
Auto Scaling

API Reference

API Version 2011-01-01



Auto Scaling: API Reference

Copyright © 2014 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

The following are trademarks of Amazon Web Services, Inc.: Amazon, Amazon Web Services Design, AWS, Amazon CloudFront, Cloudfront, Amazon DevPay, DynamoDB, ElastiCache, Amazon EC2, Amazon Elastic Compute Cloud, Amazon Glacier, Kindle, Kindle Fire, AWS Marketplace Design, Mechanical Turk, Amazon Redshift, Amazon Route 53, Amazon S3, Amazon VPC. In addition, Amazon.com graphics, logos, page headers, button icons, scripts, and service names are trademarks, or trade dress of Amazon in the U.S. and/or other countries. Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon.

All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Welcome	1
Actions	2
AttachInstances	4
Description	4
Request Parameters	4
CompleteLifecycleAction	5
Description	5
Request Parameters	5
CreateAutoScalingGroup	7
Description	7
Request Parameters	7
Errors	10
Examples	10
CreateLaunchConfiguration	11
Description	11
Request Parameters	11
Errors	14
Examples	14
CreateOrUpdateTags	16
Description	16
Request Parameters	16
Errors	16
Examples	17
DeleteAutoScalingGroup	18
Description	18
Request Parameters	18
Errors	18
Examples	18
DeleteLaunchConfiguration	20
Description	20
Request Parameters	20
Errors	20
Examples	20
DeleteLifecycleHook	21
Description	21
Request Parameters	21
DeleteNotificationConfiguration	22
Description	22
Request Parameters	22
DeletePolicy	23
Description	23
Request Parameters	23
DeleteScheduledAction	24
Description	24
Request Parameters	24
DeleteTags	25
Description	25
Request Parameters	25
DescribeAccountLimits	26
Description	26
Response Elements	26
Examples	26
DescribeAdjustmentTypes	27
Description	27
Response Elements	27

Examples	27
DescribeAutoScalingGroups	28
Description	28
Request Parameters	28
Response Elements	28
Errors	28
Examples	29
DescribeAutoScalingInstances	31
Description	31
Request Parameters	31
Response Elements	31
Errors	32
Examples	32
DescribeAutoScalingNotificationTypes	33
Description	33
Response Elements	33
DescribeLaunchConfigurations	34
Description	34
Request Parameters	34
Response Elements	34
Errors	34
Examples	35
DescribeLifecycleHookTypes	36
Description	36
Response Elements	36
DescribeLifecycleHooks	37
Description	37
Request Parameters	37
Response Elements	37
DescribeMetricCollectionTypes	38
Description	38
Response Elements	38
Examples	38
DescribeNotificationConfigurations	40
Description	40
Request Parameters	40
Response Elements	40
Errors	40
DescribePolicies	41
Description	41
Request Parameters	41
Response Elements	41
Errors	42
Examples	42
DescribeScalingActivities	44
Description	44
Request Parameters	44
Response Elements	44
Errors	45
Examples	45
DescribeScalingProcessTypes	47
Description	47
Response Elements	47
Examples	47
DescribeScheduledActions	49
Description	49
Request Parameters	49
Response Elements	50

Errors	50
DescribeTags	51
Description	51
Request Parameters	51
Response Elements	51
Errors	52
Examples	52
DescribeTerminationPolicyTypes	53
Description	53
Response Elements	53
Examples	53
DetachInstances	54
Description	54
Request Parameters	54
Response Elements	54
Examples	55
DisableMetricsCollection	56
Description	56
Request Parameters	56
EnableMetricsCollection	57
Description	57
Request Parameters	57
EnterStandby	58
Description	58
Request Parameters	58
Response Elements	58
Examples	58
ExecutePolicy	60
Description	60
Request Parameters	60
Errors	60
ExitStandby	61
Description	61
Request Parameters	61
Response Elements	61
Examples	61
PutLifecycleHook	63
Description	63
Request Parameters	63
Errors	65
Examples	65
PutNotificationConfiguration	66
Description	66
Request Parameters	66
Errors	66
PutScalingPolicy	67
Description	67
Request Parameters	67
Response Elements	68
Errors	68
Examples	68
PutScheduledUpdateGroupAction	70
Description	70
Request Parameters	70
Errors	71
Examples	72
RecordLifecycleActionHeartbeat	73
Description	73

Request Parameters	73
ResumeProcesses	74
Description	74
Request Parameters	74
SetDesiredCapacity	75
Description	75
Request Parameters	75
Errors	75
Examples	75
SetInstanceHealth	77
Description	77
Request Parameters	77
SuspendProcesses	78
Description	78
Request Parameters	78
TerminateInstanceInAutoScalingGroup	79
Description	79
Request Parameters	79
Response Elements	79
Errors	79
UpdateAutoScalingGroup	80
Description	80
Request Parameters	80
Errors	82
Examples	82
Data Types	84
Activity	85
Description	85
Contents	85
AdjustmentType	87
Description	87
Contents	87
Alarm	87
Description	87
Contents	87
AutoScalingGroup	87
Description	87
Contents	88
AutoScalingInstanceDetails	90
Description	90
Contents	90
BlockDeviceMapping	91
Description	91
Contents	91
CompleteLifecycleActionResult	92
Description	92
Contents	92
DeleteLifecycleHookResult	92
Description	92
Contents	92
DescribeAccountLimitsResult	92
Description	92
Contents	93
DescribeAdjustmentTypesResult	93
Description	93
Contents	93
DescribeAutoScalingGroupsResult	93
Description	93

Contents	93
DescribeAutoScalingInstancesResult	94
Description	94
Contents	94
DescribeAutoScalingNotificationTypesResult	94
Description	94
Contents	94
DescribeLaunchConfigurationsResult	94
Description	94
Contents	95
DescribeLifecycleHookTypesResult	95
Description	95
Contents	95
DescribeLifecycleHooksResult	95
Description	95
Contents	95
DescribeMetricCollectionTypesResult	96
Description	96
Contents	96
DescribeNotificationConfigurationsResult	96
Description	96
Contents	96
DescribePoliciesResult	97
Description	97
Contents	97
DescribeScalingActivitiesResult	97
Description	97
Contents	97
DescribeScalingProcessTypesResult	98
Description	98
Contents	98
DescribeScheduledActionsResult	98
Description	98
Contents	98
DescribeTagsResult	98
Description	98
Contents	98
DescribeTerminationPolicyTypesResult	99
Description	99
Contents	99
DetachInstancesResult	99
Description	99
Contents	99
Ebs	99
Description	99
Contents	100
EnabledMetric	101
Description	101
Contents	101
EnterStandbyResult	101
Description	101
Contents	101
ExitStandbyResult	101
Description	101
Contents	102
Filter	102
Description	102
Contents	102

Instance	102
Description	102
Contents	102
InstanceMonitoring	103
Description	103
Contents	103
LaunchConfiguration	104
Description	104
Contents	104
LifecycleHook	106
Description	106
Contents	106
MetricCollectionType	108
Description	108
Contents	108
MetricGranularityType	108
Description	108
Contents	108
NotificationConfiguration	108
Description	108
Contents	109
ProcessType	109
Description	109
Contents	110
PutLifecycleHookResult	111
Description	111
Contents	111
PutScalingPolicyResult	111
Description	111
Contents	111
RecordLifecycleActionHeartbeatResult	111
Description	111
Contents	111
ScalingPolicy	111
Description	111
Contents	112
ScheduledUpdateGroupAction	113
Description	113
Contents	113
SuspendedProcess	114
Description	114
Contents	114
Tag	115
Description	115
Contents	115
TagDescription	115
Description	115
Contents	116
TerminateInstanceInAutoScalingGroupResult	116
Description	116
Contents	116
Common Parameters	117
.....	117
Common Parameters for Signature V4 Signing	119
.....	119
Common Errors	121
.....	121

Welcome

Auto Scaling is a web service designed to automatically launch or terminate Amazon Elastic Compute Cloud (Amazon EC2) instances based on user-defined policies, schedules, and health checks. This service is used in conjunction with Amazon CloudWatch and Elastic Load Balancing services.

Auto Scaling provides APIs that you can call by submitting a Query Request. Query requests are HTTP or HTTPS requests that use the HTTP verbs GET or POST and a Query parameter named *Action* or *Operation* that specifies the API you are calling. Action is used throughout this documentation, although Operation is also supported for backward compatibility with other Amazon Web Services (AWS) Query APIs.

Calling the API using a Query request is the most direct way to access the web service, but requires that your application handle low-level details such as generating the hash to sign the request and error handling. The benefit of calling the service using a Query request is that you are assured of having access to the complete functionality of the API. For information about signing a query request, see [Use Query Requests to Call Auto Scaling APIs](#)

This guide provides detailed information about Auto Scaling actions, data types, parameters, and errors. For detailed information about Auto Scaling features and their associated API actions, go to the [Auto Scaling Developer Guide](#).

This reference is based on the current WSDL, which is available at:

<http://autoscaling.amazonaws.com/doc/2011-01-01/AutoScaling.wsdl>

Endpoints

The examples in this guide assume that your instances are launched in the US East (Northern Virginia) region and use us-east-1 as the endpoint.

You can set up your Auto Scaling infrastructure in other AWS regions. For information about this product's regions and endpoints, see [Regions and Endpoints](#) in the Amazon Web Services General Reference.

This document was last updated on July 30, 2014.

Actions

The following actions are supported:

- [AttachInstances](#) (p. 4)
- [CompleteLifecycleAction](#) (p. 5)
- [CreateAutoScalingGroup](#) (p. 7)
- [CreateLaunchConfiguration](#) (p. 11)
- [CreateOrUpdateTags](#) (p. 16)
- [DeleteAutoScalingGroup](#) (p. 18)
- [DeleteLaunchConfiguration](#) (p. 20)
- [DeleteLifecycleHook](#) (p. 21)
- [DeleteNotificationConfiguration](#) (p. 22)
- [DeletePolicy](#) (p. 23)
- [DeleteScheduledAction](#) (p. 24)
- [DeleteTags](#) (p. 25)
- [DescribeAccountLimits](#) (p. 26)
- [DescribeAdjustmentTypes](#) (p. 27)
- [DescribeAutoScalingGroups](#) (p. 28)
- [DescribeAutoScalingInstances](#) (p. 31)
- [DescribeAutoScalingNotificationTypes](#) (p. 33)
- [DescribeLaunchConfigurations](#) (p. 34)
- [DescribeLifecycleHookTypes](#) (p. 36)
- [DescribeLifecycleHooks](#) (p. 37)
- [DescribeMetricCollectionTypes](#) (p. 38)
- [DescribeNotificationConfigurations](#) (p. 40)
- [DescribePolicies](#) (p. 41)
- [DescribeScalingActivities](#) (p. 44)
- [DescribeScalingProcessTypes](#) (p. 47)
- [DescribeScheduledActions](#) (p. 49)
- [DescribeTags](#) (p. 51)
- [DescribeTerminationPolicyTypes](#) (p. 53)
- [DetachInstances](#) (p. 54)
- [DisableMetricsCollection](#) (p. 56)

- [EnableMetricsCollection](#) (p. 57)
- [EnterStandby](#) (p. 58)
- [ExecutePolicy](#) (p. 60)
- [ExitStandby](#) (p. 61)
- [PutLifecycleHook](#) (p. 63)
- [PutNotificationConfiguration](#) (p. 66)
- [PutScalingPolicy](#) (p. 67)
- [PutScheduledUpdateGroupAction](#) (p. 70)
- [RecordLifecycleActionHeartbeat](#) (p. 73)
- [ResumeProcesses](#) (p. 74)
- [SetDesiredCapacity](#) (p. 75)
- [SetInstanceHealth](#) (p. 77)
- [SuspendProcesses](#) (p. 78)
- [TerminateInstanceInAutoScalingGroup](#) (p. 79)
- [UpdateAutoScalingGroup](#) (p. 80)

AttachInstances

Description

Attaches one or more Amazon EC2 instances to an existing Auto Scaling group. After the instance(s) is attached, it becomes a part of the Auto Scaling group.

For more information, see [Attach Amazon EC2 Instances to Your Existing Auto Scaling Group](#) in the *Auto Scaling Developer Guide*.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

The name of the Auto Scaling group to which to attach the specified instance(s).

Type: String

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

Required: Yes

InstanceIds.member.N

One or more IDs of the Amazon EC2 instances to attach to the specified Auto Scaling group. You must specify at least one instance ID.

Type: String list

Length constraints: Minimum length of 1. Maximum length of 16.

Required: No

CompleteLifecycleAction

Description

Completes the lifecycle action for the associated token initiated under the given lifecycle hook with the specified result.

This operation is a part of the basic sequence for adding a lifecycle hook to an Auto Scaling group:

1. Create a notification target. A target can be either an Amazon SQS queue or an Amazon SNS topic.
2. Create an IAM role. This role allows Auto Scaling to publish lifecycle notifications to the designated SQS queue or SNS topic.
3. Create the lifecycle hook. You can create a hook that acts when instances launch or when instances terminate.
4. If necessary, record the lifecycle action heartbeat to keep the instance in a pending state.
5. **Complete the lifecycle action.**

To learn more, see [Auto Scaling Pending State](#) and [Auto Scaling Terminating State](#).

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

The name of the Auto Scaling group to which the lifecycle hook belongs.

Type: String

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

Required: Yes

LifecycleActionResult

The action the Auto Scaling group should take. The value for 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: Minimum length of 36. Maximum length of 36.

Required: Yes

LifecycleHookName

The name of the lifecycle hook.

Type: String

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

Required: Yes

CreateAutoScalingGroup

Description

Creates a new Auto Scaling group with the specified name and other attributes. When the creation request is completed, the Auto Scaling group is ready to be used in other calls.

Note

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

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

The name of the Auto Scaling group.

Type: String

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

Required: Yes

AvailabilityZones.member.N

A list of Availability Zones for the Auto Scaling group. This is required unless you have specified subnets.

Type: String list

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

Length constraints: Minimum of 1 item(s) in the list.

Required: No

DefaultCooldown

The amount of time, in seconds, between a successful scaling activity and the succeeding scaling activity.

If a `DefaultCooldown` period is not specified, Auto Scaling uses the default value of 300 as the default cool down period for the Auto Scaling group. For more information, see [Cooldown Period](#)

Type: Integer

Required: No

DesiredCapacity

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

Type: Integer

Required: No

HealthCheckGracePeriod

Length of time in seconds after a new Amazon EC2 instance comes into service that Auto Scaling starts checking its health. During this time any health check failure for the that instance is ignored.

This is required if you are adding `ELB` health check. Frequently, new instances need to warm up, briefly, before they can pass a health check. To provide ample warm-up time, set the health check grace period of the group to match the expected startup period of your application.

For more information, see [Add an Elastic Load Balancing Health Check](#).

Type: Integer

Required: No

HealthCheckType

The service you want the health checks from, Amazon EC2 or Elastic Load Balancer. Valid values are `EC2` or `ELB`.

By default, the Auto Scaling health check uses the results of Amazon EC2 instance status checks to determine the health of an instance. For more information, see [Health Check](#).

Type: String

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

Required: No

InstanceId

The ID of the Amazon EC2 instance you want to use to create the Auto Scaling group. Use this attribute if you want to create an Auto Scaling group using an EC2 instance instead of a launch configuration.

When you use an instance to create an Auto Scaling group, a new launch configuration is first created and then associated with the Auto Scaling group. The new launch configuration derives all its attributes from the instance that is used to create the Auto Scaling group, with the exception of `BlockDeviceMapping`.

For more information, see [Create an Auto Scaling Group Using EC2 Instance](#) in the *Auto Scaling Developer Guide*.

Type: String

Length constraints: Minimum length of 1. Maximum length of 16.

Required: No

LaunchConfigurationName

The name of an existing launch configuration to use to launch new instances. Use this attribute if you want to create an Auto Scaling group using an existing launch configuration instead of an EC2 instance.

Type: String

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

Required: No

LoadBalancerNames.member.N

A list of existing Elastic Load Balancing load balancers to use. The load balancers must be associated with the AWS account.

For information on using load balancers, see [Load Balance Your Auto Scaling Group](#) in the *Auto Scaling Developer Guide*.

Type: String list

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

Required: No

MaxSize

The maximum size of the Auto Scaling group.

Type: Integer

Required: Yes

MinSize

The minimum size of the Auto Scaling group.

Type: Integer

Required: Yes

PlacementGroup

Physical location of an existing cluster placement group into which you want to launch your instances. For information about cluster placement group, see [Using Cluster Instances](#)

Type: String

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

Required: No

Tags.member.N

The tag to be created or updated. Each tag should be defined by its resource type, resource ID, key, value, and a propagate flag. Valid values: `key=value`, `value=value`, `propagate=true` or `false`. Value and propagate are optional parameters.

For information about using tags, see [Tag Your Auto Scaling Groups and Amazon EC2 Instances in the Auto Scaling Developer Guide](#).

Type: [Tag \(p. 115\)](#) list

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 on configuring a termination policy for your Auto Scaling group, see [Instance Termination Policy for Your Auto Scaling Group](#) in the *Auto Scaling Developer Guide*.

Type: String list

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

Required: No

VPCZoneIdentifier

A comma-separated list of subnet identifiers of Amazon Virtual Private Clouds (Amazon VPCs).

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

For information on launching your Auto Scaling group into Amazon VPC subnets, see [Auto Scaling in Amazon Virtual Private Cloud](#) in the *Auto Scaling Developer Guide*.

Type: String

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

Required: No

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

AlreadyExists

The named Auto Scaling group or launch configuration already exists.

HTTP Status Code: 400

LimitExceeded

The quota for capacity groups or launch configurations for this customer has already been reached.

HTTP Status Code: 400

Examples

Sample Request

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

Sample Response

```
<CreateAutoScalingGroupResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-
01-01/">
  <ResponseMetadata>
    <RequestId>8d798a29-f083-11e1-bdfb-cb223EXAMPLE</RequestId>
  </ResponseMetadata>
</CreateAutoScalingGroupResponse>
```

CreateLaunchConfiguration

Description

Creates a new launch configuration. The launch configuration name must be unique within the scope of the client's AWS account. The maximum limit of launch configurations, which by default is 100, must not yet have been met; otherwise, the call will fail. When created, the new launch configuration is available for immediate use.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AssociatePublicIpAddress

Used for Auto Scaling groups that launch instances into an Amazon Virtual Private Cloud (Amazon VPC). Specifies whether to assign a public IP address to each instance launched in a Amazon VPC. For more information, see [Auto Scaling in Amazon Virtual Private Cloud](#).

Note

If you specify a value for this parameter, be sure to specify at least one VPC subnet using the *VPCZoneIdentifier* parameter when you create your Auto Scaling group.

Default: If the instance is launched into a default subnet in a default VPC, the default is `true`. If the instance is launched into a nondefault subnet in a VPC, the default is `false`. For information about default VPC and VPC platforms, see [Supported Platforms](#).

Type: Boolean

Required: No

BlockDeviceMappings.member.N

A list of mappings that specify how block devices are exposed to the instance. Each mapping is made up of a *VirtualName*, a *DeviceName*, and an *ebs* data structure that contains information about the associated Elastic Block Storage volume. For more information about Amazon EC2 BlockDeviceMappings, go to [Block Device Mapping](#) in the Amazon EC2 product documentation.

Type: [BlockDeviceMapping \(p. 91\)](#) list

Required: No

EbsOptimized

Whether the instance is optimized for EBS I/O. The optimization provides dedicated throughput to Amazon EBS and an optimized configuration stack to provide optimal EBS I/O performance. This optimization is not available with all instance types. Additional usage charges apply when using an EBS Optimized instance. By default the instance is not optimized for EBS I/O. For information about EBS-optimized instances, go to [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.

Amazon 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 Amazon EC2 instances to securely access other AWS resources. For information on launching EC2 instances with an IAM role, go to [Launching Auto Scaling Instances With an IAM Role](#) in the *Auto Scaling Developer Guide*.

Type: String

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

Required: No

ImageId

Unique ID of the Amazon Machine Image (AMI) you want to use to launch your EC2 instances. For information about finding Amazon EC2 AMIs, see [Finding a Suitable AMI](#) in the *Amazon Elastic Compute Cloud User Guide*.

Type: String

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

Required: No

InstanceId

The ID of the Amazon EC2 instance you want to use to create the launch configuration. Use this attribute if you want the launch configuration to derive its attributes from an EC2 instance.

When you use an instance to create a launch configuration, all you need to specify is the `InstanceId`. The new launch configuration, by default, derives all the attributes from the specified instance with the exception of `BlockDeviceMapping`.

If you want to create a launch configuration with `BlockDeviceMapping` or override any other instance attributes, specify them as part of the same request.

For more information on using an InstanceID to create a launch configuration, see [Create a Launch Configuration Using an Amazon EC2 Instance](#) in the *Auto Scaling Developer Guide*.

Type: String

Length constraints: Minimum length of 1. Maximum length of 16.

Required: No

InstanceMonitoring

Enables detailed monitoring if it is disabled. Detailed monitoring is enabled by default.

When detailed monitoring is enabled, Amazon Cloudwatch will generate metrics every minute and your account will be charged a fee. When you disable detailed monitoring, by specifying `False`, Cloudwatch will generate metrics every 5 minutes. For more information, see [Monitor Your Auto Scaling Instances](#). For information about Amazon CloudWatch, see the [Amazon CloudWatch Developer Guide](#).

Type: [InstanceMonitoring](#) (p. 103)

Required: No

InstanceType

The instance type of the Amazon EC2 instance. For information about available Amazon EC2 instance types, see [Available Instance Types](#) in the *Amazon Elastic Cloud Compute User Guide*.

Type: String

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

Required: No

KernelId

The ID of the kernel associated with the Amazon EC2 AMI.

Type: String

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

Required: No

KeyName

The name of the Amazon EC2 key pair. For more information, see [Getting a Key Pair](#) in the *Amazon Elastic Compute Cloud User Guide*.

Type: String

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

Required: No

LaunchConfigurationName

The name of the launch configuration to create.

Type: String

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

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 in a VPC.

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

If you specify a value for this parameter, be sure to specify at least one VPC subnet using the `VP-CZoneIdentifier` parameter when you create your Auto Scaling group.

For more information, see [Auto Scaling in Amazon Virtual Private Cloud](#) in the *Auto Scaling Developer Guide*.

Valid values: `default` | `dedicated`

Type: String

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

Required: No

RamdiskId

The ID of the RAM disk associated with the Amazon EC2 AMI.

Type: String

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

Required: No

SecurityGroups.member.N

The security groups with which to associate Amazon EC2 or Amazon VPC instances.

If your instances are launched in EC2, you can either specify Amazon EC2 security group names or the security group IDs. For more information about Amazon EC2 security groups, see [Using Security Groups](#) in the *Amazon Elastic Compute Cloud User Guide*.

If your instances are launched within VPC, specify Amazon VPC security group IDs. For more information about Amazon VPC security groups, see [Security Groups](#) in the *Amazon Virtual Private Cloud User Guide*.

Type: String list

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 on launching Spot Instances, see [Using Auto Scaling to Launch Spot Instances](#) in the *Auto Scaling Developer 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 Amazon EC2 instances. For more information about Amazon EC2 user data, see [User Data Retrieval](#) in the *Amazon Elastic Compute Cloud User Guide*.

Note

At this time, Auto Scaling launch configurations don't support compressed (e.g. zipped) user data files.

Type: String

Length constraints: Minimum length of 0. Maximum length of 21847.

Required: No

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

AlreadyExists

The named Auto Scaling group or launch configuration already exists.

HTTP Status Code: 400

LimitExceeded

The quota for capacity groups or launch configurations for this customer has already been reached.

HTTP Status Code: 400

Examples

Sample Request

```
https://autoscaling.amazonaws.com/?LaunchConfigurationName=my-test-lc
&AssociatePublicIpAddress=true
&PlacementTenancy=dedicated
&ImageId=ami-0078da69
&InstanceType=m1.small
```

```
&Action=CreateLaunchConfiguration  
&AUTHPARAMS
```

Sample Response

```
<CreateLaunchConfigurationResponse xmlns="http://autoscaling.amazon  
aws.com/doc/2011-01-01/">  
<ResponseMetadata>  
  <RequestId>7c6e177f-f082-11e1-ac58-3714bEXAMPLE</RequestId>  
</ResponseMetadata>  
</CreateLaunchConfigurationResponse>
```

CreateOrUpdateTags

Description

Creates new tags or updates existing tags for an Auto Scaling group.

Note

A tag's definition is composed of a resource ID, resource type, key and value, and the propagate flag. Value and the propagate flag are optional parameters. See the Request Parameters for more information.

For information on creating tags for your Auto Scaling group, see [Tag Your Auto Scaling Groups and Amazon EC2 Instances](#).

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

Tags.member.N

The tag to be created or updated. Each tag should be defined by its resource type, resource ID, key, value, and a propagate flag. The resource type and resource ID identify the type and name of resource for which the tag is created. Currently, `auto-scaling-group` is the only supported resource type. The valid value for the resource ID is *groupname*.

The `PropagateAtLaunch` flag defines whether the new tag will be applied to instances launched by the Auto Scaling group. Valid values are `true` or `false`. However, instances that are already running will not get the new or updated tag. Likewise, when you modify a tag, the updated version will be applied only to new instances launched by the Auto Scaling group after the change. Running instances that had the previous version of the tag will continue to have the older tag.

When you create a tag and a tag of the same name already exists, the operation overwrites the previous tag definition, but you will not get an error message.

Type: [Tag \(p. 115\)](#) list

Required: Yes

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

AlreadyExists

The named Auto Scaling group or launch configuration already exists.

HTTP Status Code: 400

LimitExceeded

The quota for capacity groups or launch configurations for this customer has already been reached.

HTTP Status Code: 400

Examples

Sample Request

```
https://autoscaling.amazonaws.com/?Tags.member.1.ResourceId=my-test-asg
&Tags.member.1.ResourceType=auto-scaling-group
&Tags.member.1.Key=version
&Tags.member.1.Value=1.0
&Tags.member.1.PropagateAtLaunch=true
&Version=2011-01-01
&Action=CreateOrUpdateTags
&AUTHPARAMS
```

Sample Response

```
<CreateOrUpdateTagsResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-
01-01/">
  <ResponseMetadata>
    <RequestId>b0203919-bf1b-11e2-8a01-13263EXAMPLE</RequestId>
  </ResponseMetadata>
</CreateOrUpdateTagsResponse>
```

DeleteAutoScalingGroup

Description

Deletes the specified Auto Scaling group if the group has no instances and no scaling activities in progress.

Note

To remove all instances before calling [DeleteAutoScalingGroup](#) (p. 18), you can call [UpdateAutoScalingGroup](#) (p. 80) to set the minimum and maximum size of the AutoScalingGroup to zero.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters](#) (p. 117).

AutoScalingGroupName

The name of the Auto Scaling group to delete.

Type: String

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

Required: Yes

ForceDelete

Starting with API version 2011-01-01, specifies that the Auto Scaling 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. 121).

ResourceInUse

This is returned when you cannot delete a launch configuration or Auto Scaling group because it is being used.

HTTP Status Code: 400

ScalingActivityInProgress

You cannot delete an Auto Scaling group while there are scaling activities in progress for that group.

HTTP Status Code: 400

Examples

Sample Request

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

```
&Action=DeleteAutoScalingGroup  
&AUTHPARAMS
```

Sample Response

```
<DeleteAutoScalingGroupResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">  
  <ResponseMetadata>  
    <RequestId>70a76d42-9665-11e2-9fdf-211deEXAMPLE</RequestId>  
  </ResponseMetadata>  
</DeleteAutoScalingGroupResponse>
```

DeleteLaunchConfiguration

Description

Deletes the specified [LaunchConfiguration](#) (p. 104).

The specified 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 common parameters that all actions use, see [Common Parameters](#) (p. 117).

LaunchConfigurationName

The name of the launch configuration.

Type: String

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

Required: Yes

Errors

For information about the errors that are common to all actions, see [Common Errors](#) (p. 121).

ResourceInUse

This is returned when you cannot delete a launch configuration or Auto Scaling group because it is being used.

HTTP Status Code: 400

Examples

Sample Request

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

Sample Response

```
<DeleteLaunchConfigurationResponse xmlns="http://autoscaling.amazon
aws.com/doc/2011-01-01/">
  <ResponseMetadata>
    <RequestId>7347261f-97df-11e2-8756-35eEXAMPLE</RequestId>
  </ResponseMetadata>
</DeleteLaunchConfigurationResponse>
```

DeleteLifecycleHook

Description

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 common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

The name of the Auto Scaling group to which the lifecycle hook belongs.

Type: String

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

Required: Yes

LifecycleHookName

The name of the lifecycle hook.

Type: String

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

Required: Yes

DeleteNotificationConfiguration

Description

Deletes notifications created by [PutNotificationConfiguration](#) (p. 66).

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters](#) (p. 117).

AutoScalingGroupName

The name of the Auto Scaling group.

Type: String

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

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.

Required: Yes

DeletePolicy

Description

Deletes a policy created by [PutScalingPolicy](#) (p. 67).

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters](#) (p. 117).

AutoScalingGroupName

The name of the Auto Scaling group.

Type: String

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

Required: No

PolicyName

The name or PolicyARN of the policy you want to delete.

Type: String

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

Required: Yes

DeleteScheduledAction

Description

Deletes a scheduled action previously created using the [PutScheduledUpdateGroupAction](#) (p. 70).

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters](#) (p. 117).

AutoScalingGroupName

The name of the Auto Scaling group.

Type: String

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

Required: No

ScheduledActionName

The name of the action you want to delete.

Type: String

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

Required: Yes

DeleteTags

Description

Removes the specified tags or a set of tags from a set of resources.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

Tags.member.N

Each tag should be defined by its resource type, resource ID, key, value, and a propagate flag. Valid values are: Resource type = *auto-scaling-group*, Resource ID = *AutoScalingGroupName*, key=*value*, value=*value*, propagate=*true* or *false*.

Type: [Tag \(p. 115\)](#) list

Required: Yes

DescribeAccountLimits

Description

Returns the limits for the Auto Scaling resources currently allowed for your AWS account.

Your AWS account comes with default limits on resources for Auto Scaling. There is a default limit of 20 Auto Scaling groups and 100 launch configurations per region.

If you reach the limits for the number of Auto Scaling groups or the launch configurations, you can go to the [Support Center](#) and place a request to raise the limits.

Response Elements

The following elements are returned in a structure named `DescribeAccountLimitsResult`.

MaxNumberOfAutoScalingGroups

The maximum number of Auto Scaling groups allowed for your AWS account.

Type: Integer

MaxNumberOfLaunchConfigurations

The maximum number of launch configurations allowed for your AWS account.

Type: Integer

Examples

Sample Request

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

Sample Response

```
<DescribeAccountLimitsResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <DescribeAccountLimitsResult>
    <MaxNumberOfLaunchConfigurations>100</MaxNumberOfLaunchConfigurations>
    <MaxNumberOfAutoScalingGroups>20</MaxNumberOfAutoScalingGroups>
  </DescribeAccountLimitsResult>
  <ResponseMetadata>
    <RequestId>a32bd184-519d-11e3-a8a4-c1c467cbcc3b</RequestId>
  </ResponseMetadata>
</DescribeAccountLimitsResponse>
```

DescribeAdjustmentTypes

Description

Returns policy adjustment types for use in the [PutScalingPolicy](#) (p. 67) action.

Response Elements

The following element is returned in a structure named `DescribeAdjustmentTypesResult`.

AdjustmentTypes

A list of specific policy adjustment types.

Type: [AdjustmentType](#) (p. 87) list

Examples

Sample Request

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

Sample Response

```
<DescribeAdjustmentTypesResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <DescribeAdjustmentTypesResult>
    <AdjustmentTypes>
      <member>
        <AdjustmentType>ChangeInCapacity</AdjustmentType>
      </member>
      <member>
        <AdjustmentType>ExactCapacity</AdjustmentType>
      </member>
      <member>
        <AdjustmentType>PercentChangeInCapacity</AdjustmentType>
      </member>
    </AdjustmentTypes>
  </DescribeAdjustmentTypesResult>
  <ResponseMetadata>
    <RequestId>cc5f0337-b694-11e2-afc0-6544dEXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeAdjustmentTypesResponse>
```

DescribeAutoScalingGroups

Description

Returns a full description of each Auto Scaling group in the given list. This includes all Amazon EC2 instances that are members of the group. If a list of names is not provided, the service returns the full details of all Auto Scaling groups.

This action supports pagination by returning a token if there are more pages to retrieve. To get the next page, call this action again with the returned token as the `NextToken` parameter.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupNames.member.N

A list of Auto Scaling group names.

Type: String list

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

Required: No

MaxRecords

The maximum number of records to return.

Type: Integer

Required: No

NextToken

A string that marks the start of the next batch of returned results.

Type: String

Required: No

Response Elements

The following elements are returned in a structure named `DescribeAutoScalingGroupsResult`.

AutoScalingGroups

A list of Auto Scaling groups.

Type: [AutoScalingGroup \(p. 87\)](#) list

NextToken

A string that marks the start of the next batch of returned results.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

InvalidNextToken

The NextToken value is invalid.

HTTP Status Code: 400

Examples

Sample Request

```
https://autoscaling.amazonaws.com/?AutoScalingGroupNames.member.1=my-test-asg-lbs
&MaxRecords=20
&Version=2011-01-01
&Action=DescribeAutoScalingGroups
&AUTHPARAMS
```

Sample Response

```
<DescribeAutoScalingGroupsResponse xmlns="http://autoscaling.amazon
aws.com/doc/2011-01-01/">
  <DescribeAutoScalingGroupsResult>
    <AutoScalingGroups>
      <member>
        <Tags/>
        <SuspendedProcesses/>
        <AutoScalingGroupName>my-test-asg-lbs</AutoScalingGroupName>
        <HealthCheckType>ELB</HealthCheckType>
        <CreatedTime>2013-05-06T17:47:15.107Z</CreatedTime>
        <EnabledMetrics/>
        <LaunchConfigurationName>my-test-lc</LaunchConfigurationName>
        <Instances/>
        <DesiredCapacity>2</DesiredCapacity>
        <AvailabilityZones>
          <member>us-east-1b</member>
          <member>us-east-1a</member>
        </AvailabilityZones>
        <LoadBalancerNames>
          <member>my-test-asg-loadbalancer</member>
        </LoadBalancerNames>
        <MinSize>2</MinSize>
        <VPCZoneIdentifier/>
        <HealthCheckGracePeriod>120</HealthCheckGracePeriod>
        <DefaultCooldown>300</DefaultCooldown>
        <AutoScalingGroupARN>arn:aws:autoscaling:us-east-1:803981987763:auto
ScalingGroup:ca861182-c8f9-4ca7-bleb-cd35505f5ebb
:autoScalingGroupName/my-test-asg-lbs</AutoScalingGroupARN>
        <TerminationPolicies>
          <member>Default</member>
        </TerminationPolicies>
        <MaxSize>10</MaxSize>
      </member>
    </AutoScalingGroups>
  </DescribeAutoScalingGroupsResult>
  <ResponseMetadata>
```

Auto Scaling API Reference Examples

```
<RequestId>0f02a07d-b677-11e2-9eb0-dd50EXAMPLE</RequestId>  
</ResponseMetadata>  
</DescribeAutoScalingGroupsResponse>
```

DescribeAutoScalingInstances

Description

Returns a description of each Auto Scaling instance in the `InstanceIds` list. If a list is not provided, the service returns the full details of all instances up to a maximum of 50. By default, the service returns a list of 20 items.

This action supports pagination by returning a token if there are more pages to retrieve. To get the next page, call this action again with the returned token as the `NextToken` parameter.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

InstanceIds.member.N

The list of Auto Scaling instances to describe. If this list is omitted, all auto scaling instances are described. The list of requested instances cannot contain more than 50 items. If unknown instances are requested, they are ignored with no error.

Type: String list

Length constraints: Minimum length of 1. Maximum length of 16.

Required: No

MaxRecords

The maximum number of Auto Scaling instances to be described with each call.

Type: Integer

Required: No

NextToken

The token returned by a previous call to indicate that there is more data available.

Type: String

Required: No

Response Elements

The following elements are returned in a structure named `DescribeAutoScalingInstancesResult`.

AutoScalingInstances

A list of Auto Scaling instances.

Type: [AutoScalingInstanceDetails \(p. 90\)](#) list

NextToken

A string that marks the start of the next batch of returned results.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

InvalidNextToken

The `NextToken` value is invalid.

HTTP Status Code: 400

Examples

Sample Request

```
https://autoscaling.amazonaws.com/?MaxRecords=20
&InstanceIds.member.1=i-78e0d40b
&Version=2011-01-01
&Action=DescribeAutoScalingInstances
&AUTHPARAMS
```

Sample Response

```
<DescribeAutoScalingInstancesResponse xmlns="http://autoscaling.amazon
aws.com/doc/2011-01-01/">
  <DescribeAutoScalingInstancesResult>
    <AutoScalingInstances>
      <member>
        <HealthStatus>Healthy</HealthStatus>
        <AutoScalingGroupName>my-test-asg</AutoScalingGroupName>
        <AvailabilityZone>us-east-1e</AvailabilityZone>
        <InstanceId>i-78e0d40b</InstanceId>
        <LaunchConfigurationName>my-test-lc</LaunchConfigurationName>
        <LifecycleState>InService</LifecycleState>
      </member>
    </AutoScalingInstances>
  </DescribeAutoScalingInstancesResult>
  <ResponseMetadata>
    <RequestId>df992dc3-b72f-11e2-81e1-750aa6EXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeAutoScalingInstancesResponse>
```


DescribeAutoScalingNotificationTypes

Description

Returns a list of all notification types that are supported by Auto Scaling.

Response Elements

The following element is returned in a structure named `DescribeAutoScalingNotificationTypes-Result`.

AutoScalingNotificationTypes

Returns a list of all notification types supported by Auto Scaling. They are:

- `autoscaling:EC2_INSTANCE_LAUNCH`
- `autoscaling:EC2_INSTANCE_LAUNCH_ERROR`
- `autoscaling:EC2_INSTANCE_TERMINATE`
- `autoscaling:EC2_INSTANCE_TERMINATE_ERROR`
- `autoscaling:TEST_NOTIFICATION`

Type: String list

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

DescribeLaunchConfigurations

Description

Returns a full description of the launch configurations, or the specified launch configurations, if they exist.

If no name is specified, then the full details of all launch configurations are returned.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

LaunchConfigurationNames.member.N

A list of launch configuration names.

Type: String list

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

Required: No

MaxRecords

The maximum number of launch configurations. The default is 100.

Type: Integer

Required: No

NextToken

A string that marks the start of the next batch of returned results.

Type: String

Required: No

Response Elements

The following elements are returned in a structure named `DescribeLaunchConfigurationsResult`.

LaunchConfigurations

A list of launch configurations.

Type: [LaunchConfiguration \(p. 104\)](#) list

NextToken

A string that marks the start of the next batch of returned results.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

InvalidNextToken

The `NextToken` value is invalid.

HTTP Status Code: 400

Examples

Sample Request

```
https://autoscaling.amazonaws.com/?LaunchConfigurationNames.member.1=my-test-
lc
&MaxRecords=20
&Version=2011-01-01
&Action=DescribeLaunchConfigurations
&AUTHPARAMS
```

Sample Response

```
<DescribeLaunchConfigurationsResponse xmlns="http://autoscaling.amazon
aws.com/doc/2011-01-01/">
  <DescribeLaunchConfigurationsResult>
    <LaunchConfigurations>
      <member>
        <AssociatePublicIpAddress>true</AssociatePublicIpAddress>
        <SecurityGroups/>
        <PlacementTenancy>dedicated</PlacementTenancy>
        <CreatedTime>2013-01-21T23:04:42.200Z</CreatedTime>
        <KernelId/>
        <LaunchConfigurationName>my-test-lc</LaunchConfigurationName>
        <UserData/>
        <InstanceType>m1.small</InstanceType>
        <LaunchConfigurationARN>arn:aws:autoscaling:us-east-
1:803981987763:launchConfiguration:
9dbbbf87-6141-428a-a409-0752edbe6cad:launchConfigurationName/my-test-
lc</LaunchConfigurationARN>
        <BlockDeviceMappings/>
        <ImageId>ami-514ac838</ImageId>
        <KeyName/>
        <RamdiskId/>
        <InstanceMonitoring>
          <Enabled>true</Enabled>
        </InstanceMonitoring>
        <EbsOptimized>>false</EbsOptimized>
      </member>
    </LaunchConfigurations>
  </DescribeLaunchConfigurationsResult>
  <ResponseMetadata>
    <RequestId>d05a22f8-b690-11e2-bf8e-2113fEXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeLaunchConfigurationsResponse>
```

DescribeLifecycleHookTypes

Description

Describes the available types of lifecycle hooks.

Response Elements

The following element is returned in a structure named `DescribeLifecycleHookTypesResult`.

LifecycleHookTypes

Returns a list of all notification types supported by Auto Scaling. They are:

- `autoscaling:EC2_INSTANCE_LAUNCHING`
- `autoscaling:EC2_INSTANCE_TERMINATING`

Type: String list

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

DescribeLifecycleHooks

Description

Describes the lifecycle hooks that currently belong to the specified Auto Scaling group.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

The name of one or more Auto Scaling groups.

Type: String

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

Required: Yes

LifecycleHookNames.member.N

The name of one or more lifecycle hooks.

Type: String list

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

Required: No

Response Elements

The following element is returned in a structure named `DescribeLifecycleHooksResult`.

LifecycleHooks

A list describing the lifecycle hooks that belong to the specified Auto Scaling group.

Type: [LifecycleHook \(p. 106\)](#) list

DescribeMetricCollectionTypes

Description

Returns a list of metrics and a corresponding list of granularities for each metric.

Note

The `GroupStandbyInstances` metric is not returned by default. You must explicitly request it when calling [EnableMetricsCollection](#) (p. 57).

Response Elements

The following elements are returned in a structure named `DescribeMetricCollectionTypesResult`.

Granularities

A list of granularities for the listed Metrics.

Type: [MetricGranularityType](#) (p. 108) list

Metrics

The list of Metrics collected. The following metrics are supported:

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

Note

The `GroupStandbyInstances` metric is not returned by default. You must explicitly request it when calling [EnableMetricsCollection](#) (p. 57).

Type: [MetricCollectionType](#) (p. 108) list

Examples

Sample Request

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

Sample Response

```
<DescribeMetricCollectionTypesResponse xmlns="http://autoscaling.amazonaws.com/2011-01-01/">
  <DescribeMetricCollectionTypesResult>
```

Auto Scaling API Reference Examples

```
<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>GroupStandbyInstances</Metric>
  </member>
  <member>
    <Metric>GroupTerminatingInstances</Metric>
  </member>
  <member>
    <Metric>GroupTotalInstances</Metric>
  </member>
</Metrics>
<Granularities>
  <member>
    <Granularity>1Minute</Granularity>
  </member>
</Granularities>
</DescribeMetricCollectionTypesResult>
<ResponseMetadata>
  <RequestId>07f3fea2-bf3c-11e2-9b6f-f3cdbb80c073</RequestId>
</ResponseMetadata>
</DescribeMetricCollectionTypesResponse>
```

DescribeNotificationConfigurations

Description

Returns a list of notification actions associated with Auto Scaling groups for specified events.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupNames.member.N

The name of the Auto Scaling group.

Type: String list

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

Required: No

MaxRecords

Maximum number of records to be returned.

Type: Integer

Required: No

NextToken

A string that is used to mark the start of the next batch of returned results for pagination.

Type: String

Required: No

Response Elements

The following elements are returned in a structure named `DescribeNotificationConfiguration-
sResult`.

NextToken

A string that is used to mark the start of the next batch of returned results for pagination.

Type: String

NotificationConfigurations

The list of notification configurations.

Type: [NotificationConfiguration \(p. 108\)](#) list

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

InvalidNextToken

The `NextToken` value is invalid.

HTTP Status Code: 400

DescribePolicies

Description

Returns descriptions of what each policy does. This action supports pagination. If the response includes a token, there are more records available. To get the additional records, repeat the request with the response token as the `NextToken` parameter.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

The name of the Auto Scaling group.

Type: String

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

Required: No

MaxRecords

The maximum number of policies that will be described with each call.

Type: Integer

Required: No

NextToken

A string that is used to mark the start of the next batch of returned results for pagination.

Type: String

Required: No

PolicyNames.member.N

A list of policy names or policy ARNs to be described. If this list is omitted, all policy names are described. If an auto scaling group name is provided, the results are limited to that group. The list of requested policy names cannot contain more than 50 items. If unknown policy names are requested, they are ignored with no error.

Type: String list

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

Required: No

Response Elements

The following elements are returned in a structure named `DescribePoliciesResult`.

NextToken

A string that marks the start of the next batch of returned results.

Type: String

ScalingPolicies

A list of scaling policies.

Type: [ScalingPolicy](#) (p. 111) list

Errors

For information about the errors that are common to all actions, see [Common Errors](#) (p. 121).

InvalidNextToken

The `NextToken` value is invalid.

HTTP Status Code: 400

Examples

Sample Request

```
https://autoscaling.amazonaws.com/?AutoScalingGroupName=my-test-asg
&MaxRecords=20
&Version=2011-01-01
&Action=DescribePolicies
&AUTHPARAMS
```

Sample Response

```
<DescribePoliciesResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <DescribePoliciesResult>
    <ScalingPolicies>
      <member>
        <PolicyARN>arn:aws:autoscaling:us-east-1:803981987763:scalingPolicy:c322761b-3172-4d56-9a21-0ed9d6161d67:autoScalingGroupName/my-test-asg:policyName/MyScaleDownPolicy</PolicyARN>
        <AdjustmentType>ChangeInCapacity</AdjustmentType>
        <ScalingAdjustment>-1</ScalingAdjustment>
        <PolicyName>MyScaleDownPolicy</PolicyName>
        <AutoScalingGroupName>my-test-asg</AutoScalingGroupName>
        <Cooldown>60</Cooldown>
        <Alarms>
          <member>
            <AlarmName>TestQueue</AlarmName>
            <AlarmARN>arn:aws:cloudwatch:us-east-1:803981987763:alarm:TestQueue</AlarmARN>
          </member>
        </Alarms>
      </member>
      <member>
        <PolicyARN>arn:aws:autoscaling:us-east-1:803981987763:scalingPolicy:c55a5cdd-9be0-435b-b60b-a8dd313159f5:autoScalingGroupName/my-test-asg:policyName/MyScaleUpPolicy</PolicyARN>
        <AdjustmentType>ChangeInCapacity</AdjustmentType>
        <ScalingAdjustment>1</ScalingAdjustment>
        <PolicyName>MyScaleUpPolicy</PolicyName>
        <AutoScalingGroupName>my-test-asg</AutoScalingGroupName>
        <Cooldown>60</Cooldown>
      </member>
    </ScalingPolicies>
  </DescribePoliciesResult>
</DescribePoliciesResponse>
```

Auto Scaling API Reference Examples

```
<Alarms>
  <member>
    <AlarmName>TestQueue</AlarmName>
    <AlarmARN>arn:aws:cloudwatch:us-east-
1:803981987763:alarm:TestQueue</AlarmARN>
  </member>
</Alarms>
</member>
</ScalingPolicies>
</DescribePoliciesResult>
<ResponseMetadata>
  <RequestId>ec3bffad-b739-11e2-b38d-15fbEXAMPLE</RequestId>
</ResponseMetadata>
</DescribePoliciesResponse>
```

DescribeScalingActivities

Description

Returns the scaling activities for the specified Auto Scaling group.

If the specified `ActivityIds` list is empty, all the activities from the past six weeks are returned. Activities are sorted by the start time. Activities still in progress appear first on the list.

This action supports pagination. If the response includes a token, there are more records available. To get the additional records, repeat the request with the response token as the `NextToken` parameter.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

ActivityIds.member.N

A list containing the activity IDs of the desired scaling activities. If this list is omitted, all activities are described. If an `AutoScalingGroupName` is provided, 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: String list

Required: No

AutoScalingGroupName

The name of the [AutoScalingGroup \(p. 87\)](#).

Type: String

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

Required: No

MaxRecords

The maximum number of scaling activities to return.

Type: Integer

Required: No

NextToken

A string that marks the start of the next batch of returned results for pagination.

Type: String

Required: No

Response Elements

The following elements are returned in a structure named `DescribeScalingActivitiesResult`.

Activities

A list of the requested scaling activities.

Type: [Activity \(p. 85\)](#) list

NextToken

Acts as a paging mechanism for large result sets. Set to a non-empty string if there are additional results waiting to be returned. Pass this in to subsequent calls to return additional results.

Type: String

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

InvalidNextToken

The `NextToken` value is invalid.

HTTP Status Code: 400

Examples

Sample Request

```
https://autoscaling.amazonaws.com/?AutoScalingGroupName=my-test-asg
&MaxRecords=20
&Version=2011-01-01
&Action=DescribeScalingActivities
&AUTHPARAMS
```

Sample Response

```
<DescribeScalingActivitiesResponse xmlns="http://ec2.amazonaws.com/doc/2011-01-01/">
  <DescribeScalingActivitiesResult>
    <Activities>
      <member>
        <StatusCode>Failed</StatusCode>
        <Progress>0</Progress>
        <ActivityId>063308ae-aa22-4a9b-94f4-9faeEXAMPLE</ActivityId>
        <StartTime>2012-04-12T17:32:07.882Z</StartTime>
        <AutoScalingGroupName>my-test-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>7a641adc-84c5-11e1-a8a5-217ebEXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeScalingActivitiesResponse>
```

Auto Scaling API Reference Examples

```
</ResponseMetadata>  
</DescribeScalingActivitiesResponse>
```

DescribeScalingProcessTypes

Description

Returns scaling process types for use in the [ResumeProcesses](#) (p. 74) and [SuspendProcesses](#) (p. 78) actions.

Response Elements

The following element is returned in a structure named `DescribeScalingProcessTypesResult`.

Processes

A list of [ProcessType](#) (p. 109) names.

Type: [ProcessType](#) (p. 109) list

Examples

Sample Request

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

Sample Response

```
<DescribeScalingProcessTypesResponse xmlns="http://autoscaling.amazon
aws.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>
    </Processes>
  </DescribeScalingProcessTypesResult>
</DescribeScalingProcessTypesResponse>
```

Auto Scaling API Reference Examples

```
<member>
  <ProcessName>Terminate</ProcessName>
</member>
</Processes>
</DescribeScalingProcessTypesResult>
<ResponseMetadata>
  <RequestId>27f2eacc-b73f-11e2-ad99-c7aba3a9c963</RequestId>
</ResponseMetadata>
</DescribeScalingProcessTypesResponse>
```


DescribeScheduledActions

Description

Lists all the actions scheduled for your Auto Scaling group that haven't been executed. To see a list of actions already executed, see the activity record returned in [DescribeScalingActivities](#) (p. 44).

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters](#) (p. 117).

AutoScalingGroupName

The name of the Auto Scaling group.

Type: String

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

Required: No

EndTime

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

Type: DateTime

Required: No

MaxRecords

The maximum number of scheduled actions to return.

Type: Integer

Required: No

NextToken

A string that marks the start of the next batch of returned results.

Type: String

Required: No

ScheduledActionNames.member.N

A list of scheduled actions to be described. If this list is omitted, all scheduled actions are described. The list of requested scheduled actions cannot contain more than 50 items. If an auto scaling group name is provided, the results are limited to that group. If unknown scheduled actions are requested, they are ignored with no error.

Type: String list

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

Required: No

StartTime

The earliest scheduled start time to return. If scheduled action names are provided, this field will be ignored.

Type: DateTime

Required: No

Response Elements

The following elements are returned in a structure named `DescribeScheduledActionsResult`.

NextToken

A string that marks the start of the next batch of returned results.

Type: String

ScheduledUpdateGroupActions

A list of scheduled actions designed to update an Auto Scaling group.

Type: [ScheduledUpdateGroupAction \(p. 113\)](#) list

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

InvalidNextToken

The `NextToken` value is invalid.

HTTP Status Code: 400

DescribeTags

Description

Lists the Auto Scaling group tags.

You can use filters to limit results when describing tags. For example, you can query for tags of a particular 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 your filters. If there's no match, no special message is returned.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

Filters.member.N

The value of the filter type used to identify the tags to be returned. For example, you can filter so that tags are returned according to Auto Scaling group, the key and value, or whether the new tag will be applied to instances launched after the tag is created (PropagateAtLaunch).

Type: [Filter \(p. 102\)](#) list

Required: No

MaxRecords

The maximum number of records to return.

Type: Integer

Required: No

NextToken

A string that marks the start of the next batch of returned results.

Type: String

Required: No

Response Elements

The following elements are returned in a structure named `DescribeTagsResult`.

NextToken

A string used to mark the start of the next batch of returned results.

Type: String

Tags

The list of tags.

Type: [TagDescription \(p. 115\)](#) list

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

InvalidNextToken

The `NextToken` value is invalid.

HTTP Status Code: 400

Examples

Sample Request

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

Sample Response

```
<DescribeTagsResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <DescribeTagsResult>
    <Tags>
      <member>
        <ResourceId>my-test-asg</ResourceId>
        <PropagateAtLaunch>true</PropagateAtLaunch>
        <Value>1.0</Value>
        <Key>version</Key>
        <ResourceType>auto-scaling-group</ResourceType>
      </member>
    </Tags>
  </DescribeTagsResult>
  <ResponseMetadata>
    <RequestId>086265fd-bf3e-11e2-85fc-fbb1EXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeTagsResponse>
```

DescribeTerminationPolicyTypes

Description

Returns a list of all termination policies supported by Auto Scaling.

Response Elements

The following element is returned in a structure named `DescribeTerminationPolicyTypesResult`.

TerminationPolicyTypes

Termination policies supported by Auto Scaling. They are: `OldestInstance`, `OldestLaunchConfiguration`, `NewestInstance`, `ClosestToNextInstanceHour`, `Default`

Type: String list

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

Examples

Sample Request

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

Sample Response

```
<DescribeTerminationPolicyTypesResponse xmlns="http://autoscaling.amazon
aws.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>d9a05827-b735-11e2-a40c-c79a5EXAMPLE</RequestId>
  </ResponseMetadata>
</DescribeTerminationPolicyTypesResponse>
```

DetachInstances

Description

Using `DetachInstances`, you can remove an instance from an Auto Scaling group. After the instances are detached, you can manage them independently from the rest of the Auto Scaling group.

To learn more about detaching instances, see [Detach Amazon EC2 Instances From Your Auto Scaling Group](#).

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

The name of the Auto Scaling group from which to detach instances.

Type: String

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

Required: Yes

InstanceIds.member.N

A list of instances to detach from the Auto Scaling group. You must specify at least one instance ID.

Type: String list

Length constraints: Minimum length of 1. Maximum length of 16.

Required: No

ShouldDecrementDesiredCapacity

Specifies if the detached instance should decrement the desired capacity value for the Auto Scaling group. If set to `True`, 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 in a structure named `DetachInstancesResult`.

Activities

A list describing the activities related to detaching the instances from the Auto Scaling group.

Type: [Activity \(p. 85\)](#) list

Examples

Sample Request

```
https://autoscaling.amazonaws.com/?AutoScalingGroupName=my-asg&ShouldDecrementDesiredCapacity=true&InstanceIds.member.1=i-5f2e8a0d&Version=2011-01-01&Action=DetachInstances&SignatureVersion=2&SignatureMethod=HmacSHA256&Timestamp=2014-06-14T00%3A07%3A29.962Z&AUTHPARAMS
```

Sample Response

```
<DetachInstancesResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <DetachInstancesResult>
    <Activities>
      <member>
        <ActivityId>e54ff599-bf05-4076-8b95-a0f090ed90bb</ActivityId>
        <Progress>50</Progress>
        <StatusCode>InProgress</StatusCode>
        <StartTime>2014-06-14T00:07:30.280Z</StartTime>
        <Cause>At 2014-06-14T00:07:30Z instance i-5f2e8a0d was detached in response to a user request, shrinking the capacity from 4 to 3.</Cause>
        <AutoScalingGroupName>my-asg</AutoScalingGroupName>
        <Details>{"Availability Zone":"us-east-1a"}</Details>
        <Description>Detaching EC2 instance: i-5f2e8a0d</Description>
      </member>
    </Activities>
  </DetachInstancesResult>
  <ResponseMetadata>
    <RequestId>e04f3b11-f357-11e3-a434-7f10009d5849</RequestId>
  </ResponseMetadata>
</DetachInstancesResponse>
```

DisableMetricsCollection

Description

Disables monitoring of group metrics for the Auto Scaling group specified in `AutoScalingGroupName`. You can specify the list of affected metrics with the `Metrics` parameter.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

The name or ARN of the Auto Scaling Group.

Type: String

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

Required: Yes

Metrics.member.N

The list of metrics to disable. If no metrics are specified, all metrics are disabled. The following metrics are supported:

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

Type: String list

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

Required: No

EnableMetricsCollection

Description

Enables monitoring of group metrics for the Auto Scaling group specified in `AutoScalingGroupName`. You can specify the list of enabled metrics with the `Metrics` parameter.

Auto Scaling metrics collection can be turned on only if the `InstanceMonitoring` flag, in the Auto Scaling group's launch configuration, is set to `True`.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

The name or ARN of the Auto Scaling group.

Type: String

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

Required: Yes

Granularity

The granularity to associate with the metrics to collect. Currently, the only legal granularity is "1Minute".

Type: String

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

Required: Yes

Metrics.member.N

The list of metrics to collect. If no metrics are specified, all metrics are enabled. The following metrics are supported:

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

Note

The `GroupStandbyInstances` metric is not returned by default. You must explicitly request it when calling [EnableMetricsCollection \(p. 57\)](#).

Type: String list

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

Required: No

EnterStandby

Description

Move instances in an Auto Scaling group into a Standby mode.

To learn more about how to put instances into a Standby mode, see [Auto Scaling InService State](#).

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

The name of the Auto Scaling group from which to move instances into Standby mode.

Type: String

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

Required: Yes

InstanceIds.member.N

The instances to move into Standby mode. You must specify at least one instance ID.

Type: String list

Length constraints: Minimum length of 1. Maximum length of 16.

Required: No

ShouldDecrementDesiredCapacity

Specifies whether the instances moved to Standby mode count as part of the Auto Scaling group's desired capacity. If set, the desired capacity for the Auto Scaling group decrements by the number of instances moved to Standby mode.

Type: Boolean

Required: Yes

Response Elements

The following element is returned in a structure named `EnterStandbyResult`.

Activities

A list describing the activities related to moving instances into Standby mode.

Type: [Activity \(p. 85\)](#) list

Examples

Sample Request

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

```
EnterStandby&SignatureVersion=2&SignatureMethod=HmacSHA256&Timestamp=2014-06-13T22%3A35%3A50.567Z&AUTHPARAMS
```

Sample Response

```
<EnterStandbyResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <EnterStandbyResult>
    <Activities>
      <member>
        <ActivityId>462b4bc3-ad3b-4e67-a58d-96cd00f02f9e</ActivityId>
        <Progress>50</Progress>
        <StatusCode>InProgress</StatusCode>
        <StartTime>2014-06-13T22:35:50.884Z</StartTime>
        <Cause>At 2014-06-13T22:35:50Z instance i-5b73d709 was moved to standby
in response to a user request, shrinking the capacity from 4 to 3.</Cause>
        <AutoScalingGroupName>my-asg</AutoScalingGroupName>
        <Details>{"Availability Zone":"us-east-1a"}</Details>
        <Description>Moving EC2 instance to Standby: i-5b73d709</Description>
      </member>
    </Activities>
  </EnterStandbyResult>
  <ResponseMetadata>
    <RequestId>126f2f31-f34b-11e3-bc51-b35178f0274f</RequestId>
  </ResponseMetadata>
</EnterStandbyResponse>
```

ExecutePolicy

Description

Executes the specified policy.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

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

Type: String

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

Required: No

HonorCooldown

Set to `True` if you want Auto Scaling to wait for the cooldown period associated with the Auto Scaling group to complete before executing the policy.

Set to `False` if you want Auto Scaling to circumvent the cooldown period associated with the Auto Scaling group and execute the policy before the cooldown period ends.

For information about cooldown period, see [Cooldown Period](#) in the *Auto Scaling Developer Guide*.

Type: Boolean

Required: No

PolicyName

The name or ARN of the policy you want to run.

Type: String

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

Required: Yes

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

ScalingActivityInProgress

You cannot delete an Auto Scaling group while there are scaling activities in progress for that group.

HTTP Status Code: 400

ExitStandby

Description

Move an instance out of Standby mode.

To learn more about how to put instances that are in a Standby mode back into service, see [Auto Scaling InService State](#).

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

The name of the Auto Scaling group from which to move instances out of Standby mode.

Type: String

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

Required: Yes

InstanceIds.member.N

A list of instances to move out of Standby mode. You must specify at least one instance ID.

Type: String list

Length constraints: Minimum length of 1. Maximum length of 16.

Required: No

Response Elements

The following element is returned in a structure named `ExitStandbyResult`.

Activities

A list describing the activities related to moving instances out of Standby mode.

Type: [Activity \(p. 85\)](#) list

Examples

Sample Request

```
https://autoscaling.amazonaws.com/?InstanceIds.member.1=i-5b73d709&AutoScalingGroupName=my-asg&Version=2011-01-01&Action=ExitStandby&SignatureVersion=2&SignatureMet
hod=HmacSHA256&Timestamp=2014-06-13T22%3A43%3A53.182Z&AUTHPARAMS
```

Sample Response

```
<ExitStandbyResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <ExitStandbyResult>
    <Activities>
      <member>
        <ActivityId>dca4efcf-eea6-4844-8064-cab1fecdlaa2</ActivityId>
        <Progress>30</Progress>
        <StatusCode>PreInService</StatusCode>
        <StartTime>2014-06-13T22:43:53.523Z</StartTime>
        <Cause>At 2014-06-13T22:43:53Z instance i-5b73d709 was moved out of
standby in response to a user request, increasing the capacity from 3 to
4.</Cause>
        <AutoScalingGroupName>my-asg</AutoScalingGroupName>
        <Details>{"Availability Zone":"us-east-1a"}</Details>
        <Description>Moving EC2 instance out of Standby: i-5b73d709</Description>
      </member>
    </Activities>
  </ExitStandbyResult>
  <ResponseMetadata>
    <RequestId>321a11c8-f34c-11e3-a434-7f10009d5849</RequestId>
  </ResponseMetadata>
</ExitStandbyResponse>
```

PutLifecycleHook

Description

Creates or updates a lifecycle hook for an 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 operation is a part of the basic sequence for adding a lifecycle hook to an Auto Scaling group:

1. Create a notification target. A target can be either an Amazon SQS queue or an Amazon SNS topic.
2. Create an IAM role. This role allows Auto Scaling to publish lifecycle notifications to the designated SQS queue or SNS topic.
3. **Create the lifecycle hook. You can create a hook that acts when instances launch or when instances terminate.**
4. If necessary, record the lifecycle action heartbeat to keep the instance in a pending state.
5. Complete the lifecycle action.

To learn more, see [Auto Scaling Pending State](#) and [Auto Scaling Terminating State](#).

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

The name of the Auto Scaling group to which you want to assign the lifecycle hook.

Type: String

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

Required: Yes

DefaultResult

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

Type: String

Required: No

HeartbeatTimeout

Defines the amount of time, in seconds, that can elapse before the lifecycle hook times out. When the lifecycle hook times out, Auto Scaling performs the action defined in the `DefaultResult` parameter. You can prevent the lifecycle hook from timing out by calling [RecordLifecycleActionHeartbeat \(p. 73\)](#). The default value for this parameter is 3600 seconds (1 hour).

Type: Integer

Required: No

LifecycleHookName

The name of the lifecycle hook.

Type: String

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

Required: Yes

LifecycleTransition

The Amazon EC2 instance state to which you want to attach the lifecycle hook. See [DescribeLifecycleHookTypes](#) (p. 36) for a list of available lifecycle hook types.

Note

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.

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 ARN target can be either an SQS queue or an SNS topic.

Note

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

The notification message sent to the target will include:

- **LifecycleActionToken**. The Lifecycle action token.
- **AccountId**. The user account ID.
- **AutoScalingGroupName**. The name of the Auto Scaling group.
- **LifecycleHookName**. The lifecycle hook name.
- **EC2InstanceId**. The EC2 instance ID.
- **LifecycleTransition**. The lifecycle transition.
- **NotificationMetadata**. The notification metadata.

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 call this operation, a test message is sent to the notification target. This test message contains an additional key/value pair: `Event:autoscaling:TEST_NOTIFICATION`.

Type: String

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

Required: No

RoleARN

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

Note

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.

Required: No

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

LimitExceeded

The quota for capacity groups or launch configurations for this customer has already been reached.

HTTP Status Code: 400

Examples

Sample Request

```
http://autoscaling.amazonaws.com/?RoleARN=arn%3Aaws%3Aiam%3A%3A896650972448%3Arole%2FAutoScaling&AutoScalingGroupName=my-asg&LifecycleHookName=ReadyForSoftwareInst
all&NotificationTargetARN=arn%3Aaws%3Asqs%3Aus-east-1%3A896650972448%3Alifecyclehookqueue&LifecycleTransition=autoscaling%3AEC2_INSTANCE_LAUNCHING&Version=2011-01-01&Action=PutLifecycleHook&SignatureVersion=2&SignatureMethod=HmacSHA256&Timestamp=2014-06-17T17%3A30%3A36.125Z&AUTHPARAMS
```

Sample Response

```
<PutLifecycleHookResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <PutLifecycleHookResult/>
  <ResponseMetadata>
    <RequestId>1952f458-f645-11e3-bc51-b35178f0274f</RequestId>
  </ResponseMetadata>
</PutLifecycleHookResponse>
```

PutNotificationConfiguration

Description

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

For more information see [Get Email Notifications When Your Auto Scaling Group Changes](#)

A new `PutNotificationConfiguration` overwrites an existing configuration.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

The name of the Auto Scaling group.

Type: String

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

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 \(p. 33\)](#).

Type: String list

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

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.

Required: Yes

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

LimitExceeded

The quota for capacity groups or launch configurations for this customer has already been reached.

HTTP Status Code: 400

PutScalingPolicy

Description

Creates or updates a policy for an Auto Scaling group. To update an existing policy, use the existing policy name and set the parameter(s) you want to change. Any existing parameter not changed in an update to an existing policy is not changed in this update request.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AdjustmentType

Specifies whether the `ScalingAdjustment` is an absolute number or a percentage of the current capacity. Valid values are `ChangeInCapacity`, `ExactCapacity`, and `PercentChangeInCapacity`.

For more information about the adjustment types supported by Auto Scaling, see [Scale Based on Demand](#).

Type: String

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

Required: Yes

AutoScalingGroupName

The name or ARN of the Auto Scaling group.

Type: String

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

Required: Yes

Cooldown

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

For more information, see [Cooldown Period](#)

Type: Integer

Required: No

MinAdjustmentStep

Used with `AdjustmentType` with the value `PercentChangeInCapacity`, the scaling policy changes the `DesiredCapacity` of the Auto Scaling group by at least the number of instances specified in the value.

You will get a `ValidationError` if you use `MinAdjustmentStep` on a policy with an `AdjustmentType` other than `PercentChangeInCapacity`.

Type: Integer

Required: No

PolicyName

The name of the policy you want to create or update.

Type: String

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

Required: Yes

ScalingAdjustment

The number of instances by which to scale. `AdjustmentType` determines the interpretation of this number (e.g., as an absolute number or as a percentage of the existing Auto Scaling group size). A positive increment adds to the current capacity and a negative value removes from the current capacity.

Type: Integer

Required: Yes

Response Elements

The following element is returned in a structure named `PutScalingPolicyResult`.

PolicyARN

A policy's Amazon Resource Name (ARN).

Type: String

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

LimitExceeded

The quota for capacity groups or launch configurations for this customer has already been reached.

HTTP Status Code: 400

Examples

Sample Request

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

Sample Response

```
<PutScalingPolicyResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <PutScalingPolicyResult>
    <PolicyARN>arn:aws:autoscaling:us-east-1:803981987763:scalingPolicy:b0dcf5e8-02e6-4e31-9719-0675d0dc31ae:autoScalingGroupName/my-test-asg:policyName/my-scaleout-policy</PolicyARN>
```

Auto Scaling API Reference Examples

```
</PutScalingPolicyResult>  
<ResponseMetadata>  
  <RequestId>3cfc6fef-c08b-11e2-a697-2922EXAMPLE</RequestId>  
</ResponseMetadata>  
</PutScalingPolicyResponse>
```

PutScheduledUpdateGroupAction

Description

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 in the affected Auto Scaling group.

For information on creating or updating a scheduled action for your Auto Scaling group, see [Scale Based on a Schedule](#).

Note

Auto Scaling supports the date and time expressed in "YYYY-MM-DDThh:mm:ssZ" format in UTC/GMT only.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

The name or ARN of the Auto Scaling group.

Type: String

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

Required: Yes

DesiredCapacity

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

Type: Integer

Required: No

EndTime

The time for this action to end.

Type: DateTime

Required: No

MaxSize

The maximum size for the Auto Scaling group.

Type: Integer

Required: No

MinSize

The minimum size for the new Auto Scaling group.

Type: Integer

Required: No

Recurrence

The time when recurring future actions will start. Start time is specified by the user following the Unix cron syntax format. For information about cron syntax, go to [Wikipedia, The Free Encyclopedia](#).

When `StartTime` and `EndTime` are specified with `Recurrence`, they form the boundaries of when the recurring action will start and stop.

Type: String

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

Required: No

ScheduledActionName

The name of this scaling action.

Type: String

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

Required: Yes

StartTime

The time for this action to start, as in `--start-time 2010-06-01T00:00:00Z`.

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

When `StartTime` and `EndTime` are specified with `Recurrence`, they form the boundaries of when the recurring action will start and stop.

Type: DateTime

Required: No

Time

Time is deprecated.

The time for this action to start. `Time` is an alias for `StartTime` and can be specified instead of `StartTime`, or vice versa. If both `Time` and `StartTime` are specified, their values should be identical. Otherwise, `PutScheduledUpdateGroupAction` will return an error.

Type: DateTime

Required: No

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

AlreadyExists

The named Auto Scaling group or launch configuration already exists.

HTTP Status Code: 400

LimitExceeded

The quota for capacity groups or launch configurations for this customer has already been reached.

HTTP Status Code: 400

Examples

Schedule based on a specific date and time

Sample Request

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

Sample Response

```
<PutScheduledUpdateGroupActionResponse xmlns="http://autoscaling.amazon
aws.com/doc/2011-01-01/">
  <ResponseMetadata>
    <RequestId>3bc8c9bc-6a62-11e2-8a51-4b8a1EXAMPLE</RequestId>
  </ResponseMetadata>
</PutScheduledUpdateGroupActionResponse>
```

Recurring Schedule

Sample Request

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

Sample Response

```
<PutScheduledUpdateGroupActionResponse xmlns="http://autoscaling.amazon
aws.com/doc/2011-01-01/">
  <ResponseMetadata>
    <RequestId>3bc8c9bc-6a62-11e2-8a51-4b8a1EXAMPLE</RequestId>
  </ResponseMetadata>
</PutScheduledUpdateGroupActionResponse>
```


RecordLifecycleActionHeartbeat

Description

Records a heartbeat for the lifecycle action associated with a specific token. This extends the timeout by the length of time defined by the `HeartbeatTimeout` parameter of the [PutLifecycleHook \(p. 63\)](#) operation.

This operation is a part of the basic sequence for adding a lifecycle hook to an Auto Scaling group:

1. Create a notification target. A target can be either an Amazon SQS queue or an Amazon SNS topic.
2. Create an IAM role. This role allows Auto Scaling to publish lifecycle notifications to the designated SQS queue or SNS topic.
3. Create the lifecycle hook. You can create a hook that acts when instances launch or when instances terminate.
4. **If necessary, record the lifecycle action heartbeat to keep the instance in a pending state.**
5. Complete the lifecycle action.

To learn more, see [Auto Scaling Pending State](#) and [Auto Scaling Terminating State](#).

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

The name of the Auto Scaling group to which the hook belongs.

Type: String

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

Required: Yes

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: Minimum length of 36. Maximum length of 36.

Required: Yes

LifecycleHookName

The name of the lifecycle hook.

Type: String

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

Required: Yes

ResumeProcesses

Description

Resumes all suspended Auto Scaling processes for an Auto Scaling group. For information on suspending and resuming Auto Scaling process, see [Suspend and Resume Auto Scaling Process](#).

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

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

Type: String

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

Required: Yes

ScalingProcesses.member.N

The processes that you want to suspend or resume, which can include one or more of the following:

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

To suspend all process types, omit this parameter.

Type: String list

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

Required: No

SetDesiredCapacity

Description

Sets the desired size of the specified [AutoScalingGroup](#) (p. 87).

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters](#) (p. 117).

AutoScalingGroupName

The name of the Auto Scaling group.

Type: String

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

Required: Yes

DesiredCapacity

The new capacity setting for the Auto Scaling group.

Type: Integer

Required: Yes

HonorCooldown

By default, `SetDesiredCapacity` overrides any cooldown period associated with the Auto Scaling group. Set to `True` if you want Auto Scaling to wait for the cooldown period associated with the Auto Scaling group to complete before initiating a scaling activity to set your Auto Scaling group to the new capacity setting.

Type: Boolean

Required: No

Errors

For information about the errors that are common to all actions, see [Common Errors](#) (p. 121).

ScalingActivityInProgress

You cannot delete an Auto Scaling group while there are scaling activities in progress for that group.

HTTP Status Code: 400

Examples

Sample Request

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

```
&Action=SetDesiredCapacity  
&AUTHPARAMS
```

Sample Response

```
<SetDesiredCapacityResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">  
  <ResponseMetadata>  
    <RequestId>9fb7e2db-6998-11e2-a985-57c82EXAMPLE</RequestId>  
  </ResponseMetadata>  
</SetDesiredCapacityResponse>
```

SetInstanceHealth

Description

Sets the health status of a specified instance that belongs to any of your Auto Scaling groups.

For more information, see [Configure Health Checks for Your Auto Scaling group](#).

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

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.

Required: Yes

InstanceId

The identifier of the Amazon EC2 instance.

Type: String

Length constraints: Minimum length of 1. Maximum length of 16.

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 `HealthCheckGracePeriod` parameter description in the [CreateAutoScalingGroup \(p. 7\)](#) action.

Type: Boolean

Required: No

SuspendProcesses

Description

Suspends Auto Scaling processes for an Auto Scaling group. To suspend specific process types, specify them by name with the `ScalingProcesses.member.N` parameter. To suspend all process types, omit the `ScalingProcesses.member.N` parameter.

Important

Suspending either of the two primary process types, `Launch` or `Terminate`, can prevent other process types from functioning properly.

To resume processes that have been suspended, use [ResumeProcesses \(p. 74\)](#) For more information on suspending and resuming Auto Scaling process, see [Suspend and Resume Auto Scaling Process](#).

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

AutoScalingGroupName

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

Type: String

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

Required: Yes

ScalingProcesses.member.N

The processes that you want to suspend or resume, which can include one or more of the following:

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

To suspend all process types, omit this parameter.

Type: String list

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

Required: No

TerminateInstanceInAutoScalingGroup

Description

Terminates the specified instance. Optionally, the desired group size can be adjusted.

Note

This call simply registers a termination request. The termination of the instance cannot happen immediately.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters \(p. 117\)](#).

InstanceId

The ID of the Amazon EC2 instance to be terminated.

Type: String

Length constraints: Minimum length of 1. Maximum length of 16.

Required: Yes

ShouldDecrementDesiredCapacity

Specifies whether (*true*) or not (*false*) terminating this instance should also decrement the size of the [AutoScalingGroup \(p. 87\)](#).

Type: Boolean

Required: Yes

Response Elements

The following element is returned in a structure named `TerminateInstanceInAutoScalingGroupResult`.

Activity

A scaling Activity.

Type: [Activity \(p. 85\)](#)

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

ScalingActivityInProgress

You cannot delete an Auto Scaling group while there are scaling activities in progress for that group.

HTTP Status Code: 400

UpdateAutoScalingGroup

Description

Updates the configuration for the specified [AutoScalingGroup](#) (p. 87).

Note

To update an Auto Scaling group with a launch configuration that has the `InstanceMonitoring` flag set to `False`, you must first ensure that collection of group metrics is disabled. Otherwise, calls to [UpdateAutoScalingGroup](#) (p. 80) will fail. If you have previously enabled group metrics collection, you can disable collection of all group metrics by calling [DisableMetricsCollection](#) (p. 56).

The new settings are registered upon the completion of this call. Any launch configuration settings take effect on any triggers after this call returns. Scaling activities that are currently in progress aren't affected.

Note

- If a new value is specified for *MinSize* without specifying the value for *DesiredCapacity*, and if the new *MinSize* is larger than the current size of the Auto Scaling Group, there will be an implicit call to [SetDesiredCapacity](#) (p. 75) to set the group to the new *MinSize*.
- If a new value is specified for *MaxSize* without specifying the value for *DesiredCapacity*, and the new *MaxSize* is smaller than the current size of the Auto Scaling Group, there will be an implicit call to [SetDesiredCapacity](#) (p. 75) to set the group to the new *MaxSize*.
- All other optional parameters are left unchanged if not passed in the request.

Request Parameters

For information about the common parameters that all actions use, see [Common Parameters](#) (p. 117).

AutoScalingGroupName

The name of the Auto Scaling group.

Type: String

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

Required: Yes

AvailabilityZones.member.N

Availability Zones for the group.

Type: String list

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

Length constraints: Minimum of 1 item(s) in the list.

Required: No

DefaultCooldown

The amount of time, in seconds, after a scaling activity completes before any further scaling activities can start. For more information, see [Cooldown Period](#).

Type: Integer

Required: No

DesiredCapacity

The desired capacity for the Auto Scaling group.

Type: Integer

Required: No

HealthCheckGracePeriod

The length of time that Auto Scaling waits before checking an instance's health status. The grace period begins when the instance passes System Status and the Instance Status checks from Amazon EC2. For more information, see [DescribeInstanceStatus](#).

Type: Integer

Required: No

HealthCheckType

The type of health check for the instances in the Auto Scaling group. The health check type can either be `EC2` for Amazon EC2 or `ELB` for Elastic Load Balancing.

Type: String

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

Required: No

LaunchConfigurationName

The name of the launch configuration.

Type: String

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

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

PlacementGroup

The name of the cluster placement group, if applicable. For more information, go to [Using Cluster Instances](#) in the Amazon EC2 User Guide.

Type: String

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

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 on creating a termination policy for your Auto Scaling group, go to [Instance Termination Policy for Your Auto Scaling Group](#) in the *Auto Scaling Developer Guide*.

Type: String list

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

Required: No

VPCZoneIdentifier

The subnet identifier for the Amazon VPC connection, if applicable. 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 on creating your Auto Scaling group in Amazon VPC by specifying subnets, see [Launch Auto Scaling Instances into Amazon VPC](#) in the the *Auto Scaling Developer Guide*.

Type: String

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

Required: No

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 121\)](#).

ScalingActivityInProgress

You cannot delete an Auto Scaling group while there are scaling activities in progress for that group.

HTTP Status Code: 400

Examples

Update existing Auto Scaling group with ELB health check

Sample Request

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

Sample Response

```
<UpdateAutoScalingGroupResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <ResponseMetadata>
    <RequestId>adafead0-ab8a-11e2-ba13-ab0ccEXAMPLE</RequestId>
  </ResponseMetadata>
</UpdateAutoScalingGroupResponse>
```

Update existing Auto Scaling group with a new Availability Zone

Sample Request

```
https://autoscaling.amazonaws.com/?AutoScalingGroupName=my-test-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
&Action=UpdateAutoScalingGroup
&AUTHPARAMS
```

Sample Response

```
<UpdateAutoScalingGroupResponse xmlns="http://autoscaling.amazonaws.com/doc/2011-01-01/">
  <ResponseMetadata>
    <RequestId>adafead0-ab8a-11e2-ba13-ab0ccEXAMPLE</RequestId>
  </ResponseMetadata>
</UpdateAutoScalingGroupResponse>
```

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 the response is not guaranteed. Applications should not assume a particular order.

The following data types are supported:

- [Activity](#) (p. 85)
- [AdjustmentType](#) (p. 87)
- [Alarm](#) (p. 87)
- [AutoScalingGroup](#) (p. 87)
- [AutoScalingInstanceDetails](#) (p. 90)
- [BlockDeviceMapping](#) (p. 91)
- [CompleteLifecycleActionResult](#) (p. 92)
- [DeleteLifecycleHookResult](#) (p. 92)
- [DescribeAccountLimitsResult](#) (p. 92)
- [DescribeAdjustmentTypesResult](#) (p. 93)
- [DescribeAutoScalingGroupsResult](#) (p. 93)
- [DescribeAutoScalingInstancesResult](#) (p. 94)
- [DescribeAutoScalingNotificationTypesResult](#) (p. 94)
- [DescribeLaunchConfigurationsResult](#) (p. 94)
- [DescribeLifecycleHookTypesResult](#) (p. 95)
- [DescribeLifecycleHooksResult](#) (p. 95)
- [DescribeMetricCollectionTypesResult](#) (p. 96)
- [DescribeNotificationConfigurationsResult](#) (p. 96)
- [DescribePoliciesResult](#) (p. 97)
- [DescribeScalingActivitiesResult](#) (p. 97)
- [DescribeScalingProcessTypesResult](#) (p. 98)
- [DescribeScheduledActionsResult](#) (p. 98)
- [DescribeTagsResult](#) (p. 98)
- [DescribeTerminationPolicyTypesResult](#) (p. 99)

- [DetachInstancesResult](#) (p. 99)
- [Ebs](#) (p. 99)
- [EnabledMetric](#) (p. 101)
- [EnterStandbyResult](#) (p. 101)
- [ExitStandbyResult](#) (p. 101)
- [Filter](#) (p. 102)
- [Instance](#) (p. 102)
- [InstanceMonitoring](#) (p. 103)
- [LaunchConfiguration](#) (p. 104)
- [LifecycleHook](#) (p. 106)
- [MetricCollectionType](#) (p. 108)
- [MetricGranularityType](#) (p. 108)
- [NotificationConfiguration](#) (p. 108)
- [ProcessType](#) (p. 109)
- [PutLifecycleHookResult](#) (p. 111)
- [PutScalingPolicyResult](#) (p. 111)
- [RecordLifecycleActionHeartbeatResult](#) (p. 111)
- [ScalingPolicy](#) (p. 111)
- [ScheduledUpdateGroupAction](#) (p. 113)
- [SuspendedProcess](#) (p. 114)
- [Tag](#) (p. 115)
- [TagDescription](#) (p. 115)
- [TerminateInstanceInAutoScalingGroupResult](#) (p. 116)

Activity

Description

A scaling Activity is a long-running process that represents a change to your AutoScalingGroup, such as changing the size of the group. It can also be a process to replace an instance, or a process to perform any other long-running operations supported by the API.

Contents

ActivityId

Specifies the ID of the activity.

Type: String

Required: Yes

AutoScalingGroupName

The name of the Auto Scaling group.

Type: String

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

Required: Yes

Cause

Contains the reason the activity was begun.

Type: String

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

Required: Yes

Description

Contains a friendly, more verbose description of the scaling activity.

Type: String

Required: No

Details

Contains details of the scaling activity.

Type: String

Required: No

EndTime

Provides the end time of this activity.

Type: DateTime

Required: No

Progress

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

Type: Integer

Required: No

StartTime

Provides the start time of this activity.

Type: DateTime

Required: Yes

StatusCode

Contains the current status of the activity.

Type: String

Valid Values: `WaitingForSpotInstanceRequestId` | `WaitingForSpotInstanceId` | `WaitingForInstanceId` | `PreInService` | `InProgress` | `WaitingForELBConnection-Draining` | `MidLifecycleAction` | `Successful` | `Failed` | `Cancelled`

Required: Yes

StatusMessage

Contains a friendly, more verbose description of the activity status.

Type: String

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

Required: No

AdjustmentType

Description

Specifies whether the [PutScalingPolicy](#) (p. 67) `ScalingAdjustment` parameter is an absolute number or a percentage of the current capacity.

Contents

AdjustmentType

A policy adjustment type. Valid values are `ChangeInCapacity`, `ExactCapacity`, and `PercentChangeInCapacity`.

Type: String

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

Required: No

Alarm

Description

The Alarm data type.

Contents

AlarmARN

The Amazon Resource Name (ARN) of the alarm.

Type: String

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

Required: No

AlarmName

The name of the alarm.

Type: String

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

Required: No

AutoScalingGroup

Description

The AutoScalingGroup data type.

Contents

AutoScalingGroupARN

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

Type: String

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

Required: No

AutoScalingGroupName

Specifies the name of the group.

Type: String

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

Required: Yes

AvailabilityZones

Contains a list of Availability Zones for the group.

Type: String list

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

Length constraints: Minimum of 1 item(s) in the list.

Required: Yes

CreatedTime

Specifies the date and time the Auto Scaling group was created.

Type: DateTime

Required: Yes

DefaultCooldown

The number of seconds after a scaling activity completes before any further scaling activities can start.

Type: Integer

Required: Yes

DesiredCapacity

Specifies the desired capacity for the Auto Scaling group.

Type: Integer

Required: Yes

EnabledMetrics

A list of metrics enabled for this Auto Scaling group.

Type: [EnabledMetric \(p. 101\)](#) list

Required: No

HealthCheckGracePeriod

The length of time that Auto Scaling waits before checking an instance's health status. The grace period begins when an instance comes into service.

Type: Integer

Required: No

HealthCheckType

The service of interest for the health status check, either "EC2" for Amazon EC2 or "ELB" for Elastic Load Balancing.

Type: String

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

Required: Yes

Instances

Provides a summary list of Amazon EC2 instances.

Type: [Instance \(p. 102\)](#) list

Required: No

LaunchConfigurationName

Specifies the name of the associated [LaunchConfiguration \(p. 104\)](#).

Type: String

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

Required: Yes

LoadBalancerNames

A list of load balancers associated with this Auto Scaling group.

Type: String list

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

Required: No

MaxSize

Contains the maximum size of the Auto Scaling group.

Type: Integer

Required: Yes

MinSize

Contains the minimum size of the Auto Scaling group.

Type: Integer

Required: Yes

PlacementGroup

The name of the cluster placement group, if applicable. For more information, go to [Using Cluster Instances](#) in the Amazon EC2 User Guide.

Type: String

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

Required: No

Status

The current state of the Auto Scaling group when a [DeleteAutoScalingGroup \(p. 18\)](#) action is in progress.

Type: String

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

Required: No

SuspendedProcesses

Suspended processes associated with this Auto Scaling group.

Type: [SuspendedProcess](#) (p. 114) list

Required: No

Tags

A list of tags for the Auto Scaling group.

Type: [TagDescription](#) (p. 115) list

Required: No

TerminationPolicies

A standalone termination policy or a list of termination policies for this Auto Scaling group.

Type: String list

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

Required: No

VPCZoneIdentifier

The subnet identifier for the Amazon VPC connection, if applicable. 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`.

Type: String

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

Required: No

AutoScalingInstanceDetails

Description

The `AutoScalingInstanceDetails` data type.

Contents

AutoScalingGroupName

The name of the Auto Scaling group associated with this instance.

Type: String

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

Required: Yes

AvailabilityZone

The Availability Zone in which this instance resides.

Type: String

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

Required: Yes

HealthStatus

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

Type: String

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

Required: Yes

InstanceId

The instance ID of the Amazon EC2 instance.

Type: String

Length constraints: Minimum length of 1. Maximum length of 16.

Required: Yes

LaunchConfigurationName

The launch configuration associated with this instance.

Type: String

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

Required: Yes

LifecycleState

The life cycle state of this instance. for more information, see [Instance Lifecycle State](#) in the *Auto Scaling Developer Guide*.

Type: String

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

Required: Yes

BlockDeviceMapping

Description

The `BlockDeviceMapping` data type.

Contents

DeviceName

The name of the device within Amazon EC2 (for example, `/dev/sdh` or `xvdh`).

Type: String

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

Required: Yes

Ebs

The Elastic Block Storage volume information.

Type: [Ebs \(p. 99\)](#)

Required: No

NoDevice

Suppresses the device mapping.

Note

If `NoDevice` is set to `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 virtual name associated with the device.

Type: String

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

Required: No

CompleteLifecycleActionResult

Description

The output of the [CompleteLifecycleAction \(p. 5\)](#).

Contents

DeleteLifecycleHookResult

Description

The output of the [DeleteLifecycleHook \(p. 21\)](#) action.

Contents

DescribeAccountLimitsResult

Description

The output of the [DescribeAccountLimitsResult \(p. 92\)](#) action.

Contents

MaxNumberOfAutoScalingGroups

The maximum number of Auto Scaling groups allowed for your AWS account.

Type: Integer

Required: No

MaxNumberOfLaunchConfigurations

The maximum number of launch configurations allowed for your AWS account.

Type: Integer

Required: No

DescribeAdjustmentTypesResult

Description

The output of the [DescribeAdjustmentTypes](#) (p. 27) action.

Contents

AdjustmentTypes

A list of specific policy adjustment types.

Type: [AdjustmentType](#) (p. 87) list

Required: No

DescribeAutoScalingGroupsResult

Description

The `AutoScalingGroupsType` data type.

Contents

AutoScalingGroups

A list of Auto Scaling groups.

Type: [AutoScalingGroup](#) (p. 87) list

Required: Yes

NextToken

A string that marks the start of the next batch of returned results.

Type: String

Required: No

DescribeAutoScalingInstancesResult

Description

The `AutoScalingInstancesType` data type.

Contents

AutoScalingInstances

A list of Auto Scaling instances.

Type: [AutoScalingInstanceDetails](#) (p. 90) list

Required: No

NextToken

A string that marks the start of the next batch of returned results.

Type: String

Required: No

DescribeAutoScalingNotificationTypesResult

Description

The `AutoScalingNotificationTypes` data type.

Contents

AutoScalingNotificationTypes

Returns a list of all notification types supported by Auto Scaling. They are:

- `autoscaling:EC2_INSTANCE_LAUNCH`
- `autoscaling:EC2_INSTANCE_LAUNCH_ERROR`
- `autoscaling:EC2_INSTANCE_TERMINATE`
- `autoscaling:EC2_INSTANCE_TERMINATE_ERROR`
- `autoscaling:TEST_NOTIFICATION`

Type: String list

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

Required: No

DescribeLaunchConfigurationsResult

Description

The `LaunchConfigurationsType` data type.

Contents

LaunchConfigurations

A list of launch configurations.

Type: [LaunchConfiguration \(p. 104\)](#) list

Required: Yes

NextToken

A string that marks the start of the next batch of returned results.

Type: String

Required: No

DescribeLifecycleHookTypesResult

Description

Contents

LifecycleHookTypes

Returns a list of all notification types supported by Auto Scaling. They are:

- `autoscaling:EC2_INSTANCE_LAUNCHING`
- `autoscaling:EC2_INSTANCE_TERMINATING`

Type: String list

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

Required: No

DescribeLifecycleHooksResult

Description

The output of the [DescribeLifecycleHooks \(p. 37\)](#) action.

Contents

LifecycleHooks

A list describing the lifecycle hooks that belong to the specified Auto Scaling group.

Type: [LifecycleHook \(p. 106\)](#) list

Required: No

DescribeMetricCollectionTypesResult

Description

The output of the [DescribeMetricCollectionTypes](#) (p. 38) action.

Contents

Granularities

A list of granularities for the listed Metrics.

Type: [MetricGranularityType](#) (p. 108) list

Required: No

Metrics

The list of Metrics collected. The following metrics are supported:

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

Note

The `GroupStandbyInstances` metric is not returned by default. You must explicitly request it when calling [EnableMetricsCollection](#) (p. 57).

Type: [MetricCollectionType](#) (p. 108) list

Required: No

DescribeNotificationConfigurationsResult

Description

The output of the [DescribeNotificationConfigurations](#) (p. 40) action.

Contents

NextToken

A string that is used to mark the start of the next batch of returned results for pagination.

Type: String

Required: No

NotificationConfigurations

The list of notification configurations.

Type: [NotificationConfiguration](#) (p. 108) list

Required: Yes

DescribePoliciesResult

Description

The `PoliciesType` data type.

Contents

NextToken

A string that marks the start of the next batch of returned results.

Type: String

Required: No

ScalingPolicies

A list of scaling policies.

Type: [ScalingPolicy](#) (p. 111) list

Required: No

DescribeScalingActivitiesResult

Description

The output for the [DescribeScalingActivities](#) (p. 44) action.

Contents

Activities

A list of the requested scaling activities.

Type: [Activity](#) (p. 85) list

Required: Yes

NextToken

Acts as a paging mechanism for large result sets. Set to a non-empty string if there are additional results waiting to be returned. Pass this in to subsequent calls to return additional results.

Type: String

Required: No

DescribeScalingProcessTypesResult

Description

The output of the [DescribeScalingProcessTypes](#) (p. 47) action.

Contents

Processes

A list of [ProcessType](#) (p. 109) names.

Type: [ProcessType](#) (p. 109) list

Required: No

DescribeScheduledActionsResult

Description

A scaling action that is scheduled for a future time and date. An action can be scheduled up to thirty days in advance.

Starting with API version 2011-01-01, you can use `recurrence` to specify that a scaling action occurs regularly on a schedule.

Contents

NextToken

A string that marks the start of the next batch of returned results.

Type: String

Required: No

ScheduledUpdateGroupActions

A list of scheduled actions designed to update an Auto Scaling group.

Type: [ScheduledUpdateGroupAction](#) (p. 113) list

Required: No

DescribeTagsResult

Description

Contents

NextToken

A string used to mark the start of the next batch of returned results.

Type: String

Required: No

Tags

The list of tags.

Type: [TagDescription \(p. 115\)](#) list

Required: No

DescribeTerminationPolicyTypesResult

Description

The `TerminationPolicyTypes` data type.

Contents

TerminationPolicyTypes

Termination policies supported by Auto Scaling. They are: `OldestInstance`, `OldestLaunchConfiguration`, `NewestInstance`, `ClosestToNextInstanceHour`, `Default`

Type: String list

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

Required: No

DetachInstancesResult

Description

The output of the [DetachInstances \(p. 54\)](#) action.

Contents

Activities

A list describing the activities related to detaching the instances from the Auto Scaling group.

Type: [Activity \(p. 85\)](#) list

Required: No

Ebs

Description

The Ebs data type.

Contents

DeleteOnTermination

Indicates whether to delete the volume on instance termination.

Default: `true`

Type: Boolean

Required: No

lops

The number of I/O operations per second (IOPS) that the volume supports.

The maximum ratio of IOPS to volume size is 30.0

Valid Values: Range is 100 to 4000.

Default: None.

Type: Integer

Required: No

SnapshotId

The snapshot ID.

Type: String

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

Required: No

VolumeSize

The volume size, in gigabytes.

Valid values: If the volume type is `io1`, the minimum size of the volume is 10.

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

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

Type: Integer

Required: No

VolumeType

The volume type.

Valid values: `standard` | `io1`

Default: `standard`

Type: String

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

Required: No

EnabledMetric

Description

The `EnabledMetric` data type.

Contents

Granularity

The granularity of the enabled metric.

Type: String

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

Required: No

Metric

The name of the enabled metric.

Type: String

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

Required: No

EnterStandbyResult

Description

The output of the [EnterStandby \(p. 58\)](#) action.

Contents

Activities

A list describing the activities related to moving instances into Standby mode.

Type: [Activity \(p. 85\)](#) list

Required: No

ExitStandbyResult

Description

The output of the [ExitStandby \(p. 61\)](#) action.

Contents

Activities

A list describing the activities related to moving instances out of Standby mode.

Type: [Activity \(p. 85\)](#) list

Required: No

Filter

Description

The `Filter` data type.

Contents

Name

The name of the filter. Valid Name values are: "auto-scaling-group", "key", "value", and "propagate-at-launch".

Type: String

Required: No

Values

The value of the filter.

Type: String list

Required: No

Instance

Description

The `Instance` data type.

Contents

AvailabilityZone

Availability Zones associated with this instance.

Type: String

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

Required: Yes

HealthStatus

The instance's health status.

Type: String

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

Required: Yes

InstanceId

Specifies the ID of the Amazon EC2 instance.

Type: String

Length constraints: Minimum length of 1. Maximum length of 16.

Required: Yes

LaunchConfigurationName

The launch configuration associated with this instance.

Type: String

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

Required: Yes

LifecycleState

Contains a description of the current *lifecycle* state.

Note

The `Quarantined` lifecycle state is currently 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

InstanceMonitoring

Description

The `InstanceMonitoring` data type.

Contents

Enabled

If `True`, instance monitoring is enabled.

Type: Boolean

Required: No

LaunchConfiguration

Description

The `LaunchConfiguration` data type.

Contents

AssociatePublicIpAddress

Specifies whether the instance is associated with a public IP address (`true`) or not (`false`).

Type: Boolean

Required: No

BlockDeviceMappings

Specifies how block devices are exposed to the instance. Each mapping is made up of a *virtualName* and a *deviceName*.

Type: [BlockDeviceMapping](#) (p. 91) list

Required: No

CreatedTime

Provides the creation date and time for this launch configuration.

Type: DateTime

Required: Yes

EbsOptimized

Specifies whether the instance is optimized for EBS I/O (`true`) or not (`false`).

Type: Boolean

Required: No

IamInstanceProfile

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

Type: String

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

Required: No

ImageId

Provides the unique ID of the *Amazon Machine Image* (AMI) that was assigned during registration.

Type: String

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

Required: Yes

InstanceMonitoring

Controls whether instances in this group are launched with detailed monitoring or not.

Type: [InstanceMonitoring](#) (p. 103)

Required: No

InstanceType

Specifies the instance type of the Amazon EC2 instance.

Type: String

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

Required: Yes

KernelId

Provides the ID of the kernel associated with the Amazon EC2 AMI.

Type: String

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

Required: No

KeyName

Provides the name of the Amazon EC2 key pair.

Type: String

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

Required: No

LaunchConfigurationARN

The launch configuration's Amazon Resource Name (ARN).

Type: String

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

Required: No

LaunchConfigurationName

Specifies the name of the launch configuration.

Type: String

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

Required: Yes

PlacementTenancy

Specifies the tenancy of the instance. It can be either `default` or `dedicated`. An instance with `dedicated` tenancy runs in an isolated, single-tenant hardware and it can only be launched in a VPC.

Type: String

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

Required: No

RamdiskId

Provides ID of the RAM disk associated with the Amazon EC2 AMI.

Type: String

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

Required: No

SecurityGroups

A description of the security groups to associate with the Amazon EC2 instances.

Type: String list

Required: No

SpotPrice

Specifies 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 launched Amazon EC2 instances.

Type: String

Length constraints: Minimum length of 0. Maximum length of 21847.

Required: No

LifecycleHook

Description

A lifecycle hook tells Auto Scaling that you want to perform an action when an instance launches or terminates. When you have a lifecycle hook in place, the Auto Scaling group will either:

- Pause the instance after it launches, but before it is put into service
- Pause the instance as it terminates, but before it is fully terminated

To learn more, see [Auto Scaling Pending State](#) and [Auto Scaling Terminating State](#).

Contents

AutoScalingGroupName

The name of the Auto Scaling group to which the lifecycle action belongs.

Type: String

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

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 value for this parameter can be either `CONTINUE` or `ABANDON`. The default value for this parameter is `CONTINUE`.

Type: String

Required: No

GlobalTimeout

The maximum length of time an instance can remain in a `Pending:Wait` or `Terminating:Wait` state. Currently, this value is set at 48 hours.

Type: Integer

Required: No

HeartbeatTimeout

Defines the amount of time that can elapse before the lifecycle hook times out. When the lifecycle hook times out, Auto Scaling performs the action defined in the `DefaultResult` parameter. You can prevent the lifecycle hook from timing out by calling [RecordLifecycleActionHeartbeat](#) (p. 73).

Type: Integer

Required: No

LifecycleHookName

The name of the lifecycle action hook.

Type: String

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

Required: No

LifecycleTransition

The Amazon EC2 instance state to which you want to attach the lifecycle hook. See [DescribeLifecycleHooks](#) (p. 37) for a list of available lifecycle hook types.

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.

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 ARN target can be either an SQS queue or an SNS topic. The notification message sent to the target will include:

- Lifecycle action token
- User account ID
- Name of the Auto Scaling group
- Lifecycle hook name
- EC2 instance ID
- Lifecycle transition
- Notification metadata

Type: String

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

Required: No

RoleARN

The ARN of the Amazon 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.

Required: No

MetricCollectionType

Description

The MetricCollectionType data type.

Contents

Metric

Type: String

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

Required: No

MetricGranularityType

Description

The MetricGranularityType data type.

Contents

Granularity

The granularity of a Metric.

Type: String

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

Required: No

NotificationConfiguration

Description

The NotificationConfiguration data type.

Contents

AutoScalingGroupName

Specifies the Auto Scaling group name.

Type: String

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

Required: No

NotificationType

The types of events for an action to start.

Type: String

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

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.

Required: No

ProcessType

Description

There are two primary Auto Scaling process types--`Launch` and `Terminate`. The `Launch` process creates a new Amazon EC2 instance for an Auto Scaling group, and the `Terminate` process removes an existing Amazon EC2 instance.

The remaining Auto Scaling process types relate to specific Auto Scaling features:

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

Important

If you suspend `Launch` or `Terminate`, all other process types are affected to varying degrees. The following descriptions discuss how each process type is affected by a suspension of `Launch` or `Terminate`.

The `AddToLoadBalancer` process type adds instances to the load balancer when the instances are launched. If you suspend this process, Auto Scaling will launch the instances but will not add them to the load balancer. If you resume the `AddToLoadBalancer` process, Auto Scaling will also resume adding new instances to the load balancer when they are launched. However, Auto Scaling will not add running

instances that were launched while the process was suspended; those instances must be added manually using the [RegisterInstancesWithLoadBalancer](#) call in the *Elastic Load Balancing API Reference*.

The `AlarmNotification` process type accepts notifications from Amazon CloudWatch alarms that are associated with the Auto Scaling group. If you suspend the `AlarmNotification` process type, Auto Scaling will not automatically execute scaling policies that would be triggered by alarms.

Although the `AlarmNotification` process type is not directly affected by a suspension of `Launch` or `Terminate`, alarm notifications are often used to signal that a change in the size of the Auto Scaling group is warranted. If you suspend `Launch` or `Terminate`, Auto Scaling might not be able to implement the alarm's associated policy.

The `AZRebalance` process type seeks to maintain a balanced number of instances across Availability Zones within a Region. If you remove an Availability Zone from your Auto Scaling group or an Availability Zone otherwise becomes unhealthy or unavailable, Auto Scaling launches new instances in an unaffected Availability Zone before terminating the unhealthy or unavailable instances. When the unhealthy Availability Zone returns to a healthy state, Auto Scaling automatically redistributes the application instances evenly across all of the designated Availability Zones.

Important

If you call [SuspendProcesses \(p. 78\)](#) on the `launch` process type, the `AZRebalance` process will neither launch new instances nor terminate existing instances. This is because the `AZRebalance` process terminates existing instances only after launching the replacement instances.

If you call [SuspendProcesses \(p. 78\)](#) on the `terminate` process type, the `AZRebalance` process can cause your Auto Scaling group to grow up to ten percent larger than the maximum size. This is because Auto Scaling allows groups to temporarily grow larger than the maximum size during rebalancing activities. If Auto Scaling cannot terminate instances, your Auto Scaling group could remain up to ten percent larger than the maximum size until you resume the `terminate` process type.

The `HealthCheck` process type checks the health of the instances. Auto Scaling marks an instance as unhealthy if Amazon EC2 or Elastic Load Balancing informs Auto Scaling that the instance is unhealthy. The `HealthCheck` process can override the health status of an instance that you set with [SetInstanceHealth \(p. 77\)](#).

The `ReplaceUnhealthy` process type terminates instances that are marked as unhealthy and subsequently creates new instances to replace them. This process calls both of the primary process types—first `Terminate` and then `Launch`.

Important

The `HealthCheck` process type works in conjunction with the `ReplaceUnhealthy` process type to provide health check functionality. If you suspend either `Launch` or `Terminate`, the `ReplaceUnhealthy` process type will not function properly.

The `ScheduledActions` process type performs scheduled actions that you create with [PutScheduledUpdateGroupAction \(p. 70\)](#). Scheduled actions often involve launching new instances or terminating existing instances. If you suspend either `Launch` or `Terminate`, your scheduled actions might not function as expected.

Contents

ProcessName

The name of a process.

Type: String

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

Required: Yes

PutLifecycleHookResult

Description

The output of the [PutLifecycleHook](#) (p. 63) action.

Contents

PutScalingPolicyResult

Description

The `PolicyARNType` data type.

Contents

PolicyARN

A policy's Amazon Resource Name (ARN).

Type: String

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

Required: No

RecordLifecycleActionHeartbeatResult

Description

The output of the [RecordLifecycleActionHeartbeat](#) (p. 73) action.

Contents

ScalingPolicy

Description

The `ScalingPolicy` data type.

Contents

AdjustmentType

Specifies whether the `ScalingAdjustment` is an absolute number or a percentage of the current capacity. Valid values are `ChangeInCapacity`, `ExactCapacity`, and `PercentChangeInCapacity`.

Type: String

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

Required: No

Alarms

A list of CloudWatch Alarms related to the policy.

Type: [Alarm \(p. 87\)](#) list

Required: No

AutoScalingGroupName

The name of the Auto Scaling group associated with this scaling policy.

Type: String

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

Required: No

Cooldown

The amount of time, in seconds, after a scaling activity completes before any further trigger-related scaling activities can start.

Type: Integer

Required: No

MinAdjustmentStep

Changes the `DesiredCapacity` of the Auto Scaling group by at least the specified number of instances.

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.

Required: No

PolicyName

The name of the scaling policy.

Type: String

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

Required: No

ScalingAdjustment

The number associated with the specified adjustment type. A positive value adds to the current capacity and a negative value removes from the current capacity.

Type: Integer

Required: No

ScheduledUpdateGroupAction

Description

This data type stores information about a scheduled update to an Auto Scaling group.

Contents

AutoScalingGroupName

The name of the Auto Scaling group to be updated.

Type: String

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

Required: No

DesiredCapacity

The number of instances you prefer to maintain in your Auto Scaling group.

Type: Integer

Required: No

EndTime

The time that the action is scheduled to end. This value can be up to one month in the future.

Type: DateTime

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

Recurrence

The regular schedule that an action occurs.

Type: String

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

Required: No

ScheduledActionARN

The Amazon Resource Name (ARN) of this scheduled action.

Type: String

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

Required: No

ScheduledActionName

The name of this scheduled action.

Type: String

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

Required: No

StartTime

The time that the action is scheduled to begin. This value 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: DateTime

Required: No

Time

`Time` is deprecated.

The time that the action is scheduled to begin. `Time` is an alias for `StartTime`.

Type: DateTime

Required: No

SuspendedProcess

Description

An Auto Scaling process that has been suspended. For more information, see [ProcessType \(p. 109\)](#).

Contents

ProcessName

The name of the suspended process.

Type: String

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

Required: No

SuspensionReason

The reason that the process was suspended.

Type: String

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

Required: No

Tag

Description

The tag applied to an Auto Scaling group.

Contents

Key

The key of the tag.

Type: String

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

Required: Yes

PropagateAtLaunch

Specifies whether the new tag will be applied to instances launched after the tag is created. The same behavior applies to updates: If you change a tag, the changed tag will be applied to all instances launched after you made the change.

Type: Boolean

Required: No

ResourceId

The name of the Auto Scaling group.

Type: String

Required: No

ResourceType

The kind of resource to which the tag is applied. Currently, Auto Scaling supports the `auto-scaling-group` resource type.

Type: String

Required: No

Value

The value of the tag.

Type: String

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

Required: No

TagDescription

Description

The tag applied to an Auto Scaling group.

Contents

Key

The key of the tag.

Type: String

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

Required: No

PropagateAtLaunch

Specifies whether the new tag will be applied to instances launched after the tag is created. The same behavior applies to updates: If you change a tag, the changed tag will be applied to all instances launched after you made the change.

Type: Boolean

Required: No

ResourceId

The name of the Auto Scaling group.

Type: String

Required: No

ResourceType

The kind of resource to which the tag is applied. Currently, Auto Scaling supports the `auto-scaling-group` resource type.

Type: String

Required: No

Value

The value of the tag.

Type: String

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

Required: No

TerminateInstanceInAutoScalingGroupResult

Description

The output for the [TerminateInstanceInAutoScalingGroup](#) (p. 79) action.

Contents

Activity

A scaling Activity.

Type: [Activity](#) (p. 85)

Required: No

Common Parameters

This section lists the request parameters that all actions use. Any action-specific parameters are listed in the topic for the action.

Action

The action to be performed.

Default: None

Type: string

Required: Yes

AuthParams

The parameters that are required to authenticate a Conditional request. Contains:

- AWSAccessKeyID
- SignatureVersion
- Timestamp
- Signature

Default: None

Required: Conditional

AWSAccessKeyID

The access key ID that corresponds to the secret access key that you used to sign the request.

Default: None

Type: string

Required: Yes

Expires

The date and time when the request signature expires, expressed in the format YYYY-MM-DDThh:mm:ssZ, as specified in the ISO 8601 standard.

Condition: Requests must include either *Timestamp* or *Expires*, but not both.

Default: None

Type: string

Required: Conditional

SecurityToken

The temporary security token that was obtained through a call to AWS Security Token Service. For a list of services that support AWS Security Token Service, go to [Using Temporary Security Credentials to Access AWS](#) in **Using Temporary Security Credentials**.

Default: None

Type: string

Required: No

Signature

The digital signature that you created for the request. For information about generating a signature, go to the service's developer documentation.

Default: None

Type: string

Required: Yes

SignatureMethod

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

Default: None

Type: string

Valid Values: HmacSHA256 | HmacSHA1

Required: Yes

SignatureVersion

The signature version you use to sign the request. Set this to the value that is recommended for your service.

Default: None

Type: string

Required: Yes

Timestamp

The date and time when the request was signed, expressed in the format YYYY-MM-DDThh:mm:ssZ, as specified in the ISO 8601 standard.

Condition: Requests must include either *Timestamp* or *Expires*, but not both.

Default: None

Type: string

Required: Conditional

Version

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

Default: None

Type: string

Required: Yes

Common Parameters for Signature V4 Signing

The following table lists the parameters that all actions use for signing Signature Version 4 requests. Any action-specific parameters are listed in the topic for that action. To view sample requests, see [Examples of Signed Signature Version 4 Requests](#) or [Signature Version 4 Test Suite](#) 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 (YYYYMM-DD'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. For a list of services that support AWS Security Token Service, go to [Using Temporary Security Credentials to Access AWS](#) in *Using Temporary Security Credentials*.

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 common errors that all actions return. Any action-specific errors are listed in the topic for the action.

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

InvalidQueryStringParameter

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

Throttling

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