](#).

ConcurrentUpdateException

Concurrent updates caused an exception, for example, if you request an update to an Application Auto Scaling resource that already has a pending update.

HTTP Status Code: 400

InternalServiceException

The service encountered an internal error.

HTTP Status Code: 400

InvalidNextTokenException

The next token supplied was invalid.

HTTP Status Code: 400

ValidationException

An exception was thrown for a validation issue. Review the available parameters for the API request.

HTTP Status Code: 400

Example

If you plan to create requests manually, you must replace the Authorization header contents in the examples (AUTHPARAMS) with a signature. For more information, see [Signature Version 4 Signing Process](#) in the *Amazon Web Services General Reference*. If you plan to use the [AWS CLI](#) or one of the [AWS SDKs](#), these tools sign the requests for you.

Example

The following example describes the scalable targets for the ecs service namespace.

Sample Request

```
POST / HTTP/1.1
Host: autoscaling.us-west-2.amazonaws.com
Accept-Encoding: identity
Content-Length: 27
X-Amz-Target: AnyScaleFrontendService.DescribeScalableTargets
X-Amz-Date: 20160506T184921Z
User-Agent: aws-cli/1.10.23 Python/2.7.11 Darwin/15.4.0 botocore/1.4.8
Content-Type: application/x-amz-json-1.1
Authorization: AUTHPARAMS

{
  "ServiceNamespace": "ecs"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: 3f10dab0-13bb-11e6-a873-676fff004c09
Content-Type: application/x-amz-json-1.1
Content-Length: 272
Date: Fri, 06 May 2016 18:49:21 GMT
```

```
{
  "ScalableTargets": [
    {
      "CreationTime": 1462558906.199,
      "MaxCapacity": 10,
      "MinCapacity": 1,
      "ResourceId": "service/default/web-app",
      "RoleARN": "arn:aws:iam::012345678910:role/ApplicationAutoscalingECSRole",
      "ScalableDimension": "ecs:service:DesiredCount",
      "ServiceNamespace": "ecs"
    }
  ]
}
```

See Also

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

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DescribeScalingActivities

Provides descriptive information about the scaling activities in the specified namespace from the previous six weeks.

You can filter the results using the `ResourceId` and `ScalableDimension` parameters.

Scaling activities are triggered by CloudWatch alarms that are associated with scaling policies. To view the scaling policies for a service namespace, see [DescribeScalingPolicies \(p. 25\)](#). To create a scaling policy or update an existing one, see [PutScalingPolicy \(p. 37\)](#).

Request Syntax

```
{  
  "MaxResults": number,  
  "NextToken": "string",  
  "ResourceId": "string",  
  "ScalableDimension": "string",  
  "ServiceNamespace": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 80\)](#).

The request accepts the following data in JSON format.

MaxResults (p. 19)

The maximum number of scalable targets. This value can be between 1 and 50. The default value is 50.

If this parameter is used, the operation returns up to `MaxResults` results at a time, along with a `NextToken` value. To get the next set of results, include the `NextToken` value in a subsequent call. If this parameter is not used, the operation returns up to 50 results and a `NextToken` value, if applicable.

Type: Integer

Required: No

NextToken (p. 19)

The token for the next set of results.

Type: String

Pattern: `[\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDFFF\x\n\t]*`

Required: No

ResourceId (p. 19)

The identifier of the resource associated with the scaling activity. This string consists of the resource type and unique identifier. If you specify a scalable dimension, you must also specify a resource ID.

- ECS service - The resource type is `service` and the unique identifier is the cluster name and service name. Example: `service/default/sample-webapp`.

- Spot fleet request - The resource type is `spot-fleet-request` and the unique identifier is the Spot fleet request ID. Example: `spot-fleet-request/sfr-73fbd2ce-aa30-494c-8788-1cee4EXAMPLE`.
- EMR cluster - The resource type is `instancegroup` and the unique identifier is the cluster ID and instance group ID. Example: `instancegroup/j-2EEZNYKUA1NTV/ig-1791Y4E1L8YIO`.
- AppStream 2.0 fleet - The resource type is `fleet` and the unique identifier is the fleet name. Example: `fleet/sample-fleet`.
- DynamoDB table - The resource type is `table` and the unique identifier is the resource ID. Example: `table/my-table`.
- DynamoDB global secondary index - The resource type is `index` and the unique identifier is the resource ID. Example: `table/my-table/index/my-table-index`.
- Aurora DB cluster - The resource type is `cluster` and the unique identifier is the cluster name. Example: `cluster:my-db-cluster`.
- Amazon SageMaker endpoint variants - The resource type is `variant` and the unique identifier is the resource ID. Example: `endpoint/my-end-point/variant/KMeansClustering`.
- Custom resources are not supported with a resource type. This parameter must specify the `OutputValue` from the CloudFormation template stack used to access the resources. The unique identifier is defined by the service provider. More information is available in our [GitHub repository](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1600.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\uD800\uDC00-\uDBFF\uDF00-\uDFFF\r\n\t]*`

Required: No

ScalableDimension (p. 19)

The scalable dimension. This string consists of the service namespace, resource type, and scaling property. If you specify a scalable dimension, you must also specify a resource ID.

- `ecs:service:DesiredCount` - The desired task count of an ECS service.
- `ec2:spot-fleet-request:TargetCapacity` - The target capacity of a Spot fleet request.
- `elasticmapreduce:instancegroup:InstanceCount` - The instance count of an EMR Instance Group.
- `appstream:fleet:DesiredCapacity` - The desired capacity of an AppStream 2.0 fleet.
- `dynamodb:table:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB table.
- `dynamodb:table:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB table.
- `dynamodb:index:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB global secondary index.
- `dynamodb:index:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB global secondary index.
- `rds:cluster:ReadReplicaCount` - The count of Aurora Replicas in an Aurora DB cluster. Available for Aurora MySQL-compatible edition and Aurora PostgreSQL-compatible edition.
- `sagemaker:variant:DesiredInstanceCount` - The number of EC2 instances for an Amazon SageMaker model endpoint variant.
- `custom-resource:ResourceType:Property` - The scalable dimension for a custom resource provided by your own application or service.

Type: String

Valid Values: `ecs:service:DesiredCount` | `ec2:spot-fleet-request:TargetCapacity` | `elasticmapreduce:instancegroup:InstanceCount`

```
| appstream:fleet:DesiredCapacity | dynamodb:table:ReadCapacityUnits  
| dynamodb:table:WriteCapacityUnits | dynamodb:index:ReadCapacityUnits  
| dynamodb:index:WriteCapacityUnits | rds:cluster:ReadReplicaCount  
| sagemaker:variant:DesiredInstanceCount | custom-  
resource:ResourceType:Property
```

Required: No

ServiceNamespace (p. 19)

The namespace of the AWS service that provides the resource or `custom-resource` for a resource provided by your own application or service. For more information, see [AWS Service Namespaces](#) in the *Amazon Web Services General Reference*.

Type: String

Valid Values: `ecs` | `elasticmapreduce` | `ec2` | `appstream` | `dynamodb` | `rds` | `sagemaker` | `custom-resource`

Required: Yes

Response Syntax

```
{  
  "NextToken": "string",  
  "ScalingActivities": [  
    {  
      "ActivityId": "string",  
      "Cause": "string",  
      "Description": "string",  
      "Details": "string",  
      "EndTime": number,  
      "ResourceId": "string",  
      "ScalableDimension": "string",  
      "ServiceNamespace": "string",  
      "StartTime": number,  
      "StatusCode": "string",  
      "StatusMessage": "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 (p. 21)

The token required to get the next set of results. This value is `null` if there are no more results to return.

Type: String

Pattern: `[\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDF00-\uDFFF\r\n\t]*`

ScalingActivities (p. 21)

A list of scaling activity objects.

Type: Array of [ScalingActivity](#) (p. 64) objects

Errors

For information about the errors that are common to all actions, see [Common Errors](#) (p. 82).

ConcurrentUpdateException

Concurrent updates caused an exception, for example, if you request an update to an Application Auto Scaling resource that already has a pending update.

HTTP Status Code: 400

InternalServiceException

The service encountered an internal error.

HTTP Status Code: 400

InvalidNextTokenException

The next token supplied was invalid.

HTTP Status Code: 400

ValidationException

An exception was thrown for a validation issue. Review the available parameters for the API request.

HTTP Status Code: 400

Example

If you plan to create requests manually, you must replace the Authorization header contents in the examples (AUTHPARAMS) with a signature. For more information, see [Signature Version 4 Signing Process](#) in the *Amazon Web Services General Reference*. If you plan to use the [AWS CLI](#) or one of the [AWS SDKs](#), these tools sign the requests for you.

Example

The following example describes the scaling activities for an Amazon ECS service called `web-app` that is running in the `default` cluster, and it limits the returned results to 2.

Sample Request

```
POST / HTTP/1.1
Host: autoscaling.us-west-2.amazonaws.com
Accept-Encoding: identity
Content-Length: 117
X-Amz-Target: AnyScaleFrontendService.DescribeScalingActivities
X-Amz-Date: 20160506T224112Z
User-Agent: aws-cli/1.10.26 Python/2.7.11 Darwin/15.4.0 botocore/1.4.8
Content-Type: application/x-amz-json-1.1
Authorization: AUTHPARAMS

{
  "ResourceId": "service/default/web-app",
  "ServiceNamespace": "ecs",
  "ScalableDimension": "ecs:service:DesiredCount"
```

```
}

```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: a2704130-13db-11e6-9fca-039a3edb2541
Content-Type: application/x-amz-json-1.1
Content-Length: 1784
Date: Fri, 06 May 2016 22:41:12 GMT

{
  "ScalingActivities": [
    {
      "ActivityId": "0b812df9-a093-4074-9064-8a8f6c0521f5",
      "Cause": "monitor alarm Step-Scaling-AlarmHigh-ECS:service/default/web-app in state
ALARM triggered policy my-scaleout-policy",
      "Description": "Setting desired count to 3.",
      "ResourceId": "service/default/web-app",
      "ScalableDimension": "ecs:service:DesiredCount",
      "ServiceNamespace": "ecs",
      "StartTime": 1462568034.684,
      "StatusCode": "Pending"
    },
    {
      "ActivityId": "4d759079-a31f-4d0c-8468-504c56e2eecf",
      "Cause": "monitor alarm Step-Scaling-AlarmHigh-ECS:service/default/web-app in state
ALARM triggered policy my-scaleout-policy",
      "Description": "Setting desired count to 3.",
      "EndTime": 1462574276.686,
      "ResourceId": "service/default/web-app",
      "ScalableDimension": "ecs:service:DesiredCount",
      "ServiceNamespace": "ecs",
      "StartTime": 1462574194.658,
      "StatusCode": "Successful",
      "StatusMessage": "Successfully set desired count to 3. Change successfully fulfilled
by ecs."
    },
    {
      "ActivityId": "90aff0eb-dd6a-443c-889b-b809e78061c1",
      "Cause": "monitor alarm Step-Scaling-AlarmHigh-ECS:service/default/web-app in state
ALARM triggered policy my-scaleout-policy",
      "Description": "Setting desired count to 9.",
      "EndTime": 1462574333.492,
      "ResourceId": "service/default/web-app",
      "ScalableDimension": "ecs:service:DesiredCount",
      "ServiceNamespace": "ecs",
      "StartTime": 1462574254.223,
      "StatusCode": "Successful",
      "StatusMessage": "Successfully set desired count to 9. Change successfully fulfilled
by ecs."
    },
    {
      "ActivityId": "ee381679-5079-46b5-ac1a-418253981efd",
      "Cause": "monitor alarm Step-Scaling-AlarmHigh-ECS:service/default/web-app in state
ALARM triggered policy my-scaleout-policy",
      "Description": "Setting desired count to 10.",
      "ResourceId": "service/default/web-app",
      "ScalableDimension": "ecs:service:DesiredCount",
      "ServiceNamespace": "ecs",
      "StartTime": 1462574434.077,
      "StatusCode": "InProgress",
      "StatusMessage": "Successfully set desired count to 10. Waiting for change to be
fulfilled by ecs."
    }
  ]
}
```

```
]
}
```

See Also

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

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DescribeScalingPolicies

Describes the scaling policies for the specified service namespace.

You can filter the results using the `ResourceId`, `ScalableDimension`, and `PolicyNames` parameters.

To create a scaling policy or update an existing one, see [PutScalingPolicy \(p. 37\)](#). If you are no longer using a scaling policy, you can delete it using [DeleteScalingPolicy \(p. 3\)](#).

Request Syntax

```
{  
  "MaxResults": number,  
  "NextToken": "string",  
  "PolicyNames": [ "string" ],  
  "ResourceId": "string",  
  "ScalableDimension": "string",  
  "ServiceNamespace": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 80\)](#).

The request accepts the following data in JSON format.

MaxResults (p. 25)

The maximum number of scalable targets. This value can be between 1 and 50. The default value is 50.

If this parameter is used, the operation returns up to `MaxResults` results at a time, along with a `NextToken` value. To get the next set of results, include the `NextToken` value in a subsequent call. If this parameter is not used, the operation returns up to 50 results and a `NextToken` value, if applicable.

Type: Integer

Required: No

NextToken (p. 25)

The token for the next set of results.

Type: String

Pattern: `[\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*`

Required: No

PolicyNames (p. 25)

The names of the scaling policies to describe.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 1600.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*`

- `custom-resource:ResourceType:Property` - The scalable dimension for a custom resource provided by your own application or service.

Type: String

Valid Values: `ecs:service:DesiredCount` | `ec2:spot-fleet-request:TargetCapacity` | `elasticmapreduce:instancegroup:InstanceCount` | `appstream:fleet:DesiredCapacity` | `dynamodb:table:ReadCapacityUnits` | `dynamodb:table:WriteCapacityUnits` | `dynamodb:index:ReadCapacityUnits` | `dynamodb:index:WriteCapacityUnits` | `rds:cluster:ReadReplicaCount` | `sagemaker:variant:DesiredInstanceCount` | `custom-resource:ResourceType:Property`

Required: No

ServiceNamespace (p. 25)

The namespace of the AWS service that provides the resource or `custom-resource` for a resource provided by your own application or service. For more information, see [AWS Service Namespaces](#) in the *Amazon Web Services General Reference*.

Type: String

Valid Values: `ecs` | `elasticmapreduce` | `ec2` | `appstream` | `dynamodb` | `rds` | `sagemaker` | `custom-resource`

Required: Yes

Response Syntax

```
{
  "NextToken": "string",
  "ScalingPolicies": [
    {
      "Alarms": [
        {
          "AlarmARN": "string",
          "AlarmName": "string"
        }
      ],
      "CreationTime": number,
      "PolicyARN": "string",
      "PolicyName": "string",
      "PolicyType": "string",
      "ResourceId": "string",
      "ScalableDimension": "string",
      "ServiceNamespace": "string",
      "StepScalingPolicyConfiguration": {
        "AdjustmentType": "string",
        "Cooldown": number,
        "MetricAggregationType": "string",
        "MinAdjustmentMagnitude": number,
        "StepAdjustments": [
          {
            "MetricIntervalLowerBound": number,
            "MetricIntervalUpperBound": number,
            "ScalingAdjustment": number
          }
        ]
      }
    },
    "TargetTrackingScalingPolicyConfiguration": {
```



```
    "CustomizedMetricSpecification": {
      "Dimensions": [
        {
          "Name": "string",
          "Value": "string"
        }
      ],
      "MetricName": "string",
      "Namespace": "string",
      "Statistic": "string",
      "Unit": "string"
    },
    "DisableScaleIn": boolean,
    "PredefinedMetricSpecification": {
      "PredefinedMetricType": "string",
      "ResourceLabel": "string"
    },
    "ScaleInCooldown": number,
    "ScaleOutCooldown": number,
    "TargetValue": 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.

NextToken (p. 27)

The token required to get the next set of results. This value is null if there are no more results to return.

Type: String

Pattern: [\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*

ScalingPolicies (p. 27)

Information about the scaling policies.

Type: Array of [ScalingPolicy \(p. 67\)](#) objects

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 82\)](#).

ConcurrentUpdateException

Concurrent updates caused an exception, for example, if you request an update to an Application Auto Scaling resource that already has a pending update.

HTTP Status Code: 400

FailedResourceAccessException

Failed access to resources caused an exception. This exception is thrown when Application Auto Scaling is unable to retrieve the alarms associated with a scaling policy due to a client error,

for example, if the role ARN specified for a scalable target does not have permission to call the CloudWatch [DescribeAlarms](#) on your behalf.

HTTP Status Code: 400

InternalServerErrorException

The service encountered an internal error.

HTTP Status Code: 400

InvalidNextTokenException

The next token supplied was invalid.

HTTP Status Code: 400

ValidationException

An exception was thrown for a validation issue. Review the available parameters for the API request.

HTTP Status Code: 400

Example

If you plan to create requests manually, you must replace the Authorization header contents in the examples (AUTHPARAMS) with a signature. For more information, see [Signature Version 4 Signing Process](#) in the *Amazon Web Services General Reference*. If you plan to use the [AWS CLI](#) or one of the [AWS SDKs](#), these tools sign the requests for you.

Example

The following example describes the scaling policies for the `ecs` service namespace.

Sample Request

```
POST / HTTP/1.1
Host: autoscaling.us-west-2.amazonaws.com
Accept-Encoding: identity
Content-Length: 27
X-Amz-Target: AnyScaleFrontendService.DescribeScalingPolicies
X-Amz-Date: 20160506T194435Z
User-Agent: aws-cli/1.10.23 Python/2.7.11 Darwin/15.4.0 botocore/1.4.8
Content-Type: application/x-amz-json-1.1
Authorization: AUTHPARAMS

{
  "ServiceNamespace": "ecs"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: f662c515-13c2-11e6-add4-41b78770ca43
Content-Type: application/x-amz-json-1.1
Content-Length: 1393
Date: Fri, 06 May 2016 19:44:35 GMT

{
```

Application Auto Scaling API Reference Example

```
"ScalingPolicies": [
  {
    "Alarms": [
      {
        "AlarmARN": "arn:aws:cloudwatch:us-west-2:012345678910:alarm:step-
scaling-alarmhigh-ecs:service/default/web-app",
        "AlarmName": "Step-Scaling-AlarmHigh-ECS:service/default/web-app"
      }
    ],
    "CreationTime": 1462561899.23,
    "PolicyARN": "arn:aws:autoscaling:us-
west-2:012345678910:scalingPolicy:ac542982-cbeb-4294-891c-a5a941dfa787:resource/ecs/
service/default/web-app:policyName/my-scaleout-policy",
    "PolicyName": "my-scaleout-policy",
    "PolicyType": "StepScaling",
    "ResourceId": "service/default/web-app",
    "ScalableDimension": "ecs:service:DesiredCount",
    "ServiceNamespace": "ecs",
    "StepScalingPolicyConfiguration": {
      "AdjustmentType": "PercentChangeInCapacity",
      "Cooldown": 60,
      "MetricAggregationType": "Average",
      "StepAdjustments": [
        {
          "MetricIntervalLowerBound": 0,
          "ScalingAdjustment": 200
        }
      ]
    }
  },
  {
    "Alarms": [
      {
        "AlarmARN": "arn:aws:cloudwatch:us-west-2:012345678910:alarm:step-
scaling-alarmlow-ecs:service/default/web-app",
        "AlarmName": "Step-Scaling-AlarmLow-ECS:service/default/web-app"
      }
    ],
    "CreationTime": 1462562575.099,
    "PolicyARN": "arn:aws:autoscaling:us-
west-2:012345678910:scalingPolicy:6d8972f3-efc8-437c-92d1-6270f29a66e7:resource/ecs/
service/default/web-app:policyName/my-scalein-policy",
    "PolicyName": "my-scalein-policy",
    "PolicyType": "StepScaling",
    "ResourceId": "service/default/web-app",
    "ScalableDimension": "ecs:service:DesiredCount",
    "ServiceNamespace": "ecs",
    "StepScalingPolicyConfiguration": {
      "AdjustmentType": "PercentChangeInCapacity",
      "Cooldown": 120,
      "MetricAggregationType": "Average",
      "MinAdjustmentMagnitude": 1,
      "StepAdjustments": [
        {
          "MetricIntervalLowerBound": -15,
          "MetricIntervalUpperBound": 0,
          "ScalingAdjustment": -25,
        },
        {
          "MetricIntervalUpperBound": -15,
          "ScalingAdjustment": -50
        }
      ]
    }
  }
]
```

}

See Also

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

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DescribeScheduledActions

Describes the scheduled actions for the specified service namespace.

You can filter the results using the `ResourceId`, `ScalableDimension`, and `ScheduledActionNames` parameters.

To create a scheduled action or update an existing one, see [PutScheduledAction \(p. 45\)](#). If you are no longer using a scheduled action, you can delete it using [DeleteScheduledAction \(p. 7\)](#).

Request Syntax

```
{
  "MaxResults": number,
  "NextToken": "string",
  "ResourceId": "string",
  "ScalableDimension": "string",
  "ScheduledActionNames": [ "string" ],
  "ServiceNamespace": "string"
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 80\)](#).

The request accepts the following data in JSON format.

MaxResults (p. 32)

The maximum number of scheduled action results. This value can be between 1 and 50. The default value is 50.

If this parameter is used, the operation returns up to `MaxResults` results at a time, along with a `NextToken` value. To get the next set of results, include the `NextToken` value in a subsequent call. If this parameter is not used, the operation returns up to 50 results and a `NextToken` value, if applicable.

Type: Integer

Required: No

NextToken (p. 32)

The token for the next set of results.

Type: String

Pattern: `[\u0020-\u007F\u00E000-\u00FFFD\u0800\uDC00-\uDBFF\uDFFF\r\n\t]*`

Required: No

ResourceId (p. 32)

The identifier of the resource associated with the scheduled action. This string consists of the resource type and unique identifier. If you specify a scalable dimension, you must also specify a resource ID.

- ECS service - The resource type is `service` and the unique identifier is the cluster name and service name. Example: `service/default/sample-webapp`.


```
| dynamodb:table:WriteCapacityUnits | dynamodb:index:ReadCapacityUnits  
| dynamodb:index:WriteCapacityUnits | rds:cluster:ReadReplicaCount  
| sagemaker:variant:DesiredInstanceCount | custom-  
resource:ResourceType:Property
```

Required: No

ScheduledActionNames (p. 32)

The names of the scheduled actions to describe.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 1600.

Pattern: [\\u0020-\\uD7FF\\uE000-\\uFFFD\\uD800\\uDC00-\\uDBFF\\uDFFF\\r\\n\\t]*

Required: No

ServiceNamespace (p. 32)

The namespace of the AWS service that provides the resource or `custom-resource` for a resource provided by your own application or service. For more information, see [AWS Service Namespaces](#) in the *Amazon Web Services General Reference*.

Type: String

Valid Values: `ecs` | `elasticmapreduce` | `ec2` | `appstream` | `dynamodb` | `rds` | `sagemaker` | `custom-resource`

Required: Yes

Response Syntax

```
{  
  "NextToken": "string",  
  "ScheduledActions": [  
    {  
      "CreationTime": number,  
      "EndTime": number,  
      "ResourceId": "string",  
      "ScalableDimension": "string",  
      "ScalableTargetAction": {  
        "MaxCapacity": number,  
        "MinCapacity": number  
      },  
      "Schedule": "string",  
      "ScheduledActionARN": "string",  
      "ScheduledActionName": "string",  
      "ServiceNamespace": "string",  
      "StartTime": 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.

[NextToken \(p. 34\)](#)

The token required to get the next set of results. This value is null if there are no more results to return.

Type: String

Pattern: [\u0020-\uD7FF\uE000-\uFFFD\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*

[ScheduledActions \(p. 34\)](#)

Information about the scheduled actions.

Type: Array of [ScheduledAction \(p. 70\)](#) objects

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 82\)](#).

ConcurrentUpdateException

Concurrent updates caused an exception, for example, if you request an update to an Application Auto Scaling resource that already has a pending update.

HTTP Status Code: 400

InternalServiceException

The service encountered an internal error.

HTTP Status Code: 400

InvalidNextTokenException

The next token supplied was invalid.

HTTP Status Code: 400

ValidationException

An exception was thrown for a validation issue. Review the available parameters for the API 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](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

PutScalingPolicy

Creates or updates a policy for an Application Auto Scaling scalable target.

Each scalable target is identified by a service namespace, resource ID, and scalable dimension. A scaling policy applies to the scalable target identified by those three attributes. You cannot create a scaling policy until you have registered the resource as a scalable target using [RegisterScalableTarget](#) (p. 50).

To update a policy, specify its policy name and the parameters that you want to change. Any parameters that you don't specify are not changed by this update request.

You can view the scaling policies for a service namespace using [DescribeScalingPolicies](#) (p. 25). If you are no longer using a scaling policy, you can delete it using [DeleteScalingPolicy](#) (p. 3).

Multiple scaling policies can be in force at the same time for the same scalable target. You can have one or more target tracking scaling policies, one or more step scaling policies, or both. However, there is a chance that multiple policies could conflict, instructing the scalable target to scale out or in at the same time. Application Auto Scaling gives precedence to the policy that provides the largest capacity for both scale in and scale out. For example, if one policy increases capacity by 3, another policy increases capacity by 200 percent, and the current capacity is 10, Application Auto Scaling uses the policy with the highest calculated capacity (200% of 10 = 20) and scales out to 30.

Learn more about how to work with scaling policies in the [Application Auto Scaling User Guide](#).

Request Syntax

```
{
  "PolicyName": "string",
  "PolicyType": "string",
  "ResourceId": "string",
  "ScalableDimension": "string",
  "ServiceNamespace": "string",
  "StepScalingPolicyConfiguration": {
    "AdjustmentType": "string",
    "Cooldown": number,
    "MetricAggregationType": "string",
    "MinAdjustmentMagnitude": number,
    "StepAdjustments": [
      {
        "MetricIntervalLowerBound": number,
        "MetricIntervalUpperBound": number,
        "ScalingAdjustment": number
      }
    ]
  },
  "TargetTrackingScalingPolicyConfiguration": {
    "CustomizedMetricSpecification": {
      "Dimensions": [
        {
          "Name": "string",
          "Value": "string"
        }
      ],
      "MetricName": "string",
      "Namespace": "string",
      "Statistic": "string",
      "Unit": "string"
    },
    "DisableScaleIn": boolean,
    "PredefinedMetricSpecification": {
      "PredefinedMetricType": "string",
```

```
    "ResourceLabel": "string"  
  },  
  "ScaleInCooldown": number,  
  "ScaleOutCooldown": number,  
  "TargetValue": number  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 80\)](#).

The request accepts the following data in JSON format.

PolicyName (p. 37)

The name of the scaling policy.

Type: String

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

Pattern: `\p{Print}+`

Required: Yes

PolicyType (p. 37)

The policy type. This parameter is required if you are creating a scaling policy.

For information on which services do not support `StepScaling` or `TargetTrackingScaling`, see the information about **Limits** in [Step Scaling Policies](#) and [Target Tracking Scaling Policies](#) in the *Application Auto Scaling User Guide*.

Type: String

Valid Values: `StepScaling` | `TargetTrackingScaling`

Required: No

ResourceId (p. 37)

The identifier of the resource associated with the scaling policy. This string consists of the resource type and unique identifier.

- ECS service - The resource type is `service` and the unique identifier is the cluster name and service name. Example: `service/default/sample-webapp`.
- Spot fleet request - The resource type is `spot-fleet-request` and the unique identifier is the Spot fleet request ID. Example: `spot-fleet-request/sfr-73fbd2ce-aa30-494c-8788-1cee4EXAMPLE`.
- EMR cluster - The resource type is `instancegroup` and the unique identifier is the cluster ID and instance group ID. Example: `instancegroup/j-2EEZNYKUA1NTV/ig-1791Y4E1L8YI0`.
- AppStream 2.0 fleet - The resource type is `fleet` and the unique identifier is the fleet name. Example: `fleet/sample-fleet`.
- DynamoDB table - The resource type is `table` and the unique identifier is the resource ID. Example: `table/my-table`.
- DynamoDB global secondary index - The resource type is `index` and the unique identifier is the resource ID. Example: `table/my-table/index/my-table-index`.
- Aurora DB cluster - The resource type is `cluster` and the unique identifier is the cluster name. Example: `cluster:my-db-cluster`.

- Amazon SageMaker endpoint variants - The resource type is `variant` and the unique identifier is the resource ID. Example: `endpoint/my-end-point/variant/KMeansClustering`.
- Custom resources are not supported with a resource type. This parameter must specify the `OutputValue` from the CloudFormation template stack used to access the resources. The unique identifier is defined by the service provider. More information is available in our [GitHub repository](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1600.

Pattern: `[\u0020-\u007F\u00E0-\u00FF\u0100-\u017F\u0180-\u01FF\u0200-\u02FF\u0300-\u037F\u0380-\u03FF\u0400-\u047F\u0480-\u04FF\u0500-\u057F\u0580-\u05FF\u0600-\u06FF\u0700-\u077F\u0780-\u07FF\u0800-\u087F\u0880-\u08FF\u0900-\u097F\u0980-\u09FF\u0A00-\u0A7F\u0A80-\u0AFF\u0B00-\u0B7F\u0B80-\u0BFF\u0C00-\u0C7F\u0C80-\u0CFF\u0D00-\u0D7F\u0D80-\u0DBF\u0DC0-\u0DBF\u0DF0-\u0DF7\u0E00-\u0E7F\u0E80-\u0EFF\u0F00-\u0F7F\u0F80-\u0FFF\u1000-\u107F\u1080-\u10FF\u1100-\u117F\u1180-\u11FF\u1200-\u127F\u1280-\u12FF\u1300-\u137F\u1380-\u13FF\u1400-\u147F\u1480-\u14FF\u1500-\u157F\u1580-\u15FF\u1600-\u167F\u1680-\u16FF\u1700-\u177F\u1780-\u17FF\u1800-\u187F\u1880-\u18FF\u1900-\u197F\u1980-\u19FF\u1A00-\u1A7F\u1A80-\u1AFF\u1B00-\u1B7F\u1B80-\u1BFF\u1C00-\u1C7F\u1C80-\u1CFF\u1D00-\u1D7F\u1D80-\u1DBF\u1E00-\u1E7F\u1E80-\u1EFF\u1F00-\u1F7F\u1F80-\u1FFF\u2000-\u207F\u2080-\u20FF\u2100-\u217F\u2180-\u21FF\u2200-\u227F\u2280-\u22FF\u2300-\u237F\u2380-\u23FF\u2400-\u247F\u2480-\u24FF\u2500-\u257F\u2580-\u25FF\u2600-\u267F\u2680-\u26FF\u2700-\u277F\u2780-\u27FF\u2800-\u287F\u2880-\u28FF\u2900-\u297F\u2980-\u29FF\u2A00-\u2A7F\u2A80-\u2AFF\u2B00-\u2B7F\u2B80-\u2BFF\u2C00-\u2C7F\u2C80-\u2CFF\u2D00-\u2D7F\u2D80-\u2DBF\u2E00-\u2E7F\u2E80-\u2EFF\u2F00-\u2F7F\u2F80-\u2FFF\u3000-\u307F\u3080-\u30FF\u3100-\u317F\u3180-\u31FF\u3200-\u327F\u3280-\u32FF\u3300-\u337F\u3380-\u33FF\u3400-\u347F\u3480-\u34FF\u3500-\u357F\u3580-\u35FF\u3600-\u367F\u3680-\u36FF\u3700-\u377F\u3780-\u37FF\u3800-\u387F\u3880-\u38FF\u3900-\u397F\u3980-\u39FF\u3A00-\u3A7F\u3A80-\u3AFF\u3B00-\u3B7F\u3B80-\u3BFF\u3C00-\u3C7F\u3C80-\u3CFF\u3D00-\u3D7F\u3D80-\u3DBF\u3E00-\u3E7F\u3E80-\u3EFF\u3F00-\u3F7F\u3F80-\u3FFF\u4000-\u407F\u4080-\u40FF\u4100-\u417F\u4180-\u41FF\u4200-\u427F\u4280-\u42FF\u4300-\u437F\u4380-\u43FF\u4400-\u447F\u4480-\u44FF\u4500-\u457F\u4580-\u45FF\u4600-\u467F\u4680-\u46FF\u4700-\u477F\u4780-\u47FF\u4800-\u487F\u4880-\u48FF\u4900-\u497F\u4980-\u49FF\u4A00-\u4A7F\u4A80-\u4AFF\u4B00-\u4B7F\u4B80-\u4BFF\u4C00-\u4C7F\u4C80-\u4CFF\u4D00-\u4D7F\u4D80-\u4DBF\u4E00-\u4E7F\u4E80-\u4EFF\u4F00-\u4F7F\u4F80-\u4FFF\u5000-\u507F\u5080-\u50FF\u5100-\u517F\u5180-\u51FF\u5200-\u527F\u5280-\u52FF\u5300-\u537F\u5380-\u53FF\u5400-\u547F\u5480-\u54FF\u5500-\u557F\u5580-\u55FF\u5600-\u567F\u5680-\u56FF\u5700-\u577F\u5780-\u57FF\u5800-\u587F\u5880-\u58FF\u5900-\u597F\u5980-\u59FF\u5A00-\u5A7F\u5A80-\u5AFF\u5B00-\u5B7F\u5B80-\u5BFF\u5C00-\u5C7F\u5C80-\u5CFF\u5D00-\u5D7F\u5D80-\u5DBF\u5E00-\u5E7F\u5E80-\u5EFF\u5F00-\u5F7F\u5F80-\u5FFF\u6000-\u607F\u6080-\u60FF\u6100-\u617F\u6180-\u61FF\u6200-\u627F\u6280-\u62FF\u6300-\u637F\u6380-\u63FF\u6400-\u647F\u6480-\u64FF\u6500-\u657F\u6580-\u65FF\u6600-\u667F\u6680-\u66FF\u6700-\u677F\u6780-\u67FF\u6800-\u687F\u6880-\u68FF\u6900-\u697F\u6980-\u69FF\u6A00-\u6A7F\u6A80-\u6AFF\u6B00-\u6B7F\u6B80-\u6BFF\u6C00-\u6C7F\u6C80-\u6CFF\u6D00-\u6D7F\u6D80-\u6DBF\u6E00-\u6E7F\u6E80-\u6EFF\u6F00-\u6F7F\u6F80-\u6FFF\u7000-\u707F\u7080-\u70FF\u7100-\u717F\u7180-\u71FF\u7200-\u727F\u7280-\u72FF\u7300-\u737F\u7380-\u73FF\u7400-\u747F\u7480-\u74FF\u7500-\u757F\u7580-\u75FF\u7600-\u767F\u7680-\u76FF\u7700-\u777F\u7780-\u77FF\u7800-\u787F\u7880-\u78FF\u7900-\u797F\u7980-\u79FF\u7A00-\u7A7F\u7A80-\u7AFF\u7B00-\u7B7F\u7B80-\u7BFF\u7C00-\u7C7F\u7C80-\u7CFF\u7D00-\u7D7F\u7D80-\u7DBF\u7E00-\u7E7F\u7E80-\u7EFF\u7F00-\u7F7F\u7F80-\u7FFF\u8000-\u807F\u8080-\u80FF\u8100-\u817F\u8180-\u81FF\u8200-\u827F\u8280-\u82FF\u8300-\u837F\u8380-\u83FF\u8400-\u847F\u8480-\u84FF\u8500-\u857F\u8580-\u85FF\u8600-\u867F\u8680-\u86FF\u8700-\u877F\u8780-\u87FF\u8800-\u887F\u8880-\u88FF\u8900-\u897F\u8980-\u89FF\u8A00-\u8A7F\u8A80-\u8AFF\u8B00-\u8B7F\u8B80-\u8BFF\u8C00-\u8C7F\u8C80-\u8CFF\u8D00-\u8D7F\u8D80-\u8DBF\u8E00-\u8E7F\u8E80-\u8EFF\u8F00-\u8F7F\u8F80-\u8FFF\u9000-\u907F\u9080-\u90FF\u9100-\u917F\u9180-\u91FF\u9200-\u927F\u9280-\u92FF\u9300-\u937F\u9380-\u93FF\u9400-\u947F\u9480-\u94FF\u9500-\u957F\u9580-\u95FF\u9600-\u967F\u9680-\u96FF\u9700-\u977F\u9780-\u97FF\u9800-\u987F\u9880-\u98FF\u9900-\u997F\u9980-\u99FF\u9A00-\u9A7F\u9A80-\u9AFF\u9B00-\u9B7F\u9B80-\u9BFF\u9C00-\u9C7F\u9C80-\u9CFF\u9D00-\u9D7F\u9D80-\u9DBF\u9E00-\u9E7F\u9E80-\u9EFF\u9F00-\u9F7F\u9F80-\u9FFF\uA000-\uA07F\uA080-\uA0FF\uA100-\uA17F\uA180-\uA1FF\uA200-\uA27F\uA280-\uA2FF\uA300-\uA37F\uA380-\uA3FF\uA400-\uA47F\uA480-\uA4FF\uA500-\uA57F\uA580-\uA5FF\uA600-\uA67F\uA680-\uA6FF\uA700-\uA77F\uA780-\uA7FF\uA800-\uA87F\uA880-\uA8FF\uA900-\uA97F\uA980-\uA9FF\uAA00-\uAA7F\uAA80-\uAAFF\uAB00-\uAB7F\uAB80-\uABFF\uAC00-\uAC7F\uAC80-\uACFF\uAD00-\uAD7F\uAD80-\uADB\uAEE00-\uAEE7F\uAEE80-\uAEEFF\uAF00-\uAF7F\uAF80-\uAFFF\uB000-\uB07F\uB080-\uB0FF\uB100-\uB17F\uB180-\uB1FF\uB200-\uB27F\uB280-\uB2FF\uB300-\uB37F\uB380-\uB3FF\uB400-\uB47F\uB480-\uB4FF\uB500-\uB57F\uB580-\uB5FF\uB600-\uB67F\uB680-\uB6FF\uB700-\uB77F\uB780-\uB7FF\uB800-\uB87F\uB880-\uB8FF\uB900-\uB97F\uB980-\uB9FF\uBA00-\uBA7F\uBA80-\uBAFF\uBB00-\uBB7F\uBB80-\uBBFF\uBC00-\uBC7F\uBC80-\uBCFF\uBD00-\uBD7F\uBD80-\uBDB\uBE00-\uBE7F\uBE80-\uBEFF\uBF00-\uBF7F\uBF80-\uBFFF\uC000-\uC07F\uC080-\uC0FF\uC100-\uC17F\uC180-\uC1FF\uC200-\uC27F\uC280-\uC2FF\uC300-\uC37F\uC380-\uC3FF\uC400-\uC47F\uC480-\uC4FF\uC500-\uC57F\uC580-\uC5FF\uC600-\uC67F\uC680-\uC6FF\uC700-\uC77F\uC780-\uC7FF\uC800-\uC87F\uC880-\uC8FF\uC900-\uC97F\uC980-\uC9FF\uCA00-\uCA7F\uCA80-\uCAFF\uCB00-\uCB7F\uCB80-\uCBFF\uCC00-\uCC7F\uCC80-\uCCFF\uCD00-\uCD7F\uCD80-\uCDFF\uCE00-\uCE7F\uCE80-\uCEFF\uCF00-\uCF7F\uCF80-\uCFFF\uD000-\uD07F\uD080-\uD0FF\uD100-\uD17F\uD180-\uD1FF\uD200-\uD27F\uD280-\uD2FF\uD300-\uD37F\uD380-\uD3FF\uD400-\uD47F\uD480-\uD4FF\uD500-\uD57F\uD580-\uD5FF\uD600-\uD67F\uD680-\uD6FF\uD700-\uD77F\uD780-\uD7FF\uD800-\uD87F\uD880-\uD8FF\uD900-\uD97F\uD980-\uD9FF\uDA00-\uDA7F\uDA80-\uDAFF\uDB00-\uDB7F\uDB80-\uDBFF\uDC00-\uDC7F\uDC80-\uDCFF\uDD00-\uDD7F\uDD80-\uDDFF\uDE00-\uDE7F\uDE80-\uDEFF\uDF00-\uDF7F\uDF80-\uDFFF\uE000-\uE07F\uE080-\uE0FF\uE100-\uE17F\uE180-\uE1FF\uE200-\uE27F\uE280-\uE2FF\uE300-\uE37F\uE380-\uE3FF\uE400-\uE47F\uE480-\uE4FF\uE500-\uE57F\uE580-\uE5FF\uE600-\uE67F\uE680-\uE6FF\uE700-\uE77F\uE780-\uE7FF\uE800-\uE87F\uE880-\uE8FF\uE900-\uE97F\uE980-\uE9FF\uEA00-\uEA7F\uEA80-\uEAFF\uEB00-\uEB7F\uEB80-\uEBFF\uEC00-\uEC7F\uEC80-\uECFF\uED00-\uED7F\uED80-\uEDFF\uEE00-\uEE7F\uEE80-\uEEFF\uEF00-\uEF7F\uEF80-\uEFFF\uF000-\uF07F\uF080-\uF0FF\uF100-\uF17F\uF180-\uF1FF\uF200-\uF27F\uF280-\uF2FF\uF300-\uF37F\uF380-\uF3FF\uF400-\uF47F\uF480-\uF4FF\uF500-\uF57F\uF580-\uF5FF\uF600-\uF67F\uF680-\uF6FF\uF700-\uF77F\uF780-\uF7FF\uF800-\uF87F\uF880-\uF8FF\uF900-\uF97F\uF980-\uF9FF\uFA00-\uFA7F\uFA80-\uFAFF\uFB00-\uFB7F\uFB80-\uFBFF\uFC00-\uFC7F\uFC80-\uFCFF\uFD00-\uFD7F\uFD80-\uFDFF\uFE00-\uFE7F\uFE80-\uFEFF\uFF00-\uFF7F\uFF80-\uFFFF]`*

Required: Yes

ScalableDimension (p. 37)

The scalable dimension. This string consists of the service namespace, resource type, and scaling property.

- `ecs:service:DesiredCount` - The desired task count of an ECS service.
- `ec2:spot-fleet-request:TargetCapacity` - The target capacity of a Spot fleet request.
- `elasticmapreduce:instancegroup:InstanceCount` - The instance count of an EMR Instance Group.
- `appstream:fleet:DesiredCapacity` - The desired capacity of an AppStream 2.0 fleet.
- `dynamodb:table:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB table.
- `dynamodb:table:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB table.
- `dynamodb:index:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB global secondary index.
- `dynamodb:index:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB global secondary index.
- `rds:cluster:ReadReplicaCount` - The count of Aurora Replicas in an Aurora DB cluster. Available for Aurora MySQL-compatible edition and Aurora PostgreSQL-compatible edition.
- `sagemaker:variant:DesiredInstanceCount` - The number of EC2 instances for an Amazon SageMaker model endpoint variant.
- `custom-resource:ResourceType:Property` - The scalable dimension for a custom resource provided by your own application or service.

Type: String

Valid Values: `ecs:service:DesiredCount | ec2:spot-fleet-request:TargetCapacity | elasticmapreduce:instancegroup:InstanceCount | appstream:fleet:DesiredCapacity | dynamodb:table:ReadCapacityUnits | dynamodb:table:WriteCapacityUnits | dynamodb:index:ReadCapacityUnits | dynamodb:index:WriteCapacityUnits | rds:cluster:ReadReplicaCount | sagemaker:variant:DesiredInstanceCount | custom-resource:ResourceType:Property`

Required: Yes

ServiceNamespace (p. 37)

The namespace of the AWS service that provides the resource or `custom-resource` for a resource provided by your own application or service. For more information, see [AWS Service Namespaces](#) in the *Amazon Web Services General Reference*.

Type: String

Valid Values: `ecs | elasticmapreduce | ec2 | appstream | dynamodb | rds | sagemaker | custom-resource`

Required: Yes

[StepScalingPolicyConfiguration \(p. 37\)](#)

A step scaling policy.

This parameter is required if you are creating a policy and the policy type is `StepScaling`.

Type: [StepScalingPolicyConfiguration \(p. 76\)](#) object

Required: No

[TargetTrackingScalingPolicyConfiguration \(p. 37\)](#)

A target tracking scaling policy. Includes support for predefined or customized metrics.

This parameter is required if you are creating a policy and the policy type is `TargetTrackingScaling`.

Type: [TargetTrackingScalingPolicyConfiguration \(p. 78\)](#) object

Required: No

Response Syntax

```
{
  "Alarms": [
    {
      "AlarmARN": "string",
      "AlarmName": "string"
    }
  ],
  "PolicyARN": "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.

[Alarms \(p. 40\)](#)

The CloudWatch alarms created for the target tracking scaling policy.

Type: Array of [Alarm \(p. 55\)](#) objects

[PolicyARN \(p. 40\)](#)

The Amazon Resource Name (ARN) of the resulting scaling policy.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1600.

Pattern: [\u0020-\uD7FF\uE000-\uFFFD\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 82\)](#).

ConcurrentUpdateException

Concurrent updates caused an exception, for example, if you request an update to an Application Auto Scaling resource that already has a pending update.

HTTP Status Code: 400

FailedResourceAccessException

Failed access to resources caused an exception. This exception is thrown when Application Auto Scaling is unable to retrieve the alarms associated with a scaling policy due to a client error, for example, if the role ARN specified for a scalable target does not have permission to call the CloudWatch [DescribeAlarms](#) on your behalf.

HTTP Status Code: 400

InternalServiceException

The service encountered an internal error.

HTTP Status Code: 400

LimitExceededException

A per-account resource limit is exceeded. For more information, see [Application Auto Scaling Limits](#).

HTTP Status Code: 400

ObjectNotFoundException

The specified object could not be found. For any operation that depends on the existence of a scalable target, this exception is thrown if the scalable target with the specified service namespace, resource ID, and scalable dimension does not exist. For any operation that deletes or deregisters a resource, this exception is thrown if the resource cannot be found.

HTTP Status Code: 400

ValidationException

An exception was thrown for a validation issue. Review the available parameters for the API request.

HTTP Status Code: 400

Examples

If you plan to create requests manually, you must replace the Authorization header contents in the examples (AUTHPARAMS) with a signature. For more information, see [Signature Version 4 Signing Process](#) in the *Amazon Web Services General Reference*. If you plan to use the [AWS CLI](#) or one of the [AWS SDKs](#), these tools sign the requests for you.

Example of a target tracking scaling policy

The following example applies a target tracking scaling policy to an Amazon ECS service called `web-app` in the `default` cluster. The policy keeps the average CPU utilization of the service at 75 percent, with scale-out and scale-in cooldown periods of 60 seconds. The output contains the ARNs and names of the two CloudWatch alarms created on your behalf.

Sample Request

```
POST / HTTP/1.1
Host: autoscaling.us-west-2.amazonaws.com
Accept-Encoding: identity
```

```
Content-Length: 392
X-Amz-Target: AnyScaleFrontendService.PutScalingPolicy
X-Amz-Date: 20160506T191044Z
User-Agent: aws-cli/1.10.23 Python/2.7.11 Darwin/15.4.0 botocore/1.4.8
Content-Type: application/x-amz-json-1.1
Authorization: AUTHPARAMS
```

```
{
  "PolicyName": "cpu75-target-tracking-scaling-policy",
  "PolicyType": "TargetTrackingScaling",
  "TargetTrackingScalingPolicyConfiguration": {
    "TargetValue": 75.0,
    "PredefinedMetricSpecification": {
      "PredefinedMetricType": "ECSServiceAverageCPUUtilization"
    },
    "ScaleOutCooldown": 60,
    "ScaleInCooldown": 60
  },
  "ServiceNamespace": "ecs",
  "ScalableDimension": "ecs:service:DesiredCount",
  "ResourceId": "service/default/web-app"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: 4a0f8f18-cb5f-11e0-8364-37acb4b5a1b2
Content-Type: application/x-amz-json-1.1
Content-Length: 314
Date: Fri, 06 May 2016 19:10:44 GMT

{
  "PolicyARN": "arn:aws:autoscaling:us-west-2:012345678910:scalingPolicy:6d8972f3-efc8-437c-92d1-6270f29a66e7:resource/ecs/service/default/web-app:policyName/cpu75-target-tracking-scaling-policy",
  "Alarms": [
    {
      "AlarmARN": "arn:aws:cloudwatch:us-west-2:012345678910:alarm:TargetTracking-service/default/web-app-AlarmHigh-d4f0770c-b46e-434a-a60f-3b36d653feca",
      "AlarmName": "TargetTracking-service/default/web-app-AlarmHigh-d4f0770c-b46e-434a-a60f-3b36d653feca"
    },
    {
      "AlarmARN": "arn:aws:cloudwatch:us-west-2:012345678910:alarm:TargetTracking-service/default/web-app-AlarmLow-1b437334-d19b-4a63-a812-6c67aaf2910d",
      "AlarmName": "TargetTracking-service/default/web-app-AlarmLow-1b437334-d19b-4a63-a812-6c67aaf2910d"
    }
  ]
}
```

Example of a step scaling policy for scale out

The following example applies a step scaling policy to an Amazon ECS service called `web-app` in the default cluster. The policy increases the desired count of the service by 200%, with a cooldown period of 60 seconds. The output includes the ARN for the policy, which you need to create the CloudWatch alarm.

Sample Request

```
POST / HTTP/1.1
```

```
Host: autoscaling.us-west-2.amazonaws.com
Accept-Encoding: identity
Content-Length: 358
X-Amz-Target: AnyScaleFrontendService.PutScalingPolicy
X-Amz-Date: 20160506T191138Z
User-Agent: aws-cli/1.10.23 Python/2.7.11 Darwin/15.4.0 botocore/1.4.8
Content-Type: application/x-amz-json-1.1
Authorization: AUTHPARAMS

{
  "PolicyName": "my-scaleout-policy",
  "PolicyType": "StepScaling",
  "StepScalingPolicyConfiguration": {
    "AdjustmentType": "PercentChangeInCapacity",
    "Cooldown": 60,
    "MetricAggregationType": "Average",
    "StepAdjustments": [
      {
        "ScalingAdjustment": 200,
        "MetricIntervalLowerBound": 0
      }
    ]
  },
  "ServiceNamespace": "ecs",
  "ScalableDimension": "ecs:service:DesiredCount",
  "ResourceId": "service/default/web-app"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: 5ec6d08e-17ce-1e165a468-73cad4b5ce16
Content-Type: application/x-amz-json-1.1
Content-Length: 175
Date: Fri, 06 May 2016 19:11:38 GMT

{
  "PolicyARN": "arn:aws:autoscaling:us-west-2:012345678910:scalingPolicy:ac542982-cbeb-4294-891c-a5a941dfa787:resource/ecs/service/default/web-app:policyName/my-scaleout-policy"
}
```

Example of a step scaling policy for scale in

The following example applies a step scaling policy to the same Amazon ECS service as in the above example. The policy has two step adjustments that decrease the desired count of the service by 25% or 50%, depending on the size of the alarm breach, with a cooldown period of 120 seconds. The output includes the ARN for the policy, which you need to create the CloudWatch alarm.

Sample Request

```
POST / HTTP/1.1
Host: autoscaling.us-west-2.amazonaws.com
Accept-Encoding: identity
Content-Length: 431
X-Amz-Target: AnyScaleFrontendService.PutScalingPolicy
X-Amz-Date: 20160506T191152Z
User-Agent: aws-cli/1.10.23 Python/2.7.11 Darwin/15.4.0 botocore/1.4.8
Content-Type: application/x-amz-json-1.1
Authorization: AUTHPARAMS
```



```
{
  "PolicyName": "my-scalein-policy",
  "PolicyType": "StepScaling",
  "StepScalingPolicyConfiguration": {
    "AdjustmentType": "PercentChangeInCapacity",
    "Cooldown": 120,
    "MetricAggregationType": "Average",
    "MinAdjustmentMagnitude": 1,
    "StepAdjustments": [
      {
        "ScalingAdjustment": -25,
        "MetricIntervalLowerBound": -15,
        "MetricIntervalUpperBound": 0
      },
      {
        "ScalingAdjustment": -50,
        "MetricIntervalUpperBound": -15
      }
    ]
  },
  "ServiceNamespace": "ecs",
  "ScalableDimension": "ecs:service:DesiredCount",
  "ResourceId": "service/default/web-app"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: 5a64c9e1-3cfe-11e74bfad-8d1c65ec6d08
Content-Type: application/x-amz-json-1.1
Content-Length: 174
Date: Fri, 06 May 2016 19:11:52 GMT

{
  "PolicyARN": "arn:aws:autoscaling:us-west-2:012345678910:scalingPolicy:6d8972f3-efc8-437c-92d1-6270f29a66e7:resource/ecs/service/default/web-app:policyName/my-scalein-policy"
}
```

See Also

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

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

PutScheduledAction

Creates or updates a scheduled action for an Application Auto Scaling scalable target.

Each scalable target is identified by a service namespace, resource ID, and scalable dimension. A scheduled action applies to the scalable target identified by those three attributes. You cannot create a scheduled action until you have registered the resource as a scalable target using [RegisterScalableTarget](#) (p. 50).

To update an action, specify its name and the parameters that you want to change. If you don't specify start and end times, the old values are deleted. Any other parameters that you don't specify are not changed by this update request.

You can view the scheduled actions using [DescribeScheduledActions](#) (p. 32). If you are no longer using a scheduled action, you can delete it using [DeleteScheduledAction](#) (p. 7).

Learn more about how to work with scheduled actions in the [Application Auto Scaling User Guide](#).

Request Syntax

```
{
  "EndTime": number,
  "ResourceId": "string",
  "ScalableDimension": "string",
  "ScalableTargetAction": {
    "MaxCapacity": number,
    "MinCapacity": number
  },
  "Schedule": "string",
  "ScheduledActionName": "string",
  "ServiceNamespace": "string",
  "StartTime": number
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters](#) (p. 80).

The request accepts the following data in JSON format.

EndTime (p. 45)

The date and time for the scheduled action to end.

Type: Timestamp

Required: No

ResourceId (p. 45)

The identifier of the resource associated with the scheduled action. This string consists of the resource type and unique identifier.

- ECS service - The resource type is `service` and the unique identifier is the cluster name and service name. Example: `service/default/sample-webapp`.
- Spot fleet request - The resource type is `spot-fleet-request` and the unique identifier is the Spot fleet request ID. Example: `spot-fleet-request/sfr-73fbd2ce-aa30-494c-8788-1cee4EXAMPLE`.


```
| sagemaker:variant:DesiredInstanceCount | custom-  
resource:ResourceType:Property
```

Required: Yes

ScalableTargetAction (p. 45)

The new minimum and maximum capacity. You can set both values or just one. During the scheduled time, if the current capacity is below the minimum capacity, Application Auto Scaling scales out to the minimum capacity. If the current capacity is above the maximum capacity, Application Auto Scaling scales in to the maximum capacity.

Type: [ScalableTargetAction \(p. 63\)](#) object

Required: No

Schedule (p. 45)

The schedule for this action. The following formats are supported:

- At expressions - "at(*yyyy-mm-ddThh:mm:ss*)"
- Rate expressions - "rate(*value unit*)"
- Cron expressions - "cron(*fields*)"

At expressions are useful for one-time schedules. Specify the time, in UTC.

For rate expressions, *value* is a positive integer and *unit* is minute | minutes | hour | hours | day | days.

For more information about cron expressions, see [Cron Expressions](#) in the *Amazon CloudWatch Events User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1600.

```
Pattern: [\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDFFF\x\n\t]*
```

Required: No

ScheduledActionName (p. 45)

The name of the scheduled action.

Type: String

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

```
Pattern: (?!(^[ ]+.*)|(.*([\u0000-\u001f]|[\u007f-\u009f]|[:/])+.*)|(.*[ ]+  
$))).+
```

Required: Yes

ServiceNamespace (p. 45)

The namespace of the AWS service that provides the resource or `custom-resource` for a resource provided by your own application or service. For more information, see [AWS Service Namespaces](#) in the *Amazon Web Services General Reference*.

Type: String

Valid Values: ecs | elasticmapreduce | ec2 | appstream | dynamodb | rds | sagemaker | custom-resource

Required: Yes

StartTime (p. 45)

The date and time for the scheduled action to start.

Type: Timestamp

Required: No

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 Errors](#) (p. 82).

ConcurrentUpdateException

Concurrent updates caused an exception, for example, if you request an update to an Application Auto Scaling resource that already has a pending update.

HTTP Status Code: 400

InternalServiceException

The service encountered an internal error.

HTTP Status Code: 400

LimitExceededException

A per-account resource limit is exceeded. For more information, see [Application Auto Scaling Limits](#).

HTTP Status Code: 400

ObjectNotFoundException

The specified object could not be found. For any operation that depends on the existence of a scalable target, this exception is thrown if the scalable target with the specified service namespace, resource ID, and scalable dimension does not exist. For any operation that deletes or deregisters a resource, this exception is thrown if the resource cannot be found.

HTTP Status Code: 400

ValidationException

An exception was thrown for a validation issue. Review the available parameters for the API 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](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)

- [AWS SDK for Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

RegisterScalableTarget

Registers or updates a scalable target. A scalable target is a resource that Application Auto Scaling can scale out and scale in. Each scalable target has a resource ID, scalable dimension, and namespace, as well as values for minimum and maximum capacity.

After you register a scalable target, you do not need to register it again to use other Application Auto Scaling operations. To see which resources have been registered, use [DescribeScalableTargets \(p. 14\)](#). You can also view the scaling policies for a service namespace using [DescribeScalableTargets \(p. 14\)](#).

If you no longer need a scalable target, you can deregister it using [DeregisterScalableTarget \(p. 10\)](#).

Request Syntax

```
{  
  "MaxCapacity": number,  
  "MinCapacity": number,  
  "ResourceId": "string",  
  "RoleARN": "string",  
  "ScalableDimension": "string",  
  "ServiceNamespace": "string"  
}
```

Request Parameters

For information about the parameters that are common to all actions, see [Common Parameters \(p. 80\)](#).

The request accepts the following data in JSON format.

MaxCapacity (p. 50)

The maximum value to scale to in response to a scale-out event. This parameter is required to register a scalable target.

Type: Integer

Required: No

MinCapacity (p. 50)

The minimum value to scale to in response to a scale-in event. This parameter is required to register a scalable target.

Type: Integer

Required: No

ResourceId (p. 50)

The identifier of the resource associated with the scalable target. This string consists of the resource type and unique identifier.

- ECS service - The resource type is `service` and the unique identifier is the cluster name and service name. Example: `service/default/sample-webapp`.
- Spot fleet request - The resource type is `spot-fleet-request` and the unique identifier is the Spot fleet request ID. Example: `spot-fleet-request/sfr-73fbd2ce-aa30-494c-8788-1cee4EXAMPLE`.

- `dynamodb:index:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB global secondary index.
- `rds:cluster:ReadReplicaCount` - The count of Aurora Replicas in an Aurora DB cluster. Available for Aurora MySQL-compatible edition and Aurora PostgreSQL-compatible edition.
- `sagemaker:variant:DesiredInstanceCount` - The number of EC2 instances for an Amazon SageMaker model endpoint variant.
- `custom-resource:ResourceType:Property` - The scalable dimension for a custom resource provided by your own application or service.

Type: String

Valid Values: `ecs:service:DesiredCount` | `ec2:spot-fleet-request:TargetCapacity` | `elasticmapreduce:instancegroup:InstanceCount` | `appstream:fleet:DesiredCapacity` | `dynamodb:table:ReadCapacityUnits` | `dynamodb:table:WriteCapacityUnits` | `dynamodb:index:ReadCapacityUnits` | `dynamodb:index:WriteCapacityUnits` | `rds:cluster:ReadReplicaCount` | `sagemaker:variant:DesiredInstanceCount` | `custom-resource:ResourceType:Property`

Required: Yes

[ServiceNamespace \(p. 50\)](#)

The namespace of the AWS service that provides the resource or `custom-resource` for a resource provided by your own application or service. For more information, see [AWS Service Namespaces](#) in the *Amazon Web Services General Reference*.

Type: String

Valid Values: `ecs` | `elasticmapreduce` | `ec2` | `appstream` | `dynamodb` | `rds` | `sagemaker` | `custom-resource`

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 Errors \(p. 82\)](#).

ConcurrentUpdateException

Concurrent updates caused an exception, for example, if you request an update to an Application Auto Scaling resource that already has a pending update.

HTTP Status Code: 400

InternalServiceException

The service encountered an internal error.

HTTP Status Code: 400

LimitExceededException

A per-account resource limit is exceeded. For more information, see [Application Auto Scaling Limits](#).

HTTP Status Code: 400

ValidationException

An exception was thrown for a validation issue. Review the available parameters for the API request.

HTTP Status Code: 400

Example

If you plan to create requests manually, you must replace the Authorization header contents in the examples (AUTHPARAMS) with a signature. For more information, see [Signature Version 4 Signing Process](#) in the *Amazon Web Services General Reference*. If you plan to use the [AWS CLI](#) or one of the [AWS SDKs](#), these tools sign the requests for you.

Example

The following example registers an Amazon ECS service with Application Auto Scaling.

Sample Request

```
POST / HTTP/1.1
Host: autoscaling.us-west-2.amazonaws.com
Accept-Encoding: identity
Content-Length: 229
X-Amz-Target: AnyScaleFrontendService.RegisterScalableTarget
X-Amz-Date: 20160506T182145Z
User-Agent: aws-cli/1.10.23 Python/2.7.11 Darwin/15.4.0 botocore/1.4.8
Content-Type: application/x-amz-json-1.1
Authorization: AUTHPARAMS

{
  "ScalableDimension": "ecs:service:DesiredCount",
  "ResourceId": "service/default/web-app",
  "RoleARN": "arn:aws:iam::012345678910:role/ApplicationAutoscalingECSRole",
  "MinCapacity": 1,
  "ServiceNamespace": "ecs",
  "MaxCapacity": 10
}
```

See Also

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

- [AWS Command Line Interface](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

Data Types

The Application Auto Scaling API contains several data types that various actions use. This section describes each data type in detail.

Note

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

The following data types are supported:

- [Alarm](#) (p. 55)
- [CustomizedMetricSpecification](#) (p. 56)
- [MetricDimension](#) (p. 58)
- [PredefinedMetricSpecification](#) (p. 59)
- [ScalableTarget](#) (p. 60)
- [ScalableTargetAction](#) (p. 63)
- [ScalingActivity](#) (p. 64)
- [ScalingPolicy](#) (p. 67)
- [ScheduledAction](#) (p. 70)
- [StepAdjustment](#) (p. 74)
- [StepScalingPolicyConfiguration](#) (p. 76)
- [TargetTrackingScalingPolicyConfiguration](#) (p. 78)

Alarm

Represents a CloudWatch alarm associated with a scaling policy.

Contents

AlarmARN

The Amazon Resource Name (ARN) of the alarm.

Type: String

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*`

Required: Yes

AlarmName

The name of the alarm.

Type: String

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*`

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 Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

CustomizedMetricSpecification

Represents a CloudWatch metric of your choosing for a target tracking scaling policy to use with Application Auto Scaling.

To create your customized metric specification:

- Add values for each required parameter from CloudWatch. You can use an existing metric, or a new metric that you create. To use your own metric, you must first publish the metric to CloudWatch. For more information, see [Publish Custom Metrics](#) in the *Amazon CloudWatch User Guide*.
- Choose a metric that changes proportionally with capacity. The value of the metric should increase or decrease in inverse proportion to the number of capacity units. That is, the value of the metric should decrease when capacity increases.

For more information about CloudWatch, see [Amazon CloudWatch Concepts](#).

Contents

Dimensions

The dimensions of the metric.

Conditional: If you published your metric with dimensions, you must specify the same dimensions in your scaling policy.

Type: Array of [MetricDimension \(p. 58\)](#) objects

Required: No

MetricName

The name of the metric.

Type: String

Required: Yes

Namespace

The namespace of the metric.

Type: String

Required: Yes

Statistic

The statistic of the metric.

Type: String

Valid Values: `Average` | `Minimum` | `Maximum` | `SampleCount` | `Sum`

Required: Yes

Unit

The unit of the metric.

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 Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

MetricDimension

Describes the dimension names and values associated with a metric.

Contents

Name

The name of the dimension.

Type: String

Required: Yes

Value

The value of the dimension.

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 Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

PredefinedMetricSpecification

Represents a predefined metric for a target tracking scaling policy to use with Application Auto Scaling.

Contents

PredefinedMetricType

The metric type. The `ALBRequestCountPerTarget` metric type applies only to Spot fleet requests and ECS services.

Type: String

Valid Values: `DynamoDBReadCapacityUtilization` | `DynamoDBWriteCapacityUtilization` | `ALBRequestCountPerTarget` | `RDSReaderAverageCPUUtilization` | `RDSReaderAverageDatabaseConnections` | `EC2SpotFleetRequestAverageCPUUtilization` | `EC2SpotFleetRequestAverageNetworkIn` | `EC2SpotFleetRequestAverageNetworkOut` | `SageMakerVariantInvocationsPerInstance` | `ECSServiceAverageCPUUtilization` | `ECSServiceAverageMemoryUtilization`

Required: Yes

ResourceLabel

Identifies the resource associated with the metric type. You can't specify a resource label unless the metric type is `ALBRequestCountPerTarget` and there is a target group attached to the Spot fleet request or ECS service.

The format is `app/<load-balancer-name>/<load-balancer-id>/targetgroup/<target-group-name>/<target-group-id>`, where:

- `app/<load-balancer-name>/<load-balancer-id>` is the final portion of the load balancer ARN
- `targetgroup/<target-group-name>/<target-group-id>` is the final portion of the target group ARN.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1023.

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 Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

ScalableTarget

Represents a scalable target.

Contents

CreationTime

The Unix timestamp for when the scalable target was created.

Type: Timestamp

Required: Yes

MaxCapacity

The maximum value to scale to in response to a scale-out event.

Type: Integer

Required: Yes

MinCapacity

The minimum value to scale to in response to a scale-in event.

Type: Integer

Required: Yes

ResourceId

The identifier of the resource associated with the scalable target. This string consists of the resource type and unique identifier.

- ECS service - The resource type is `service` and the unique identifier is the cluster name and service name. Example: `service/default/sample-webapp`.
- Spot fleet request - The resource type is `spot-fleet-request` and the unique identifier is the Spot fleet request ID. Example: `spot-fleet-request/sfr-73fbd2ce-aa30-494c-8788-1cee4EXAMPLE`.
- EMR cluster - The resource type is `instancegroup` and the unique identifier is the cluster ID and instance group ID. Example: `instancegroup/j-2EEZNYKUA1NTV/ig-1791Y4E1L8YI0`.
- AppStream 2.0 fleet - The resource type is `fleet` and the unique identifier is the fleet name. Example: `fleet/sample-fleet`.
- DynamoDB table - The resource type is `table` and the unique identifier is the resource ID. Example: `table/my-table`.
- DynamoDB global secondary index - The resource type is `index` and the unique identifier is the resource ID. Example: `table/my-table/index/my-table-index`.
- Aurora DB cluster - The resource type is `cluster` and the unique identifier is the cluster name. Example: `cluster:my-db-cluster`.
- Amazon SageMaker endpoint variants - The resource type is `variant` and the unique identifier is the resource ID. Example: `endpoint/my-end-point/variant/KMeansClustering`.
- Custom resources are not supported with a resource type. This parameter must specify the `OutputValue` from the CloudFormation template stack used to access the resources. The unique identifier is defined by the service provider. More information is available in our [GitHub repository](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1600.

Pattern: [\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*

Required: Yes

RoleARN

The ARN of an IAM role that allows Application Auto Scaling to modify the scalable target on your behalf.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1600.

Pattern: [\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*

Required: Yes

ScalableDimension

The scalable dimension associated with the scalable target. This string consists of the service namespace, resource type, and scaling property.

- `ecs:service:DesiredCount` - The desired task count of an ECS service.
- `ec2:spot-fleet-request:TargetCapacity` - The target capacity of a Spot fleet request.
- `elasticmapreduce:instancegroup:InstanceCount` - The instance count of an EMR Instance Group.
- `appstream:fleet:DesiredCapacity` - The desired capacity of an AppStream 2.0 fleet.
- `dynamodb:table:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB table.
- `dynamodb:table:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB table.
- `dynamodb:index:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB global secondary index.
- `dynamodb:index:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB global secondary index.
- `rds:cluster:ReadReplicaCount` - The count of Aurora Replicas in an Aurora DB cluster. Available for Aurora MySQL-compatible edition and Aurora PostgreSQL-compatible edition.
- `sagemaker:variant:DesiredInstanceCount` - The number of EC2 instances for an Amazon SageMaker model endpoint variant.
- `custom-resource:ResourceType:Property` - The scalable dimension for a custom resource provided by your own application or service.

Type: String

Valid Values: `ecs:service:DesiredCount` | `ec2:spot-fleet-request:TargetCapacity` | `elasticmapreduce:instancegroup:InstanceCount` | `appstream:fleet:DesiredCapacity` | `dynamodb:table:ReadCapacityUnits` | `dynamodb:table:WriteCapacityUnits` | `dynamodb:index:ReadCapacityUnits` | `dynamodb:index:WriteCapacityUnits` | `rds:cluster:ReadReplicaCount` | `sagemaker:variant:DesiredInstanceCount` | `custom-resource:ResourceType:Property`

Required: Yes

ServiceNamespace

The namespace of the AWS service that provides the resource or `custom-resource` for a resource provided by your own application or service. For more information, see [AWS Service Namespaces](#) in the *Amazon Web Services General Reference*.

Type: String

Valid Values: `ecs` | `elasticmapreduce` | `ec2` | `appstream` | `dynamodb` | `rds` | `sagemaker` | `custom-resource`

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 Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

ScalableTargetAction

Represents the minimum and maximum capacity for a scheduled action.

Contents

MaxCapacity

The maximum capacity.

Type: Integer

Required: No

MinCapacity

The minimum capacity.

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 Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

ScalingActivity

Represents a scaling activity.

Contents

ActivityId

The unique identifier of the scaling activity.

Type: String

Pattern: [\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*

Required: Yes

Cause

A simple description of what caused the scaling activity to happen.

Type: String

Pattern: [\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*

Required: Yes

Description

A simple description of what action the scaling activity intends to accomplish.

Type: String

Pattern: [\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*

Required: Yes

Details

The details about the scaling activity.

Type: String

Pattern: [\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*

Required: No

EndTime

The Unix timestamp for when the scaling activity ended.

Type: Timestamp

Required: No

ResourceId

The identifier of the resource associated with the scaling activity. This string consists of the resource type and unique identifier.

- ECS service - The resource type is `service` and the unique identifier is the cluster name and service name. Example: `service/default/sample-webapp`.
- Spot fleet request - The resource type is `spot-fleet-request` and the unique identifier is the Spot fleet request ID. Example: `spot-fleet-request/sfr-73fbd2ce-aa30-494c-8788-1cee4EXAMPLE`.

- EMR cluster - The resource type is `instancegroup` and the unique identifier is the cluster ID and instance group ID. Example: `instancegroup/j-2EEZNYKUA1NTV/ig-1791Y4E1L8YIO`.
- AppStream 2.0 fleet - The resource type is `fleet` and the unique identifier is the fleet name. Example: `fleet/sample-fleet`.
- DynamoDB table - The resource type is `table` and the unique identifier is the resource ID. Example: `table/my-table`.
- DynamoDB global secondary index - The resource type is `index` and the unique identifier is the resource ID. Example: `table/my-table/index/my-table-index`.
- Aurora DB cluster - The resource type is `cluster` and the unique identifier is the cluster name. Example: `cluster:my-db-cluster`.
- Amazon SageMaker endpoint variants - The resource type is `variant` and the unique identifier is the resource ID. Example: `endpoint/my-end-point/variant/KMeansClustering`.
- Custom resources are not supported with a resource type. This parameter must specify the `OutputValue` from the CloudFormation template stack used to access the resources. The unique identifier is defined by the service provider. More information is available in our [GitHub repository](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1600.

Pattern: `[\u0020-\u007F\u00E0-\u00FF\u0100-\u017F\u0180-\u01FF\u0200-\u02FF\u0300-\u036F\u0370-\u03FF\u0400-\u047F\u0480-\u04FF\u0500-\u056F\u0570-\u05FF\u0600-\u06FF\u0700-\u077F\u0780-\u07FF\u0800-\u087F\u0880-\u08FF\u0900-\u097F\u0980-\u09FF\u0A00-\u0A7F\u0A80-\u0AFF\u0B00-\u0B7F\u0B80-\u0BFF\u0C00-\u0C7F\u0C80-\u0CFF\u0D00-\u0D7F\u0D80-\u0DBF\u0DC0-\u0DBF\u0E00-\u0E7F\u0E80-\u0EFF\u0F00-\u0F7F\u0F80-\u0FFF\u1000-\u107F\u1080-\u10FF\u1100-\u117F\u1180-\u11FF\u1200-\u127F\u1280-\u12FF\u1300-\u137F\u1380-\u13FF\u1400-\u147F\u1480-\u14FF\u1500-\u157F\u1580-\u15FF\u1600-\u167F\u1680-\u16FF\u1700-\u177F\u1780-\u17FF\u1800-\u187F\u1880-\u18FF\u1900-\u197F\u1980-\u19FF\u1A00-\u1A7F\u1A80-\u1AFF\u1B00-\u1B7F\u1B80-\u1BFF\u1C00-\u1C7F\u1C80-\u1CFF\u1D00-\u1D7F\u1D80-\u1DBF\u1E00-\u1E7F\u1E80-\u1EFF\u1F00-\u1F7F\u1F80-\u1FFF\u2000-\u206F\u2070-\u209F\u20A0-\u20FF\u2100-\u214F\u2150-\u218F\u2190-\u21FF\u2200-\u223F\u2240-\u227F\u2280-\u22BF\u22C0-\u22FF\u2300-\u232F\u2330-\u236F\u2370-\u239F\u23A0-\u23FF\u2400-\u243F\u2440-\u247F\u2480-\u24BF\u24C0-\u24FF\u2500-\u254F\u2550-\u258F\u2590-\u25FF\u2600-\u264F\u2650-\u268F\u2690-\u26FF\u2700-\u273F\u2740-\u277F\u2780-\u27BF\u27C0-\u27FF\u2800-\u283F\u2840-\u287F\u2880-\u28BF\u28C0-\u28FF\u2900-\u293F\u2940-\u297F\u2980-\u29BF\u29C0-\u29FF\u2A00-\u2A3F\u2A40-\u2A7F\u2A80-\u2ABF\u2AC0-\u2AFF\u2B00-\u2B3F\u2B40-\u2B7F\u2B80-\u2BBF\u2BC0-\u2BFF\u2C00-\u2C3F\u2C40-\u2C7F\u2C80-\u2CBF\u2CC0-\u2CFF\u2D00-\u2D3F\u2D40-\u2D7F\u2D80-\u2DBF\u2DC0-\u2DFF\u2E00-\u2E3F\u2E40-\u2E7F\u2E80-\u2EBF\u2EC0-\u2EFF\u2F00-\u2F3F\u2F40-\u2F7F\u2F80-\u2FBF\u2FC0-\u2FFF\u3000-\u303F\u3040-\u307F\u3080-\u30BF\u30C0-\u30FF\u3100-\u313F\u3140-\u317F\u3180-\u31BF\u31C0-\u31FF\u3200-\u323F\u3240-\u327F\u3280-\u32BF\u32C0-\u32FF\u3300-\u333F\u3340-\u337F\u3380-\u33BF\u33C0-\u33FF\u3400-\u343F\u3440-\u347F\u3480-\u34BF\u34C0-\u34FF\u3500-\u353F\u3540-\u357F\u3580-\u35BF\u35C0-\u35FF\u3600-\u363F\u3640-\u367F\u3680-\u36BF\u36C0-\u36FF\u3700-\u373F\u3740-\u377F\u3780-\u37BF\u37C0-\u37FF\u3800-\u383F\u3840-\u387F\u3880-\u38BF\u38C0-\u38FF\u3900-\u393F\u3940-\u397F\u3980-\u39BF\u39C0-\u39FF\u3A00-\u3A3F\u3A40-\u3A7F\u3A80-\u3ABF\u3AC0-\u3AFF\u3B00-\u3B3F\u3B40-\u3B7F\u3B80-\u3BBF\u3BC0-\u3BFF\u3C00-\u3C3F\u3C40-\u3C7F\u3C80-\u3CBF\u3CC0-\u3CFF\u3D00-\u3D3F\u3D40-\u3D7F\u3D80-\u3DBF\u3DC0-\u3DFF\u3E00-\u3E3F\u3E40-\u3E7F\u3E80-\u3EBF\u3EC0-\u3EFF\u3F00-\u3F3F\u3F40-\u3F7F\u3F80-\u3FBF\u3FC0-\u3FFF\u4000-\u403F\u4040-\u407F\u4080-\u40BF\u40C0-\u40FF\u4100-\u413F\u4140-\u417F\u4180-\u41BF\u41C0-\u41FF\u4200-\u423F\u4240-\u427F\u4280-\u42BF\u42C0-\u42FF\u4300-\u433F\u4340-\u437F\u4380-\u43BF\u43C0-\u43FF\u4400-\u443F\u4440-\u447F\u4480-\u44BF\u44C0-\u44FF\u4500-\u453F\u4540-\u457F\u4580-\u45BF\u45C0-\u45FF\u4600-\u463F\u4640-\u467F\u4680-\u46BF\u46C0-\u46FF\u4700-\u473F\u4740-\u477F\u4780-\u47BF\u47C0-\u47FF\u4800-\u483F\u4840-\u487F\u4880-\u48BF\u48C0-\u48FF\u4900-\u493F\u4940-\u497F\u4980-\u49BF\u49C0-\u49FF\u4A00-\u4A3F\u4A40-\u4A7F\u4A80-\u4ABF\u4AC0-\u4AFF\u4B00-\u4B3F\u4B40-\u4B7F\u4B80-\u4BBF\u4BC0-\u4BFF\u4C00-\u4C3F\u4C40-\u4C7F\u4C80-\u4CBF\u4CC0-\u4CFF\u4D00-\u4D3F\u4D40-\u4D7F\u4D80-\u4DBF\u4DC0-\u4DFF\u4E00-\u4E3F\u4E40-\u4E7F\u4E80-\u4EBF\u4EC0-\u4EFF\u4F00-\u4F3F\u4F40-\u4F7F\u4F80-\u4FBF\u4FC0-\u4FFF\u5000-\u503F\u5040-\u507F\u5080-\u50BF\u50C0-\u50FF\u5100-\u513F\u5140-\u517F\u5180-\u51BF\u51C0-\u51FF\u5200-\u523F\u5240-\u527F\u5280-\u52BF\u52C0-\u52FF\u5300-\u533F\u5340-\u537F\u5380-\u53BF\u53C0-\u53FF\u5400-\u543F\u5440-\u547F\u5480-\u54BF\u54C0-\u54FF\u5500-\u553F\u5540-\u557F\u5580-\u55BF\u55C0-\u55FF\u5600-\u563F\u5640-\u567F\u5680-\u56BF\u56C0-\u56FF\u5700-\u573F\u5740-\u577F\u5780-\u57BF\u57C0-\u57FF\u5800-\u583F\u5840-\u587F\u5880-\u58BF\u58C0-\u58FF\u5900-\u593F\u5940-\u597F\u5980-\u59BF\u59C0-\u59FF\u5A00-\u5A3F\u5A40-\u5A7F\u5A80-\u5ABF\u5AC0-\u5AFF\u5B00-\u5B3F\u5B40-\u5B7F\u5B80-\u5BBF\u5BC0-\u5BFF\u5C00-\u5C3F\u5C40-\u5C7F\u5C80-\u5CBF\u5CC0-\u5CFF\u5D00-\u5D3F\u5D40-\u5D7F\u5D80-\u5DBF\u5DC0-\u5DFF\u5E00-\u5E3F\u5E40-\u5E7F\u5E80-\u5EBF\u5EC0-\u5EFF\u5F00-\u5F3F\u5F40-\u5F7F\u5F80-\u5FBF\u5FC0-\u5FFF\u6000-\u603F\u6040-\u607F\u6080-\u60BF\u60C0-\u60FF\u6100-\u613F\u6140-\u617F\u6180-\u61BF\u61C0-\u61FF\u6200-\u623F\u6240-\u627F\u6280-\u62BF\u62C0-\u62FF\u6300-\u633F\u6340-\u637F\u6380-\u63BF\u63C0-\u63FF\u6400-\u643F\u6440-\u647F\u6480-\u64BF\u64C0-\u64FF\u6500-\u653F\u6540-\u657F\u6580-\u65BF\u65C0-\u65FF\u6600-\u663F\u6640-\u667F\u6680-\u66BF\u66C0-\u66FF\u6700-\u673F\u6740-\u677F\u6780-\u67BF\u67C0-\u67FF\u6800-\u683F\u6840-\u687F\u6880-\u68BF\u68C0-\u68FF\u6900-\u693F\u6940-\u697F\u6980-\u69BF\u69C0-\u69FF\u6A00-\u6A3F\u6A40-\u6A7F\u6A80-\u6ABF\u6AC0-\u6AFF\u6B00-\u6B3F\u6B40-\u6B7F\u6B80-\u6BBF\u6BC0-\u6BFF\u6C00-\u6C3F\u6C40-\u6C7F\u6C80-\u6CBF\u6CC0-\u6CFF\u6D00-\u6D3F\u6D40-\u6D7F\u6D80-\u6DBF\u6DC0-\u6DFF\u6E00-\u6E3F\u6E40-\u6E7F\u6E80-\u6EBF\u6EC0-\u6EFF\u6F00-\u6F3F\u6F40-\u6F7F\u6F80-\u6FBF\u6FC0-\u6FFF\u7000-\u703F\u7040-\u707F\u7080-\u70BF\u70C0-\u70FF\u7100-\u713F\u7140-\u717F\u7180-\u71BF\u71C0-\u71FF\u7200-\u723F\u7240-\u727F\u7280-\u72BF\u72C0-\u72FF\u7300-\u733F\u7340-\u737F\u7380-\u73BF\u73C0-\u73FF\u7400-\u743F\u7440-\u747F\u7480-\u74BF\u74C0-\u74FF\u7500-\u753F\u7540-\u757F\u7580-\u75BF\u75C0-\u75FF\u7600-\u763F\u7640-\u767F\u7680-\u76BF\u76C0-\u76FF\u7700-\u773F\u7740-\u777F\u7780-\u77BF\u77C0-\u77FF\u7800-\u783F\u7840-\u787F\u7880-\u78BF\u78C0-\u78FF\u7900-\u793F\u7940-\u797F\u7980-\u79BF\u79C0-\u79FF\u7A00-\u7A3F\u7A40-\u7A7F\u7A80-\u7ABF\u7AC0-\u7AFF\u7B00-\u7B3F\u7B40-\u7B7F\u7B80-\u7BBF\u7BC0-\u7BFF\u7C00-\u7C3F\u7C40-\u7C7F\u7C80-\u7CBF\u7CC0-\u7CFF\u7D00-\u7D3F\u7D40-\u7D7F\u7D80-\u7DBF\u7DC0-\u7DFF\u7E00-\u7E3F\u7E40-\u7E7F\u7E80-\u7EBF\u7EC0-\u7EFF\u7F00-\u7F3F\u7F40-\u7F7F\u7F80-\u7FBF\u7FC0-\u7FFF\u8000-\u803F\u8040-\u807F\u8080-\u80BF\u80C0-\u80FF\u8100-\u813F\u8140-\u817F\u8180-\u81BF\u81C0-\u81FF\u8200-\u823F\u8240-\u827F\u8280-\u82BF\u82C0-\u82FF\u8300-\u833F\u8340-\u837F\u8380-\u83BF\u83C0-\u83FF\u8400-\u843F\u8440-\u847F\u8480-\u84BF\u84C0-\u84FF\u8500-\u853F\u8540-\u857F\u8580-\u85BF\u85C0-\u85FF\u8600-\u863F\u8640-\u867F\u8680-\u86BF\u86C0-\u86FF\u8700-\u873F\u8740-\u877F\u8780-\u87BF\u87C0-\u87FF\u8800-\u883F\u8840-\u887F\u8880-\u88BF\u88C0-\u88FF\u8900-\u893F\u8940-\u897F\u8980-\u89BF\u89C0-\u89FF\u8A00-\u8A3F\u8A40-\u8A7F\u8A80-\u8ABF\u8AC0-\u8AFF\u8B00-\u8B3F\u8B40-\u8B7F\u8B80-\u8BBF\u8BC0-\u8BFF\u8C00-\u8C3F\u8C40-\u8C7F\u8C80-\u8CBF\u8CC0-\u8CFF\u8D00-\u8D3F\u8D40-\u8D7F\u8D80-\u8DBF\u8DC0-\u8DFF\u8E00-\u8E3F\u8E40-\u8E7F\u8E80-\u8EBF\u8EC0-\u8EFF\u8F00-\u8F3F\u8F40-\u8F7F\u8F80-\u8FBF\u8FC0-\u8FFF\u9000-\u903F\u9040-\u907F\u9080-\u90BF\u90C0-\u90FF\u9100-\u913F\u9140-\u917F\u9180-\u91BF\u91C0-\u91FF\u9200-\u923F\u9240-\u927F\u9280-\u92BF\u92C0-\u92FF\u9300-\u933F\u9340-\u937F\u9380-\u93BF\u93C0-\u93FF\u9400-\u943F\u9440-\u947F\u9480-\u94BF\u94C0-\u94FF\u9500-\u953F\u9540-\u957F\u9580-\u95BF\u95C0-\u95FF\u9600-\u963F\u9640-\u967F\u9680-\u96BF\u96C0-\u96FF\u9700-\u973F\u9740-\u977F\u9780-\u97BF\u97C0-\u97FF\u9800-\u983F\u9840-\u987F\u9880-\u98BF\u98C0-\u98FF\u9900-\u993F\u9940-\u997F\u9980-\u99BF\u99C0-\u99FF\u9A00-\u9A3F\u9A40-\u9A7F\u9A80-\u9ABF\u9AC0-\u9AFF\u9B00-\u9B3F\u9B40-\u9B7F\u9B80-\u9BBF\u9BC0-\u9BFF\u9C00-\u9C3F\u9C40-\u9C7F\u9C80-\u9CBF\u9CC0-\u9CFF\u9D00-\u9D3F\u9D40-\u9D7F\u9D80-\u9DBF\u9DC0-\u9DFF\u9E00-\u9E3F\u9E40-\u9E7F\u9E80-\u9EBF\u9EC0-\u9EFF\u9F00-\u9F3F\u9F40-\u9F7F\u9F80-\u9FBF\u9FC0-\u9FFF\uA000-\uA03F\uA040-\uA07F\uA080-\uA0BF\uA0C0-\uA0FF\uA100-\uA13F\uA140-\uA17F\uA180-\uA1BF\uA1C0-\uA1FF\uA200-\uA23F\uA240-\uA27F\uA280-\uA2BF\uA2C0-\uA2FF\uA300-\uA33F\uA340-\uA37F\uA380-\uA3BF\uA3C0-\uA3FF\uA400-\uA43F\uA440-\uA47F\uA480-\uA4BF\uA4C0-\uA4FF\uA500-\uA53F\uA540-\uA57F\uA580-\uA5BF\uA5C0-\uA5FF\uA600-\uA63F\uA640-\uA67F\uA680-\uA6BF\uA6C0-\uA6FF\uA700-\uA73F\uA740-\uA77F\uA780-\uA7BF\uA7C0-\uA7FF\uA800-\uA83F\uA840-\uA87F\uA880-\uA8BF\uA8C0-\uA8FF\uA900-\uA93F\uA940-\uA97F\uA980-\uA9BF\uA9C0-\uA9FF\uAA00-\uAA3F\uAA40-\uAA7F\uAA80-\uAABF\uAAC0-\uAAFF\uAB00-\uAB3F\uAB40-\uAB7F\uAB80-\uABBF\uABC0-\uABFF\uAC00-\uAC3F\uAC40-\uAC7F\uAC80-\uACBF\uACC0-\uACFF\uAD00-\uAD3F\uAD40-\uAD7F\uAD80-\uADBFB\uADC0-\uADFF\uAE00-\uAE3F\uAE40-\uAE7F\uAE80-\uAEBF\uAEC0-\uAEFF\uAF00-\uAF3F\uAF40-\uAF7F\uAF80-\uAFBF\uAFC0-\uAFFF\uB000-\uB03F\uB040-\uB07F\uB080-\uB0BF\uB0C0-\uB0FF\uB100-\uB13F\uB140-\uB17F\uB180-\uB1BF\uB1C0-\uB1FF\uB200-\uB23F\uB240-\uB27F\uB280-\uB2BF\uB2C0-\uB2FF\uB300-\uB33F\uB340-\uB37F\uB380-\uB3BF\uB3C0-\uB3FF\uB400-\uB43F\uB440-\uB47F\uB480-\uB4BF\uB4C0-\uB4FF\uB500-\uB53F\uB540-\uB57F\uB580-\uB5BF\uB5C0-\uB5FF\uB600-\uB63F\uB640-\uB67F\uB680-\uB6BF\uB6C0-\uB6FF\uB700-\uB73F\uB740-\uB77F\uB780-\uB7BF\uB7C0-\uB7FF\uB800-\uB83F\uB840-\uB87F\uB880-\uB8BF\uB8C0-\uB8FF\uB900-\uB93F\uB940-\uB97F\uB980-\uB9BF\uB9C0-\uB9FF\uBA00-\uBA3F\uBA40-\uBA7F\uBA80-\uBABF\uBAC0-\uBAFF\uBB00-\uBB3F\uBB40-\uBB7F\uBB80-\uBBBF\uBBC0-\uBBFF\uBC00-\uBC3F\uBC40-\uBC7F\uBC80-\uBCBF\uBCC0-\uBCFF\uBD00-\uBD3F\uBD40-\uBD7F\uBD80-\uDBFB\uBDC0-\uBDBF\uBE00-\uBE3F\uBE40-\uBE7F\uBE80-\uEBFB\uBEC0-\uBEFF\uBF00-\uBF3F\uBF40-\uBF7F\uBF80-\uFBFB\uBFC0-\uBDBF\uC000-\uC03F\uC040-\uC07F\uC080-\u00BF\uC0C0-\u00FF\uC100-\u013F\uC140-\u017F\uC180-\u01BF\uC1C0-\u01FF\uC200-\u023F\uC240-\u027F\uC280-\u02BF\uC2C0-\u02FF\uC300-\u033F\uC340-\u037F\uC380-\u03BF\uC3C0-\u03FF\uC400-\u043F\uC440-\u047F\uC480-\u04BF\uC4C0-\u04FF\uC500-\u053F\uC540-\u057F\uC580-\u05BF\uC5C0-\u05FF\uC600-\u063F\uC640-\u067F\uC680-\u06BF\uC6C0-\u06FF\uC700-\u073F\uC740-\u077F\uC780-\u07BF\uC7C0-\u07FF\uC800-\u083F\uC840-\u087F\uC880-\u08BF\uC8C0-\u08FF\uC900-\u093F\uC940-\u097F\uC980-\u09BF\uC9C0-\u09FF\uCA00-\u0A3F\uCA40-\u0A7F\uCA80-\u0ABF\uCAC0-\u0AFF\uCB00-\u0B3F\uCB40-\u0B7F\uCB80-\u0BBF\uCBC0-\u0BFF\uCC00-\u0C3F\uCC40-\u0C7F\uCC80-\u0CBF\uCCC0-\u0CFF\uCD00-\u0D3F\uCD40-\u0D7F\uCD80-\u0DBFB\uCDC0-\u0DFF\uCE00-\u0E3F\uCE40-\u0E7F\uCE80-\u0EBFB\uCEC0-\u0EFF\uCF00-\u0F3F\uCF40-\u0F7F\uCF80-\u0FBFB\uCFC0-\u0DBF\uD000-\uD03F\uD040-\uD07F\uD080-\uD0BF\uD0C0-\uD0FF\uD100-\uD13F\uD140-\uD17F\uD180-\uD1BF\uD1C0-\uD1FF\uD200-\uD23F\uD240-\uD27F\uD280-\uD2BF\uD2C0-\uD2FF\uD300-\uD33F\uD340-\uD37F\uD380-\uD3BF\uD3C0-\uD3FF\uD400-\uD43F\uD440-\uD47F\uD480-\uD4BF\uD4C0-\uD4FF\uD500-\uD53F\uD540-\uD57F\uD580-\uD5BF\uD5C0-\uD5FF\uD600-\uD63F\uD640-\uD67F\uD680-\uD6BF\uD6C0-\uD6FF\uD700-\uD73F\uD740-\uD77F\uD780-\uD7BF\uD7C0-\uD7FF\uD800-\uD83F\uD840-\uD87F\uD880-\uD8BF\uD8C0-\uD8FF\uD900-\uD93F\uD940-\uD97F\uD980-\uD9BF\uD9C0-\uD9FF\uDA00-\uDA3F\uDA40-\uDA7F\uDA80-\uDABF\uDAC0-\uDAFF\uDB00-\uDB3F\uDB40-\uDB7F\uDB80-\uDBBF\uDBC0-\uDBFF\uDC00-\uDC3F\uDC40-\uDC7F\uDC80-\uDCBF\uDCC0-\uDCFF\uDD00-\uDD3F\uDD40-\uDD7F\uDD80-\uDDBF\uDDC0-\uDDFF\uDE00-\uDE3F\uDE40-\uDE7F\uDE80-\uDEBF\uDEC0-\uDEFF\uDF00-\uDF3F\uDF40-\uDF7F\uDF80-\uDFBF\uDFC0-\uDBFF\uE000-\uE03F\uE040-\uE07F\uE080-\uE0BF\uE0C0-\uE0FF\uE100-\uE13F\uE140-\uE17F\uE180-\uE1BF\uE1C0-\uE1FF\uE200-\uE23F\uE240-\uE27F\uE280-\uE2BF\uE2C0-\uE2FF\uE300-\uE33F\uE340-\uE37F\uE380-\uE3BF\uE3C0-\uE3FF\uE400-\uE43F\uE440-\uE47F\uE480-\uE4BF\uE4C0-\uE4FF\uE500-\uE53F\uE540-\uE57F\uE580-\uE5BF\uE5C0-\uE5FF\uE600-\uE63F\uE640-\uE67F\uE680-\uE6BF\uE6C0-\uE6FF\uE700-\uE73F\uE740-\uE77F\uE780-\uE7BF\uE7C0-\uE7FF\uE800-\uE83F\uE840-\uE87F\uE880-\uE8BF\uE8C0-\uE8FF\uE900-\uE93F\uE940-\uE97F\uE980-\uE9BF\uE9C0-\uE9FF\uEA00-\uEA3F\uEA40-\uEA7F\uEA80-\uEABF\uEAC0-\uEAFF\uEB00-\uEB3F\uEB40-\uEB7F\uEB80-\uEBBF\uEBC0-\uEBFF\uEC00-\uEC3F\uEC40-\uEC7F\uEC80-\uECBF\uECC0-\uECFF\uED00-\uED3F\uED40-\uED7F\uED80-\uEDBF\uEDC0-\uEDFF\uEE00-\uEE3F\uEE40-\uEE7F\uEE80-\uEEBF\uEEC0-\uEEFF\uEF00-\uEF3F\uEF40-\uEF7F\uEF80-\uEFBF\uEFC0-\uEDBF\uF000-\uF03F\uF040-\uF07F\uF080-\uF0BF\uF0C0-\uF0FF\uF100-\uF13F\uF140-\uF17F\uF180-\uF1BF\uF1C0-\uF1FF\uF200-\uF23F\uF240-\uF27F\uF280-\uF2BF\uF2C0-\uF2FF\uF300-\uF33F\uF340-\uF37F\uF380-\uF3BF\uF3C0-\uF3FF\uF400-\uF43F\uF440-\uF47F\uF480-\uF4BF\uF4C0-\uF4FF\uF500-\uF53F\uF540-\uF57F\uF580-\uF5BF\uF5C0-\uF5FF\uF600-\uF63F\uF640-\uF67F\uF680-\uF6BF\uF6C0-\uF6FF\uF700-\uF73F\uF740-\uF77F\uF780-\uF7BF\uF7C0-\uF7FF\uF800-\uF83F\uF840-\uF87F\uF880-\uF8BF\uF8C0-\uF8FF\uF900-\uF93F\uF940-\uF97F\uF980-\uF9BF\uF9C0-\uF9FF\uFA00-\uFA3F\uFA40-\uFA7F\uFA80-\uFABF\uFAC0-\uFAFF\uFB00-\uFB3F\uFB40-\uFB7F\uFB80-\uFBBF\uFBC0-\uFBFF\uFC00-\uFC3F\uFC40-\uFC7F\uFC80-\uFCBF\uFCC0-\uFCFF\uFD00-\uFD3F\uFD40-\uFD7F\uFD80-\uFDBFB\uFDC0-\uFDBF\uFE00-\uFE3F\uFE40-\uFE7F\uFE80-\uFEBF\uFEC0-\uFEFF\uFF00-\uFF3F\uFF40-\uFF7F\uFF80-\uFFBF\uFFC0-\uDBFF`

Required: Yes

ScalableDimension

The scalable dimension. This string consists of the service namespace, resource type, and scaling property.

- `ecs:service:DesiredCount` - The desired task count of an ECS service.
- `ec2:spot-fleet-request:TargetCapacity` - The target capacity of a Spot fleet request.
- `elasticmapreduce:instancegroup:InstanceCount` - The instance count of an EMR Instance Group.
- `appstream:fleet:DesiredCapacity` - The desired capacity of an AppStream 2.0 fleet.
- `dynamodb:table:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB table.
- `dynamodb:table:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB table.
- `dynamodb:index:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB global secondary index.
- `dynamodb:index:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB global secondary index.
- `rds:cluster:ReadReplicaCount` - The count of Aurora Replicas in an Aurora DB cluster. Available for Aurora MySQL-compatible edition and Aurora PostgreSQL-compatible edition.
- `sagemaker:variant:DesiredInstanceCount` - The number of EC2 instances for an Amazon SageMaker model endpoint variant.
- `custom-resource:ResourceType:Property` - The scalable dimension for a custom resource provided by your own application or service.

Type: String

Valid Values: `ecs:service:DesiredCount | ec2:spot-fleet-request:TargetCapacity | elasticmapreduce:instancegroup:InstanceCount | appstream:fleet:DesiredCapacity | dynamodb:table:ReadCapacityUnits | dynamodb:table:WriteCapacityUnits | dynamodb:index:ReadCapacityUnits | dynamodb:index:WriteCapacityUnits | rds:cluster:ReadReplicaCount`

| `sagemaker:variant:DesiredInstanceCount` | `custom-resource:ResourceType:Property`

Required: Yes

ServiceNamespace

The namespace of the AWS service that provides the resource or `custom-resource` for a resource provided by your own application or service. For more information, see [AWS Service Namespaces](#) in the *Amazon Web Services General Reference*.

Type: String

Valid Values: `ecs` | `elasticmapreduce` | `ec2` | `appstream` | `dynamodb` | `rds` | `sagemaker` | `custom-resource`

Required: Yes

StartTime

The Unix timestamp for when the scaling activity began.

Type: Timestamp

Required: Yes

StatusCode

Indicates the status of the scaling activity.

Type: String

Valid Values: `Pending` | `InProgress` | `Successful` | `Overridden` | `Unfulfilled` | `Failed`

Required: Yes

StatusMessage

A simple message about the current status of the scaling activity.

Type: String

Pattern: `[\u0020-\uD7FF\uE000-\uFFFF\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*`

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 Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

ScalingPolicy

Represents a scaling policy to use with Application Auto Scaling.

Contents

Alarms

The CloudWatch alarms associated with the scaling policy.

Type: Array of [Alarm \(p. 55\)](#) objects

Required: No

CreationTime

The Unix timestamp for when the scaling policy was created.

Type: Timestamp

Required: Yes

PolicyARN

The Amazon Resource Name (ARN) of the scaling policy.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1600.

Pattern: `[\u0020-\u007F\u00E000-\u00FFFD\u00800\u0DC00-\u0DBFF\u0DFFF\r\n\t]*`

Required: Yes

PolicyName

The name of the scaling policy.

Type: String

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

Pattern: `\p{Print}+`

Required: Yes

PolicyType

The scaling policy type.

Type: String

Valid Values: `StepScaling` | `TargetTrackingScaling`

Required: Yes

ResourceId

The identifier of the resource associated with the scaling policy. This string consists of the resource type and unique identifier.

- ECS service - The resource type is `service` and the unique identifier is the cluster name and service name. Example: `service/default/sample-webapp`.

- Spot fleet request - The resource type is `spot-fleet-request` and the unique identifier is the Spot fleet request ID. Example: `spot-fleet-request/sfr-73fbd2ce-aa30-494c-8788-1cee4EXAMPLE`.
- EMR cluster - The resource type is `instancegroup` and the unique identifier is the cluster ID and instance group ID. Example: `instancegroup/j-2EEZNYKUA1NTV/ig-1791Y4E1L8YI0`.
- AppStream 2.0 fleet - The resource type is `fleet` and the unique identifier is the fleet name. Example: `fleet/sample-fleet`.
- DynamoDB table - The resource type is `table` and the unique identifier is the resource ID. Example: `table/my-table`.
- DynamoDB global secondary index - The resource type is `index` and the unique identifier is the resource ID. Example: `table/my-table/index/my-table-index`.
- Aurora DB cluster - The resource type is `cluster` and the unique identifier is the cluster name. Example: `cluster:my-db-cluster`.
- Amazon SageMaker endpoint variants - The resource type is `variant` and the unique identifier is the resource ID. Example: `endpoint/my-end-point/variant/KMeansClustering`.
- Custom resources are not supported with a resource type. This parameter must specify the `OutputValue` from the CloudFormation template stack used to access the resources. The unique identifier is defined by the service provider. More information is available in our [GitHub repository](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1600.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\uD800\uDC00-\uDBFF\uDF00-\uDFFF\x\n\t]*`

Required: Yes

ScalableDimension

The scalable dimension. This string consists of the service namespace, resource type, and scaling property.

- `ecs:service:DesiredCount` - The desired task count of an ECS service.
- `ec2:spot-fleet-request:TargetCapacity` - The target capacity of a Spot fleet request.
- `elasticmapreduce:instancegroup:InstanceCount` - The instance count of an EMR Instance Group.
- `appstream:fleet:DesiredCapacity` - The desired capacity of an AppStream 2.0 fleet.
- `dynamodb:table:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB table.
- `dynamodb:table:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB table.
- `dynamodb:index:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB global secondary index.
- `dynamodb:index:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB global secondary index.
- `rds:cluster:ReadReplicaCount` - The count of Aurora Replicas in an Aurora DB cluster. Available for Aurora MySQL-compatible edition and Aurora PostgreSQL-compatible edition.
- `sagemaker:variant:DesiredInstanceCount` - The number of EC2 instances for an Amazon SageMaker model endpoint variant.
- `custom-resource:ResourceType:Property` - The scalable dimension for a custom resource provided by your own application or service.

Type: String

Valid Values: `ecs:service:DesiredCount` | `ec2:spot-fleet-request:TargetCapacity` | `elasticmapreduce:instancegroup:InstanceCount` | `appstream:fleet:DesiredCapacity` | `dynamodb:table:ReadCapacityUnits`

```
| dynamodb:table:WriteCapacityUnits | dynamodb:index:ReadCapacityUnits  
| dynamodb:index:WriteCapacityUnits | rds:cluster:ReadReplicaCount  
| sagemaker:variant:DesiredInstanceCount | custom-  
resource:ResourceType:Property
```

Required: Yes

ServiceNamespace

The namespace of the AWS service that provides the resource or `custom-resource` for a resource provided by your own application or service. For more information, see [AWS Service Namespaces](#) in the *Amazon Web Services General Reference*.

Type: String

Valid Values: `ecs` | `elasticmapreduce` | `ec2` | `appstream` | `dynamodb` | `rds` | `sagemaker` | `custom-resource`

Required: Yes

StepScalingPolicyConfiguration

A step scaling policy.

Type: [StepScalingPolicyConfiguration \(p. 76\)](#) object

Required: No

TargetTrackingScalingPolicyConfiguration

A target tracking scaling policy.

Type: [TargetTrackingScalingPolicyConfiguration \(p. 78\)](#) 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 Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

ScheduledAction

Represents a scheduled action.

Contents

CreationTime

The date and time that the scheduled action was created.

Type: Timestamp

Required: Yes

EndTime

The date and time that the action is scheduled to end.

Type: Timestamp

Required: No

ResourceId

The identifier of the resource associated with the scaling policy. This string consists of the resource type and unique identifier.

- ECS service - The resource type is `service` and the unique identifier is the cluster name and service name. Example: `service/default/sample-webapp`.
- Spot fleet request - The resource type is `spot-fleet-request` and the unique identifier is the Spot fleet request ID. Example: `spot-fleet-request/sfr-73fbd2ce-aa30-494c-8788-1cee4EXAMPLE`.
- EMR cluster - The resource type is `instancegroup` and the unique identifier is the cluster ID and instance group ID. Example: `instancegroup/j-2EEZNYKUA1NTV/ig-1791Y4E1L8YI0`.
- AppStream 2.0 fleet - The resource type is `fleet` and the unique identifier is the fleet name. Example: `fleet/sample-fleet`.
- DynamoDB table - The resource type is `table` and the unique identifier is the resource ID. Example: `table/my-table`.
- DynamoDB global secondary index - The resource type is `index` and the unique identifier is the resource ID. Example: `table/my-table/index/my-table-index`.
- Aurora DB cluster - The resource type is `cluster` and the unique identifier is the cluster name. Example: `cluster:my-db-cluster`.
- Amazon SageMaker endpoint variants - The resource type is `variant` and the unique identifier is the resource ID. Example: `endpoint/my-end-point/variant/KMeansClustering`.
- Custom resources are not supported with a resource type. This parameter must specify the `OutputValue` from the CloudFormation template stack used to access the resources. The unique identifier is defined by the service provider. More information is available in our [GitHub repository](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1600.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\uD800\uDC00-\uDBFF\uDFFF\r\n\t]*`

Required: Yes

ScalableDimension

The scalable dimension. This string consists of the service namespace, resource type, and scaling property.

- `ecs:service:DesiredCount` - The desired task count of an ECS service.
- `ec2:spot-fleet-request:TargetCapacity` - The target capacity of a Spot fleet request.
- `elasticmapreduce:instancegroup:InstanceCount` - The instance count of an EMR Instance Group.
- `appstream:fleet:DesiredCapacity` - The desired capacity of an AppStream 2.0 fleet.
- `dynamodb:table:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB table.
- `dynamodb:table:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB table.
- `dynamodb:index:ReadCapacityUnits` - The provisioned read capacity for a DynamoDB global secondary index.
- `dynamodb:index:WriteCapacityUnits` - The provisioned write capacity for a DynamoDB global secondary index.
- `rds:cluster:ReadReplicaCount` - The count of Aurora Replicas in an Aurora DB cluster. Available for Aurora MySQL-compatible edition and Aurora PostgreSQL-compatible edition.
- `sagemaker:variant:DesiredInstanceCount` - The number of EC2 instances for an Amazon SageMaker model endpoint variant.
- `custom-resource:ResourceType:Property` - The scalable dimension for a custom resource provided by your own application or service.

Type: String

Valid Values: `ecs:service:DesiredCount` | `ec2:spot-fleet-request:TargetCapacity` | `elasticmapreduce:instancegroup:InstanceCount` | `appstream:fleet:DesiredCapacity` | `dynamodb:table:ReadCapacityUnits` | `dynamodb:table:WriteCapacityUnits` | `dynamodb:index:ReadCapacityUnits` | `dynamodb:index:WriteCapacityUnits` | `rds:cluster:ReadReplicaCount` | `sagemaker:variant:DesiredInstanceCount` | `custom-resource:ResourceType:Property`

Required: No

ScalableTargetAction

The new minimum and maximum capacity. You can set both values or just one. During the scheduled time, if the current capacity is below the minimum capacity, Application Auto Scaling scales out to the minimum capacity. If the current capacity is above the maximum capacity, Application Auto Scaling scales in to the maximum capacity.

Type: [ScalableTargetAction](#) (p. 63) object

Required: No

Schedule

The schedule for this action. The following formats are supported:

- At expressions - `"at(yyyy-mm-ddThh:mm:ss)"`
- Rate expressions - `"rate(value unit)"`
- Cron expressions - `"cron(fields)"`

At expressions are useful for one-time schedules. Specify the time, in UTC.

For rate expressions, *value* is a positive integer and *unit* is `minute` | `minutes` | `hour` | `hours` | `day` | `days`.

For more information about cron expressions, see [Cron Expressions](#) in the *Amazon CloudWatch Events User Guide*.

Type: String

StepAdjustment

Represents a step adjustment for a [StepScalingPolicyConfiguration](#) (p. 76). Describes an adjustment based on the difference between the value of the aggregated CloudWatch metric and the breach threshold that you've defined for the alarm.

For the following examples, suppose that you have an alarm with a breach threshold of 50:

- To trigger the adjustment when the metric is greater than or equal to 50 and less than 60, specify a lower bound of 0 and an upper bound of 10.
- To trigger the adjustment when the metric is greater than 40 and less than or equal to 50, specify a lower bound of -10 and an upper bound of 0.

There are a few rules for the step adjustments for your step policy:

- The ranges of your step adjustments can't overlap or have a gap.
- At most one step adjustment can have a null lower bound. If one step adjustment has a negative lower bound, then there must be a step adjustment with a null lower bound.
- At most one step adjustment can have a null upper bound. If one step adjustment has a positive upper bound, then there must be a step adjustment with a null upper bound.
- The upper and lower bound can't be null in the same step adjustment.

Contents

MetricIntervalLowerBound

The lower bound for the difference between the alarm threshold and the CloudWatch metric. If the metric value is above the breach threshold, the lower bound is inclusive (the metric must be greater than or equal to the threshold plus the lower bound). Otherwise, it is exclusive (the metric must be greater than the threshold plus the lower bound). A null value indicates negative infinity.

Type: Double

Required: No

MetricIntervalUpperBound

The upper bound for the difference between the alarm threshold and the CloudWatch metric. If the metric value is above the breach threshold, the upper bound is exclusive (the metric must be less than the threshold plus the upper bound). Otherwise, it is inclusive (the metric must be less than or equal to the threshold plus the upper bound). A null value indicates positive infinity.

The upper bound must be greater than the lower bound.

Type: Double

Required: No

ScalingAdjustment

The amount by which to scale, based on the specified adjustment type. A positive value adds to the current scalable dimension while a negative number removes from the current scalable dimension.

Type: Integer

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 Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

StepScalingPolicyConfiguration

Represents a step scaling policy configuration to use with Application Auto Scaling.

Contents

AdjustmentType

The adjustment type, which specifies how the `ScalingAdjustment` parameter in a [StepAdjustment](#) (p. 74) is interpreted.

Type: String

Valid Values: `ChangeInCapacity` | `PercentChangeInCapacity` | `ExactCapacity`

Required: No

Cooldown

The amount of time, in seconds, after a scaling activity completes where previous trigger-related scaling activities can influence future scaling events.

For scale-out policies, while the cooldown period is in effect, the capacity that has been added by the previous scale-out event that initiated the cooldown is calculated as part of the desired capacity for the next scale out. The intention is to continuously (but not excessively) scale out. For example, an alarm triggers a step scaling policy to scale out an Amazon ECS service by 2 tasks, the scaling activity completes successfully, and a cooldown period of 5 minutes starts. During the cooldown period, if the alarm triggers the same policy again but at a more aggressive step adjustment to scale out the service by 3 tasks, the 2 tasks that were added in the previous scale-out event are considered part of that capacity and only 1 additional task is added to the desired count.

For scale-in policies, the cooldown period is used to block subsequent scale-in requests until it has expired. The intention is to scale in conservatively to protect your application's availability. However, if another alarm triggers a scale-out policy during the cooldown period after a scale-in, Application Auto Scaling scales out your scalable target immediately.

Type: Integer

Required: No

MetricAggregationType

The aggregation type for the CloudWatch metrics. Valid values are `Minimum`, `Maximum`, and `Average`. If the aggregation type is null, the value is treated as `Average`.

Type: String

Valid Values: `Average` | `Minimum` | `Maximum`

Required: No

MinAdjustmentMagnitude

The minimum number to adjust your scalable dimension as a result of a scaling activity. If the adjustment type is `PercentChangeInCapacity`, the scaling policy changes the scalable dimension of the scalable target by this amount.

For example, suppose that you create a step scaling policy to scale out an Amazon ECS service by 25 percent and you specify a `MinAdjustmentMagnitude` of 2. If the service has 4 tasks and the scaling policy is performed, 25 percent of 4 is 1. However, because you specified a `MinAdjustmentMagnitude` of 2, Application Auto Scaling scales out the service by 2 tasks.

Type: Integer

Required: No

StepAdjustments

A set of adjustments that enable you to scale based on the size of the alarm breach.

Type: Array of [StepAdjustment \(p. 74\)](#) 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 Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

TargetTrackingScalingPolicyConfiguration

Represents a target tracking scaling policy configuration to use with Application Auto Scaling.

Contents

CustomizedMetricSpecification

A customized metric. You can specify either a predefined metric or a customized metric.

Type: [CustomizedMetricSpecification \(p. 56\)](#) object

Required: No

DisableScaleIn

Indicates whether scale in by the target tracking scaling policy is disabled. If the value is `true`, scale in is disabled and the target tracking scaling policy won't remove capacity from the scalable resource. Otherwise, scale in is enabled and the target tracking scaling policy can remove capacity from the scalable resource. The default value is `false`.

Type: Boolean

Required: No

PredefinedMetricSpecification

A predefined metric. You can specify either a predefined metric or a customized metric.

Type: [PredefinedMetricSpecification \(p. 59\)](#) object

Required: No

ScaleInCooldown

The amount of time, in seconds, after a scale-in activity completes before another scale in activity can start.

The cooldown period is used to block subsequent scale-in requests until it has expired. The intention is to scale in conservatively to protect your application's availability. However, if another alarm triggers a scale-out policy during the cooldown period after a scale-in, Application Auto Scaling scales out your scalable target immediately.

Type: Integer

Required: No

ScaleOutCooldown

The amount of time, in seconds, after a scale-out activity completes before another scale-out activity can start.

While the cooldown period is in effect, the capacity that has been added by the previous scale-out event that initiated the cooldown is calculated as part of the desired capacity for the next scale out. The intention is to continuously (but not excessively) scale out.

Type: Integer

Required: No

TargetValue

The target value for the metric. The range is 8.515920e-109 to 1.174271e+108 (Base 10) or 2e-360 to 2e360 (Base 2).

Type: Double

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 Go](#)
- [AWS SDK for Go - Pilot](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

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 [Signature Version 4 Signing Process](#) in the *Amazon Web Services General Reference*.

Action

The action to be performed.

Type: string

Required: Yes

Version

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

Type: string

Required: Yes

X-Amz-Algorithm

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

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

Type: string

Valid Values: `AWS4-HMAC-SHA256`

Required: Conditional

X-Amz-Credential

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

For more information, see [Task 2: Create a String to Sign for Signature Version 4](#) in the *Amazon Web Services General Reference*.

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'THHMMSS'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 [Handling Dates in Signature Version 4](#) in the *Amazon Web Services General Reference*.

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 Security Token Service, go to [AWS Services That Work with IAM](#) in the *IAM User Guide*.

Condition: If you're using temporary security credentials from the AWS Security Token Service, 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 [Task 1: Create a Canonical Request For Signature Version 4](#) in the *Amazon Web Services General Reference*.

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

Type: string

Required: Conditional

Common Errors

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

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

IncompleteSignature

The request signature does not conform to AWS standards.

HTTP Status Code: 400

InternalFailure

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

HTTP Status Code: 500

InvalidAction

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

HTTP Status Code: 400

InvalidClientTokenId

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

HTTP Status Code: 403

InvalidParameterCombination

Parameters that must not be used together were used together.

HTTP Status Code: 400

InvalidParameterValue

An invalid or out-of-range value was supplied for the input parameter.

HTTP Status Code: 400

InvalidQueryParameter

The AWS query string is malformed or does not adhere to AWS standards.

HTTP Status Code: 400

MalformedQueryString

The query string contains a syntax error.

HTTP Status Code: 404

MissingAction

The request is missing an action or a required parameter.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 403

MissingParameter

A required parameter for the specified action is not supplied.

HTTP Status Code: 400

OptInRequired

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

HTTP Status Code: 403

RequestExpired

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

HTTP Status Code: 400

ServiceUnavailable

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

HTTP Status Code: 503

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

ValidationError

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

HTTP Status Code: 400

Logging Application Auto Scaling API Calls with AWS CloudTrail

Application Auto Scaling is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service using the Application Auto Scaling API. CloudTrail captures all API calls for Application Auto Scaling as events. The calls captured include calls from the AWS Management Console and code calls to the Application Auto Scaling API. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Application Auto Scaling. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine the request that was made to Application Auto Scaling, the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, see the [AWS CloudTrail User Guide](#).

Application Auto Scaling Information in CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When Application Auto Scaling activity occurs, that activity is recorded in a CloudTrail event along with other AWS service events in **Event history**. You can view, search, and download recent events in your AWS account. For more information, see [Viewing Events with CloudTrail Event History](#).

For an ongoing record of events in your AWS account, including events for Application Auto Scaling, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all AWS Regions. The trail logs events from all Regions in the AWS partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- [Overview for Creating a Trail](#)
- [CloudTrail Supported Services and Integrations](#)
- [Configuring Amazon SNS Notifications for CloudTrail](#)
- [Receiving CloudTrail Log Files from Multiple Regions and Receiving CloudTrail Log Files from Multiple Accounts](#)

All Application Auto Scaling actions are logged by CloudTrail and are documented in the [Application Auto Scaling API Reference](#). For example, calls to the `PutScalingPolicy`, `DeleteScalingPolicy`, and `DescribeScalingPolicies` actions generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or AWS Identity and Access Management (IAM) user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see the [CloudTrail userIdentity Element](#).

Understanding Application Auto Scaling Log File Entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the `DescribeScalableTargets` action.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "Root",
    "principalId": "123456789012",
    "arn": "arn:aws:iam::123456789012:root",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "sessionContext": {
      "attributes": {
        "mfaAuthenticated": "false",
        "creationDate": "2018-08-21T17:05:42Z"
      }
    }
  },
  "eventTime": "2018-08-16T23:20:32Z",
  "eventSource": "autoscaling.amazonaws.com",
  "eventName": "DescribeScalableTargets",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "72.21.196.68",
  "userAgent": "EC2 Spot Console",
  "requestParameters": {
    "serviceNamespace": "ec2",
    "scalableDimension": "ec2:spot-fleet-request:TargetCapacity",
    "resourceIds": [
      "spot-fleet-request/sfr-05ceaf79-3ba2-405d-e87b-612857f1357a"
    ]
  },
  "responseElements": null,
  "additionalEventData": {
    "service": "application-autoscaling"
  },
  "requestID": "0737e2ea-fb2d-11e3-bfd8-99133058e7bb",
  "eventID": "3fcfb182-98f8-4744-bd45-b38835ab61cb",
  "eventType": "AwsApiCall",
  "recipientAccountId": "123456789012"
}
```