Amazon EC2 Auto Scaling

API Reference

API Version 2011-01-01

Amazon EC2 Auto Scaling: API Reference

Table of Contents

Welcome

Amazon EC2 Auto Scaling is designed to automatically launch or terminate EC2 instances based on user-defined policies, schedules, and health checks. Use this service in conjunction with the AWS Auto Scaling, Amazon CloudWatch, and Elastic Load Balancing services.

This document was last published on February 14, 2018.

List of Actions by Function

Account Limits

- DescribeAccountLimits (p. 42)

Auto Scaling Groups

- AttachLoadBalancers (p. 8)
- AttachLoadBalancerTargetGroups (p. 10)
- CreateAutoScalingGroup (p. 15)
- DeleteAutoScalingGroup (p. 28)
- DescribeAutoScalingGroups (p. 46)
- DescribeLoadBalancers (p. 61)
- DescribeLoadBalancerTargetGroups (p. 64)
- DetachLoadBalancerTargetGroups (p. 92)
- DetachLoadBalancers (p. 90)
- UpdateAutoScalingGroup (p. 132)

Auto Scaling Instances

- AttachInstances (p. 6)
- DescribeAutoScalingInstances (p. 49)
- DetachInstances (p. 87)
- SetInstanceHealth (p. 124)
- SetInstanceProtection (p. 126)
- TerminateInstanceInAutoScalingGroup (p. 130)

Launch Configurations

- CreateLaunchConfiguration (p. 20)
- DeleteLaunchConfiguration (p. 30)
- DescribeLaunchConfigurations (p. 54)

Lifecycle Hooks

- CompleteLifecycleAction (p. 12)
- DeleteLifecycleHook (p. 32)
- DescribeLifecycleHooks (p. 57)
- DescribeLifecycleHookTypes (p. 59)
- PutLifecycleHook (p. 105)
- RecordLifecycleActionHeartbeat (p. 118)

Monitoring

- DeleteNotificationConfiguration (p. 34)

- DescribeAutoScalingNotificationTypes (p. 52)
- DescribeMetricCollectionTypes (p. 66)
- DescribeNotificationConfigurations (p. 68)
- DisableMetricsCollection (p. 94)
- EnableMetricsCollection (p. 96)
- PutNotificationConfiguration (p. 109)

Scaling

- DeletePolicy (p. 36)
- DescribeAdjustmentTypes (p. 44)
- DescribePolicies (p. 71)
- DescribeScalingActivities (p. 74)
- DescribeScalingProcessTypes (p. 77)
- DescribeTerminationPolicyTypes (p. 85)
- ExecutePolicy (p. 101)
- PutScalingPolicy (p. 111)
- ResumeProcesses (p. 120)
- SetDesiredCapacity (p. 122)
- SuspendProcesses (p. 128)

Scheduled Scaling

- DeleteScheduledAction (p. 38)
- DescribeScheduledActions (p. 79)
- PutScheduledUpdateGroupAction (p. 115)

Standby State

- EnterStandby (p. 98)
- ExitStandby (p. 103)

Tags

- CreateOrUpdateTags (p. 26)
- DeleteTags (p. 40)
- DescribeTags (p. 82)

Actions

The following actions are supported:

AttachInstances

Attaches one or more EC2 instances to the specified Auto Scaling group.

When you attach instances, Auto Scaling increases the desired capacity of the group by the number of instances being attached. If the number of instances being attached plus the desired capacity of the group exceeds the maximum size of the group, the operation fails.

If there is a Classic Load Balancer attached to your Auto Scaling group, the instances are also registered with the load balancer. If there are target groups attached to your Auto Scaling group, the instances are also registered with the target groups.

For more information, see Attach EC2 Instances to Your Auto Scaling Group in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

InstanceIds.member.N

The IDs of the instances. You can specify up to 20 instances.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 19.

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

Required: No

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=AttachInstances
&AutoScalingGroupName=my-asg
&InstanceIds.member.1=i-12345678
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

AttachLoadBalancers

Attaches one or more Classic Load Balancers to the specified Auto Scaling group.

To attach an Application Load Balancer instead, see AttachLoadBalancerTargetGroups (p. 10).

To describe the load balancers for an Auto Scaling group, use DescribeLoadBalancers (p. 61). To detach the load balancer from the Auto Scaling group, use DetachLoadBalancers (p. 90).

For more information, see Attach a Load Balancer to Your Auto Scaling Group in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

LoadBalancerNames.member.N

The names of the load balancers. You can specify up to 10 load balancers.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=AttachLoadBalancers
```

```
&AutoScalingGroupName=my-asg
&LoadBalancerNames.member.1=my-lb
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

AttachLoadBalancerTargetGroups

Attaches one or more target groups to the specified Auto Scaling group.

To describe the target groups for an Auto Scaling group, use DescribeLoadBalancerTargetGroups (p. 64). To detach the target group from the Auto Scaling group, use DetachLoadBalancerTargetGroups (p. 92).

For more information, see Attach a Load Balancer to Your Auto Scaling Group in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

TargetGroupARNs.member.N

The Amazon Resource Names (ARN) of the target groups. You can specify up to 10 target groups.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 511.

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

Required: Yes

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

CompleteLifecycleAction

Completes the lifecycle action for the specified token or instance with the specified result.

This step is a part of the procedure for adding a lifecycle hook to an Auto Scaling group:

1. (Optional) Create a Lambda function and a rule that allows CloudWatch Events to invoke your Lambda function when Auto Scaling launches or terminates instances.
2. (Optional) Create a notification target and an IAM role. The target can be either an Amazon SQS queue or an Amazon SNS topic. The role allows Auto Scaling to publish lifecycle notifications to the target.
3. Create the lifecycle hook. Specify whether the hook is used when the instances launch or terminate.
4. If you need more time, record the lifecycle action heartbeat to keep the instance in a pending state.
5. **If you finish before the timeout period ends, complete the lifecycle action.**

For more information, see Auto Scaling Lifecycle in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

InstanceId

The ID of the instance.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 19.

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

Required: No

LifecycleActionResult

The action for the group to take. This parameter can be either `CONTINUE` or `ABANDON`.

Type: String

Required: Yes

LifecycleActionToken

A universally unique identifier (UUID) that identifies a specific lifecycle action associated with an instance. Auto Scaling sends this token to the notification target you specified when you created the lifecycle hook.

Type: String

Length Constraints: Fixed length of 36.

Required: No

LifecycleHookName

The name of the lifecycle hook.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[A-Za-z0-9\-_\/]+`

Required: Yes

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=CompleteLifecycleAction
&AutoScalingGroupName=my-asg
&LifecycleHookName=my-launch-hook
&LifecycleActionResult=CONTINUE
&LifecycleActionToken=bcd2f1b8-9a78-44d3-8a7a-4dd07EXAMPLE
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python

- AWS SDK for Ruby V2

CreateAutoScalingGroup

Creates an Auto Scaling group with the specified name and attributes.

If you exceed your maximum limit of Auto Scaling groups, the call fails. For information about viewing this limit, see DescribeAccountLimits (p. 42). For information about updating this limit, see Auto Scaling Limits in the *Auto Scaling User Guide*.

For more information, see Auto Scaling Groups in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group. This name must be unique within the scope of your AWS account.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

AvailabilityZones.member.N

One or more Availability Zones for the group. This parameter is optional if you specify one or more subnets.

Type: Array of strings

Array Members: Minimum number of 1 item.

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

DefaultCooldown

The amount of time, in seconds, after a scaling activity completes before another scaling activity can start. The default is 300.

For more information, see Auto Scaling Cooldowns in the *Auto Scaling User Guide*.

Type: Integer

Required: No

DesiredCapacity

The number of EC2 instances that should be running in the group. This number must be greater than or equal to the minimum size of the group and less than or equal to the maximum size of the group. If you do not specify a desired capacity, the default is the minimum size of the group.

Type: Integer

Required: No

HealthCheckGracePeriod

The amount of time, in seconds, that Auto Scaling waits before checking the health status of an EC2 instance that has come into service. During this time, any health check failures for the instance are ignored. The default is 0.

This parameter is required if you are adding an `ELB` health check.

For more information, see Health Checks in the *Auto Scaling User Guide*.

Type: Integer

Required: No

HealthCheckType

The service to use for the health checks. The valid values are `EC2` and `ELB`.

By default, health checks use Amazon EC2 instance status checks to determine the health of an instance. For more information, see Health Checks in the *Auto Scaling User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

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

Required: No

InstanceId

The ID of the instance used to create a launch configuration for the group. You must specify one of the following: an EC2 instance, a launch configuration, or a launch template.

When you specify an ID of an instance, Auto Scaling creates a new launch configuration and associates it with the group. This launch configuration derives its attributes from the specified instance, with the exception of the block device mapping.

For more information, see Create an Auto Scaling Group Using an EC2 Instance in the *Auto Scaling User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 19.

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

Required: No

LaunchConfigurationName

The name of the launch configuration. You must specify one of the following: a launch configuration, a launch template, or an EC2 instance.

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: No

LaunchTemplate

The launch template to use to launch instances. You must specify one of the following: a launch template, a launch configuration, or an EC2 instance.

Type: LaunchTemplateSpecification (p. 161) object

Required: No

LifecycleHookSpecificationList.member.N

One or more lifecycle hooks.

Type: Array of LifecycleHookSpecification (p. 164) objects

Required: No

LoadBalancerNames.member.N

One or more Classic Load Balancers. To specify an Application Load Balancer, use `TargetGroupARNs` instead.

For more information, see Using a Load Balancer With an Auto Scaling Group in the *Auto Scaling User Guide*.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

MaxSize

The maximum size of the group.

Type: Integer

Required: Yes

MinSize

The minimum size of the group.

Type: Integer

Required: Yes

NewInstancesProtectedFromScaleIn

Indicates whether newly launched instances are protected from termination by Auto Scaling when scaling in.

Type: Boolean

Required: No

PlacementGroup

The name of the placement group into which you'll launch your instances, if any. For more information, see Placement Groups in the *Amazon Elastic Compute Cloud User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

Tags.member.N

One or more tags.

For more information, see Tagging Auto Scaling Groups and Instances in the *Auto Scaling User Guide*.

Type: Array of Tag (p. 182) objects

Required: No

TargetGroupARNs.member.N

The Amazon Resource Names (ARN) of the target groups.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 511.

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

Required: No

TerminationPolicies.member.N

One or more termination policies used to select the instance to terminate. These policies are executed in the order that they are listed.

For more information, see Controlling Which Instances Auto Scaling Terminates During Scale In in the *Auto Scaling User Guide*.

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

VPCZoneIdentifier

A comma-separated list of subnet identifiers for your virtual private cloud (VPC).

If you specify subnets and Availability Zones with this call, ensure that the subnets' Availability Zones match the Availability Zones specified.

For more information, see Launching Auto Scaling Instances in a VPC in the *Auto Scaling User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2047.

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

Required: No

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

AlreadyExists

You already have an Auto Scaling group or launch configuration with this name.

HTTP Status Code: 400

LimitExceeded

You have already reached a limit for your Auto Scaling resources (for example, groups, launch configurations, or lifecycle hooks). For more information, see DescribeAccountLimits (p. 42).

HTTP Status Code: 400

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=CreateAutoScalingGroup
&AutoScalingGroupName=my-asg
&AvailabilityZones.member.1=us-east-1a
&AvailabilityZones.member.2=us-east-1b
&MinSize=2
&MaxSize=10
&DesiredCapacity=2
&LoadBalancerNames.member.1=my-loadbalancer
&HealthCheckType=ELB
&HealthCheckGracePeriod=120
&LaunchConfigurationName=my-lc
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

CreateLaunchConfiguration

Creates a launch configuration.

If you exceed your maximum limit of launch configurations, the call fails. For information about viewing this limit, see DescribeAccountLimits (p. 42). For information about updating this limit, see Auto Scaling Limits in the *Auto Scaling User Guide*.

For more information, see Launch Configurations in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AssociatePublicIpAddress

Used for groups that launch instances into a virtual private cloud (VPC). Specifies whether to assign a public IP address to each instance. For more information, see Launching Auto Scaling Instances in a VPC in the *Auto Scaling User Guide*.

If you specify this parameter, be sure to specify at least one subnet when you create your group.

Default: If the instance is launched into a default subnet, the default is to assign a public IP address. If the instance is launched into a nondefault subnet, the default is not to assign a public IP address.

Type: Boolean

Required: No

BlockDeviceMappings.member.N

One or more mappings that specify how block devices are exposed to the instance. For more information, see Block Device Mapping in the *Amazon Elastic Compute Cloud User Guide*.

Type: Array of BlockDeviceMapping (p. 148) objects

Required: No

ClassicLinkVPCId

The ID of a ClassicLink-enabled VPC to link your EC2-Classic instances to. This parameter is supported only if you are launching EC2-Classic instances. For more information, see ClassicLink in the *Amazon Elastic Compute Cloud User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

ClassicLinkVPCSecurityGroups.member.N

The IDs of one or more security groups for the specified ClassicLink-enabled VPC. This parameter is required if you specify a ClassicLink-enabled VPC, and is not supported otherwise. For more information, see ClassicLink in the *Amazon Elastic Compute Cloud User Guide*.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

EbsOptimized

Indicates whether the instance is optimized for Amazon EBS I/O. By default, the instance is not optimized for EBS I/O. The optimization provides dedicated throughput to Amazon EBS and an optimized configuration stack to provide optimal I/O performance. This optimization is not available with all instance types. Additional usage charges apply. For more information, see Amazon EBS-Optimized Instances in the *Amazon Elastic Compute Cloud User Guide*.

Type: Boolean

Required: No

IamInstanceProfile

The name or the Amazon Resource Name (ARN) of the instance profile associated with the IAM role for the instance.

EC2 instances launched with an IAM role will automatically have AWS security credentials available. You can use IAM roles with Auto Scaling to automatically enable applications running on your EC2 instances to securely access other AWS resources. For more information, see Launch Auto Scaling Instances with an IAM Role in the *Auto Scaling User Guide*.

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: No

ImageId

The ID of the Amazon Machine Image (AMI) to use to launch your EC2 instances.

If you do not specify `InstanceId`, you must specify `ImageId`.

For more information, see Finding an AMI in the *Amazon Elastic Compute Cloud User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

InstanceId

The ID of the instance to use to create the launch configuration. The new launch configuration derives attributes from the instance, with the exception of the block device mapping.

If you do not specify `InstanceId`, you must specify both `ImageId` and `InstanceType`.

To create a launch configuration with a block device mapping or override any other instance attributes, specify them as part of the same request.

For more information, see Create a Launch Configuration Using an EC2 Instance in the *Auto Scaling User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 19.

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

Required: No

InstanceMonitoring

Enables detailed monitoring (`true`) or basic monitoring (`false`) for the Auto Scaling instances. The default is `true`.

Type: InstanceMonitoring (p. 156) object

Required: No

InstanceType

The instance type of the EC2 instance.

If you do not specify `InstanceId`, you must specify `InstanceType`.

For information about available instance types, see Available Instance Types in the *Amazon Elastic Compute Cloud User Guide.*

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

KernelId

The ID of the kernel associated with the AMI.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

KeyName

The name of the key pair. For more information, see Amazon EC2 Key Pairs in the *Amazon Elastic Compute Cloud User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

LaunchConfigurationName

The name of the launch configuration. This name must be unique within the scope of your AWS account.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

PlacementTenancy

The tenancy of the instance. An instance with a tenancy of `dedicated` runs on single-tenant hardware and can only be launched into a VPC.

You must set the value of this parameter to `dedicated` if want to launch Dedicated Instances into a shared tenancy VPC (VPC with instance placement tenancy attribute set to `default`).

If you specify this parameter, be sure to specify at least one subnet when you create your group.

For more information, see Launching Auto Scaling Instances in a VPC in the *Auto Scaling User Guide*.

Valid values: `default` | `dedicated`

Type: String

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

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

Required: No

RamdiskId

The ID of the RAM disk associated with the AMI.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

SecurityGroups.member.N

One or more security groups with which to associate the instances.

If your instances are launched in EC2-Classic, you can either specify security group names or the security group IDs. For more information about security groups for EC2-Classic, see Amazon EC2 Security Groups in the *Amazon Elastic Compute Cloud User Guide*.

If your instances are launched into a VPC, specify security group IDs. For more information, see Security Groups for Your VPC in the *Amazon Virtual Private Cloud User Guide*.

Type: Array of strings

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

Required: No

SpotPrice

The maximum hourly price to be paid for any Spot Instance launched to fulfill the request. Spot Instances are launched when the price you specify exceeds the current Spot market price. For more information, see Launching Spot Instances in Your Auto Scaling Group in the *Auto Scaling User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

UserData

The user data to make available to the launched EC2 instances. For more information, see Instance Metadata and User Data in the *Amazon Elastic Compute Cloud User Guide*.

Type: String

Length Constraints: Maximum length of 21847.

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

Required: No

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

AlreadyExists

You already have an Auto Scaling group or launch configuration with this name.

HTTP Status Code: 400

LimitExceeded

You have already reached a limit for your Auto Scaling resources (for example, groups, launch configurations, or lifecycle hooks). For more information, see DescribeAccountLimits (p. 42).

HTTP Status Code: 400

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=CreateLaunchConfiguration
&LaunchConfigurationName=my-lc
&ImageId=ami-12345678
&InstanceType=t2.micro
&SecurityGroups.member.1=sg-12345678
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

CreateOrUpdateTags

Creates or updates tags for the specified Auto Scaling group.

When you specify a tag with a key that already exists, the operation overwrites the previous tag definition, and you do not get an error message.

For more information, see Tagging Auto Scaling Groups and Instances in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

Tags.member.N

One or more tags.

Type: Array of Tag (p. 182) objects

Required: Yes

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

AlreadyExists

You already have an Auto Scaling group or launch configuration with this name.

HTTP Status Code: 400

LimitExceeded

You have already reached a limit for your Auto Scaling resources (for example, groups, launch configurations, or lifecycle hooks). For more information, see DescribeAccountLimits (p. 42).

HTTP Status Code: 400

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

ResourceInUse

The operation can't be performed because the resource is in use.

HTTP Status Code: 400

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=CreateOrUpdateTags
```

```
&Tags.member.1.ResourceId=my-asg
&Tags.member.1.ResourceType=auto-scaling-group
&Tags.member.1.Key=environment
&Tags.member.1.Value=test
&Tags.member.1.PropagateAtLaunch=true
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DeleteAutoScalingGroup

Deletes the specified Auto Scaling group.

If the group has instances or scaling activities in progress, you must specify the option to force the deletion in order for it to succeed.

If the group has policies, deleting the group deletes the policies, the underlying alarm actions, and any alarm that no longer has an associated action.

To remove instances from the Auto Scaling group before deleting it, call DetachInstances (p. 87) with the list of instances and the option to decrement the desired capacity so that Auto Scaling does not launch replacement instances.

To terminate all instances before deleting the Auto Scaling group, call UpdateAutoScalingGroup (p. 132) and set the minimum size and desired capacity of the Auto Scaling group to zero.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

ForceDelete

Specifies that the group will be deleted along with all instances associated with the group, without waiting for all instances to be terminated. This parameter also deletes any lifecycle actions associated with the group.

Type: Boolean

Required: No

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

ResourceInUse

The operation can't be performed because the resource is in use.

HTTP Status Code: 400

ScalingActivityInProgress

The operation can't be performed because there are scaling activities in progress.

HTTP Status Code: 400

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DeleteAutoScalingGroup
&AutoScalingGroupName=my-asg
&ForceDelete=true
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DeleteLaunchConfiguration

Deletes the specified launch configuration.

The launch configuration must not be attached to an Auto Scaling group. When this call completes, the launch configuration is no longer available for use.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

LaunchConfigurationName

The name of the launch configuration.

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

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

ResourceInUse

The operation can't be performed because the resource is in use.

HTTP Status Code: 400

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DeleteLaunchConfiguration
&LaunchConfigurationName=my-lc
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DeleteLifecycleHook

Deletes the specified lifecycle hook.

If there are any outstanding lifecycle actions, they are completed first (`ABANDON` for launching instances, `CONTINUE` for terminating instances).

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

LifecycleHookName

The name of the lifecycle hook.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[A-Za-z0-9\-_\/]+`

Required: Yes

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DeleteLifecycleHook
&AutoScalingGroupName=my-asg
&LifecycleHookName=my-hook
&Version=2011-01-01
```

```
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DeleteNotificationConfiguration

Deletes the specified notification.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

TopicARN

The Amazon Resource Name (ARN) of the Amazon Simple Notification Service (SNS) topic.

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

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DeleteNotificationConfiguration
&AutoScalingGroupName=my-asg
&TopicARN=arn:aws:sns:us-east-1:123456789012:my-sns-topic
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DeletePolicy

Deletes the specified Auto Scaling policy.

Deleting a policy deletes the underlying alarm action, but does not delete the alarm, even if it no longer has an associated action.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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: No

PolicyName

The name or Amazon Resource Name (ARN) of the 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]*`

Required: Yes

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DeletePolicy
&AutoScalingGroupName=my-asg
&PolicyName=ScaleIn
&Version=2011-01-01
```

```
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DeleteScheduledAction

Deletes the specified scheduled action.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

ScheduledActionName

The name of the action to delete.

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

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DeleteScheduledAction
&AutoScalingGroupName=my-asg
&ScheduledActionName=my-scheduled-action
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DeleteTags

Deletes the specified tags.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

Tags.member.N

One or more tags.

Type: Array of Tag (p. 182) objects

Required: Yes

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

ResourceInUse

The operation can't be performed because the resource is in use.

HTTP Status Code: 400

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DeleteTags
&Tags.member.1.ResourceId=my-asg
&Tags.member.1.ResourceType=auto-scaling-group
&Tags.member.1.Key=environment
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeAccountLimits

Describes the current Auto Scaling resource limits for your AWS account.

For information about requesting an increase in these limits, see Auto Scaling Limits in the *Auto Scaling User Guide*.

Response Elements

The following elements are returned by the service.

MaxNumberOfAutoScalingGroups

The maximum number of groups allowed for your AWS account. The default limit is 20 per region.

Type: Integer

MaxNumberOfLaunchConfigurations

The maximum number of launch configurations allowed for your AWS account. The default limit is 100 per region.

Type: Integer

NumberOfAutoScalingGroups

The current number of groups for your AWS account.

Type: Integer

NumberOfLaunchConfigurations

The current number of launch configurations for your AWS account.

Type: Integer

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DescribeAccountLimits
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DescribeAccountLimitsResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <DescribeAccountLimitsResult>
    <NumberOfLaunchConfigurations>5</NumberOfLaunchConfigurations>
    <MaxNumberOfLaunchConfigurations>100</MaxNumberOfLaunchConfigurations>
    <NumberOfAutoScalingGroups>10</NumberOfAutoScalingGroups>
    <MaxNumberOfAutoScalingGroups>20</MaxNumberOfAutoScalingGroups>
  </DescribeAccountLimitsResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeAccountLimitsResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeAdjustmentTypes

Describes the policy adjustment types for use with PutScalingPolicy (p. 111).

Response Elements

The following element is returned by the service.

AdjustmentTypes.member.N

The policy adjustment types.

Type: Array of AdjustmentType (p. 140) objects

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DescribeAdjustmentTypes
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DescribeAdjustmentTypesResponse xmlns="http://autoscaling.amazonaws.com/doc/201-01-01/">
  <DescribeAdjustmentTypesResult>
    <AdjustmentTypes>
      <member>
        <AdjustmentType>ChangeInCapacity</AdjustmentType>
      </member>
      <member>
        <AdjustmentType>ExactCapacity</AdjustmentType>
      </member>
      <member>
        <AdjustmentType>PercentChangeInCapacity</AdjustmentType>
      </member>
    </AdjustmentTypes>
  </DescribeAdjustmentTypesResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeAdjustmentTypesResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeAutoScalingGroups

Describes one or more Auto Scaling groups.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupNames.member.N

The names of the Auto Scaling groups. If you omit this parameter, all Auto Scaling groups are described.

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

MaxRecords

The maximum number of items to return with this call. The default value is 50 and the maximum value is 100.

Type: Integer

Required: No

NextToken

The token for the next set of items to return. (You received this token from a previous call.)

Type: String

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

Required: No

Response Elements

The following elements are returned by the service.

AutoScalingGroups.member.N

The groups.

Type: Array of AutoScalingGroup (p. 142) objects

NextToken

The token to use when requesting the next set of items. If there are no additional items to return, the string is empty.

Type: String

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. 189).

InvalidNextToken

The `NextToken` value is not valid.

HTTP Status Code: 400

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DescribeAutoScalingGroups
&AutoScalingGroupNames.member.1=my-asg
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DescribeAutoScalingGroupsResponse xmlns="http://autoscaling.amazonaws.com/
doc/2011-01-01/">
  <DescribeAutoScalingGroupsResult>
    <AutoScalingGroups>
      <member>
        <HealthCheckType>ELB</HealthCheckType>
        <LoadBalancerNames>
          <member>my-loadbalancer</member>
        </LoadBalancerNames>
        <Instances>
          <member>
            <LaunchConfigurationName>my-lc</LaunchConfigurationName>
            <LifecycleState>InService</LifecycleState>
            <InstanceId>i-12345678</InstanceId>
            <ProtectedFromScaleIn>false</ProtectedFromScaleIn>
            <AvailabilityZone>us-east-1c</AvailabilityZone>
          </member>
        </Instances>
        <TerminationPolicies>
          <member>Default</member>
        </TerminationPolicies>
        <DefaultCooldown>300</DefaultCooldown>
        <AutoScalingGroupARN>arn:aws:autoscaling:us-
east-1:123456789012:autoScalingGroup:12345678-1234-1234-1234-123456789012:autoScalingGroupName/
my-asg</AutoScalingGroupARN>
        <EnabledMetrics />
        <AvailabilityZones>
          <member>us-east-1b</member>
          <member>us-east-1a</member>
        </AvailabilityZones>
        <Tags>
```

```
          <member>
            <ResourceId>my-asg</ResourceId>
            <PropagateAtLaunch>true</PropagateAtLaunch>
            <Value>test</Value>
            <Key>environment</Key>
            <ResourceType>auto-scaling-group</ResourceType>
          </member>
        </Tags>
        <LaunchConfigurationName>my-lc</LaunchConfigurationName>
        <AutoScalingGroupName>my-asg</AutoScalingGroupName>
        <HealthCheckGracePeriod>300</HealthCheckGracePeriod>
        <NewInstancesProtectedFromScaleIn>false</NewInstancesProtectedFromScaleIn>
        <SuspendedProcesses />
        <CreatedTime>2015-05-06T17:47:15.107Z</CreatedTime>
        <MinSize>2</MinSize>
        <MaxSize>10</MaxSize>
        <DesiredCapacity>2</DesiredCapacity>
        <VPCZoneIdentifier>subnet-12345678,subnet-98765432</VPCZoneIdentifier>
      </member>
    </AutoScalingGroups>
  </DescribeAutoScalingGroupsResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeAutoScalingGroupsResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeAutoScalingInstances

Describes one or more Auto Scaling instances.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

InstanceIds.member.N

The instances to describe; up to 50 instance IDs. If you omit this parameter, all Auto Scaling instances are described. If you specify an ID that does not exist, it is ignored with no error.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 19.

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

Required: No

MaxRecords

The maximum number of items to return with this call. The default value is 50 and the maximum value is 50.

Type: Integer

Required: No

NextToken

The token for the next set of items to return. (You received this token from a previous call.)

Type: String

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

Required: No

Response Elements

The following elements are returned by the service.

AutoScalingInstances.member.N

The instances.

Type: Array of AutoScalingInstanceDetails (p. 146) objects

NextToken

The token to use when requesting the next set of items. If there are no additional items to return, the string is empty.

Type: String

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. 189).

InvalidNextToken

The `NextToken` value is not valid.

HTTP Status Code: 400

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DescribeAutoScalingInstances
&InstanceIds.member.1=i-12345678
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DescribeAutoScalingInstancesResponse xmlns="http://autoscaling.amazonaws.com/
doc/2011-01-01/">
  <DescribeAutoScalingInstancesResult>
    <AutoScalingInstances>
      <member>
        <LaunchConfigurationName>my-lc</LaunchConfigurationName>
        <LifecycleState>InService</LifecycleState>
        <AutoScalingGroupName>my-asg</AutoScalingGroupName>
        <InstanceId>i-12345678</InstanceId>
        <HealthStatus>HEALTHY</HealthStatus>
        <ProtectedFromScaleIn>false</ProtectedFromScaleIn>
        <AvailabilityZone>us-east-1b</AvailabilityZone>
      </member>
    </AutoScalingInstances>
  </DescribeAutoScalingInstancesResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeAutoScalingInstancesResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeAutoScalingNotificationTypes

Describes the notification types that are supported by Auto Scaling.

Response Elements

The following element is returned by the service.

AutoScalingNotificationTypes.member.N

The notification types.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

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. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?
Version=2011-01-01&Action=DescribeAutoScalingNotificationTypes
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DescribeAutoScalingNotificationTypesResponse xmlns="http://autoscaling.amazonaws.com/
doc/2011-01-01/">
  <DescribeAutoScalingNotificationTypesResult>
    <AutoScalingNotificationTypes>
      <member>autoscaling:EC2_INSTANCE_LAUNCH</member>
      <member>autoscaling:EC2_INSTANCE_LAUNCH_ERROR</member>
      <member>autoscaling:EC2_INSTANCE_TERMINATE</member>
      <member>autoscaling:EC2_INSTANCE_TERMINATE_ERROR</member>
      <member>autoscaling:TEST_NOTIFICATION</member>
    </AutoScalingNotificationTypes>
  </DescribeAutoScalingNotificationTypesResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
```

```
</DescribeAutoScalingNotificationTypesResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeLaunchConfigurations

Describes one or more launch configurations.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

LaunchConfigurationNames.member.N

The launch configuration names. If you omit this parameter, all launch configurations are described.

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

MaxRecords

The maximum number of items to return with this call. The default value is 50 and the maximum value is 100.

Type: Integer

Required: No

NextToken

The token for the next set of items to return. (You received this token from a previous call.)

Type: String

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

Required: No

Response Elements

The following elements are returned by the service.

LaunchConfigurations.member.N

The launch configurations.

Type: Array of LaunchConfiguration (p. 157) objects

NextToken

The token to use when requesting the next set of items. If there are no additional items to return, the string is empty.

Type: String

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. 189).

InvalidNextToken

The `NextToken` value is not valid.

HTTP Status Code: 400

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DescribeLaunchConfigurations
&LaunchConfigurationNames.member.1=my-lc
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DescribeLaunchConfigurationsResponse xmlns="http://autoscaling.amazonaws.com/
doc/2011-01-01/">
  <DescribeLaunchConfigurationsResult>
    <LaunchConfigurations>
      <member>
        <KernelId />
        <EbsOptimized>false</EbsOptimized>
        <RamdiskId />
        <UserData />
        <ImageId>ami-12345678</ImageId>
        <BlockDeviceMappings />
        <ClassicLinkVPCSecurityGroups />
        <InstanceType>t2.micro</InstanceType>
        <KeyName />
        <LaunchConfigurationARN>arn:aws:autoscaling:us-
east-1:123456789012:launchConfiguration:12345678-1234-1234-1234-123456789012:launchConfigurationName/
my-lc</LaunchConfigurationARN>
        <LaunchConfigurationName>my-lc</LaunchConfigurationName>
        <CreatedTime>2015-01-21T23:04:42.200Z</CreatedTime>
        <SecurityGroups>
          <member>sg-12345678</member>
        </SecurityGroups>
        <InstanceMonitoring>
          <Enabled>true</Enabled>
        </InstanceMonitoring>
      </member>
    </LaunchConfigurations>
  </DescribeLaunchConfigurationsResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
```

```
</DescribeLaunchConfigurationsResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeLifecycleHooks

Describes the lifecycle hooks for the specified Auto Scaling group.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

LifecycleHookNames.member.N

The names of one or more lifecycle hooks. If you omit this parameter, all lifecycle hooks are described.

Type: Array of strings

Array Members: Maximum number of 50 items.

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[A-Za-z0-9\-_\/]+`

Required: No

Response Elements

The following element is returned by the service.

LifecycleHooks.member.N

The lifecycle hooks for the specified group.

Type: Array of LifecycleHook (p. 162) objects

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DescribeLifecycleHooks
&AutoScalingGroupName=my-asg
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DescribeLifecycleHooksResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <DescribeLifecycleHooksResult>
    <LifecycleHooks>
      <member>
        <AutoScalingGroupName>my-asg</AutoScalingGroupName>
        <RoleARN>arn:aws:iam::1234567890:role/my-auto-scaling-role</RoleARN>
        <LifecycleTransition>autoscaling:EC2_INSTANCE_LAUNCHING</LifecycleTransition>
        <GlobalTimeout>172800</GlobalTimeout>
        <LifecycleHookName>my-launch-hook</LifecycleHookName>
        <HeartbeatTimeout>3600</HeartbeatTimeout>
        <DefaultResult>ABANDON</DefaultResult>
        <NotificationTargetARN>arn:aws:sqs:us-east-1:123456789012:my-queue</
NotificationTargetARN>
      </member>
    </LifecycleHooks>
  </DescribeLifecycleHooksResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeLifecycleHooksResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeLifecycleHookTypes

Describes the available types of lifecycle hooks.

Response Elements

The following element is returned by the service.

LifecycleHookTypes.member.N

The lifecycle hook types.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

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. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DescribeLifecycleHookTypes
&AutoScalingGroupName=my-asg
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DescribeLifecycleHookTypesResponse xmlns="http://autoscaling.amazonaws.com/
doc/2011-01-01/">
  <DescribeLifecycleHookTypesResult>
    <LifecycleHookTypes>
      <member>autoscaling:EC2_INSTANCE_LAUNCHING</member>
      <member>autoscaling:EC2_INSTANCE_TERMINATING</member>
    </LifecycleHookTypes>
  </DescribeLifecycleHookTypesResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeLifecycleHookTypesResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeLoadBalancers

Describes the load balancers for the specified Auto Scaling group.

Note that this operation describes only Classic Load Balancers. If you have Application Load Balancers, use DescribeLoadBalancerTargetGroups (p. 64) instead.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

MaxRecords

The maximum number of items to return with this call. The default value is 100 and the maximum value is 100.

Type: Integer

Required: No

NextToken

The token for the next set of items to return. (You received this token from a previous call.)

Type: String

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

Required: No

Response Elements

The following elements are returned by the service.

LoadBalancers.member.N

The load balancers.

Type: Array of LoadBalancerState (p. 166) objects

NextToken

The token to use when requesting the next set of items. If there are no additional items to return, the string is empty.

Type: String

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. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DescribeLoadBalancers
&AutoScalingGroupName=my-asg
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DescribeLoadBalancersResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <DescribeLoadBalancersResult>
    <LoadBalancers>
      <member>
        <LoadBalancerName>my-loadbalancer</LoadBalancerName>
        <State>Added</State>
      </member>
    </LoadBalancers>
  </DescribeLoadBalancersResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeLoadBalancersResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeLoadBalancerTargetGroups

Describes the target groups for the specified Auto Scaling group.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

MaxRecords

The maximum number of items to return with this call. The default value is 100 and the maximum value is 100.

Type: Integer

Required: No

NextToken

The token for the next set of items to return. (You received this token from a previous call.)

Type: String

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

Required: No

Response Elements

The following elements are returned by the service.

LoadBalancerTargetGroups.member.N

Information about the target groups.

Type: Array of LoadBalancerTargetGroupState (p. 167) objects

NextToken

The token to use when requesting the next set of items. If there are no additional items to return, the string is empty.

Type: String

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. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeMetricCollectionTypes

Describes the available CloudWatch metrics for Auto Scaling.

Note that the `GroupStandbyInstances` metric is not returned by default. You must explicitly request this metric when calling EnableMetricsCollection (p. 96).

Response Elements

The following elements are returned by the service.

Granularities.member.N

The granularities for the metrics.

Type: Array of MetricGranularityType (p. 170) objects

Metrics.member.N

One or more metrics.

Type: Array of MetricCollectionType (p. 168) objects

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Version=2011-01-01&Action=DescribeMetricCollectionTypes
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DescribeMetricCollectionTypesResponse xmlns="http://autoscaling.amazonaws.com/
doc/2011-01-01/">
  <DescribeMetricCollectionTypesResult>
    <Granularities>
      <member>
        <Granularity>1Minute</Granularity>
      </member>
    </Granularities>
    <Metrics>
      <member>
```

```
        <Metric>GroupMinSize</Metric>
      </member>
      <member>
        <Metric>GroupMaxSize</Metric>
      </member>
      <member>
        <Metric>GroupDesiredCapacity</Metric>
      </member>
      <member>
        <Metric>GroupInServiceInstances</Metric>
      </member>
      <member>
        <Metric>GroupPendingInstances</Metric>
      </member>
      <member>
        <Metric>GroupTerminatingInstances</Metric>
      </member>
      <member>
        <Metric>GroupStandbyInstances</Metric>
      </member>
      <member>
        <Metric>GroupTotalInstances</Metric>
      </member>
    </Metrics>
  </DescribeMetricCollectionTypesResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeMetricCollectionTypesResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeNotificationConfigurations

Describes the notification actions associated with the specified Auto Scaling group.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupNames.member.N

The name of the Auto Scaling group.

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

MaxRecords

The maximum number of items to return with this call. The default value is 50 and the maximum value is 100.

Type: Integer

Required: No

NextToken

The token for the next set of items to return. (You received this token from a previous call.)

Type: String

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

Required: No

Response Elements

The following elements are returned by the service.

NextToken

The token to use when requesting the next set of items. If there are no additional items to return, the string is empty.

Type: String

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

NotificationConfigurations.member.N

The notification configurations.

Type: Array of NotificationConfiguration (p. 171) objects

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

InvalidNextToken

The `NextToken` value is not valid.

HTTP Status Code: 400

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?
Version=2011-01-01&Action=DescribeNotificationConfigurations
&AutoScalingGroupNames.member.1=my-asg
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DescribeNotificationConfigurationsResponse xmlns="http://autoscaling.amazonaws.com/
doc/2011-01-01/">
  <DescribeNotificationConfigurationsResult>
    <NotificationConfigurations>
      <member>
        <AutoScalingGroupName>my-asg</AutoScalingGroupName>
        <NotificationType>autoscaling:EC2_INSTANCE_LAUNCH</NotificationType>
        <TopicARN>arn:aws:sns:us-east-1:123456789012:my-sns-topic</TopicARN>
      </member>
    </NotificationConfigurations>
  </DescribeNotificationConfigurationsResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeNotificationConfigurationsResponse>
```

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 Java
- AWS SDK for JavaScript

- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribePolicies

Describes the policies for the specified Auto Scaling group.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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: No

MaxRecords

The maximum number of items to be returned with each call. The default value is 50 and the maximum value is 100.

Type: Integer

Required: No

NextToken

The token for the next set of items to return. (You received this token from a previous call.)

Type: String

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

Required: No

PolicyNames.member.N

The names of one or more policies. If you omit this parameter, all policies are described. If an group name is provided, the results are limited to that group. This list is limited to 50 items. If you specify an unknown policy name, it is ignored with no error.

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

PolicyTypes.member.N

One or more policy types. Valid values are `SimpleScaling` and `StepScaling`.

Type: Array of strings

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

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

Required: No

Response Elements

The following elements are returned by the service.

NextToken

The token to use when requesting the next set of items. If there are no additional items to return, the string is empty.

Type: String

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

ScalingPolicies.member.N

The scaling policies.

Type: Array of ScalingPolicy (p. 174) objects

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

InvalidNextToken

The `NextToken` value is not valid.

HTTP Status Code: 400

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DescribePolicies
&AutoScalingGroupName=my-asg
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DescribePoliciesResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <DescribePoliciesResult>
    <ScalingPolicies>
      <member>
```

```
        <PolicyARN>arn:aws:autoscaling:us-
east-1:123456789012:scalingPolicy:c322761b-3172-4d56-9a21-0ed9dEXAMPLE:autoScalingGroupName/
my-asg:policyName/MyScaleDownPolicy</PolicyARN>
        <AdjustmentType>ChangeInCapacity</AdjustmentType>
        <ScalingAdjustment>-1</ScalingAdjustment>
        <PolicyName>MyScaleDownPolicy</PolicyName>
        <PolicyType>SimpleScaling</PolicyType>
        <AutoScalingGroupName>my-asg</AutoScalingGroupName>
        <Cooldown>60</Cooldown>
        <Alarms>
          <member>
            <AlarmName>TestQueue</AlarmName>
            <AlarmARN>arn:aws:cloudwatch:us-east-1:123456789012:alarm:TestQueue</AlarmARN>
          </member>
        </Alarms>
      </member>
      <member>
        <PolicyARN>arn:aws:autoscaling:us-
east-1:123456789012:scalingPolicy:c55a5cdd-9be0-435b-b60b-
a8dd3EXAMPLE:autoScalingGroupName/my-asg:policyName/MyScaleUpPolicy</PolicyARN>
        <AdjustmentType>ChangeInCapacity</AdjustmentType>
        <ScalingAdjustment>1</ScalingAdjustment>
        <PolicyName>MyScaleUpPolicy</PolicyName>
        <PolicyType>SimpleScaling</PolicyType>
        <AutoScalingGroupName>my-asg</AutoScalingGroupName>
        <Cooldown>60</Cooldown>
        <Alarms>
          <member>
            <AlarmName>TestQueue</AlarmName>
            <AlarmARN>arn:aws:cloudwatch:us-east-1:123456789012:alarm:TestQueue</AlarmARN>
          </member>
        </Alarms>
      </member>
    </ScalingPolicies>
  </DescribePoliciesResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</DescribePoliciesResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeScalingActivities

Describes one or more scaling activities for the specified Auto Scaling group.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

ActivityIds.member.N

The activity IDs of the desired scaling activities. If you omit this parameter, all activities for the past six weeks are described. If you specify an Auto Scaling group, the results are limited to that group. The list of requested activities cannot contain more than 50 items. If unknown activities are requested, they are ignored with no error.

Type: Array of strings

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

Required: No

AutoScalingGroupName

The name of the Auto Scaling group.

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: No

MaxRecords

The maximum number of items to return with this call. The default value is 100 and the maximum value is 100.

Type: Integer

Required: No

NextToken

The token for the next set of items to return. (You received this token from a previous call.)

Type: String

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

Required: No

Response Elements

The following elements are returned by the service.

Activities.member.N

The scaling activities. Activities are sorted by start time. Activities still in progress are described first.

Type: Array of Activity (p. 138) objects

NextToken

The token to use when requesting the next set of items. If there are no additional items to return, the string is empty.

Type: String

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. 189).

InvalidNextToken

The `NextToken` value is not valid.

HTTP Status Code: 400

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DescribeScalingActivities
&AutoScalingGroupName=my-asg
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DescribeScalingActivitiesResponse xmlns="http://ec2.amazonaws.com/doc/2011-01-01/">
  <DescribeScalingActivitiesResult>
    <Activities>
      <member>
        <StatusCode>Failed</StatusCode>
        <Progress>0</Progress>
        <ActivityId>12345678-1234-1234-1234-123456789012</ActivityId>
        <StartTime>2012-04-12T17:32:07.882Z</StartTime>
        <AutoScalingGroupName>my-asg</AutoScalingGroupName>
        <Cause>At 2012-04-12T17:31:30Z a user request created an AutoScalingGroup changing
 the desired capacity from 0 to 1.  At 2012-04-12T17:32:07Z an instance was started in
 response to a difference between desired and actual capacity, increasing the capacity from
 0 to 1.</Cause>
        <Details>{}</Details>
        <Description>Launching a new EC2 instance.  Status Reason: The image id
 'ami-4edb0327' does not exist. Launching EC2 instance failed.</Description>
        <EndTime>2012-04-12T17:32:08Z</EndTime>
        <StatusMessage>The image id 'ami-4edb0327' does not exist. Launching EC2 instance
 failed.</StatusMessage>
```

```
      </member>
    </Activities>
  </DescribeScalingActivitiesResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeScalingActivitiesResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeScalingProcessTypes

Describes the scaling process types for use with ResumeProcesses (p. 120) and SuspendProcesses (p. 128).

Response Elements

The following element is returned by the service.

Processes.member.N

The names of the process types.

Type: Array of ProcessType (p. 173) objects

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DescribeScalingProcessTypes
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DescribeScalingProcessTypesResponse xmlns="http://autoscaling.amazonaws.com/
doc/2011-01-01/">
  <DescribeScalingProcessTypesResult>
    <Processes>
      <member>
        <ProcessName>AZRebalance</ProcessName>
      </member>
      <member>
        <ProcessName>AddToLoadBalancer</ProcessName>
      </member>
      <member>
        <ProcessName>AlarmNotification</ProcessName>
      </member>
      <member>
        <ProcessName>HealthCheck</ProcessName>
      </member>
      <member>
        <ProcessName>Launch</ProcessName>
```

```
      </member>
      <member>
        <ProcessName>ReplaceUnhealthy</ProcessName>
      </member>
      <member>
        <ProcessName>ScheduledActions</ProcessName>
      </member>
      <member>
        <ProcessName>Terminate</ProcessName>
      </member>
    </Processes>
  </DescribeScalingProcessTypesResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeScalingProcessTypesResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeScheduledActions

Describes the actions scheduled for your Auto Scaling group that haven't run. To describe the actions that have already run, use DescribeScalingActivities (p. 74).

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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: No

EndTime

The latest scheduled start time to return. If scheduled action names are provided, this parameter is ignored.

Type: Timestamp

Required: No

MaxRecords

The maximum number of items to return with this call. The default value is 50 and the maximum value is 100.

Type: Integer

Required: No

NextToken

The token for the next set of items to return. (You received this token from a previous call.)

Type: String

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

Required: No

ScheduledActionNames.member.N

Describes one or more scheduled actions. If you omit this parameter, all scheduled actions are described. If you specify an unknown scheduled action, it is ignored with no error.

You can describe up to a maximum of 50 instances with a single call. If there are more items to return, the call returns a token. To get the next set of items, repeat the call with the returned token.

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

StartTime

The earliest scheduled start time to return. If scheduled action names are provided, this parameter is ignored.

Type: Timestamp

Required: No

Response Elements

The following elements are returned by the service.

NextToken

The token to use when requesting the next set of items. If there are no additional items to return, the string is empty.

Type: String

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

ScheduledUpdateGroupActions.member.N

The scheduled actions.

Type: Array of ScheduledUpdateGroupAction (p. 177) objects

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

InvalidNextToken

The `NextToken` value is not valid.

HTTP Status Code: 400

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeTags

Describes the specified tags.

You can use filters to limit the results. For example, you can query for the tags for a specific Auto Scaling group. You can specify multiple values for a filter. A tag must match at least one of the specified values for it to be included in the results.

You can also specify multiple filters. The result includes information for a particular tag only if it matches all the filters. If there's no match, no special message is returned.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

Filters.member.N

A filter used to scope the tags to return.

Type: Array of Filter (p. 153) objects

Required: No

MaxRecords

The maximum number of items to return with this call. The default value is 50 and the maximum value is 100.

Type: Integer

Required: No

NextToken

The token for the next set of items to return. (You received this token from a previous call.)

Type: String

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

Required: No

Response Elements

The following elements are returned by the service.

NextToken

The token to use when requesting the next set of items. If there are no additional items to return, the string is empty.

Type: String

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

Tags.member.N

One or more tags.

Type: Array of TagDescription (p. 184) objects

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

InvalidNextToken

The `NextToken` value is not valid.

HTTP Status Code: 400

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DescribeTags
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DescribeTagsResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <DescribeTagsResult>
    <Tags>
      <member>
        <ResourceId>my-asg</ResourceId>
        <PropagateAtLaunch>true</PropagateAtLaunch>
        <Value>test</Value>
        <Key>environment</Key>
        <ResourceType>auto-scaling-group</ResourceType>
      </member>
    </Tags>
  </DescribeTagsResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeTagsResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DescribeTerminationPolicyTypes

Describes the termination policies supported by Auto Scaling.

Response Elements

The following element is returned by the service.

TerminationPolicyTypes.member.N

The termination policies supported by Auto Scaling (`OldestInstance`, `OldestLaunchConfiguration`, `NewestInstance`, `ClosestToNextInstanceHour`, and `Default`).

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]*`

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DescribeTerminationPolicyTypes
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DescribeTerminationPolicyTypesResponse xmlns="http://autoscaling.amazonaws.com/
doc/2011-01-01/">
  <DescribeTerminationPolicyTypesResult>
    <TerminationPolicyTypes>
      <member>ClosestToNextInstanceHour</member>
      <member>Default</member>
      <member>NewestInstance</member>
      <member>OldestInstance</member>
      <member>OldestLaunchConfiguration</member>
    </TerminationPolicyTypes>
  </DescribeTerminationPolicyTypesResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
```

```
  </ResponseMetadata>
</DescribeTerminationPolicyTypesResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DetachInstances

Removes one or more instances from the specified Auto Scaling group.

After the instances are detached, you can manage them independent of the Auto Scaling group.

If you do not specify the option to decrement the desired capacity, Auto Scaling launches instances to replace the ones that are detached.

If there is a Classic Load Balancer attached to the Auto Scaling group, the instances are deregistered from the load balancer. If there are target groups attached to the Auto Scaling group, the instances are deregistered from the target groups.

For more information, see Detach EC2 Instances from Your Auto Scaling Group in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

InstanceIds.member.N

The IDs of the instances. You can specify up to 20 instances.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 19.

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

Required: No

ShouldDecrementDesiredCapacity

Indicates whether the Auto Scaling group decrements the desired capacity value by the number of instances detached.

Type: Boolean

Required: Yes

Response Elements

The following element is returned by the service.

Activities.member.N

The activities related to detaching the instances from the Auto Scaling group.

Type: Array of Activity (p. 138) objects

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DetachInstances
&AutoScalingGroupName=my-asg
&InstanceIds.member.1=i-12345678
&ShouldDecrementDesiredCapacity=true
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<DetachInstancesResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <DetachInstancesResult>
    <Activities>
      <member>
        <ActivityId>12345678-1234-1234-1234-123456789012</ActivityId>
        <AutoScalingGroupName>my-asg</AutoScalingGroupName>
        <Description>Detaching EC2 instance: i-12345678</Description>
        <Cause>At 2015-06-14T00:07:30Z instance i-12345678 was detached in response to a
 user request, shrinking the capacity from 4 to 3.</Cause>
        <Progress>50</Progress>
        <StartTime>2015-06-14T00:07:30.280Z</StartTime>
        <Details>{"Availability Zone":"us-east-1a","SubnetID":"subnet-12345678"}</Details>
        <StatusCode>InProgress</StatusCode>
      </member>
    </Activities>
  </DetachInstancesResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</DetachInstancesResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DetachLoadBalancers

Detaches one or more Classic Load Balancers from the specified Auto Scaling group.

Note that this operation detaches only Classic Load Balancers. If you have Application Load Balancers, use DetachLoadBalancerTargetGroups (p. 92) instead.

When you detach a load balancer, it enters the `Removing` state while deregistering the instances in the group. When all instances are deregistered, then you can no longer describe the load balancer using DescribeLoadBalancers (p. 61). Note that the instances remain running.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

LoadBalancerNames.member.N

The names of the load balancers. You can specify up to 10 load balancers.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DetachLoadBalancers
```

```
&AutoScalingGroupName=my-asg
&LoadBalancerNames.member.1=my-lb
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

DetachLoadBalancerTargetGroups

Detaches one or more target groups from the specified Auto Scaling group.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

TargetGroupARNs.member.N

The Amazon Resource Names (ARN) of the target groups. You can specify up to 10 target groups.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 511.

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

Required: Yes

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3

- AWS SDK for Python
- AWS SDK for Ruby V2

DisableMetricsCollection

Disables group metrics for the specified Auto Scaling group.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

Metrics.member.N

One or more of the following metrics. If you omit this parameter, all metrics are disabled.

- `GroupMinSize`
- `GroupMaxSize`
- `GroupDesiredCapacity`
- `GroupInServiceInstances`
- `GroupPendingInstances`
- `GroupStandbyInstances`
- `GroupTerminatingInstances`
- `GroupTotalInstances`

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=DisableMetricsCollection
&AutoScalingGroupName=my-asg
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

EnableMetricsCollection

Enables group metrics for the specified Auto Scaling group. For more information, see Monitoring Your Auto Scaling Groups and Instances in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

Granularity

The granularity to associate with the metrics to collect. The only valid value is `1Minute`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

Metrics.member.N

One or more of the following metrics. If you omit this parameter, all metrics are enabled.

- `GroupMinSize`
- `GroupMaxSize`
- `GroupDesiredCapacity`
- `GroupInServiceInstances`
- `GroupPendingInstances`
- `GroupStandbyInstances`
- `GroupTerminatingInstances`
- `GroupTotalInstances`

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=EnableMetricsCollection
&AutoScalingGroupName=my-asg
&Granularity=1Minute
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

EnterStandby

Moves the specified instances into the standby state.

For more information, see Temporarily Removing Instances from Your Auto Scaling Group in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

InstanceIds.member.N

The IDs of the instances. You can specify up to 20 instances.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 19.

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

Required: No

ShouldDecrementDesiredCapacity

Indicates whether to decrement the desired capacity of the Auto Scaling group by the number of instances moved to `Standby` mode.

Type: Boolean

Required: Yes

Response Elements

The following element is returned by the service.

Activities.member.N

The activities related to moving instances into `Standby` mode.

Type: Array of Activity (p. 138) objects

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=EnterStandby
&AutoScalingGroupName=my-asg
&InstanceIds.member.1=i-12345678
&ShouldDecrementDesiredCapacity=true
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<EnterStandbyResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <EnterStandbyResult>
    <Activities>
      <member>
        <ActivityId>12345678-1234-1234-1234-123456789012</ActivityId>
        <AutoScalingGroupName>my-asg</AutoScalingGroupName>
        <Description>Moving EC2 instance to Standby: i-12345678</Description>
        <Progress>50</Progress>
        <Cause>At 2015-06-13T22:35:50Z instance i-5b73d709 was moved to standby in response
 to a user request, shrinking the capacity from 4 to 3.</Cause>
        <StartTime>2015-06-13T22:35:50.884Z</StartTime>
        <Details>{"Availability Zone":"us-east-1a","SubnetID":"subnet-12345678"}</Details>
        <StatusCode>InProgress</StatusCode>
      </member>
    </Activities>
  </EnterStandbyResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</EnterStandbyResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

ExecutePolicy

Executes the specified policy.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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: No

BreachThreshold

The breach threshold for the alarm.

This parameter is required if the policy type is `StepScaling` and not supported otherwise.

Type: Double

Required: No

HonorCooldown

Indicates whether Auto Scaling waits for the cooldown period to complete before executing the policy.

This parameter is not supported if the policy type is `StepScaling`.

For more information, see Auto Scaling Cooldowns in the *Auto Scaling User Guide*.

Type: Boolean

Required: No

MetricValue

The metric value to compare to `BreachThreshold`. This enables you to execute a policy of type `StepScaling` and determine which step adjustment to use. For example, if the breach threshold is 50 and you want to use a step adjustment with a lower bound of 0 and an upper bound of 10, you can set the metric value to 59.

If you specify a metric value that doesn't correspond to a step adjustment for the policy, the call returns an error.

This parameter is required if the policy type is `StepScaling` and not supported otherwise.

Type: Double

Required: No

PolicyName

The name or ARN of the 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]*`

Required: Yes

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

ScalingActivityInProgress

The operation can't be performed because there are scaling activities in progress.

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

ExitStandby

Moves the specified instances out of the standby state.

For more information, see Temporarily Removing Instances from Your Auto Scaling Group in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

InstanceIds.member.N

The IDs of the instances. You can specify up to 20 instances.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 19.

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

Required: No

Response Elements

The following element is returned by the service.

Activities.member.N

The activities related to moving instances out of `Standby` mode.

Type: Array of Activity (p. 138) objects

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=ExitStandby
&AutoScalingGroupName=my-asg
&InstanceIds.member.1=i-5b73d709
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<ExitStandbyResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <ExitStandbyResult>
    <Activities>
      <member>
        <ActivityId>12345678-1234-1234-1234-123456789012</ActivityId>
        <AutoScalingGroupName>my-asg</AutoScalingGroupName>
        <Description>Moving EC2 instance out of Standby: i-12345678</Description>
        <Progress>30</Progress>
        <Cause>At 2015-06-13T22:43:53Z instance i-5b73d709 was moved out of standby in
 response to a user request, increasing the capacity from 3 to 4.</Cause>
        <StartTime>2015-06-13T22:43:53.523Z</StartTime>
        <Details>{"Availability Zone":"us-east-1a","SubnetID":"subnet-12345678"}</Details>
        <StatusCode>PreInService</StatusCode>
      </member>
    </Activities>
  </ExitStandbyResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</ExitStandbyResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

PutLifecycleHook

Creates or updates a lifecycle hook for the specified Auto Scaling Group.

A lifecycle hook tells Auto Scaling that you want to perform an action on an instance that is not actively in service; for example, either when the instance launches or before the instance terminates.

This step is a part of the procedure for adding a lifecycle hook to an Auto Scaling group:

1. (Optional) Create a Lambda function and a rule that allows CloudWatch Events to invoke your Lambda function when Auto Scaling launches or terminates instances.
2. (Optional) Create a notification target and an IAM role. The target can be either an Amazon SQS queue or an Amazon SNS topic. The role allows Auto Scaling to publish lifecycle notifications to the target.
3. **Create the lifecycle hook. Specify whether the hook is used when the instances launch or terminate.**
4. If you need more time, record the lifecycle action heartbeat to keep the instance in a pending state.
5. If you finish before the timeout period ends, complete the lifecycle action.

For more information, see Auto Scaling Lifecycle Hooks in the *Auto Scaling User Guide*.

If you exceed your maximum limit of lifecycle hooks, which by default is 50 per Auto Scaling group, the call fails. For information about updating this limit, see AWS Service Limits in the *Amazon Web Services General Reference*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

DefaultResult

Defines the action the Auto Scaling group should take when the lifecycle hook timeout elapses or if an unexpected failure occurs. This parameter can be either `CONTINUE` or `ABANDON`. The default value is `ABANDON`.

Type: String

Required: No

HeartbeatTimeout

The maximum time, in seconds, that can elapse before the lifecycle hook times out. The range is from 30 to 7200 seconds. The default is 3600 seconds (1 hour).

If the lifecycle hook times out, Auto Scaling performs the default action. You can prevent the lifecycle hook from timing out by calling RecordLifecycleActionHeartbeat (p. 118).

Type: Integer

Required: No

LifecycleHookName

The name of the lifecycle hook.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[A-Za-z0-9\-_\/]+`

Required: Yes

LifecycleTransition

The instance state to which you want to attach the lifecycle hook. For a list of lifecycle hook types, see DescribeLifecycleHookTypes (p. 59).

This parameter is required for new lifecycle hooks, but optional when updating existing hooks.

Type: String

Required: No

NotificationMetadata

Contains additional information that you want to include any time Auto Scaling sends a message to the notification target.

Type: String

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

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

Required: No

NotificationTargetARN

The ARN of the notification target that Auto Scaling will use to notify you when an instance is in the transition state for the lifecycle hook. This target can be either an SQS queue or an SNS topic. If you specify an empty string, this overrides the current ARN.

This operation uses the JSON format when sending notifications to an Amazon SQS queue, and an email key/value pair format when sending notifications to an Amazon SNS topic.

When you specify a notification target, Auto Scaling sends it a test message. Test messages contains the following additional key/value pair: `"Event": "autoscaling:TEST_NOTIFICATION"`.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1600.

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

Required: No

RoleARN

The ARN of the IAM role that allows the Auto Scaling group to publish to the specified notification target.

This parameter is required for new lifecycle hooks, but optional when updating existing hooks.

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: No

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

LimitExceeded

You have already reached a limit for your Auto Scaling resources (for example, groups, launch configurations, or lifecycle hooks). For more information, see DescribeAccountLimits (p. 42).

HTTP Status Code: 400

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
http://autoscaling.amazonaws.com/?Action=PutLifecycleHook
&LifecycleHookName=my-launch-hook
&AutoScalingGroupName=my-asg
&LifecycleTransition=autoscaling:EC2_INSTANCE_LAUNCHING
&NotificationTargetARN=arn:aws:sqs:us-east-1:123456789012:my-queue
&RoleARN=arn:aws:iam::123456789012:role/my-auto-scaling-role
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python

- AWS SDK for Ruby V2

PutNotificationConfiguration

Configures an Auto Scaling group to send notifications when specified events take place. Subscribers to the specified topic can have messages delivered to an endpoint such as a web server or an email address.

This configuration overwrites any existing configuration.

For more information see Getting SNS Notifications When Your Auto Scaling Group Scales in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

NotificationTypes.member.N

The type of event that will cause the notification to be sent. For details about notification types supported by Auto Scaling, see DescribeAutoScalingNotificationTypes (p. 52).

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

TopicARN

The Amazon Resource Name (ARN) of the Amazon Simple Notification Service (SNS) topic.

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

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

LimitExceeded

You have already reached a limit for your Auto Scaling resources (for example, groups, launch configurations, or lifecycle hooks). For more information, see DescribeAccountLimits (p. 42).

HTTP Status Code: 400

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=PutNotificationConfiguration
&AutoScalingGroupName=my-asg
&TopicARN=arn:aws:us-east-1:123456789012:my-sns-topic
&NotificationTypes.member.1=autoscaling:EC2_INSTANCE_LAUNCH
&Version=2011-01-01
&AUTHPARAMS
```

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 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 Auto Scaling group. To update an existing policy, use the existing policy name and set the parameters you want to change. Any existing parameter not changed in an update to an existing policy is not changed in this update request.

If you exceed your maximum limit of step adjustments, which by default is 20 per region, the call fails. For information about updating this limit, see AWS Service Limits in the *Amazon Web Services General Reference*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AdjustmentType

The adjustment type. The valid values are `ChangeInCapacity`, `ExactCapacity`, and `PercentChangeInCapacity`.

This parameter is supported if the policy type is `SimpleScaling` or `StepScaling`.

For more information, see Dynamic Scaling in the *Auto Scaling User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

AutoScalingGroupName

The name of the Auto Scaling group.

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

Cooldown

The amount of time, in seconds, after a scaling activity completes and before the next scaling activity can start. If this parameter is not specified, the default cooldown period for the group applies.

This parameter is supported if the policy type is `SimpleScaling`.

For more information, see Auto Scaling Cooldowns in the *Auto Scaling User Guide*.

Type: Integer

Required: No

EstimatedInstanceWarmup

The estimated time, in seconds, until a newly launched instance can contribute to the CloudWatch metrics. The default is to use the value specified for the default cooldown period for the group.

This parameter is supported if the policy type is `StepScaling` or `TargetTrackingScaling`.

Type: Integer

Required: No

MetricAggregationType

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

This parameter is supported if the policy type is `StepScaling`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

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

Required: No

MinAdjustmentMagnitude

The minimum number of instances to scale. If the value of `AdjustmentType` is `PercentChangeInCapacity`, the scaling policy changes the `DesiredCapacity` of the Auto Scaling group by at least this many instances. Otherwise, the error is `ValidationError`.

This parameter is supported if the policy type is `SimpleScaling` or `StepScaling`.

Type: Integer

Required: No

MinAdjustmentStep

This parameter has been deprecated.

Available for backward compatibility. Use `MinAdjustmentMagnitude` instead.

Type: Integer

Required: No

PolicyName

The name of the policy.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

PolicyType

The policy type. The valid values are `SimpleScaling`, `StepScaling`, and `TargetTrackingScaling`. If the policy type is null, the value is treated as `SimpleScaling`.

Type: String

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

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

Required: No

ScalingAdjustment

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

This parameter is required if the policy type is `SimpleScaling` and not supported otherwise.

Type: Integer

Required: No

StepAdjustments.member.N

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

This parameter is required if the policy type is `StepScaling` and not supported otherwise.

Type: Array of StepAdjustment (p. 179) objects

Required: No

TargetTrackingConfiguration

A target tracking policy.

This parameter is required if the policy type is `TargetTrackingScaling` and not supported otherwise.

Type: TargetTrackingConfiguration (p. 186) object

Required: No

Response Elements

The following elements are returned by the service.

Alarms.member.N

The CloudWatch alarms created for the target tracking policy.

Type: Array of Alarm (p. 141) objects

PolicyARN

The Amazon Resource Name (ARN) of the 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. 189).

LimitExceeded

You have already reached a limit for your Auto Scaling resources (for example, groups, launch configurations, or lifecycle hooks). For more information, see DescribeAccountLimits (p. 42).

HTTP Status Code: 400

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=PutScalingPolicy
&AutoScalingGroupName=my-asg
&ScalingAdjustment=30
&AdjustmentType=PercentChangeInCapacity
&PolicyName=my-scaleout-policy
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<PutScalingPolicyResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <PutScalingPolicyResult>
    <PolicyARN>arn:aws:autoscaling:us-
east-1:123456789012:scalingPolicy:b0dcf5e8-02e6-4e31-9719-0675dEXAMPLE:autoScalingGroupName/
my-asg:policyName/my-scaleout-policy</PolicyARN>
  </PutScalingPolicyResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</PutScalingPolicyResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

PutScheduledUpdateGroupAction

Creates or updates a scheduled scaling action for an Auto Scaling group. When updating a scheduled scaling action, if you leave a parameter unspecified, the corresponding value remains unchanged.

For more information, see Scheduled Scaling in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

DesiredCapacity

The number of EC2 instances that should be running in the group.

Type: Integer

Required: No

EndTime

The time for the recurring schedule to end. Auto Scaling does not perform the action after this time.

Type: Timestamp

Required: No

MaxSize

The maximum size for the Auto Scaling group.

Type: Integer

Required: No

MinSize

The minimum size for the Auto Scaling group.

Type: Integer

Required: No

Recurrence

The recurring schedule for this action, in Unix cron syntax format. For more information, see Cron in Wikipedia.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

ScheduledActionName

The name of this scaling action.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

StartTime

The time for this action to start, in "YYYY-MM-DDThh:mm:ssZ" format in UTC/GMT only (for example, `2014-06-01T00:00:00Z`).

If you specify `Recurrence` and `StartTime`, Auto Scaling performs the action at this time, and then performs the action based on the specified recurrence.

If you try to schedule your action in the past, Auto Scaling returns an error message.

Type: Timestamp

Required: No

Time

This parameter is deprecated.

Type: Timestamp

Required: No

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

AlreadyExists

You already have an Auto Scaling group or launch configuration with this name.

HTTP Status Code: 400

LimitExceeded

You have already reached a limit for your Auto Scaling resources (for example, groups, launch configurations, or lifecycle hooks). For more information, see DescribeAccountLimits (p. 42).

HTTP Status Code: 400

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Examples

Example 1: Schedule based on a specific date and time

Sample Request

```
https://autoscaling.amazonaws.com/?Action=PutScheduledUpdateGroupAction
&AutoScalingGroupName=my-asg
&ScheduledActionName=ScaleUp
&StartTime=2013-05-25T08:00:00Z
&DesiredCapacity=3
&Version=2011-01-01
&AUTHPARAMS
```

Example 2: Recurring Schedule

Sample Request

```
https://autoscaling.amazonaws.com/?Action="PutScheduledUpdateGroupAction
&AutoScalingGroupName=my-asg
&ScheduledActionName=scaleup-schedule-year
&Recurrence="30 0 1 1,6,12 *"
&DesiredCapacity=3
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

RecordLifecycleActionHeartbeat

Records a heartbeat for the lifecycle action associated with the specified token or instance. This extends the timeout by the length of time defined using PutLifecycleHook (p. 105).

This step is a part of the procedure for adding a lifecycle hook to an Auto Scaling group:

1. (Optional) Create a Lambda function and a rule that allows CloudWatch Events to invoke your Lambda function when Auto Scaling launches or terminates instances.
2. (Optional) Create a notification target and an IAM role. The target can be either an Amazon SQS queue or an Amazon SNS topic. The role allows Auto Scaling to publish lifecycle notifications to the target.
3. Create the lifecycle hook. Specify whether the hook is used when the instances launch or terminate.
4. **If you need more time, record the lifecycle action heartbeat to keep the instance in a pending state.**
5. If you finish before the timeout period ends, complete the lifecycle action.

For more information, see Auto Scaling Lifecycle in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

InstanceId

The ID of the instance.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 19.

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

Required: No

LifecycleActionToken

A token that uniquely identifies a specific lifecycle action associated with an instance. Auto Scaling sends this token to the notification target you specified when you created the lifecycle hook.

Type: String

Length Constraints: Fixed length of 36.

Required: No

LifecycleHookName

The name of the lifecycle hook.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[A-Za-z0-9\-_\/]+`

Required: Yes

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

ResumeProcesses

Resumes the specified suspended Auto Scaling processes, or all suspended process, for the specified Auto Scaling group.

For more information, see Suspending and Resuming Auto Scaling Processes in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

ScalingProcesses.member.N

One or more of the following processes. If you omit this parameter, all processes are specified.

- `Launch`
- `Terminate`
- `HealthCheck`
- `ReplaceUnhealthy`
- `AZRebalance`
- `AlarmNotification`
- `ScheduledActions`
- `AddToLoadBalancer`

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

ResourceInUse

The operation can't be performed because the resource is in use.

HTTP Status Code: 400

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=ResumeProcesses
&AutoScalingGroupName=my-asg
&ScalingProcesses.member.1=AlarmNotification
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

SetDesiredCapacity

Sets the size of the specified Auto Scaling group.

For more information about desired capacity, see What Is Auto Scaling? in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

DesiredCapacity

The number of EC2 instances that should be running in the Auto Scaling group.

Type: Integer

Required: Yes

HonorCooldown

Indicates whether Auto Scaling waits for the cooldown period to complete before initiating a scaling activity to set your Auto Scaling group to its new capacity. By default, Auto Scaling does not honor the cooldown period during manual scaling activities.

Type: Boolean

Required: No

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

ScalingActivityInProgress

The operation can't be performed because there are scaling activities in progress.

HTTP Status Code: 400

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=SetDesiredCapacity
&AutoScalingGroupName=my-asg
&HonorCooldown=false
&DesiredCapacity=2
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

SetInstanceHealth

Sets the health status of the specified instance.

For more information, see Health Checks in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

HealthStatus

The health status of the instance. Set to `Healthy` if you want the instance to remain in service. Set to `Unhealthy` if you want the instance to be out of service. Auto Scaling will terminate and replace the unhealthy instance.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

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

Required: Yes

InstanceId

The ID of the instance.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 19.

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

Required: Yes

ShouldRespectGracePeriod

If the Auto Scaling group of the specified instance has a `HealthCheckGracePeriod` specified for the group, by default, this call will respect the grace period. Set this to `False`, if you do not want the call to respect the grace period associated with the group.

For more information, see the description of the health check grace period for CreateAutoScalingGroup (p. 15).

Type: Boolean

Required: No

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=SetInstanceHealth
&InstanceId=i-12345678
&HealthStatus=Unhealthy
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

SetInstanceProtection

Updates the instance protection settings of the specified instances.

For more information, see Instance Protection in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

InstanceIds.member.N

One or more instance IDs.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 19.

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

Required: Yes

ProtectedFromScaleIn

Indicates whether the instance is protected from termination by Auto Scaling when scaling in.

Type: Boolean

Required: Yes

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

LimitExceeded

You have already reached a limit for your Auto Scaling resources (for example, groups, launch configurations, or lifecycle hooks). For more information, see DescribeAccountLimits (p. 42).

HTTP Status Code: 400

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=SetInstanceProtection
&AutoScalingGroupName=my-asg
&InstanceIds.member.1=i-12345678
&ProtectedFromScaleIn=false
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

SuspendProcesses

Suspends the specified Auto Scaling processes, or all processes, for the specified Auto Scaling group.

Note that if you suspend either the `Launch` or `Terminate` process types, it can prevent other process types from functioning properly.

To resume processes that have been suspended, use ResumeProcesses (p. 120).

For more information, see Suspending and Resuming Auto Scaling Processes in the *Auto Scaling User Guide*.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

ScalingProcesses.member.N

One or more of the following processes. If you omit this parameter, all processes are specified.

- `Launch`
- `Terminate`
- `HealthCheck`
- `ReplaceUnhealthy`
- `AZRebalance`
- `AlarmNotification`
- `ScheduledActions`
- `AddToLoadBalancer`

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

ResourceInUse

The operation can't be performed because the resource is in use.

HTTP Status Code: 400

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=SuspendProcesses
&AutoScalingGroupName=my-asg
&ScalingProcesses.member.1=AlarmNotification
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

TerminateInstanceInAutoScalingGroup

Terminates the specified instance and optionally adjusts the desired group size.

This call simply makes a termination request. The instance is not terminated immediately.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

InstanceId

The ID of the instance.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 19.

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

Required: Yes

ShouldDecrementDesiredCapacity

Indicates whether terminating the instance also decrements the size of the Auto Scaling group.

Type: Boolean

Required: Yes

Response Elements

The following element is returned by the service.

Activity

A scaling activity.

Type: Activity (p. 138) object

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

ScalingActivityInProgress

The operation can't be performed because there are scaling activities in progress.

HTTP Status Code: 400

Example

Sample Request

```
https://autoscaling.amazonaws.com/?Action=TerminateInstanceInAutoScalingGroup
&InstanceId=i-12345678
&ShouldDecrementDesiredCapacity=true
&Version=2011-01-01
&AUTHPARAMS
```

Sample Response

```
<TerminateInstanceInAutoScalingGroupResponse xmlns="http://autoscaling.amazonaws.com/
doc/2011-01-01/">
  <TerminateInstanceInAutoScalingGroupResult>
    <Activity>
      <ActivityId>12345678-1234-1234-1234-123456789012</ActivityId>
      <Description>Terminating EC2 instance: i-12345678</Description>
      <Progress>0</Progress>
      <Cause>At 2015-06-14T00:07:30Z instance i-12345678 was taken out of service in
 response to a user request, shrinking the capacity from 4 to 3.</Cause>
      <StartTime>2015-06-14T00:07:30.280Z</StartTime>
      <Details>{"Availability Zone":"us-east-1a","SubnetID":"subnet-12345678"}</Details>
      <StatusCode>InProgress</StatusCode>
    </Activity>
  </TerminateInstanceInAutoScalingGroupResult>
  <ResponseMetadata>
    <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>
  </ResponseMetadata>
</TerminateInstanceInAutoScalingGroupResponse>
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

UpdateAutoScalingGroup

Updates the configuration for the specified Auto Scaling group.

The new settings take effect on any scaling activities after this call returns. Scaling activities that are currently in progress aren't affected.

To update an Auto Scaling group with a launch configuration with `InstanceMonitoring` set to `false`, you must first disable the collection of group metrics. Otherwise, you will get an error. If you have previously enabled the collection of group metrics, you can disable it using DisableMetricsCollection (p. 94).

Note the following:

- If you specify a new value for `MinSize` without specifying a value for `DesiredCapacity`, and the new `MinSize` is larger than the current size of the group, we implicitly call SetDesiredCapacity (p. 122) to set the size of the group to the new value of `MinSize`.
- If you specify a new value for `MaxSize` without specifying a value for `DesiredCapacity`, and the new `MaxSize` is smaller than the current size of the group, we implicitly call SetDesiredCapacity (p. 122) to set the size of the group to the new value of `MaxSize`.
- All other optional parameters are left unchanged if not specified.

Request Parameters

For information about the parameters that are common to all actions, see Common Parameters (p. 187).

AutoScalingGroupName

The name of the Auto Scaling group.

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

AvailabilityZones.member.N

One or more Availability Zones for the group.

Type: Array of strings

Array Members: Minimum number of 1 item.

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

DefaultCooldown

The amount of time, in seconds, after a scaling activity completes before another scaling activity can start. The default is 300.

For more information, see Auto Scaling Cooldowns in the *Auto Scaling User Guide*.

Type: Integer

Required: No

DesiredCapacity

The number of EC2 instances that should be running in the Auto Scaling group. This number must be greater than or equal to the minimum size of the group and less than or equal to the maximum size of the group.

Type: Integer

Required: No

HealthCheckGracePeriod

The amount of time, in seconds, that Auto Scaling waits before checking the health status of an EC2 instance that has come into service. The default is 0.

For more information, see Health Checks in the *Auto Scaling User Guide*.

Type: Integer

Required: No

HealthCheckType

The service to use for the health checks. The valid values are `EC2` and `ELB`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

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

Required: No

LaunchConfigurationName

The name of the launch configuration. If you specify a launch configuration, you can't specify a launch template.

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: No

LaunchTemplate

The launch template to use to specify the updates. If you specify a launch template, you can't specify a launch configuration.

Type: LaunchTemplateSpecification (p. 161) object

Required: No

MaxSize

The maximum size of the Auto Scaling group.

Type: Integer

Required: No

MinSize

The minimum size of the Auto Scaling group.

Type: Integer

Required: No

NewInstancesProtectedFromScaleIn

Indicates whether newly launched instances are protected from termination by Auto Scaling when scaling in.

Type: Boolean

Required: No

PlacementGroup

The name of the placement group into which you'll launch your instances, if any. For more information, see Placement Groups in the *Amazon Elastic Compute Cloud User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

TerminationPolicies.member.N

A standalone termination policy or a list of termination policies used to select the instance to terminate. The policies are executed in the order that they are listed.

For more information, see Controlling Which Instances Auto Scaling Terminates During Scale In in the *Auto Scaling User Guide*.

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

VPCZoneIdentifier

The ID of the subnet, if you are launching into a VPC. You can specify several subnets in a comma-separated list.

When you specify `VPCZoneIdentifier` with `AvailabilityZones`, ensure that the subnets' Availability Zones match the values you specify for `AvailabilityZones`.

For more information, see Launching Auto Scaling Instances in a VPC in the *Auto Scaling User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2047.

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

Required: No

Errors

For information about the errors that are common to all actions, see Common Errors (p. 189).

ResourceContention

You already have a pending update to an Auto Scaling resource (for example, a group, instance, or load balancer).

HTTP Status Code: 500

ScalingActivityInProgress

The operation can't be performed because there are scaling activities in progress.

HTTP Status Code: 400

Examples

Example 1: Update existing Auto Scaling group with ELB health check

Sample Request

```
https://autoscaling.amazonaws.com/?Action=UpdateAutoScalingGroup
&HealthCheckType=ELB
&HealthCheckGracePeriod=300
&AutoScalingGroupName=my-asg
&Version=2011-01-01
&AUTHPARAMS
```

Example 2: Update existing Auto Scaling group with a new Availability Zone

Sample Request

```
https://autoscaling.amazonaws.com/?Action=UpdateAutoScalingGroup
&AutoScalingGroupName=my-asg-lbs
&AvailabilityZones.member.1=us-east-1a
&AvailabilityZones.member.2=us-east-1b
&AvailabilityZones.member.3=us-east-1c
&MinSize=3
&Version=2011-01-01
&AUTHPARAMS
```

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 Java
- AWS SDK for JavaScript
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V2

Data Types

The 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:

Activity

Describes scaling activity, which is a long-running process that represents a change to your Auto Scaling group, such as changing its size or replacing an instance.

Contents

ActivityId

The ID of the activity.

Type: String

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

Required: Yes

AutoScalingGroupName

The name of the Auto Scaling group.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

Cause

The reason the activity began.

Type: String

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

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

Required: Yes

Description

A friendly, more verbose description of the activity.

Type: String

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

Required: No

Details

The details about the activity.

Type: String

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

Required: No

EndTime

The end time of the activity.

Type: Timestamp

Required: No

Progress

A value between 0 and 100 that indicates the progress of the activity.

Type: Integer

Required: No

StartTime

The start time of the activity.

Type: Timestamp

Required: Yes

StatusCode

The current status of the activity.

Type: String

Valid Values: `PendingSpotBidPlacement | WaitingForSpotInstanceRequestId | WaitingForSpotInstanceId | WaitingForInstanceId | PreInService | InProgress | WaitingForELBConnectionDraining | MidLifecycleAction | WaitingForInstanceWarmup | Successful | Failed | Cancelled`

Required: Yes

StatusMessage

A friendly, more verbose description of the activity status.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

AdjustmentType

Describes a policy adjustment type.

For more information, see Dynamic Scaling in the *Auto Scaling User Guide*.

Contents

AdjustmentType

The policy adjustment type. The valid values are `ChangeInCapacity`, `ExactCapacity`, and `PercentChangeInCapacity`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

Alarm

Describes an alarm.

Contents

AlarmARN

The Amazon Resource Name (ARN) of the alarm.

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: No

AlarmName

The name of the alarm.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

AutoScalingGroup

Describes an Auto Scaling group.

Contents

AutoScalingGroupARN

The Amazon Resource Name (ARN) of the Auto Scaling group.

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: No

AutoScalingGroupName

The name of the Auto Scaling group.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

AvailabilityZones.member.N

One or more Availability Zones for the group.

Type: Array of strings

Array Members: Minimum number of 1 item.

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

CreatedTime

The date and time the group was created.

Type: Timestamp

Required: Yes

DefaultCooldown

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

Type: Integer

Required: Yes

DesiredCapacity

The desired size of the group.

Type: Integer

Required: Yes

EnabledMetrics.member.N

The metrics enabled for the group.

Type: Array of EnabledMetric (p. 152) objects

Required: No

HealthCheckGracePeriod

The amount of time, in seconds, that Auto Scaling waits before checking the health status of an EC2 instance that has come into service.

Type: Integer

Required: No

HealthCheckType

The service to use for the health checks. The valid values are `EC2` and `ELB`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

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

Required: Yes

Instances.member.N

The EC2 instances associated with the group.

Type: Array of Instance (p. 154) objects

Required: No

LaunchConfigurationName

The name of the associated launch configuration.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

LaunchTemplate

The launch template for the group.

Type: LaunchTemplateSpecification (p. 161) object

Required: No

LoadBalancerNames.member.N

One or more load balancers associated with the group.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

MaxSize

The maximum size of the group.

Type: Integer

Required: Yes

MinSize

The minimum size of the group.

Type: Integer

Required: Yes

NewInstancesProtectedFromScaleIn

Indicates whether newly launched instances are protected from termination by Auto Scaling when scaling in.

Type: Boolean

Required: No

PlacementGroup

The name of the placement group into which you'll launch your instances, if any. For more information, see Placement Groups in the *Amazon Elastic Compute Cloud User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

Status

The current state of the group when DeleteAutoScalingGroup (p. 28) is in progress.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

SuspendedProcesses.member.N

The suspended processes associated with the group.

Type: Array of SuspendedProcess (p. 181) objects

Required: No

Tags.member.N

The tags for the group.

Type: Array of TagDescription (p. 184) objects

Required: No

TargetGroupARNs.member.N

The Amazon Resource Names (ARN) of the target groups for your load balancer.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 511.

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

Required: No

TerminationPolicies.member.N

The termination policies for the group.

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

VPCZoneIdentifier

One or more subnet IDs, if applicable, separated by commas.

If you specify `VPCZoneIdentifier` and `AvailabilityZones`, ensure that the Availability Zones of the subnets match the values for `AvailabilityZones`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2047.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

AutoScalingInstanceDetails

Describes an EC2 instance associated with an Auto Scaling group.

Contents

AutoScalingGroupName

The name of the Auto Scaling group for the instance.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

AvailabilityZone

The Availability Zone for the instance.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

HealthStatus

The last reported health status of this instance. "Healthy" means that the instance is healthy and should remain in service. "Unhealthy" means that the instance is unhealthy and Auto Scaling should terminate and replace it.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

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

Required: Yes

InstanceId

The ID of the instance.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 19.

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

Required: Yes

LaunchConfigurationName

The launch configuration used to launch the instance. This value is not available if you attached the instance to the Auto Scaling group.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

LaunchTemplate

The launch template for the instance.

Type: LaunchTemplateSpecification (p. 161) object

Required: No

LifecycleState

The lifecycle state for the instance. For more information, see Auto Scaling Lifecycle in the *Auto Scaling User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

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

Required: Yes

ProtectedFromScaleIn

Indicates whether the instance is protected from termination by Auto Scaling when scaling in.

Type: Boolean

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 Java
- AWS SDK for Ruby V2

BlockDeviceMapping

Describes a block device mapping.

Contents

DeviceName

The device name exposed to the EC2 instance (for example, `/dev/sdh` or `xvdh`).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

Ebs

The information about the Amazon EBS volume.

Type: Ebs (p. 150) object

Required: No

NoDevice

Suppresses a device mapping.

If this parameter is true for the root device, the instance might fail the EC2 health check. Auto Scaling launches a replacement instance if the instance fails the health check.

Type: Boolean

Required: No

VirtualName

The name of the virtual device (for example, `ephemeral0`).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

CustomizedMetricSpecification

Configures a customized metric for a target tracking policy.

Contents

Dimensions.member.N

The dimensions of the metric.

Type: Array of MetricDimension (p. 169) 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 Java
- AWS SDK for Ruby V2

Ebs

Describes an Amazon EBS volume.

Contents

DeleteOnTermination

Indicates whether the volume is deleted on instance termination. The default is `true`.

Type: Boolean

Required: No

Encrypted

Indicates whether the volume should be encrypted. Encrypted EBS volumes must be attached to instances that support Amazon EBS encryption. Volumes that are created from encrypted snapshots are automatically encrypted. There is no way to create an encrypted volume from an unencrypted snapshot or an unencrypted volume from an encrypted snapshot. For more information, see Amazon EBS Encryption in the *Amazon Elastic Compute Cloud User Guide*.

Type: Boolean

Required: No

Iops

The number of I/O operations per second (IOPS) to provision for the volume.

Constraint: Required when the volume type is `io1`.

Type: Integer

Valid Range: Minimum value of 100. Maximum value of 20000.

Required: No

SnapshotId

The ID of the snapshot.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

VolumeSize

The volume size, in GiB. For `standard` volumes, specify a value from 1 to 1,024. For `io1` volumes, specify a value from 4 to 16,384. For `gp2` volumes, specify a value from 1 to 16,384. If you specify a snapshot, the volume size must be equal to or larger than the snapshot size.

Default: If you create a volume from a snapshot and you don't specify a volume size, the default is the snapshot size.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 16384.

Required: No

VolumeType

The volume type. For more information, see Amazon EBS Volume Types in the *Amazon Elastic Compute Cloud User Guide*.

Valid values: `standard` | `io1` | `gp2`

Default: `standard`

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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 Java
- AWS SDK for Ruby V2

EnabledMetric

Describes an enabled metric.

Contents

Granularity

The granularity of the metric. The only valid value is `1Minute`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

Metric

One of the following metrics:

- `GroupMinSize`
- `GroupMaxSize`
- `GroupDesiredCapacity`
- `GroupInServiceInstances`
- `GroupPendingInstances`
- `GroupStandbyInstances`
- `GroupTerminatingInstances`
- `GroupTotalInstances`

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

Filter

Describes a filter.

Contents

Name

The name of the filter. The valid values are: `"auto-scaling-group"`, `"key"`, `"value"`, and `"propagate-at-launch"`.

Type: String

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

Required: No

Values.member.N

The value of the filter.

Type: Array of strings

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

Instance

Describes an EC2 instance.

Contents

AvailabilityZone

The Availability Zone in which the instance is running.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

HealthStatus

The last reported health status of the instance. "Healthy" means that the instance is healthy and should remain in service. "Unhealthy" means that the instance is unhealthy and Auto Scaling should terminate and replace it.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

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

Required: Yes

InstanceId

The ID of the instance.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 19.

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

Required: Yes

LaunchConfigurationName

The launch configuration associated with the instance.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

LaunchTemplate

The launch template for the instance.

Type: LaunchTemplateSpecification (p. 161) object

Required: No

LifecycleState

A description of the current lifecycle state. Note that the `Quarantined` state is not used.

Type: String

Valid Values: `Pending | Pending:Wait | Pending:Proceed | Quarantined | InService | Terminating | Terminating:Wait | Terminating:Proceed | Terminated | Detaching | Detached | EnteringStandby | Standby`

Required: Yes

ProtectedFromScaleIn

Indicates whether the instance is protected from termination by Auto Scaling when scaling in.

Type: Boolean

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 Java
- AWS SDK for Ruby V2

InstanceMonitoring

Describes whether detailed monitoring is enabled for the Auto Scaling instances.

Contents

Enabled

If `true`, detailed monitoring is enabled. Otherwise, basic monitoring is enabled.

Type: Boolean

Required: No

See Also

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

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

LaunchConfiguration

Describes a launch configuration.

Contents

AssociatePublicIpAddress

[EC2-VPC] Indicates whether to assign a public IP address to each instance.

Type: Boolean

Required: No

BlockDeviceMappings.member.N

A block device mapping, which specifies the block devices for the instance.

Type: Array of BlockDeviceMapping (p. 148) objects

Required: No

ClassicLinkVPCId

The ID of a ClassicLink-enabled VPC to link your EC2-Classic instances to. This parameter can only be used if you are launching EC2-Classic instances. For more information, see ClassicLink in the *Amazon Elastic Compute Cloud User Guide*.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

ClassicLinkVPCSecurityGroups.member.N

The IDs of one or more security groups for the VPC specified in `ClassicLinkVPCId`. This parameter is required if you specify a ClassicLink-enabled VPC, and cannot be used otherwise. For more information, see ClassicLink in the *Amazon Elastic Compute Cloud User Guide*.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

CreatedTime

The creation date and time for the launch configuration.

Type: Timestamp

Required: Yes

EbsOptimized

Controls whether the instance is optimized for EBS I/O (`true`) or not (`false`).

Type: Boolean

Required: No

IamInstanceProfile

The name or Amazon Resource Name (ARN) of the instance profile associated with the IAM role for the instance.

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: No

ImageId

The ID of the Amazon Machine Image (AMI).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

InstanceMonitoring

Controls whether instances in this group are launched with detailed (`true`) or basic (`false`) monitoring.

Type: InstanceMonitoring (p. 156) object

Required: No

InstanceType

The instance type for the instances.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

KernelId

The ID of the kernel associated with the AMI.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

KeyName

The name of the key pair.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

LaunchConfigurationARN

The Amazon Resource Name (ARN) of the launch configuration.

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: No

LaunchConfigurationName

The name of the launch configuration.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

PlacementTenancy

The tenancy of the instance, either `default` or `dedicated`. An instance with `dedicated` tenancy runs in an isolated, single-tenant hardware and can only be launched into a VPC.

Type: String

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

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

Required: No

RamdiskId

The ID of the RAM disk associated with the AMI.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

SecurityGroups.member.N

The security groups to associate with the instances.

Type: Array of strings

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

Required: No

SpotPrice

The price to bid when launching Spot Instances.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

UserData

The user data available to the instances.

Type: String

Length Constraints: Maximum length of 21847.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

LaunchTemplateSpecification

Describes a launch template.

Contents

LaunchTemplateId

The ID of the launch template. You must specify either a template ID or a template name.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

LaunchTemplateName

The name of the launch template. You must specify either a template name or a template ID.

Type: String

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

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

Required: No

Version

The version number, `$Latest`, or `$Default`. If the value is `$Latest`, Auto Scaling selects the latest version of the launch template when launching instances. If the value is `$Default`, Auto Scaling selects the default version of the launch template when launching instances. The default value is `$Default`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

LifecycleHook

Describes a lifecycle hook, which tells Auto Scaling that you want to perform an action whenever it launches instances or whenever it terminates instances.

For more information, see Auto Scaling Lifecycle Hooks in the *Auto Scaling User Guide*.

Contents

AutoScalingGroupName

The name of the Auto Scaling group for the lifecycle hook.

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: No

DefaultResult

Defines the action the Auto Scaling group should take when the lifecycle hook timeout elapses or if an unexpected failure occurs. The valid values are `CONTINUE` and `ABANDON`. The default value is `CONTINUE`.

Type: String

Required: No

GlobalTimeout

The maximum time, in seconds, that an instance can remain in a `Pending:Wait` or `Terminating:Wait` state. The maximum is 172800 seconds (48 hours) or 100 times `HeartbeatTimeout`, whichever is smaller.

Type: Integer

Required: No

HeartbeatTimeout

The maximum time, in seconds, that can elapse before the lifecycle hook times out. If the lifecycle hook times out, Auto Scaling performs the default action. You can prevent the lifecycle hook from timing out by calling RecordLifecycleActionHeartbeat (p. 118).

Type: Integer

Required: No

LifecycleHookName

The name of the lifecycle hook.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[A-Za-z0-9\-_\/]+`

Required: No

LifecycleTransition

The state of the EC2 instance to which you want to attach the lifecycle hook. For a list of lifecycle hook types, see DescribeLifecycleHookTypes (p. 59).

Type: String

Required: No

NotificationMetadata

Additional information that you want to include any time Auto Scaling sends a message to the notification target.

Type: String

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

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

Required: No

NotificationTargetARN

The ARN of the target that Auto Scaling sends notifications to when an instance is in the transition state for the lifecycle hook. The notification target can be either an SQS queue or an SNS topic.

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: No

RoleARN

The ARN of the IAM role that allows the Auto Scaling group to publish to the specified notification target.

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: 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 Java
- AWS SDK for Ruby V2

LifecycleHookSpecification

Describes a lifecycle hook, which tells Auto Scaling that you want to perform an action whenever it launches instances or whenever it terminates instances.

For more information, see Auto Scaling Lifecycle Hooks in the *Auto Scaling User Guide*.

Contents

DefaultResult

Defines the action the Auto Scaling group should take when the lifecycle hook timeout elapses or if an unexpected failure occurs. The valid values are `CONTINUE` and `ABANDON`.

Type: String

Required: No

HeartbeatTimeout

The maximum time, in seconds, that can elapse before the lifecycle hook times out. If the lifecycle hook times out, Auto Scaling performs the default action. You can prevent the lifecycle hook from timing out by calling RecordLifecycleActionHeartbeat (p. 118).

Type: Integer

Required: No

LifecycleHookName

The name of the lifecycle hook.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[A-Za-z0-9\-_\/]+`

Required: Yes

LifecycleTransition

The state of the EC2 instance to which you want to attach the lifecycle hook. For a list of lifecycle hook types, see DescribeLifecycleHookTypes (p. 59).

Type: String

Required: Yes

NotificationMetadata

Additional information that you want to include any time Auto Scaling sends a message to the notification target.

Type: String

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

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

Required: No

NotificationTargetARN

The ARN of the target that Auto Scaling sends notifications to when an instance is in the transition state for the lifecycle hook. The notification target can be either an SQS queue or an SNS topic.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1600.

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

Required: No

RoleARN

The ARN of the IAM role that allows the Auto Scaling group to publish to the specified notification target.

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: 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 Java
- AWS SDK for Ruby V2

LoadBalancerState

Describes the state of a Classic Load Balancer.

If you specify a load balancer when creating the Auto Scaling group, the state of the load balancer is `InService`.

If you attach a load balancer to an existing Auto Scaling group, the initial state is `Adding`. The state transitions to `Added` after all instances in the group are registered with the load balancer. If ELB health checks are enabled for the load balancer, the state transitions to `InService` after at least one instance in the group passes the health check. If EC2 health checks are enabled instead, the load balancer remains in the `Added` state.

Contents

LoadBalancerName

The name of the load balancer.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

State

One of the following load balancer states:

- `Adding` - The instances in the group are being registered with the load balancer.
- `Added` - All instances in the group are registered with the load balancer.
- `InService` - At least one instance in the group passed an ELB health check.
- `Removing` - The instances in the group are being deregistered from the load balancer. If connection draining is enabled, Elastic Load Balancing waits for in-flight requests to complete before deregistering the instances.
- `Removed` - All instances in the group are deregistered from the load balancer.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

LoadBalancerTargetGroupState

Describes the state of a target group.

If you attach a target group to an existing Auto Scaling group, the initial state is `Adding`. The state transitions to `Added` after all Auto Scaling instances are registered with the target group. If ELB health checks are enabled, the state transitions to `InService` after at least one Auto Scaling instance passes the health check. If EC2 health checks are enabled instead, the target group remains in the `Added` state.

Contents

LoadBalancerTargetGroupARN

The Amazon Resource Name (ARN) of the target group.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 511.

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

Required: No

State

The state of the target group.

- `Adding` - The Auto Scaling instances are being registered with the target group.
- `Added` - All Auto Scaling instances are registered with the target group.
- `InService` - At least one Auto Scaling instance passed an ELB health check.
- `Removing` - The Auto Scaling instances are being deregistered from the target group. If connection draining is enabled, Elastic Load Balancing waits for in-flight requests to complete before deregistering the instances.
- `Removed` - All Auto Scaling instances are deregistered from the target group.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

MetricCollectionType

Describes a metric.

Contents

Metric

One of the following metrics:

- `GroupMinSize`
- `GroupMaxSize`
- `GroupDesiredCapacity`
- `GroupInServiceInstances`
- `GroupPendingInstances`
- `GroupStandbyInstances`
- `GroupTerminatingInstances`
- `GroupTotalInstances`

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

MetricDimension

Describes the dimension of 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 Java
- AWS SDK for Ruby V2

MetricGranularityType

Describes a granularity of a metric.

Contents

Granularity

The granularity. The only valid value is `1Minute`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

NotificationConfiguration

Describes a notification.

Contents

AutoScalingGroupName

The name of the Auto Scaling group.

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: No

NotificationType

One of the following event notification types:

- `autoscaling:EC2_INSTANCE_LAUNCH`
- `autoscaling:EC2_INSTANCE_LAUNCH_ERROR`
- `autoscaling:EC2_INSTANCE_TERMINATE`
- `autoscaling:EC2_INSTANCE_TERMINATE_ERROR`
- `autoscaling:TEST_NOTIFICATION`

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

TopicARN

The Amazon Resource Name (ARN) of the Amazon Simple Notification Service (SNS) topic.

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: 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 Java
- AWS SDK for Ruby V2

PredefinedMetricSpecification

Configures a predefined metric for a target tracking policy.

Contents

PredefinedMetricType

The metric type.

Type: String

Valid Values: `ASGAverageCPUUtilization | ASGAverageNetworkIn | ASGAverageNetworkOut | ALBRequestCountPerTarget`

Required: Yes

ResourceLabel

Identifies the resource associated with the metric type. The following predefined metrics are available:

- `ASGAverageCPUUtilization` - average CPU utilization of the Auto Scaling group
- `ASGAverageNetworkIn` - average number of bytes received on all network interfaces by the Auto Scaling group
- `ASGAverageNetworkOut` - average number of bytes sent out on all network interfaces by the Auto Scaling group
- `ALBRequestCountPerTarget` - number of requests completed per target in an Application Load Balancer target group

For predefined metric types `ASGAverageCPUUtilization`, `ASGAverageNetworkIn`, and `ASGAverageNetworkOut`, the parameter must not be specified as the resource associated with the metric type is the Auto Scaling group. For predefined metric type `ALBRequestCountPerTarget`, the parameter must be specified in the format: `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, and `targetgroup/`*`target-group-name`*`/`*`target-group-id`* is the final portion of the target group ARN. The target group must be attached to the Auto Scaling group.

Type: String

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

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

ProcessType

Describes a process type.

For more information, see Auto Scaling Processes in the *Auto Scaling User Guide*.

Contents

ProcessName

One of the following processes:

- `Launch`
- `Terminate`
- `AddToLoadBalancer`
- `AlarmNotification`
- `AZRebalance`
- `HealthCheck`
- `ReplaceUnhealthy`
- `ScheduledActions`

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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 Java
- AWS SDK for Ruby V2

ScalingPolicy

Describes a scaling policy.

Contents

AdjustmentType

The adjustment type, which specifies how `ScalingAdjustment` is interpreted. Valid values are `ChangeInCapacity`, `ExactCapacity`, and `PercentChangeInCapacity`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

Alarms.member.N

The CloudWatch alarms related to the policy.

Type: Array of Alarm (p. 141) objects

Required: No

AutoScalingGroupName

The name of the Auto Scaling group.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

Cooldown

The amount of time, in seconds, after a scaling activity completes before any further dynamic scaling activities can start.

Type: Integer

Required: No

EstimatedInstanceWarmup

The estimated time, in seconds, until a newly launched instance can contribute to the CloudWatch metrics.

Type: Integer

Required: No

MetricAggregationType

The aggregation type for the CloudWatch metrics. Valid values are `Minimum`, `Maximum`, and `Average`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 32.

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

Required: No

MinAdjustmentMagnitude

The minimum number of instances to scale. If the value of `AdjustmentType` is `PercentChangeInCapacity`, the scaling policy changes the `DesiredCapacity` of the Auto Scaling group by at least this many instances. Otherwise, the error is `ValidationError`.

Type: Integer

Required: No

MinAdjustmentStep

This member has been deprecated.

Available for backward compatibility. Use `MinAdjustmentMagnitude` instead.

Type: Integer

Required: No

PolicyARN

The Amazon Resource Name (ARN) of the 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]*`

Required: No

PolicyName

The name of the scaling policy.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

PolicyType

The policy type. Valid values are `SimpleScaling` and `StepScaling`.

Type: String

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

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

Required: No

ScalingAdjustment

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

Type: Integer

Required: No

StepAdjustments.member.N

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

Type: Array of StepAdjustment (p. 179) objects

Required: No

TargetTrackingConfiguration

A target tracking policy.

Type: TargetTrackingConfiguration (p. 186) 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 Java
- AWS SDK for Ruby V2

ScheduledUpdateGroupAction

Describes a scheduled update to an Auto Scaling group.

Contents

AutoScalingGroupName

The name of the Auto Scaling group.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

DesiredCapacity

The number of instances you prefer to maintain in the group.

Type: Integer

Required: No

EndTime

The date and time that the action is scheduled to end. This date and time can be up to one month in the future.

Type: Timestamp

Required: No

MaxSize

The maximum size of the group.

Type: Integer

Required: No

MinSize

The minimum size of the group.

Type: Integer

Required: No

Recurrence

The recurring schedule for the action.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

ScheduledActionARN

The Amazon Resource Name (ARN) of the scheduled action.

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: No

ScheduledActionName

The name of the scheduled action.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

StartTime

The date and time that the action is scheduled to begin. This date and time can be up to one month in the future.

When `StartTime` and `EndTime` are specified with `Recurrence`, they form the boundaries of when the recurring action will start and stop.

Type: Timestamp

Required: No

Time

This parameter is deprecated.

Type: Timestamp

Required: No

See Also

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

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

StepAdjustment

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:

- If you want the adjustment to be triggered 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.
- If you want the adjustment to be triggered 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 capacity while a negative number removes from the current capacity.

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 Java
- AWS SDK for Ruby V2

SuspendedProcess

Describes an Auto Scaling process that has been suspended. For more information, see ProcessType (p. 173).

Contents

ProcessName

The name of the suspended process.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

SuspensionReason

The reason that the process was suspended.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

Tag

Describes a tag for an Auto Scaling group.

Contents

Key

The tag key.

Type: String

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

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

Required: Yes

PropagateAtLaunch

Determines whether the tag is added to new instances as they are launched in the group.

Type: Boolean

Required: No

ResourceId

The name of the group.

Type: String

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

Required: No

ResourceType

The type of resource. The only supported value is `auto-scaling-group`.

Type: String

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

Required: No

Value

The tag value.

Type: String

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

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

TagDescription

Describes a tag for an Auto Scaling group.

Contents

Key

The tag key.

Type: String

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

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

Required: No

PropagateAtLaunch

Determines whether the tag is added to new instances as they are launched in the group.

Type: Boolean

Required: No

ResourceId

The name of the group.

Type: String

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

Required: No

ResourceType

The type of resource. The only supported value is `auto-scaling-group`.

Type: String

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

Required: No

Value

The tag value.

Type: String

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

Pattern: `[\u0020-\uD7FF\uE000-\uFFFD\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 Java
- AWS SDK for Ruby V2

TargetTrackingConfiguration

Represents a target tracking policy configuration.

Contents

CustomizedMetricSpecification

A customized metric.

Type: CustomizedMetricSpecification (p. 149) object

Required: No

DisableScaleIn

Indicates whether scale in by the target tracking policy is disabled. If scale in is disabled, the target tracking policy won't remove instances from the Auto Scaling group. Otherwise, the target tracking policy can remove instances from the Auto Scaling group. The default is disabled.

Type: Boolean

Required: No

PredefinedMetricSpecification

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

Type: PredefinedMetricSpecification (p. 172) object

Required: No

TargetValue

The target value for the metric.

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 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'T'HHMMSS'Z'). For example, the following date time is a valid X-Amz-Date value: `20120325T120000Z`.

Condition: X-Amz-Date is optional for all requests; it can be used to override the date used for signing requests. If the Date header is specified in the ISO 8601 basic format, X-Amz-Date is

not required. When X-Amz-Date is used, it always overrides the value of the Date header. For more information, see 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

Auto Scaling SOAP API

We have deprecated the SOAP API for Auto Scaling. As of December 4, 2017, if you make a SOAP request, you will receive the following response:

```
Client.InvalidQueryParameter: SOAP is no longer supported
```

We recommend that you use the Query API for Auto Scaling, the AWS CLI, or one of the AWS SDKs. For more information, see Accessing Auto Scaling in the *Amazon EC2 Auto Scaling User Guide*.