Auto Scaling

Product Documentation

Contents

Auto Scaling 4
- Product Introduction 4
  - Auto Scaling Overview 4
  - Strengths 8
  - Use Cases 9
  - Usage Recommendations 10
- Quick Start 11
  - Creating a Scaling Plan in 5 Minutes 11
  - Creating Launch Configuration 18
  - Creating Scaling Group 21
  - Creating Scaling Policy 22
- Operation Guide 25
  - Launch Configuration 25
    - Launch Configuration Overview 25
    - Creating a Launch Configuration 26
    - Viewing the List of Launch Configurations 28
    - Changing the Launch Configuration 29
    - Detecting Exceptions 30
    - Multi-Model Configuration 31
  - Scaling Groups 32
    - Scaling Group Overview 32
    - Creating Scaling Groups 33
    - Viewing the Scaling Group List 34
    - Modifying a Scaling Group 35
    - Modifying the CVM Bound to the Scaling Group 36
    - Combining CLBs with Scaling Groups 37
    - Deleting Scaling Groups 38
  - Scaling 39
    - Managing Scheduled Tasks 39
    - Managing Alarm Triggering Policies 41
    - Instance Health Examination 43
    - Manual Scale-out 44
    - Scaling-in 47
    - Viewing Scaling Activities 48
    - Suspending and Resuming Scaling 49
    - Exempting a Specified Instance From Scaling-in 51
    - Scaling Activity Cancellation 52
    - Scaling Activity Failed 53
    - Cooldown Period Description 54
  - Monitoring 56
    - Monitoring Alarm Metrics 56
    - Running State Examination 58
  - Create a notification 59
    - Creating a Scaling Activity Notification 59
    - Specifying Notification Recipients 60
- Troubleshooting 61
  - Scaling Group Failure 61
  - Scaling Activity Execution Failed 63
- Best Practice 66
  - Setting a Fixed Outbound Access IP for Scaling Groups 66
  - Setting Services to Start at Boot 70
- Glossary 72
- API Documentation 73
  - Auto Scaling(as) 73
    - version(2018-04-19) 73
      - API overview 73
      - CallingMethod 76
        - signaturev1 76
        - signaturev3 84
        - request 93
        - result 94
        - PublicParameter 97
      - Alarm Trigger Policy APIs 99
        - CreateScalingPolicy 99
        - DeleteScalingPolicy 100
        - DescribeScalingPolicies 102
        - ExecuteScalingPolicy 103
        - ModifyScalingPolicy 104
      - Launch Configuration APIs 105
        - CreateLaunchConfiguration 105
        - DeleteLaunchConfiguration 106
        - DescribeLaunchConfigurations 108
        - ModifyLaunchConfigurationAttributes 109
        - UpgradeLaunchConfiguration 110
      - Lifecycle Hook APIs 111
        - CompleteLifecycleAction 111
        - CreateLifecycleHook 113
        - DeleteLifecycleHook 115
        - DescribeLifecycleHooks 117
        - UpgradeLifecycleHook 118
      - Notification APIs 119
        - CreateNotificationConfiguration 119
        - DeleteNotificationConfiguration 120
        - DescribeNotificationConfigurations 122
        - ModifyNotificationConfiguration 123
      - Other APIs 124
        - DescribeAccountLimits 124
      - Scaling Group APIs 126
        - AttachInstances 126

- CreateAutoScalingGroup 128
- DeleteAutoScalingGroup 130
- DescribeAutoScalingActivities 132
- DescribeAutoScalingGroupLastActivities 133
- DescribeAutoScalingGroups 134
- DescribeAutoScalingInstances 135
- DetachInstances 136
- DisableAutoScalingGroup 138
- EnableAutoScalingGroup 140
- ModifyAutoScalingGroup 142
- ModifyDesiredCapacity 144
- ModifyLoadBalancers 146
- RemoveSelectedInstance 147
- SetInstancesProtection 149

• Scheduled Task APIs 151
- CreateScheduledAction 151
- DeleteScheduledAction 152
- DescribeScheduledActions 154
- ModifyScheduledAction 155

• Data structure 156
• ErrorCode 162

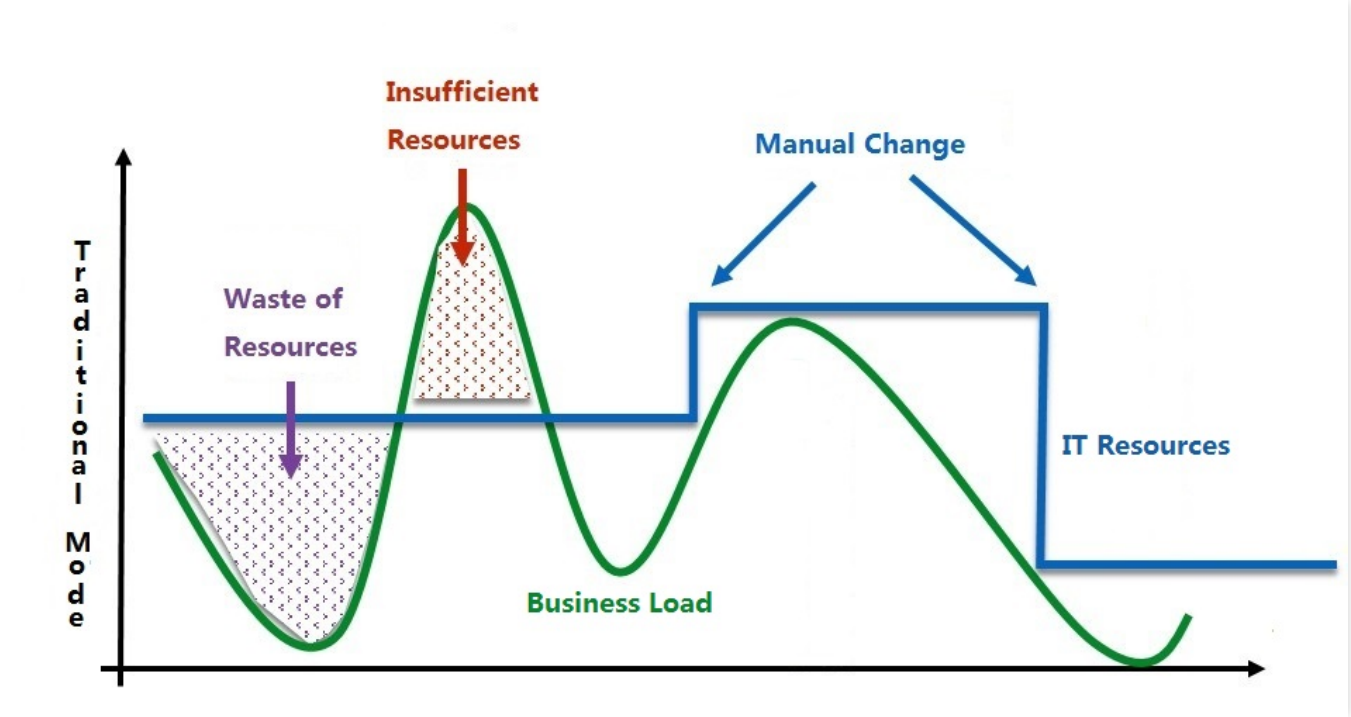
• Purchase Guide 169

Product Introduction
Auto Scaling Overview

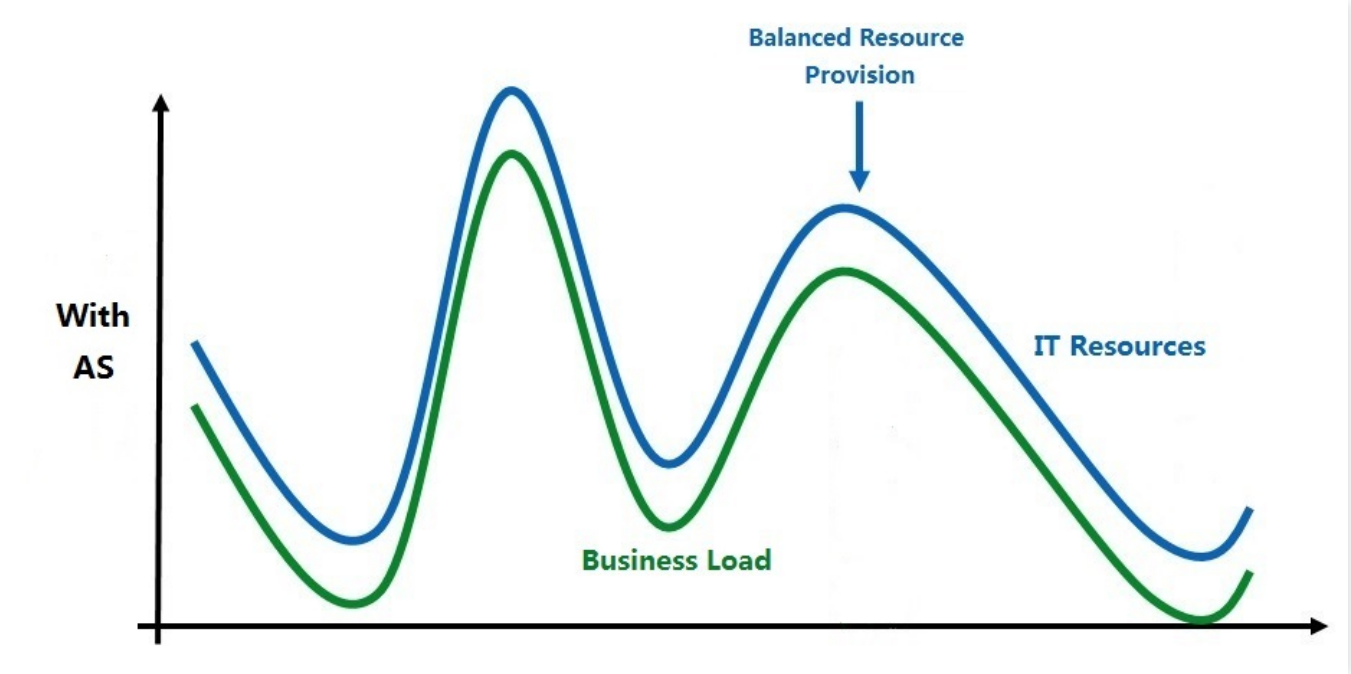
What is AS?

AS (Auto Scaling) can adjust CVM computing resources automatically based on your business needs and policies, ensuring you have the right amount of CVM instances to handle your application's load. Intelligent scaling is an important part of your Web services' cost control and resource management. When your web applications receive more traffic, you'll add more servers to handle the extra load. Similarly, you'll terminate underutilized servers when traffic to your web applications decreases.
Using AS for capacity adjustment only requires setting expansion and contraction conditions in advance. AS automatically increases the number of servers used to maintain performance when conditions are met; when demand decreases, AS reduces the number of servers based on your contraction conditions, helping you minimize costs to the greatest extent.
As shown in the comparison below, by using AS, your cluster can always maintain just the right amount of resources and stay in a healthy state. You will say goodbye to many frustrations encountered with traditional modes:

- A sudden increase in business or a CC attack causes a lack of machines, making your service unresponsive.
- The estimation of resources is based on peak traffic. As the actual traffic rarely reaches the peak, wasted resources are caused.
- Manual monitoring and frequent handling of capacity alerts lead to multiple manual changes.
  Cluster Maintenance in the Traditional Mode:



The Effect After Adopting AS:



How AS Works?

In common Web application services, your cluster usually runs multiple copies of an application to meet client traffic, for example, the frontend server cluster at the access layer, the application server cluster at the logical layer, and the backend cache server cluster. Every instance can process client requests.

These instances are similar or identical, and the quantity is usually adjustable. You can add them to one scaling group for joint management:

- You can specify the minimum number of instances in each scaling group, and AS ensures the number of instances in the group never falls below this number;
- You can specify the maximum number of instances in each scaling group, and AS ensures the number of instances in the group never exceeds this number;
- You can specify a scaling policy, and AS launches or terminates instances when the application's demand increases or decreases. There are two types of scaling policies:
  - Alarm-Triggered Scaling Policy: Dynamically scale according to specified conditions (for example, scale out when the CPU usage of the machines in the scaling group exceeds 60%.).
  - Scheduled Scaling Policy: Scale in or out at specified times (for example, scale out every night at 21:00.).
- After the policy is set, you can also set the scaling activity notifications. AS informs you via Email, SMS, and message center when scaling activities occur. You don't need to monitor your business request volume constantly. Just pay attention to the notifications from AS.
- You can also specify the number of machines required at any time or add existing machines to the scaling group for joint management.

Basic Concepts of AS

The AS product has the following basic concepts:

- Scaling Group:
  A scaling group contains a collection of CVM instances that follow the same rules and have a shared purpose. Scaling groups define attributes such as the maximum and minimum numbers of CVM instances in the group and their associated CLB instances.
- Launch Configuration:
  A launch configuration is the template for the automatic creation of CVMs. It contains the image ID, CVM instance type, system disk/data disk types and capacities, key pair, security group, etc.
  When a scaling group is created, it must be specified, and the properties of the launch configuration cannot be edited.
- Scaling Policy:
  The conditions for executing scaling actions. The trigger can be a specific time or an alarm from the cloud monitor, with actions including removing or adding CVMs.
  There are two types of scaling policies:
- Scheduled Scaling Policy:
  The automatic addition and removal of CVM instances at a specified time point can be performed on a periodic basis.
- Alarm-Triggered Scaling Policy:
  The automatic addition and removal of CVM instances based on the cloud monitor metrics such as CPU,

memory, and network traffic.

- Cooldown Period:
  Cooldown period refers to a locking duration after one scaling activity (addition or removal of CVM instances) is executed within the same scaling group. During this duration, no scaling activities are performed by the scaling group. The cooldown period can be specified from 0 to 999,999 (seconds).

Strengths

Strengths	Using Auto-scaling	Not Using Auto-scaling
Automation	Auto-scaling Instances with No Manual Intervention Required Auto-scaling dynamically and in real-time automatically creates and releases CVM instances based on business load, helping you maintain the optimal number of instances for your business needs completely without manual intervention, thereby relieving you from the burden of manual deployment. For example, you can set a scaling policy that adds new CVM instances to your scaling group when CPU usage is high, with newly added CVM instances billed by the second. Similarly, you can set a policy to remove instances from the scaling group when CPU usage is low; if you can predict your load changes, you can set scheduled tasks to plan your scaling activities. New instances can also be directly associated with the existing CLB to take on the distribution of traffic, improving service availability; you can also send alerts to administrators to help you promptly pay attention to exceptions.	Cumbersome Manual Operations Manually creating and destroying resources requires manual configuration of CLB. Manual operations are prone to errors and affect the business.
Cost Savings	Moderate Scaling Instances to Save Costs Auto-scaling helps you maintain the optimal number of instances for your business needs. When business demand increases, it automatically adds the appropriate amount of CVM instances seamlessly; when business demand decreases, it automatically reduces unnecessary CVM instances, improving device usage and saving on deployment and instance costs.	Waste Due to Resource Idle Reserve excess CVMs to prevent shortages that could affect the business.
Fault Tolerance	System Auto-detection and Timely Fault Tolerance Auto-scaling monitors the health status of instances automatically. Upon detecting any exception, it immediately replicates a healthy instance to replace the abnormal one, ensuring your application receives the expected computing capacity and safeguarding your business operations.	Inability to Tolerate Faults Timely Exceptional instances are often handled once a business interruption is detected, which affects business availability.

Use Cases

Scaling Deployment in Advance

If you know the time for scaling, you can set an Auto Scaling scheduling policy. At the corresponding time, the system automatically adds or removes the CVM instance without human effort.

Cost-Effective Handling of Business Surges

When the traffic arises, you need to prepare servers in advance to prevent server overload caused by the surge in CPU usage. When the traffic decreases, you need to reduce servers according to actual load. With the Auto Scaling policy configured in advance, the system automatically determines whether to trigger CVM parallel scaling according to the set service monitoring metrics. When the monitoring metrics reach the threshold, CVM instances are automatically added or removed, and the cloud load balancer configuration is automatically implemented. This reduces the cost and does not require customers to prepare for manual scale-out constantly.

Automatic Replacement of Unhealthy CVMs

To prevent your business from being affected by unhealthy CVMs, you need to constantly monitor the operation status of CVMs in the system and be ready to handle the problem at any time. With Auto Scaling, the system regularly performs health checks on CVMs. When the system detects an exceptional CVM instance, it automatically performs parallel scales for a new instance to replace the exceptional instance. The operation record is kept for the user to view.

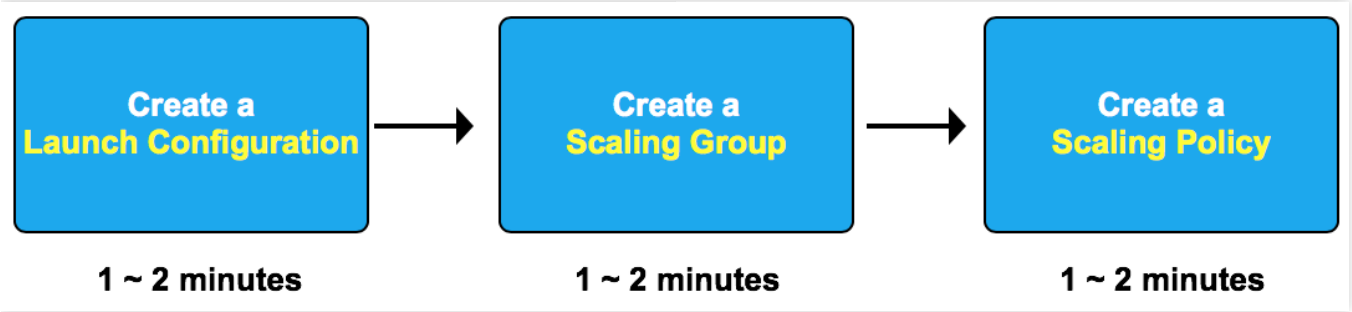
Usage Recommendations

The following table describes the usage recommendations:

Limit Type	Description
For a user in a specific region	- A maximum of 20 startup configurations can be created. - A maximum of 20 scaling groups can be created. - The maximum number of CVM instances that can be created is subject to the user's CVM quota.
For a scaling group	- There can be only one startup configuration. - A maximum of 200 CVM instances can be scaled up. - A maximum of 100 scaling policies and 10 scheduled tasks can be created. - A maximum of 5 notifications can be created. - A maximum of 10 lifecycle hooks can be created.
Others	- The number of instances in a scaling group cannot exceed the upper limit of IP addresses that the VPC subnet can provide. - AS cannot automatically scale the CPU, memory, and bandwidth of CVM instances. - Auto scaling and startup configuration are region-level concepts. They can be used to start or terminate CVM instances only when instances are in the same region. - Only TDMQ Pulsar is supported as the notification method for lifecycle hooks in a scaling group.

Quick Start
Creating a Scaling Plan in 5 Minutes

To create a complete AS scheme, you need to complete the following 3 steps:



> Note:
> The operations in the following example use the console mode.

Creating launch configuration

A launch configuration defines the configuration information for AS's CVM instances, including the CVM's image, storage, network, security groups, log-in methods, and other configuration information.

> Note:
> Creating a launch configuration service is completely free. Feel free to create.

1. Log in to the AS console.
2. In the left sidebar, click Launch Configuration to enter the Launch Configuration management page.

Selecting a region

1. Select the region for the launch configuration.
2. Click Create, and fill in the basic information of the launch configuration on the pop-up page.

Selecting Region and Model

In the 1. Select Region and Model page, fill in the Launch Configuration Name, select Region, and select Instance.

- Name the configuration.
- Select the same model of the instance that you want to bind to the scaling group.

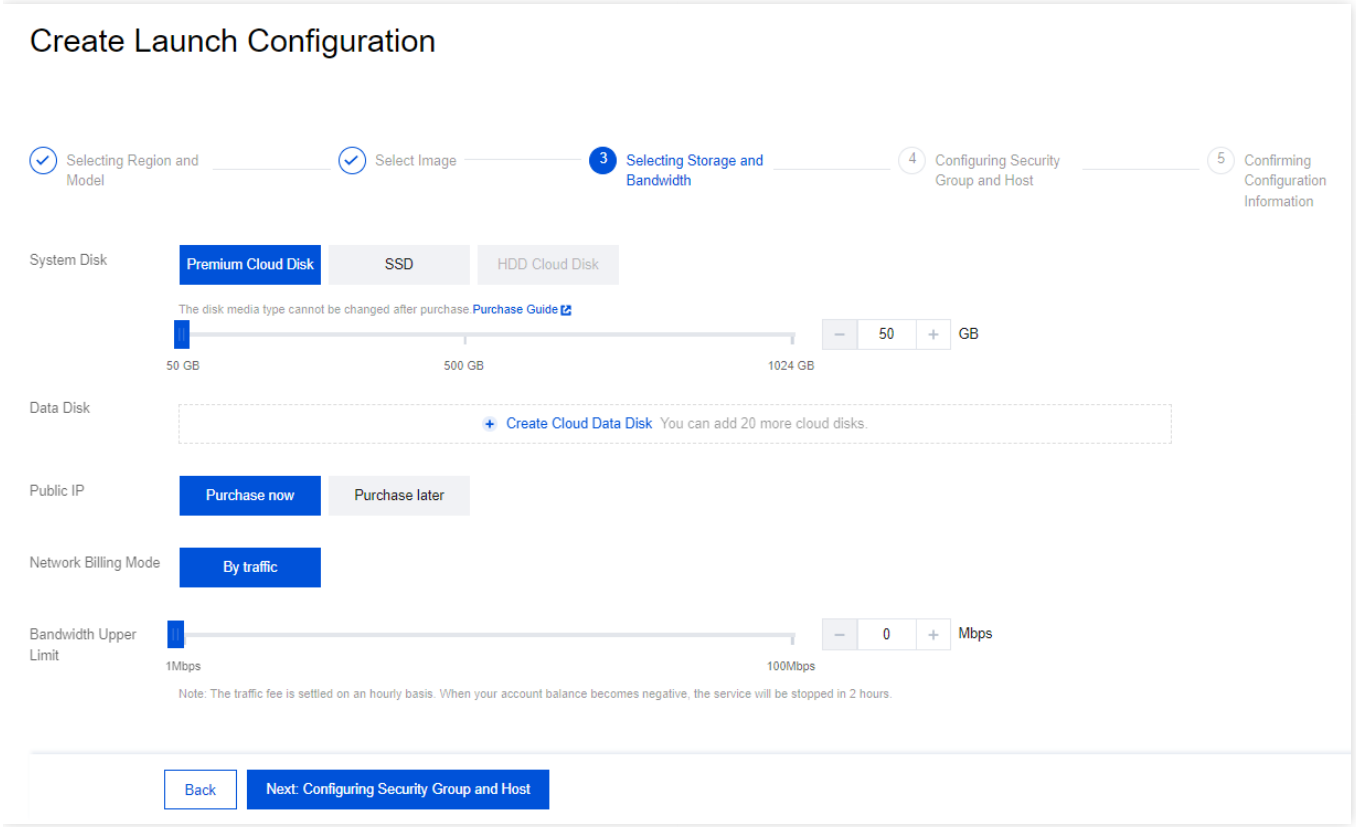
Selecting an image

A launch configuration can be created using a public image or custom image.
It is recommended to use a custom image with a pre-deployed environment for the following reasons:

- If a public image is selected, the instances scaled out have a clean OS, and you still need to deploy the application environment manually.
- If a custom image is selected, create an image from a CVM instance with a pre-deployed environment and then use this image to batch create CVM instances. The newly created CVM instances have the same software environment as that of the previous CVM instance, thus achieving the purpose of batch deployment.

Selecting Storage and Bandwidth

In the 3. Select Storage and Bandwidth page, set the disk and network.



> Note:
> If a CBS is chosen for the system disk, then the data disk can be a data disk snapshot.

For users with a large amount of data, data disks are often used to store data. When a data disk A creates a snapshot, users can use this snapshot file to quickly clone multiple disks, achieving rapid deployment of servers.

When AS automatically adds a new cloud service instance, if a data disk snapshot is specified in the launch configuration, CBS can support auto-mounting a data disk containing set data after the cloud service instance starts, meeting the need for automatic data copying.

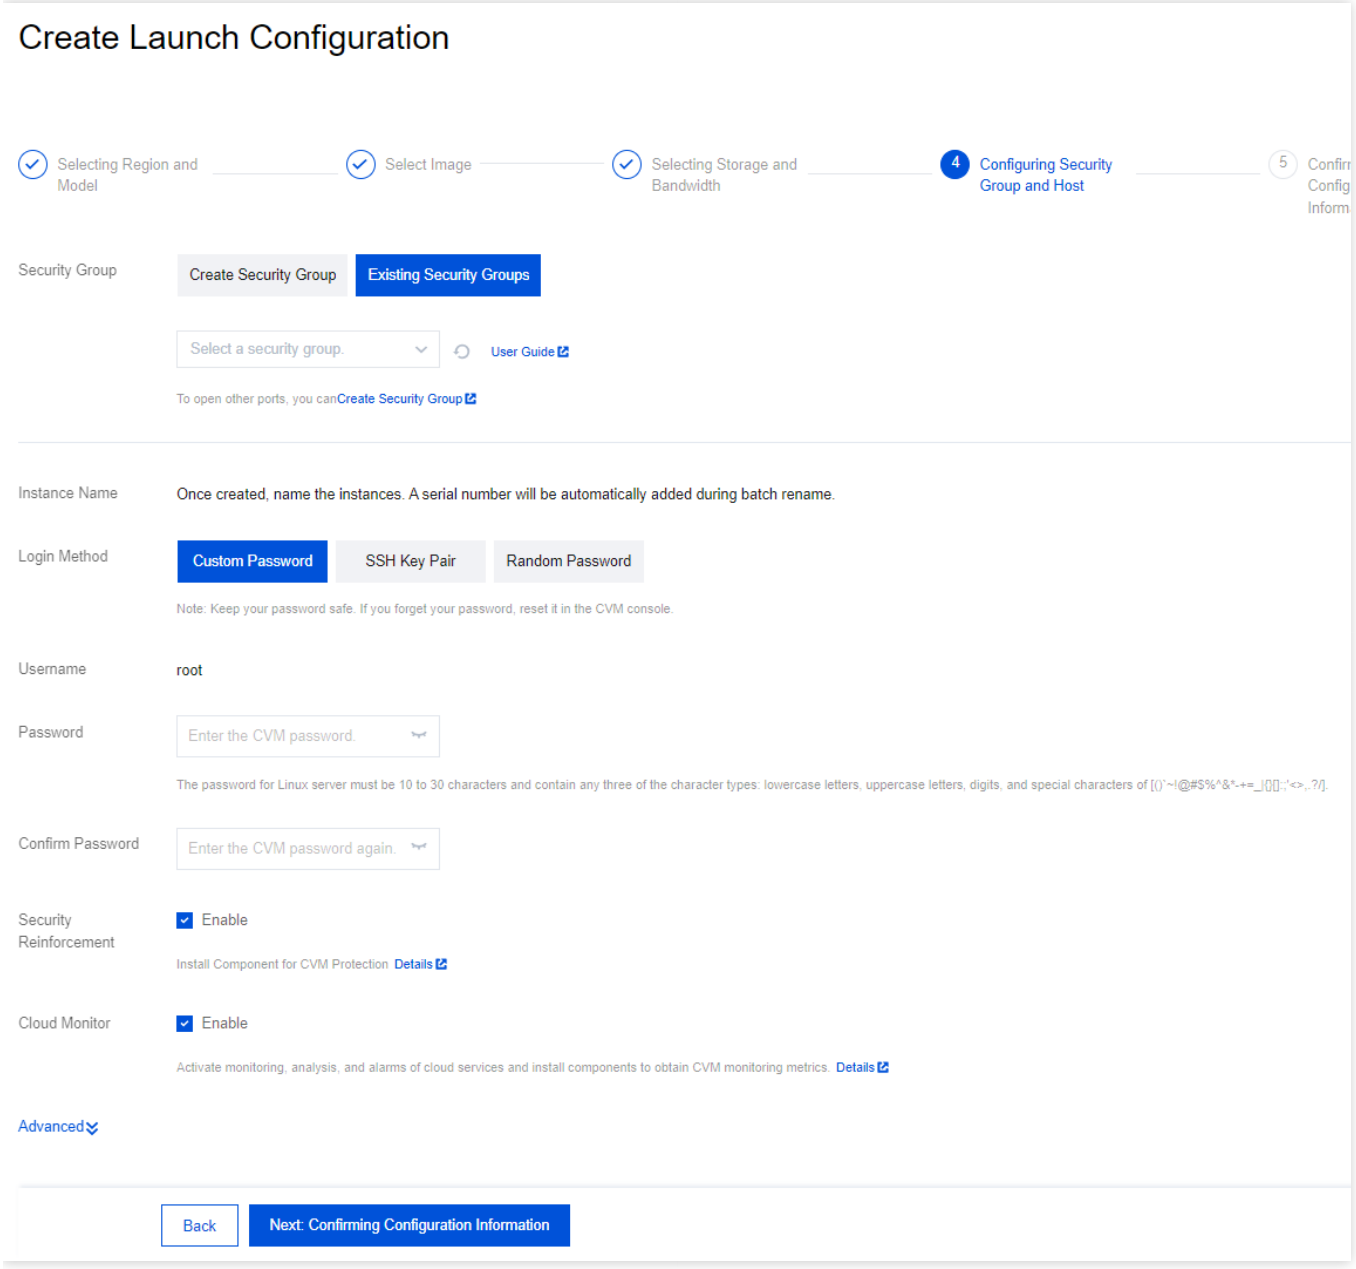
If a data disk snapshot is specified in the launch configuration, ensure the data disk can be automatically mounted correctly so that the scaling group can auto-scale successfully. You need to perform some operations on the original instance that made the data disk snapshot before setting AS to support automatic data disk mounting when a new cloud service instance is started.

> Note:
> The AS service is free. The added servers, disks, and networks will be charged based on the volume of CVM instances, disks, and networks. Prices will be shown based on your settings.

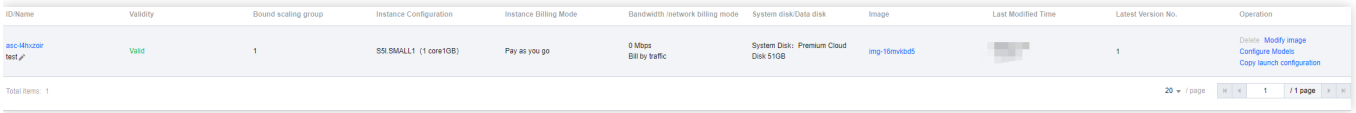
Setting information

1. On the 4. Set Security Group and CVM page, select the log-in method and security group. CVM instances added by the AS service enjoy Cloud Security and Cloud Monitor services for free by default.

   > Note:
   >
   > In the advanced settings, if the instance type of the extension you created has been set by the administrator to not support placement groups, the placement group feature will be unavailable. If you need to make a change, please contact the administrator.



2. After the configuration is completed, this entry will be displayed in the launch configuration list on the page.



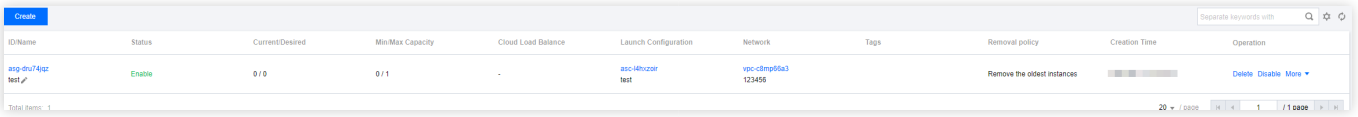
Creating scaling group

A scaling group contains a collection of CVM instances that follow the same rules and have a shared purpose.

1. Log in to the AS console.
2. In the left sidebar, click Scaling Group to enter the scaling group management page.

Creating a scaling group

1. Click Create and fill in the basic information of the scaling group on the pop-up page. * is required.
   - The current number of CVM instances in the scaling group will be maintained between the minimum and the maximum scaling group capacities.
     - The starting instance count defines the initial number of CVM instances in the scaling group.
     - If the current number of CVM instances is less than the minimum scaling number, the AS service will automatically add instances to match the minimum scaling number.
     - If the current number of CVM instances exceeds the maximum scaling number, the AS service will automatically reduce instances to match the maximum scaling number.
   - Select an existing launch configuration or create a new one.
   - Select the network and removal policies.
   - (Optional) Select an existing CLB policy to associate or create a new CLB.
2. After the configuration is completed, this entry will be displayed in the scaling group list on the page.



Adding instances (optional)

Add the instances to be bound to the instance list. After the configuration is completed, this entry will be displayed in the launch configuration list on the page, as shown below:

> Note:
> If you are unable to add or remove instances here, please check the maximum and minimum scaling numbers you have set.

Creating a scaling policy;

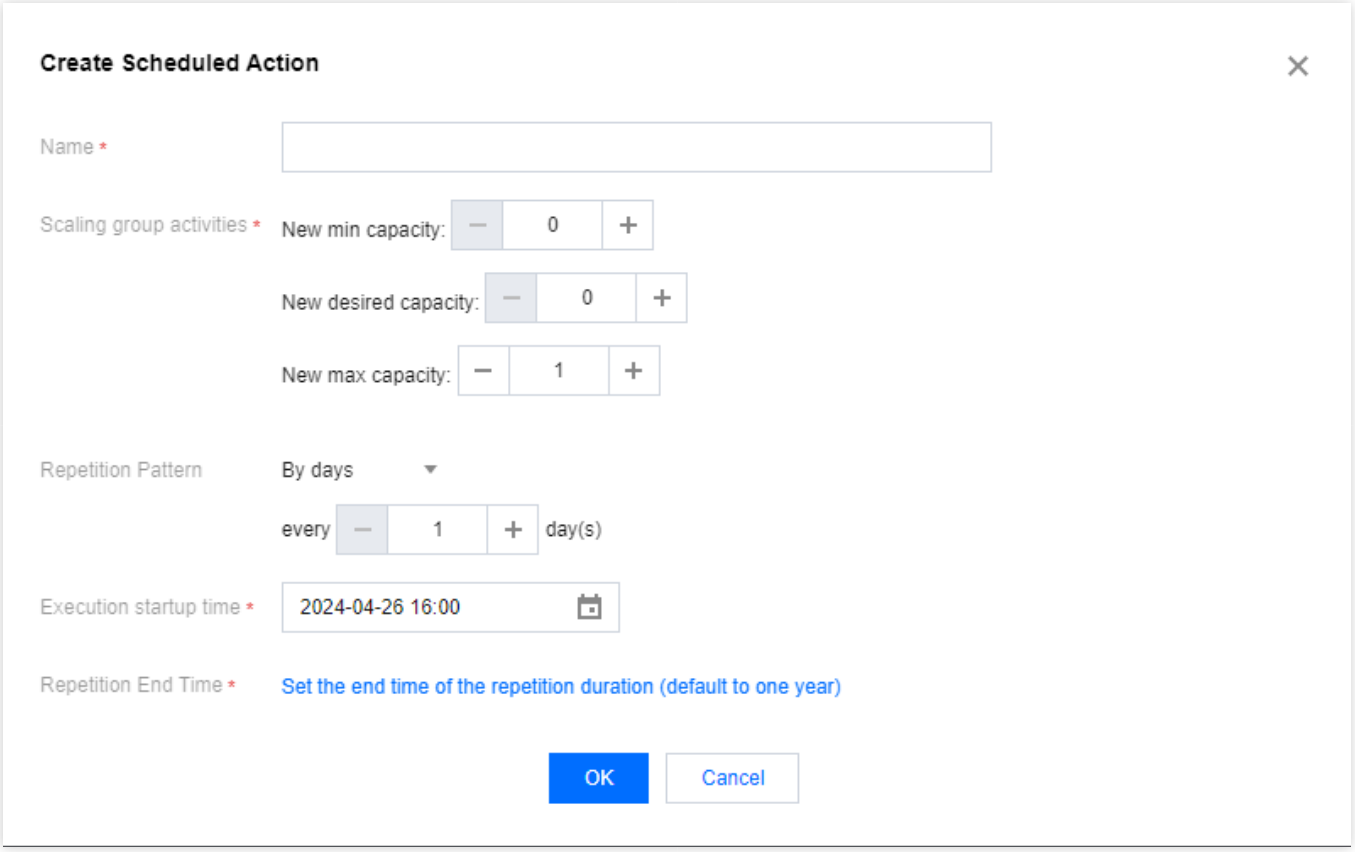
AS group adjusts the number of cloud service machines according to the scaling policies:

- Create scheduled tasks to perform scheduled scaling activities, and you can also set whether to perform them periodically.
- Create alarm-triggered scaling policies to perform scaling based on the cloud monitor metrics (such as CPU, memory usage rate, etc.).

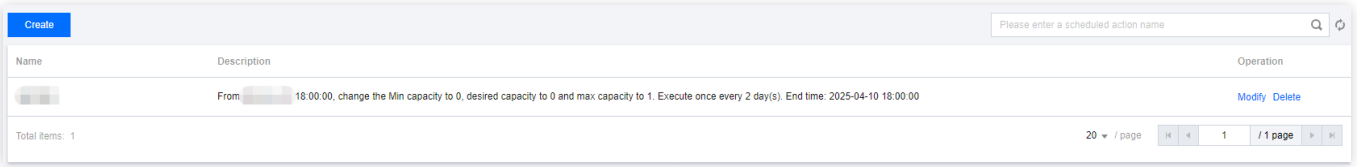
Creating scheduled task

If you can predict your load changes, you can set up scheduled tasks to plan your device scale-out activities. This feature allows you to automatically increase or decrease CVM instances regularly and periodically to flexibly respond to business load changes, enhance device usage, and save costs on deployment and instances.

1. On the Scaling Group page, click Scaling Group to enter the scaling group management page.
2. Select the Scheduled Action tab, then click Create.
3. On the Create page, specify the scheduled task name, execution time, and activities to be executed.



4. After the settings are completed, the scheduled task will be displayed in the list on the page. See the example below:

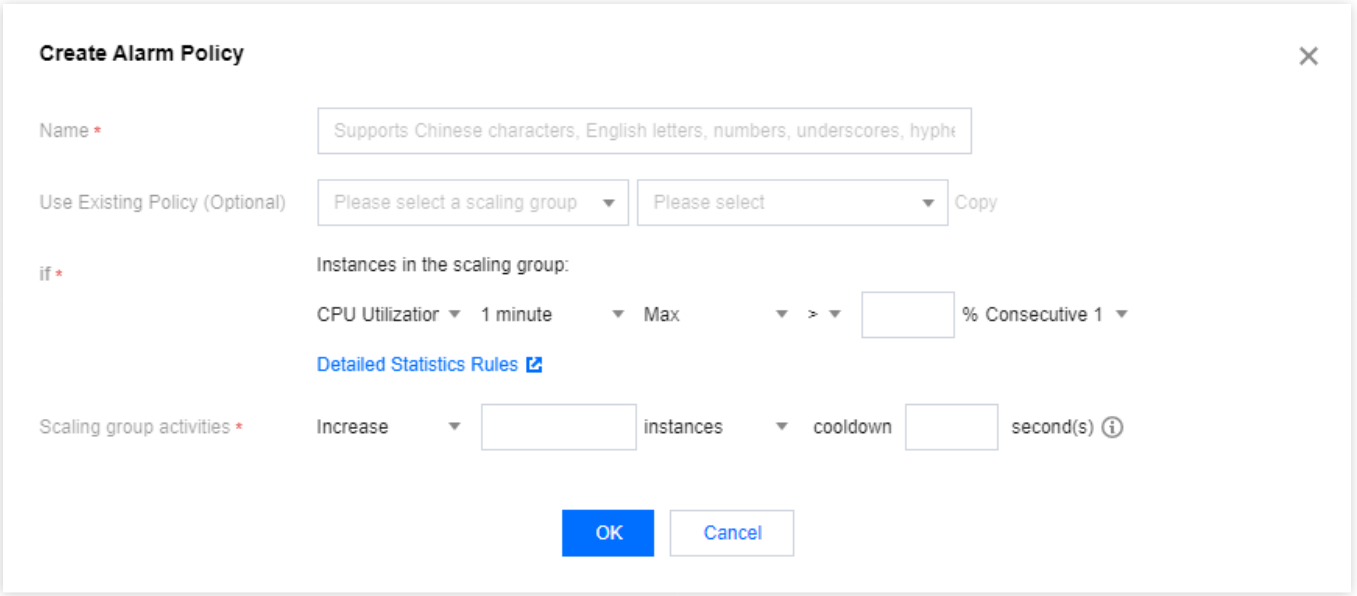


Creating an alarm policy

If you want to adjust your business deployment based on CVM metrics, you can define your own alarm-triggered scaling policy. When business loads reach the threshold, this policy automatically increases or decreases the number of CVM instances. This allows flexible adaptation to changes in business load, improving device usage and saving on deployment and instance costs.

> Note:
>
> - When a scaling group is created, a ping unreachable alarm-triggered policy is created by default, which is used to replace unhealthy CVMs.
> - Before using the Alarm-Triggered Scaling Policy, you need to install the latest version of the cloud monitor Agent in the CVM image.

1. On the Scaling Group page, click Scaling Group ID to enter the scaling group management page.
2. Select the Alarm Triggered Scaling Policy tab, and then click Create.
3. On the Create page, set the Alarm-Triggered Scaling Policy. The policy is based on the cloud monitor performance metrics (such as CPU, memory, bandwidth, etc.) and automatically increases or decreases a specified number or percentage of CVM instances for the scaling group.
   You can also copy the policy (optional), directly copying an existing scaling group's policy to the current group.



4. After the settings are completed, this alarm-triggered scaling policy will be displayed in the list on the page.

Creating Launch Configuration

A launch configuration defines the configuration information for AS's CVM instances, including the CVM's image, storage, network, security groups, log-in methods, and other configuration information.

Prerequisites

Logged in to the AS console successfully.

Directions

Selecting Region

Select the region of the cloud server you want to bind to the scaling group.

Selecting a Model

1. In the left sidebar, click Launch Configuration to enter the Launch Configuration Management page.
2. Click Create, and the Create Launch Configuration page pops up.
3. In the 1. Select the region and model tab, fill in the launch configuration name, select the availability region, and choose the model.
    - Configuration Name: Define it yourself, and fill in the configuration name.
    - Instance: Choose the same model of cloud server as that of the scaling group you want to bind to.

Selecting Image

A launch configuration can be created using a public image or custom image. It is recommended to use a custom image with a pre-deployed environment for the following reasons:

- If a public image is selected, the instances scaled out have a clean OS, and you still need to deploy the application environment manually.
- If a custom image is selected, by creating an image from a CVM instance with a pre-deployed environment and then using this image to batch create CVM instances, the newly created CVM instances have the same software environment as the previous CVM instance, thus achieving the purpose of batch deployment.
- Image: The choices of images include public images and custom images, with custom images recommended.
    - If a public image is selected, it should have the same operating system as the cloud server the scaling group wants to bind to.

    - If a custom image is selected, it needs to be made by the user through the cloud server's image.
- Operating System: You can select from operating systems such as CentOS, CoreOS, Debian, SUSE, and Windows.
- System Architecture: You can choose between 64-bit and 32-bit.
- Image Version: You can select the image version.

  > Note:
  >
  > - Differences between Public Images and Custom Images: If you select a public image, the scaled-out cloud service instances are not activated and cannot be used directly; manual deployment of the application environment is required. If you select a custom image by creating an image from a CVM instance with an already deployed environment, this image will be used as the operating system when batch-creating CVM instances. After the CVM instance is created, it will have the same software environment as the previous CVM instance, achieving the purpose of batch deployment of the software environment.

Selecting Storage and Bandwidth

In the 3. Select Storage and Bandwidth tab, set up the disk and network.

> Note:
> If a CBS is chosen for the system disk, then the data disk can be a data disk snapshot.

For users with a large amount of data, data disks are often used to store data. When a data disk snapshot is created, users can use this snapshot file to quickly clone multiple disks, achieving rapid deployment of servers.
When AS automatically adds a new cloud service instance, if a data disk snapshot is specified in the launch configuration, CBS can support auto-mounting a data disk containing set data after the cloud service instance starts, meeting the need for automatic data copying.
If a data disk snapshot is specified in the launch configuration, ensure the data disk can be automatically mounted correctly so that the scaling group can auto-scale successfully. You need to perform some operations on the original instance that made the data disk snapshot before setting AS to support automatic data disk mounting when a new cloud service instance is started.

Setting Information

1. On the 4. Set Security Group and CVM page, select the log-in method and security group.
   CVM instances added through the AS service are, by default, free to enjoy Cloud Security and Cloud Monitor services.
2. Click Next: Confirm Configuration Inoformation to check and confirm the information that has been configured.
3. Click Create Launch Configuration.

After the configuration is completed, this entry will be displayed in the launch configuration list on the page.

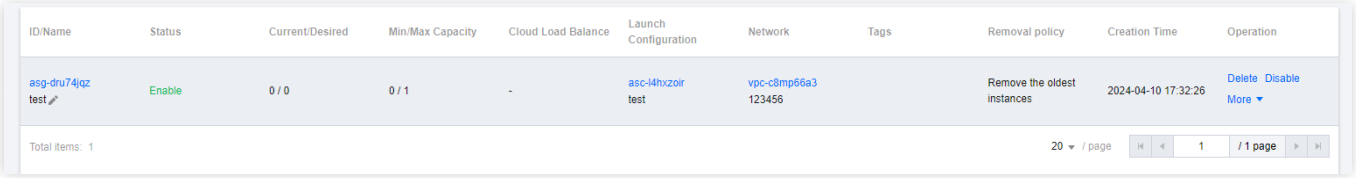
Creating Scaling Group

A scaling group contains a collection of CVM instances that follow the same rules and have a shared purpose.

Creating Scaling Groups

1. Log in to the AS console.
2. In the left sidebar, click Scaling Group to enter the scaling group management page.

3. Click Create and fill in the basic information of the scaling group on the pop-up page. * is required.
   - The number of CVM instances in the scaling group will be maintained between the minimum and the maximum scaling group capacities.
   - The starting instance count defines the initial number of CVM instances in the scaling group.
     - If the current number of CVM instances is less than the minimum scaling number, the AS service will automatically add instances to match the minimum scaling number.
     - If the current number of CVM instances exceeds the maximum scaling number, the AS service will automatically reduce instances to match the maximum scaling number.
   - Select an existing launch configuration or create a new one.
   - Select the network, availability zone, and removal policy.
   - (Optional) Select an existing CLB policy to associate or create a new CLB.
4. After the configuration is completed, this entry will be displayed in the scaling group list on the page.



Adding Instance (Optional)

1. On the Scaling Group page, click Scaling Group to enter the scaling group management page.
2. Select the Bind with Instance tab and click Add Instances.
3. In the pop-up Add Instances window, add the instances you need to bind to.
4. Click OK to complete the configuration.

   The added instance information will be displayed in the startup configuration list.

> Note:
> If you are unable to add or remove instances here, please check the maximum and minimum scaling numbers you have set.

Creating Scaling Policy

AS group adjusts the number of cloud service machines according to the scaling policies:

- Create scheduled tasks to perform scheduled scaling activities, and you can also set whether to perform them periodically.
- Create alarm-triggered scaling policies to perform scaling based on the cloud monitor metrics (such as CPU, memory usage rate, etc.).

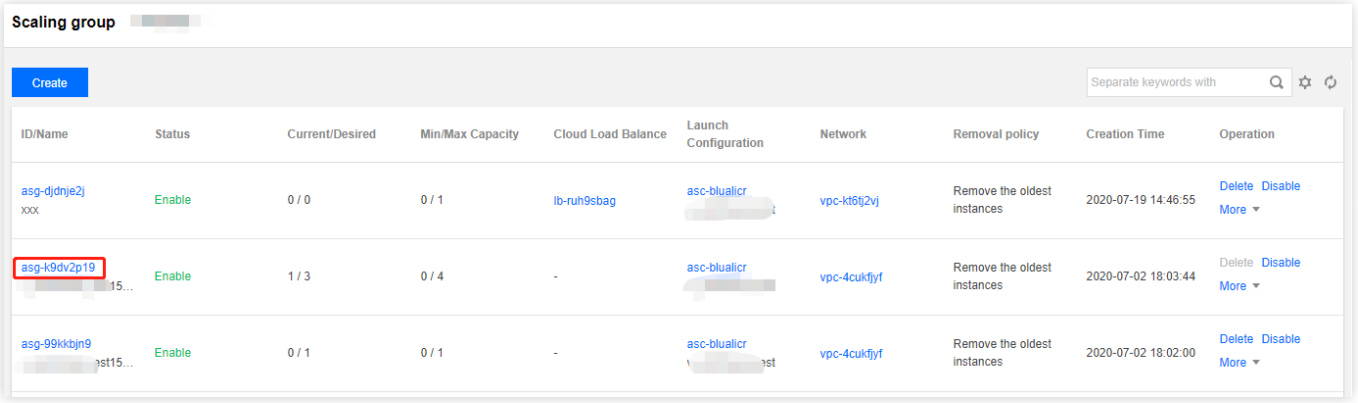
Creating Scheduled Task

Overview

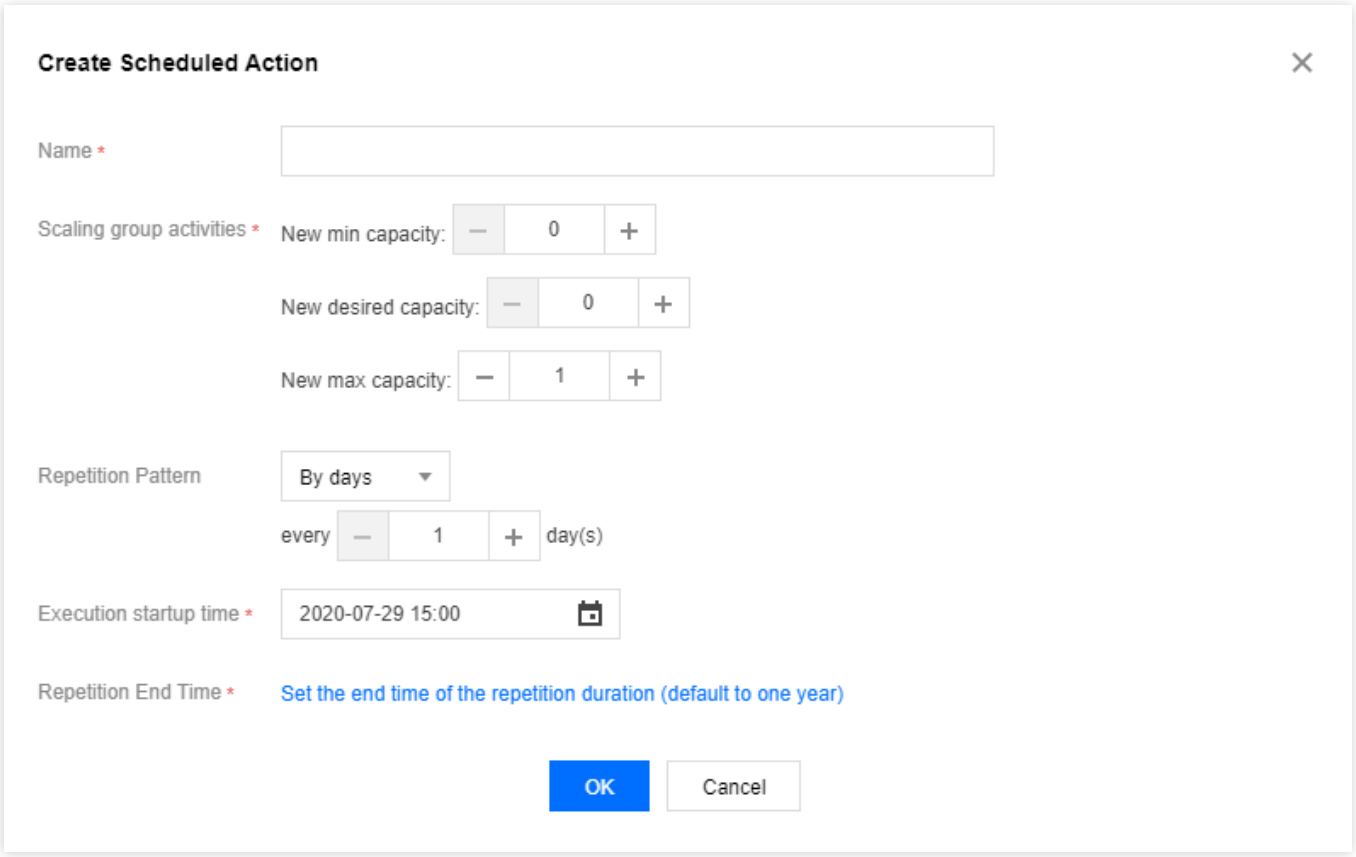
If you can predict your load changes, you can set up scheduled tasks to plan your device scale-out activities. This feature allows you to automatically increase or decrease CVM instances on a schedule or periodically to respond to business load changes flexibly, enhance device usage, and save on deployment and instance costs.

Directions

1. Log in to the AS console.
2. On the Scaling Group page, click Scaling Group.



3. In the scaling group management page, select the Scheduled Action tab, and click Create.
4. In the pop-up Create Scheduled Action window, specify the scheduled task name, execution time, and actions. You can also select Repeat to define a task for periodic scheduled execution.



5. Click OK to complete the setting. The created scheduled task will be displayed in the list on the page.

Creating Alarm-Triggered Scaling Policy

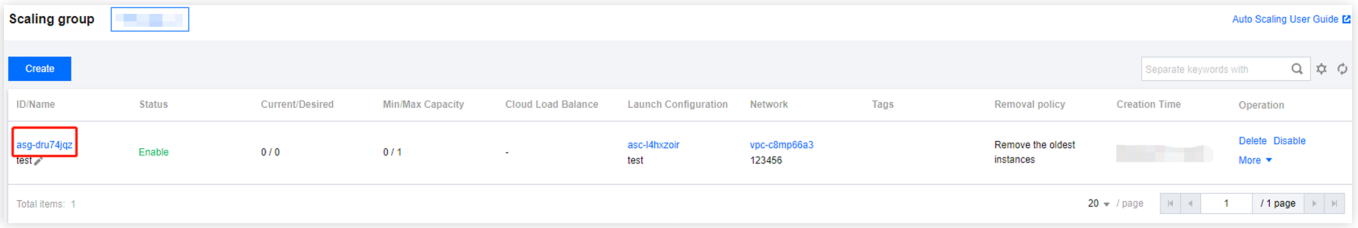
Overview

If you want to adjust your business deployment based on CVM metrics, you can define your own alarm-triggered scaling policy. When business loads reach the threshold, this policy automatically increases or decreases the number of CVM instances. This allows flexible adaptation to changes in business load, improving device usage and saving on deployment and instance costs.

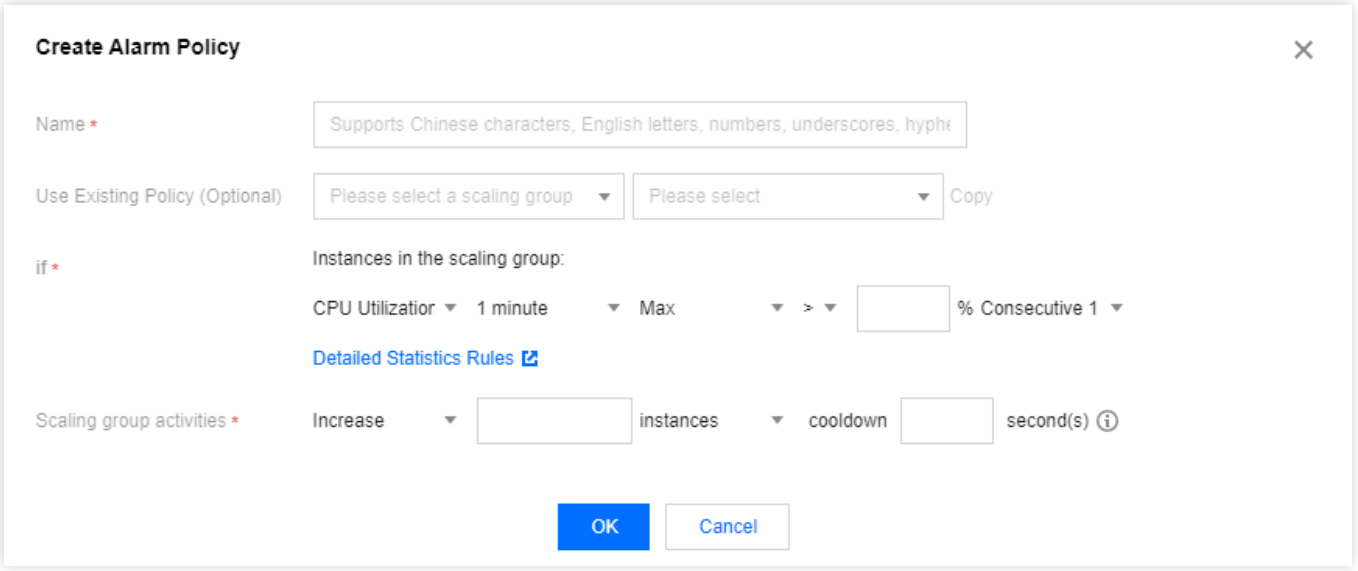
Directions

> Note:
>
> - When a scaling group is created, a ping unreachable alarm-triggered scaling policy is created by default, which is used to replace unhealthy CVMs.
> - Before using the Alarm-Triggered Scaling Policy, you need to install the latest version of the cloud monitor Agent in the CVM image.

1. Log in to the AS console.
2. On the Scaling Group page, click Scaling Group.



3. In the scaling group management page, select the Alarm Trigger Policy tab and click Create.
4. In the pop-up Create Alarm Trigger Policy window, set up the alarm policy.
   The policy is based on the cloud monitor performance metrics (such as CPU, memory, bandwidth, etc.) and automatically increases or decreases a specified number or percentage of CVM instances for the scaling group. You can also copy the policy (optional), directly copying an existing scaling group's policy to the current group.



5. Click OK to complete the setting. The created alarm-triggered scaling policy will be displayed in the list on the page.

Operation Guide
Launch Configuration
Launch Configuration Overview

The launch configuration is a template for automatically creating a CVM, which includes the image ID, CVM instance type, system disk and data disk type and capacity, key pair, security group, etc.
You must specify a launch configuration when you create a scaling group. Once a launch configuration is created, its attributes cannot be edited.

Creating a Launch Configuration

AS needs to know in advance what configuration to use to produce CVM during scaling-out. You need to specify relevant resources in advance, such as images, data disk data, instance configuration, key pairs, security groups, and block storage devices.
It is worth noting that the launch configuration is just a template, and machines are produced according to this template during automatic scaling-out. Creating a launch configuration does not produce a machine, so it is completely free. Feel free to use it.
Log in to the auto scaling console and click Launch Configuration in the sidebar.

Step 1: Select the Region

In the menu at the top of the screen, select a region for the scaling group that meets your needs.

> Note:
> Here, you must select the region where the CVM to which you want to bind the scaling group is located. Both the launch configuration and the scaling group are region attributes. For example, if NCR is selected as the region for the launch configuration, you can only bind the scaling group in NCR, and the automatically added CVM is also a CVM in NCR.

Step 2: Specify Parameters

Click Create, and follow the instructions to create a launch configuration. The steps are the same as when you purchase a CVM.

1. Fill in the Configuration Name, for example, Frontend Server Cluster Configuration A.

2. Select the model, for example, 1 core 1 GB, which means 1 core CPU and 1 GB memory.

3. Choose images: You can choose a clean public image or a custom image with a business already deployed.
   To make the machine available directly after creation, we strongly recommend that you deploy business applications in a custom image. At the same time, the business applications in the image should be set to launch together with the operating system, so that the machines scaled-out by AS can be automated.

4. Select the system disk size and data disk size.
   If you want the data disk to come with data after the machine is activated, you can specify a data disk snapshot so that the machine will come with the data in the snapshot after it is produced.

   > Note:

- As instances in a scaling group are generally stateless, we recommend you put your own data into a custom image for convenience. If the system disk is not large enough, you can apply for a larger one by submitting a ticket.
- If you want to use a data disk to store data, you need to set up automatic mounting of the data disk so that the scale-out does not require manual intervention.

5. Select the bandwidth just like when you purchase a CVM instance.

6. Set the username, password, and security group.

> Description:
>
> In Advanced settings, if the instance type of the extension you created has been set by the Administrator as unsupported for Placement Groups, then Placement Groups will be unavailable. Please contact the Administrator if changes are needed.

7. Click Next: Confirm Configuration Information.
   Create a scaling group based on the launch configuration. The launch configuration determines the instances to be created during scale-out, while the scaling group determines when to scale out.

8. Click Create Launch Configuration.

Viewing the List of Launch Configurations

The launch configuration is a template for automatically creating a CVM, which includes the image ID, CVM instance type, system disk and data disk type and capacity, key pair, and security group.

> Notes:
>
> - You must specify a launch configuration when you create a scaling group.
> - Once a launch configuration is created, its attributes cannot be edited.

Open the auto scaling console and select Launch Configuration in the sidebar to view the list.

- To view the launch configuration details, click the corresponding launch configuration ID.
- To delete a launch configuration, click Delete of the corresponding launch configuration entry.
  > Notes:
  > A launch configuration that has been bound to a scaling group cannot be deleted.

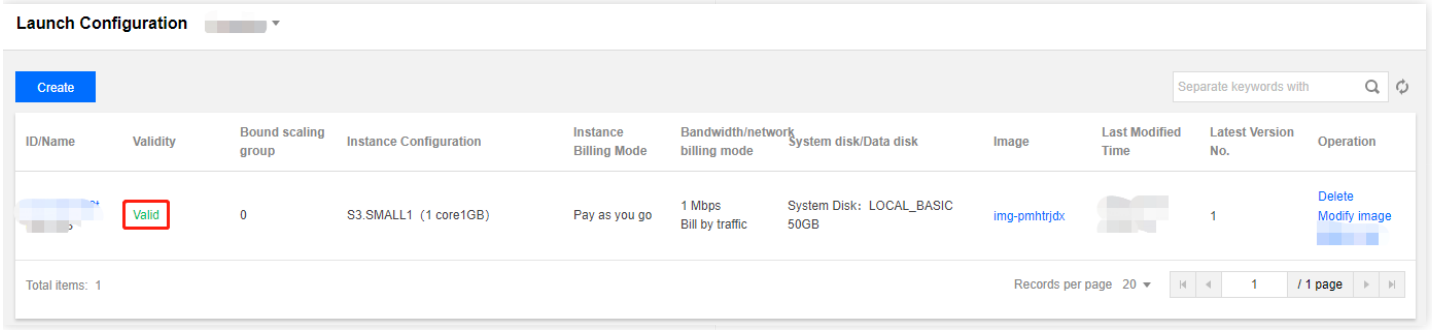
Changing the Launch Configuration

The launch configuration itself cannot be edited, but the launch configuration of the scaling group can be changed.
If you need to change the launch configuration after you upgrade the application and update the data, follow the steps below:

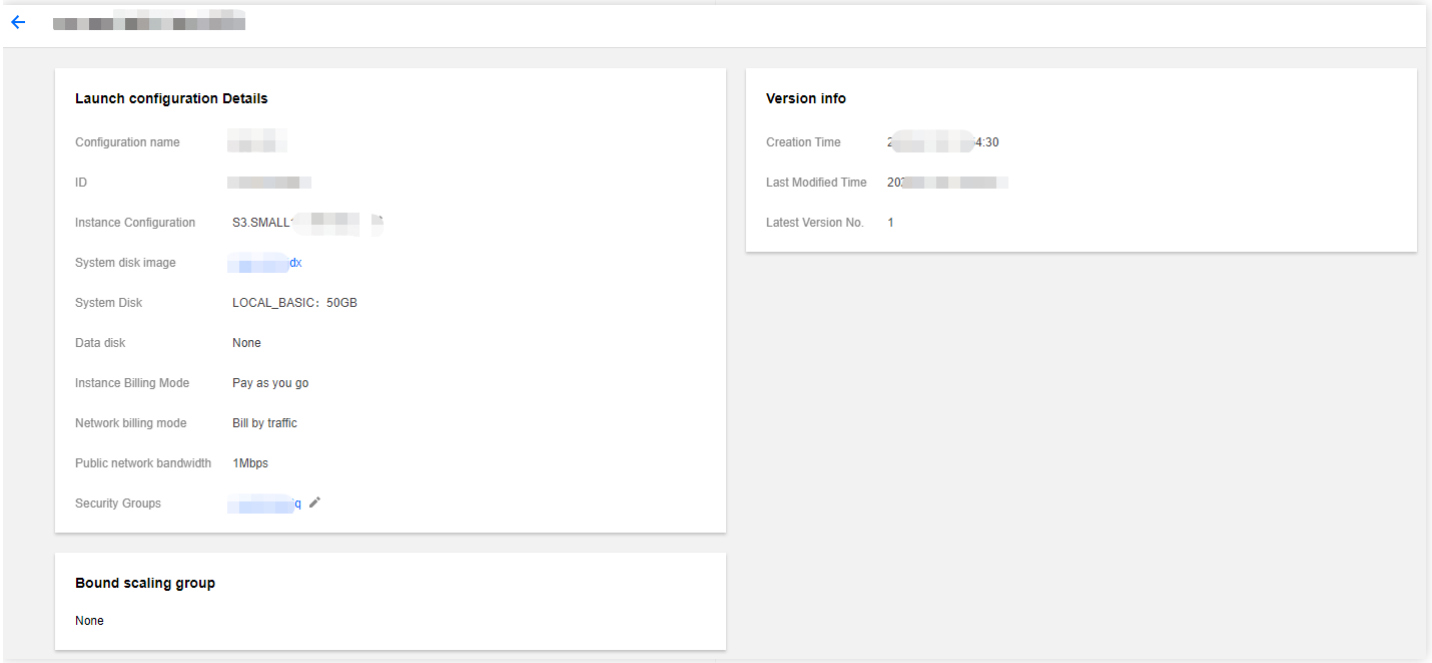
1. Create a new launch configuration.
2. To change the launch configuration of the corresponding scaling group, on the details page of the scaling group, click Edit and then point to the new launch configuration.

Detecting Exceptions

AS provides you with a very considerate service: detecting exceptions in advance.
For example, if your account balance is insufficient or you accidentally delete an image, you cannot produce the machines required for scale-out. AS will detect such anomalies in advance and issue a warning. You don't have to wait for a scale-out activity to fail before you discover the risks. This helps you stop losses in the bud.
You can check it directly in the launch configuration list. If Validity is displayed as Invalid, it means that your launch configuration has been operated incorrectly and is unavailable. When you hover the cursor over it, the reason for the failure is displayed.



You can also click the launch configuration ID to go to the launch configuration details page and view detailed issues:



Multi-Model Configuration

Operation scenarios

The multi-model configuration has the following features:

- It can effectively reduce the risk of scale-out failure caused by sold-out and instance call failure.
- We recommend a variety of models that are similar to your preferred model in terms of configuration, price, and performance. If your preferred model is sold out, we intelligently select other models with abundant inventory for you.

Operation step

Multi-Model Configuration

1. Log in to the auto scaling console and click Launch Configuration in the left sidebar.
2. On the Launch Configuration page, you can enter the multi-models configuration page in the following two ways:
    - Click Configure Models on the right side of the row where you want to configure multiple model launch configurations.
    - Click the launch configuration ID for which you want to configure multiple models. On the launch configuration details page, select the edit icon in the instance configuration row.
3. On the Multi-Model Configuration page, select any model you need and click OK.

Viewing Multi-Model Configuration

After the configuration is complete, you can view the various models and fees configured in the details page of the launch configuration.

Modifying Multi-Model Configuration

If the completed multi-model configuration does not meet your expectations, you can modify it by following the steps below.

1. To refer to multi-model configuration, enter the Multi-Model Configuration page.
2. After modifying the original model configuration, click OK to complete the modification.

Scaling Groups
Scaling Group Overview

A Scaling Group is a collection of cloud server instances that follow the same rules and target the same scenario. A scaling group defines the maximum and minimum number of CVM instances in the group, its associated load balancing instances and other attributes.

Creating Scaling Groups

Open the auto scaling console and click Scaling Group in the left sidebar.

Selecting Region

At the top of the console, select the region you wish to create.

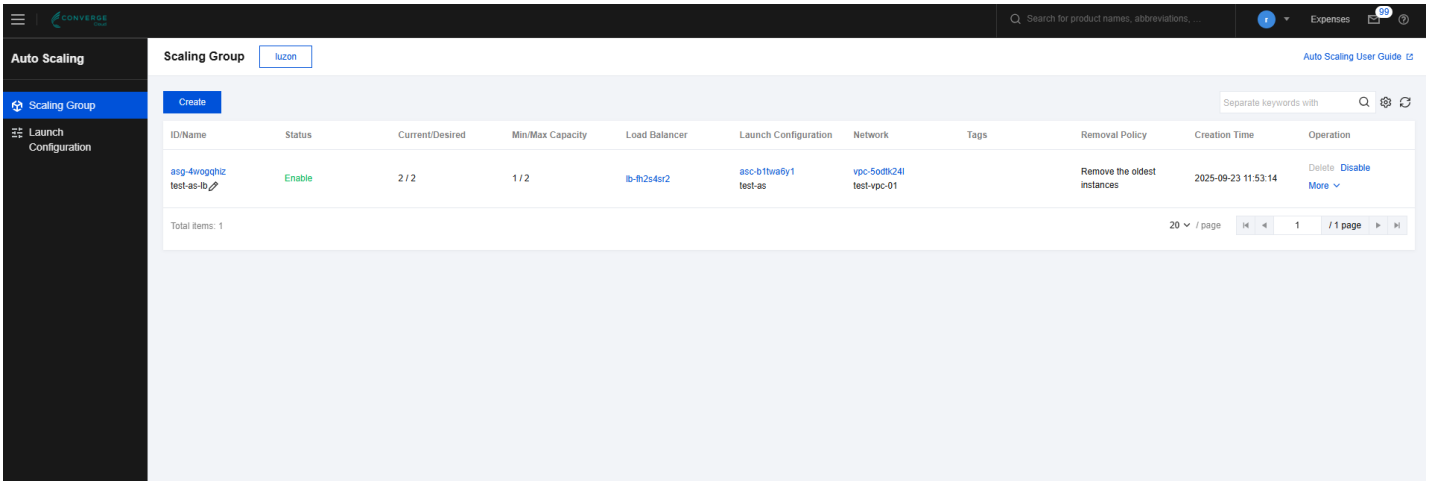
Definition Information

Click Create and define the attributes of the scaling group:

- Scaling Group Name: Used to mark this scaling group. For example, Website Logic Tier.
- Min Capacity: Specify the minimum number of instances in the scaling group.
- Intial Capacity: Specify the number of instances when the scaling group starts Automatic production. After a scaling group is created, a corresponding number of instances are generated.
- Max Capacity: Specify the maximum number of instances in the scaling group.
- Launch Configuration: Specifies the created startup configuration. When scaling-out, the machine will be scaled-out according to the startup configuration.
- Supported Networks: Selects the network attributes of the scaled-out machine, that is, whether the scaled-out machine is in the basic network or in a VPC. Unless your cluster already uses VPC, you can select Basic Network.
- Support Subnet: You can select multiple subnets, and the automatically scaled-out machines are randomly created from the subnets you checked, achieving cross-subnet disaster recovery.
- Removal Policy: When the scaling group needs to reduce instances and has multiple options, the CVM is removed based on the removal policy. Generally, just select Remove the Oldest Machine.
- CLB: If you specify a CLB, the scaled-out machines are automatically mounted to the CLB.
  The scaling group is now created. At this point, the scaling group can accommodate machines, but it does not yet have the capability for intelligent scaling. Next, we strongly recommend you to proceed with the following three actions:
- Join an existing CVM;
- Creating a scaling policy;
- Create a notification.

Viewing the Scaling Group List

Open the auto scaling console and click Scaling Group in the left sidebar to view the list.

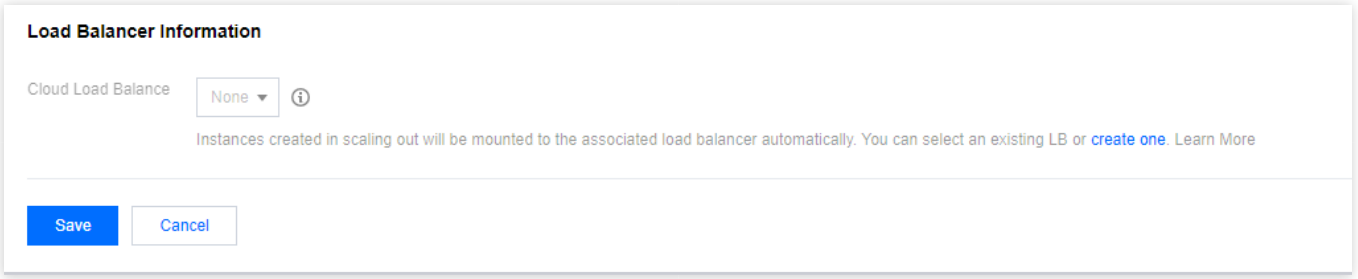


Modifying a Scaling Group

1. Open the auto scaling console and click Scaling Group in the left sidebar.
2. Select the scaling group you want to modify and click the scaling group ID/Name to enter the scaling group basic information page.
3. Click Edit to enter the editing page.

Modifying Scaling Group Details

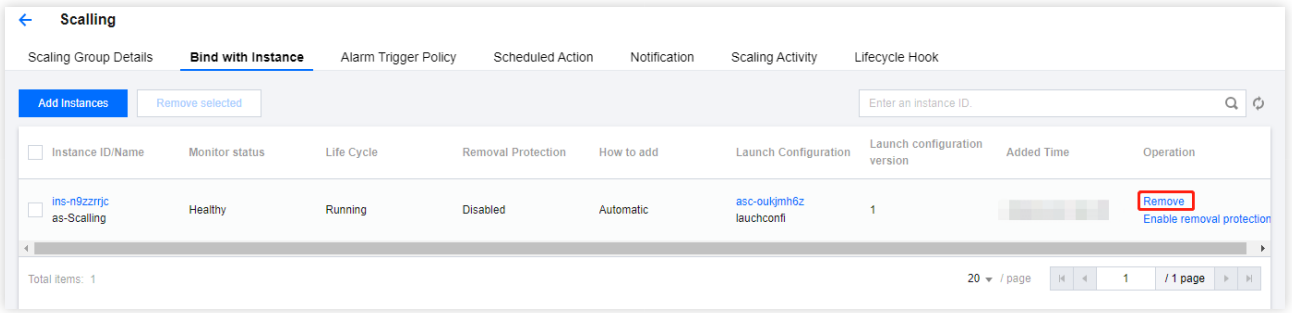
1. On the Scaling Group Details page, you can modify the scaling group name, adjust the minimum and maximum scaling values, and modify the CVM instance removal policy.
2. After the modification is completed, click Save to complete the modification of basic information.
3. If you need to modify the CLB information, click Edit on the right side of the device to enter the editing page.



4. After the configuration is completed according to actual needs, click Save to complete modifying the CLB information.

Modifying the CVM Bound to the Scaling Group

1. Open the auto scaling console and click Scaling Group in the left sidebar.
2. Select the scaling group you want to modify and click the scaling group ID/Name to enter the scaling group basic information page.
3. On this page, you can view a list of instances bound to the scaling group.
    - If you need to manually add CVM instances to the scaling group, click Add Instance, select the instances to be added (press and hold Shift to select multiple instances), and then click OK.
    - To unbind an instance, click Remove next to the corresponding instance entry.



- Automatic production machines are terminated after being removed.
- Manually added machines are not terminated after being removed. They are only removed from the scaling group and unbound from the CLB.

Combining CLBs with Scaling Groups

As the AS adds and removes CVM instances, you need to ensure that the application's traffic is allocated across all CVM instances. If you want the scaled-out machines to be under a load balance and receive traffic forwarded by the load balance without your intervention, you can assign a load balance to your machines. The CLB will serve as the single point of contact for all incoming traffic to the instances in your auto scaling group.

Adding CLBs to Scaling Groups

Scaling groups are integrated with CLB so that you can attach the CLB to the existing scaling group. After you attach the CLB, it automatically registers the instances in the group and distributes incoming traffic to them.
In the auto scaling console, select Create, and in the CLB option at the bottom of the page, select the CLB you need. If you have not created one in advance, you can click the Create link below the options to create a new CLB.

> Note:
> The CLB instance associated with the scaling group must be in the same network environment (VPC or basic network in the same region) as the scaling group.

Deleting CLBs of Scaling Groups

Click to enter the details page of the scaling group. Under the details, Click Edit to delete the corresponding CLB.
After deletion, the machines in the scaling group are also automatically unbound from the deleted CLB.

Deleting Scaling Groups

1. Open the auto scaling console and select Scaling Group in the sidebar.
2. In the scaling group list, click the Delete button next to each scaling group list.

> Note:
>
> You need to delete instances in the scaling group before you can delete the scaling group.

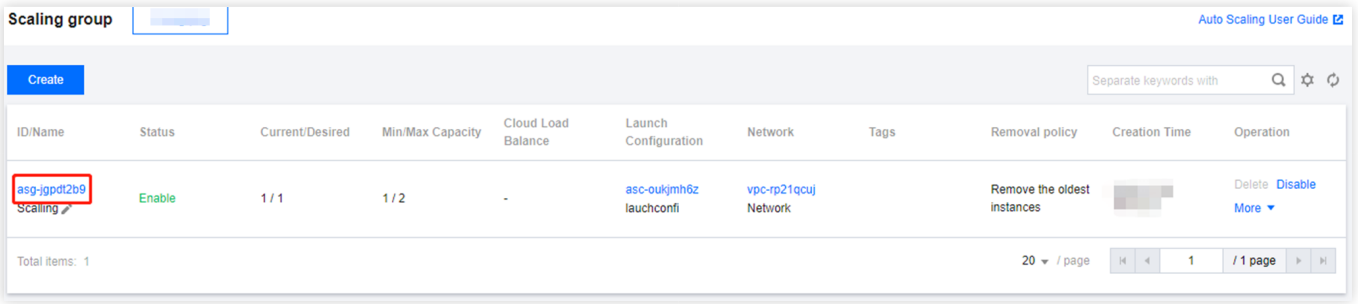
Scaling
Managing Scheduled Tasks

Introduction to Scheduled Tasks

Scheduled tasks are tasks that set a time plan so that your business can regularly scale-out or scale-in the number of CVM instances used based on predictable load changes.
For example, traffic to your web application might start to increase each Wednesday, remain high on Thursday, and then begin to decline on Friday. In this case, you can plan your scaling-out activities based on the predictable traffic patterns of your web application.
To create a scheduled scaling, specify the start time when you want the scaling to take effect and the new minimum (minimum number of instances), maximum (maximum number of instances), and desired size (desired number of instances) for the scaling. At the specified time, AS updates the number of instances in the scaling group based on these settings.
You can create a pre-scheduled action to scale only once, or you can create a pre-scheduled action to scale on a recurring schedule.

Managing Scheduled Tasks

1. Open the auto scaling console and click Scaling Group in the left sidebar.
2. Select the scaling group to be modified and click Scaling Group ID/Name to go to the scaling group basic information page.



3. Select Scheduled Action in the sidebar above to manage scheduled tasks associated with the scaling group.
   - Click Create to add a new scheduled task.
   - Select a scheduled task and click Modify. In the pop-up page, you can modify the task name, adjust the execution time, set whether to execute periodically and modify the execution activity.
   - Click Delete to delete the scheduled task.
     If you want to create a recurring task, you can specify a start time, and AS performs the action at that time according to the recurring schedule; if you specify an end time, AS no longer performs the action after that time.

Managing Alarm Triggering Policies

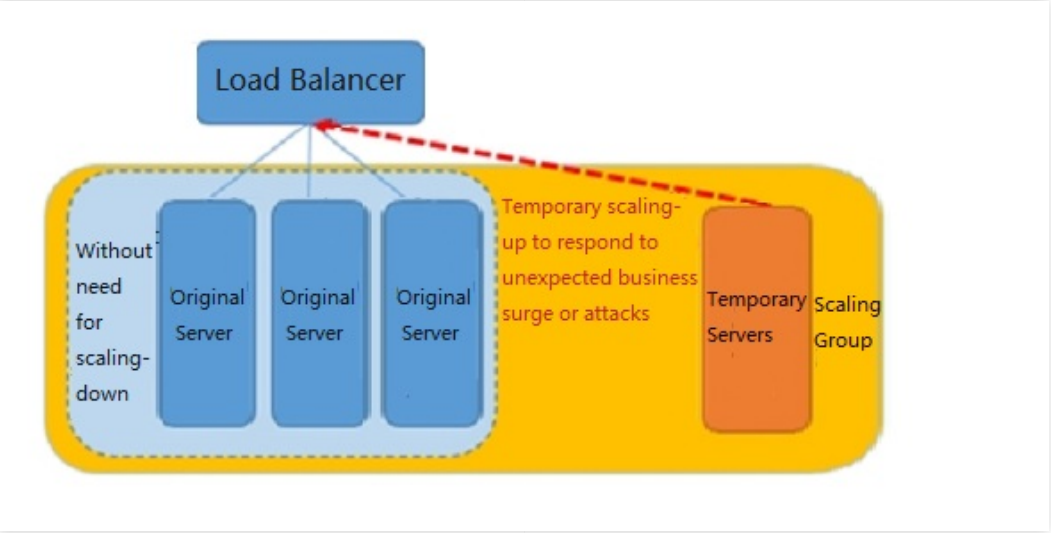
Overview

AS supports scaling out the number of instances in the scaling group dynamic based on monitored metrics. You need to define an alarm-triggering policy, that is, the monitoring indicator state that triggers scaling-out and how to scale according to demand changes.
To create an alarm policy, you need to specify conditions and actions:

- Conditional Format: A certain metric + threshold + cycle + the number of consecutive cycles that reach the threshold. That is, the metric reaches the threshold for N consecutive periods.
- Execution Actions: Send notification + increase/decrease the specified number of CVM.
  We recommend that you create two policies for each scaling group: One for scaling-out and another for scaling-in. When the business volume reaches the conditions specified by the alarm policy, the AS executes the associated policy to scale-in (by terminating instances) or scale-out (by launching instances) the scaling group.

Scenario Example

For example, you have an e-commerce website application and currently use five instances. You have done an operation activity and are worried that the traffic will be much greater than you estimated. You can set it to launch two additional instances when the load on the current instance rises to 70% and then terminate the excess instances when the load drops to 40%. You can configure a scaling group to automatically scale-out based on these conditions.

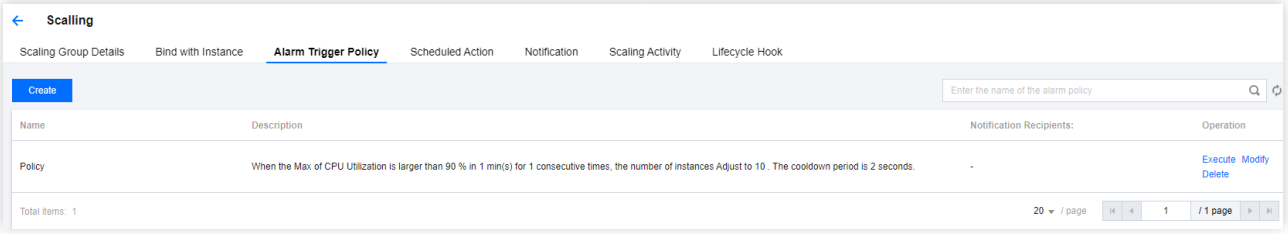


Operation step

1. Open the auto scaling console and select Scaling Group in the sidebar.
2. Select the scaling group to be modified and click Scaling Group ID/Name to go to the scaling group basic

information page.

3. Select the Alarm Triggered Policy in the sidebar above to manage the alarm trigger policy associated with the scaling group.
    - Click Create to add a new alarm trigger policy;
    - Click Delete to delete the alarm trigger policy;



Specifying a Server That is Not Affected by the Alarm Scaling Policy

Before using auto scaling, your system may already have frequently used servers, and you may not want the machines to be removed by the alarm scaling policy for the following reasons:

- Multi-purpose machine: In addition to its role in the cluster, a server may also serve other purposes. For example, during the early stages of WDS, a server could be used both as a cache/caching server and as a file server. When such cache/caching servers are placed in a scaling group, you might not want them to be removed by the alert scaling strategy.
- Store Data: This server is stateful or has data that other servers do not have. For example, the incremental data generated during the running of other servers in the cluster are all saved in this server.
- Update Image/Snapshot: Use this server regularly for mirroring and snapshots.
  Setting Method:

1. You can click the scaling group where the server is located in the scaling group list to enter the management page.
2. In the Instance List at the bottom of the management page, click Enable Removal Protection for the instance you want to set.

Instance Health Examination

If you specify the Initial Capacity when creating a scaling group, after you create the launch configuration and scaling group, the scaling group creates new CVM instances equal to the starting number of instances. At the same time, the scaling group ensures that there are running instances greater than the Min Capacity and less than the Max Capacity.

> Notes:
>
> - Min Capacity: The minimum number of instances allowed in a scaling group. When the number of CVMs in a scaling group is less than the minimum number of instances, the AS adds instances so that the current number of instances in the scaling group matches the minimum number of instances.
> - Initial Capacity: The number of CVM when the scaling group is created.
> - Max Capacity: The maximum number of instances allowed in a scaling group. When the number of CVMs in a scaling group is greater than the maximum number of instances, the AS removes instances so that the current number of instances in the scaling group matches the maximum number of instances.

To keep the instances in the scaling group running normally, AS performs regular examinations on the running state of the instances in the scaling group. If it finds an instance is unhealthy, it terminates the instance and launches a new CVM instance.

- Instance Health Examination:
  The scaling group periodically checks the instance running state to determine whether each instance is healthy. The judgment standard is whether the machine is unreachable by ping for 1 consecutive minute. If an instance is unreachable by ping for more than 1 minute, the AS marks the instance as unhealthy.
- Replacing Unhealthy Instances:
  After an unhealthy instance is marked as unhealthy, the scaling group immediately launches a new instance to replace it (except for machines with a removal protection set).

Manual Scale-out

In addition to supporting automatic scale-out based on business load, AS also supports manual intervention to achieve fast manual scaling. You can achieve scale-out effect in the following two ways:

- Adding an existing CVM instance to the scaling group;
- One-click scale-out can be achieved by modifying the expected number of instances in the scaling group.

Adding an Existing CVM Instance to the Scaling Group

Scaling groups allow you to add existing instances to an existing scaling group, enabling you to observe and manage loads together with other machines in the scaling group.

Prerequisites

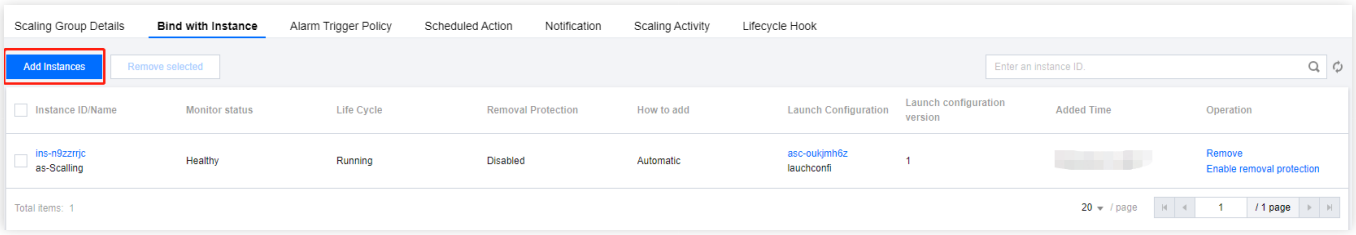
- The instance is in the running state.
- The instance and scaling group are located in the same region.
- The network attributes of the instance must be the same as those of the scaling group; that is, they must belong to the same basic network or the same VPC.

Notes

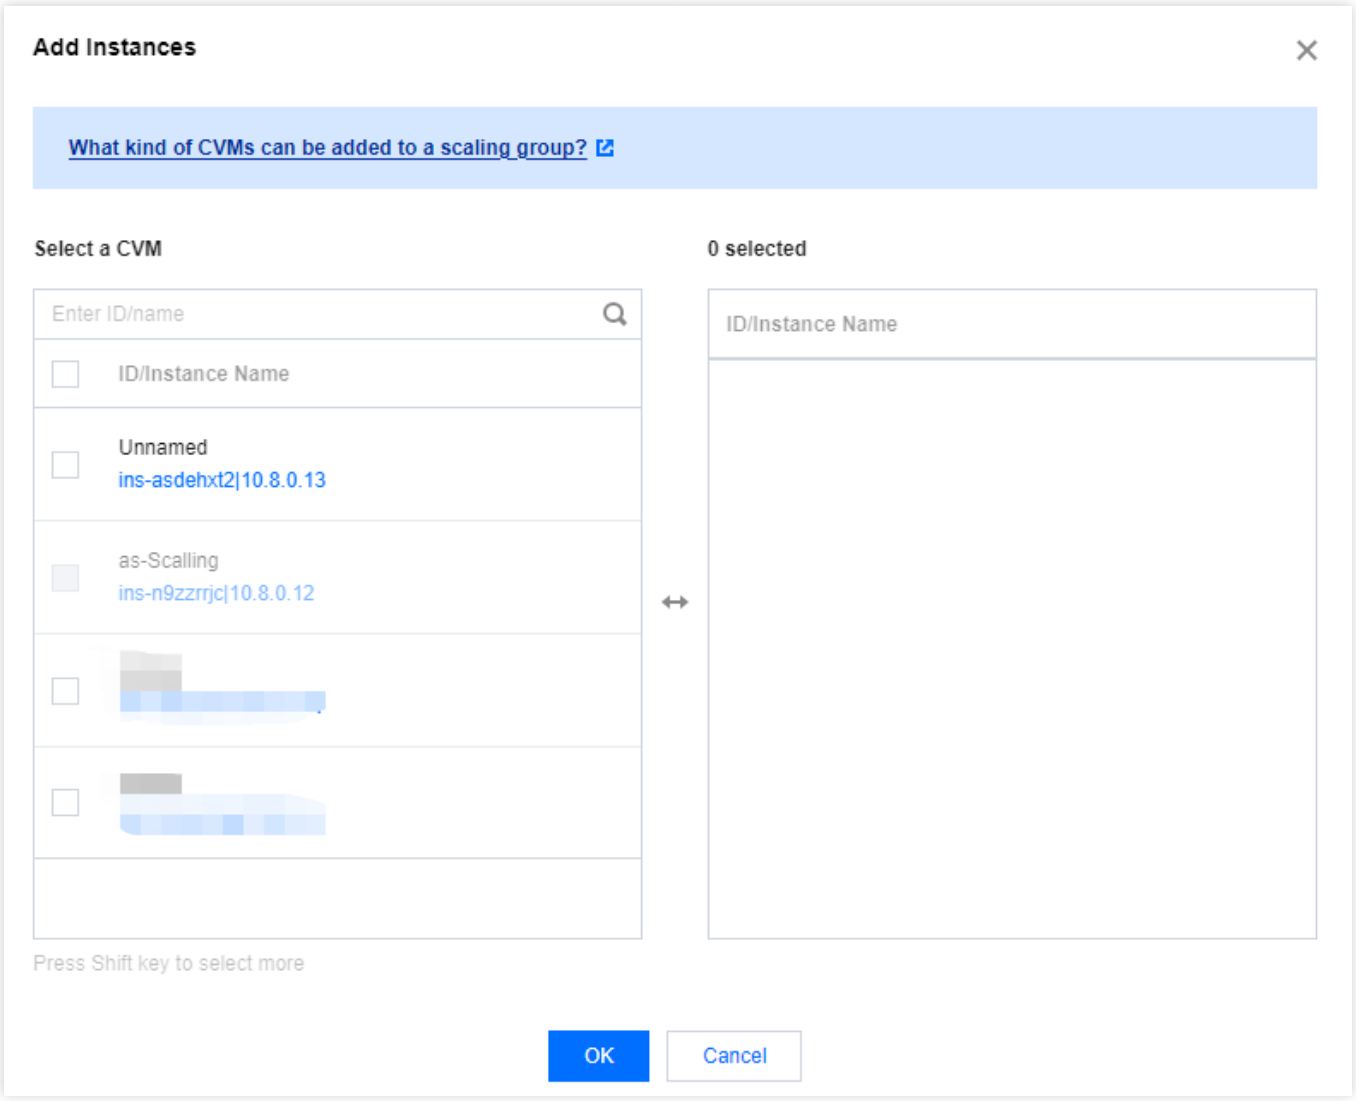
- AS will add the required capacity of the group to the number of instances to be added.
  For example, if the current desired number of instances in your scaling group is 5, and you manually add 3 instances, the desired number of instances in your scaling group will become 5 + 3 = 8. If the number of instances to be added plus the required capacity exceeds the maximum number of instances of the scaling group, the request will fail.
- The scaling group is associated with one or more CLBs. Manually added instances are automatically registered with all CLBs in the scaling group.
- When a scaling group scales-in, the automatically created machines are removed first. If there are no automatically created machines, manually added machines are removed.
- When you remove a manually added instance from a scaling group, you only remove the instance from the scaling group and CLB. This means that the instance is no longer managed by the scaling group, but your instance is not terminated.

Manually Add an Instance Using the Console

1. Log in to the scaling group console and click the scaling group ID/Name to which you want to add an instance.
2. Go to the scaling group details page and select Bind with Instance > Add Instances.



3. In the dialog box, select the corresponding instance and click OK.



One-click scale-out can be achieved by modifying the expected number of instances

Scaling-out Scenarios

If your needs match the following scenarios, you can use the console for one-click expansion. Prepare tasks such as CLB forwarding rules, machine configuration, and business deployment in advance. Even if your business

needs scaling later, you only need to modify the parameters of the scaling group with one click to quickly complete the expansion.

- The peaks and troughs of business are difficult to predict, but we are unwilling to leave the scaling entirely to the system. Business peaks and troughs are predictable.
- Your computing needs are project-specific, and the machines you use are similar each time. For example, social opinion collection, gene sequencing, weather forecasting, etc.

One-Click Scale-out in the Console

Perform the following steps to set the CVM template as the launch configuration and configure the corresponding scaling group.

1. Create a custom image.

   > Notes:
   >
   > Subsequent expansion instances deploy the environment based on this image.
   > Recommended Ideas for Creating a Custom Image: You can select an existing CVM or create a new CVM, deploy your business, set the business to launch with the operating system, and then export it as a custom image.

2. Create a launch configuration based on the custom image.
3. Create a Scaling Group.

   When creating, select the created launch configuration, and fill in the minimum scaling number, maximum scaling number, and starting number of instances according to the lower limit, upper limit, and current number of servers you need.

4. After completing the preceding steps, when the business needs to be scaled-out (for example, starting a gene sequencing task or activating a request-based machine to collect data), you can modify the scaling group configuration to increase the minimum scaling number, maximum scaling number, and expected number of instances, and AS will quickly complete the scaling-out.

Scaling-in

For each scaling group, you can control when to add instances to it (scale-out) or remove instances from it (scale-in). You can scale the size of your group manually by adding or removing instances, or you can use scaling policies to let AS automate the process.
When a scaling group automatically scales-in, you need to know which instances should be terminated first. The selection is based on the removal policy.
When scaling-in, you can prevent AS from terminating specific instances by using instance protection.

Removal Policy

When a scaling group is scaled-in, it decides which machine to remove based on the removal policy. You can choose from two removal policies:

- Remove the oldest instances: Delete the earliest added machine.
- Remove the Latest instances: Delete the most recently added machine.

Setting and Modifying Removal Policies in the Console

There are two ways to set it:

- When you create a scaling group, select the removal policy you want.
- On the scaling group details page, click Edit to modify the scaling policy.

Viewing Scaling Activities

Viewing Scaling Activities

1. Open the console and select Scaling Group in the sidebar.
2. Select the scaling group to be modified and click Scaling Group ID/Name to go to the scaling group basic information page.
3. Select Scaling Activities in the sidebar above. On this page, you can view information about scaling activities that have been executed by the scaling group based on the scaling policy.

Suspending and Resuming Scaling

Use Cases

If you need to troubleshoot configuration or other issues related to the web application (such as resetting a password after shutting down the application, upgrading a service, etc.) and want to make changes to the application without triggering the automatic scaling process, you can suspend the scaling group and resume it after completion.

Suspending a Scaling Group

Operation

1. Open the auto scaling console.
2. Click Scaling Group in the left sidebar.
3. In the operation column of the scaling group list, click Disable.
   After setting, you can see the word Disabled in the Status list.

Notes

After a scaling group is disabled, the automatically triggered activities of the scaling group are not performed, but the restrictions of the scaling group are still in effect.

- Automatically triggered activities include:
  - Alarm Scaling
  - Scheduled Task
  - Health Examination
  - Manually Causing the Expected Number of Instances to Not Match
- Scaling group limitations include:
  - Manually removing instances is not subject to the minimum number of instances.
  - If the number of manually added instances exceeds the maximum number, the addition is not allowed.
  - Manually increasing the minimum or maximum number of instances does not trigger scaling activities.

Resuming Scaling Groups

If you have completed troubleshooting or operations during the period when the scaling group activities

were suspended, you can resume the automatic scaling settings for your business.

1. Open the auto scaling console.
2. Click Scaling Group in the left sidebar.
3. In the operation column of the scaling group list, click Enable.

Exempting a Specified Instance From Scaling-in

Overview

In a scaling group, you can specify an instance that is not scaled-in during a scaling-in activity. When a scaling-in activity occurs, AS selects a sub-machine to be scaled-in from other machines.
You can enable Instance Protection for one or more scaling group instances. You can change the scaling group or instance protection settings at any time.
If all remaining instances in the scaling group are protected by scaling-in protection and a scaling-in event occurs at the same time, AS reduces the required capacity instead of removing instances.

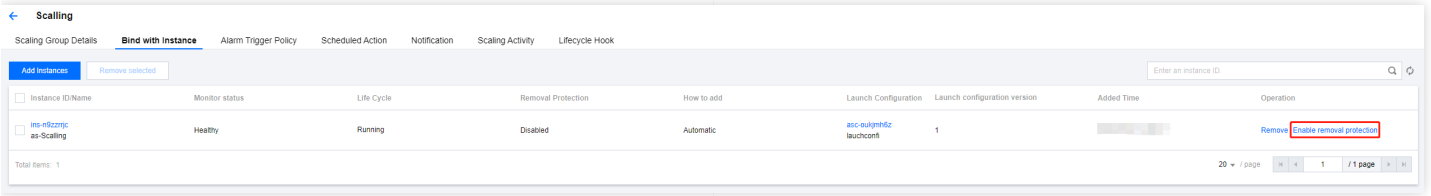
Applicable Scenario

Normally, the machines in a scaling group are stateless, and all machines can be removed at any time. However, in actual practice, the following situations apply to exempting a specific instance from being scaled-in:

- One Machine for Multiple Uses: Based on cost considerations, some machines, besides doing things in the cluster, also have other purposes, such as storing data generated in the cluster, so this machine is actually stateful.
- Avoid Misoperation: If you are worried that incorrect policy settings affect your business, you can set some machines to be exempt from being scaled-in. In this way, AS will never scale-in the machine, and the request-LB-instance path can remain unobstructed.

How to Set Up?

In the instance list of the scaling group, you can directly set:



Scaling Activity Cancellation

Scaling activity cancellation means that the scheduled task time is reached or the alarm scaling condition is met, so the scaling activity is triggered, but there is a conflict, and the scaling activity is forced to be canceled.
Cause of Conflict:

- There is ongoing scaling activity.
- The scaling group is in the cooldown period.
  Is the scaling activity retried after cancellation?
- Alarm Scaling. If the activity is canceled, it is not retried. However, if the conditions for alarm scaling continue to be met, the next alarm scaling activity is triggered.
- Scheduled Action The expected number of instances, the maximum number of scaling, and the minimum number of scaling are defined, so the scaling group keeps retrying to ensure that the actual number of instances matches the expected number of instances.

  > Note:
  > When a scaling group is suspended, the scaling group does not attempt scaling activities directly. Therefore, the Scaling Activities record does not contain any canceled scaling activities.

Scaling Activity Failed

Scaling Activity Canceled is in line with expectations. Scaling Activity Failed is not in line with expectations.

How can I view failed scaling activities?

You can configure a notification policy to be notified immediately when a scaling activity fails.

Why do failed scaling activities occur?

We have categorized the reasons for scaling activities.

Cooldown Period Description

What is the Cooldown Period?

The AS cooldown period is a configurable setting for a scaling group. By setting the cooldown period, you can ensure that AS does not launch or terminate other instances before the previous scaling activity takes effect. After a scaling group is dynamically scaled using a simple scaling policy, AS waits for the cooldown period to complete before continuing the scaling activity.
When you manually scale out a scaling group, the default is to not wait for the cooldown period, but you can set a cooldown period to override the default setting. Note that if an instance is detected to be unhealthy, AS replaces the unhealthy instance immediately without waiting for the cooldown period to complete.

Why Do We Need a Cooldown Period?

After a machine is added to the scaling group, it takes some time to reduce the load. If there is no cooldown period, the system continues to scale-out before the load drops. After the newly added machine takes over the business, if the load is too low, it scales-in.
These instances use configuration scripts to install and configure the software before the instances are put into use, so it takes about two to three minutes for the instances to be launched and put into use. (The actual time depends on many factors, such as the size of the instance and whether there are launch scripts to complete.)
Scenario Example
A spike in traffic occurs, which triggers the alarm policy. When it does, AS enables an instance to help with the increase in demand. However, there's a problem: the instance takes a couple of minutes to enable, and it will also take some time for the enabled instance to receive requests from CLB. During that time, the monitor alarm could continue to be triggered, causing AS to enable another instance each time the alarm is triggered.
However, if a cooldown period is set, after an instance is launched, AS suspends all scaling activities caused by simple scaling policies or manual scaling until the specified amount of time has passed (the default value is 60 seconds). This gives newly launched instances time to start handling application traffic.
After the cooldown period, all paused scaling operations will resume. If the alarm is triggered again, AS will launch another instance, and the cooldown period will be reactivated. However, if the added instance is sufficient to reduce the CPU utilization to a normal level, then the group will maintain its current size.

Setting Cooldown Period

The default cooldown period is 60 seconds.
If you need to modify it, follow the steps below:

- Open the scaling group details page.
- Click Alarm Triggered Policy, select the alarm scaling policy to be set, select Modify, and specify the

cooldown period (which can be set from 0 to 999,999 seconds) below the edit box.

Monitoring
Monitoring Alarm Metrics

With Cloud Monitoring, you can search statistics by an ordered set of time-series data called metrics. You can use these metrics to verify that your system is running as expected and to scale if thresholds are exceeded. You can use these metrics to verify that your system is running as expected and to scale if thresholds are exceeded.

AS Monitoring Metrics

AS currently supports the following metrics:

- CPU Uitilization
- Memory Uitilization
- Private Bandwidth In
- Public Bandwidth In
- Private Bandwidth Out
- Public Bandwidth Out
  Each metric can support the following dimensions:
- Maximum value
- Minimum value
- Average value

Metric Aggregation Method

Auto scaling is the monitoring of CVM clusters, which involves multiple CVMs and multiple monitoring data generated by these CVMs within a time period. These data are aggregated first and then operated according to the user-configured policies.
The basic statistical policy is to take values for the set monitoring items of each CVM for 1 minute in each cycle (one value is taken every minute). If the values obtained meet the set rules for multiple consecutive cycles (the number of cycles can be customized by the user), the alarm scaling behavior is triggered.
For example, there are 5 CVM instances in a scaling group, and the defined alarm scaling policy is that the maximum/minimum/average CPU usage within 5 minutes is greater than 50% and occurs 3 times. The auto scaling collects monitoring data and makes policy judgments. The steps are as follows:

1. The system takes one value for each CVM every minute and takes 25 CPU usage values in one cycle (currently set to 5 minutes).

2. The maximum/minimum/average value is configured based on the policy to determine whether it meets the alarm rules.
    - Maximum Value: If the maximum value of the 25 values exceeds the threshold (50%), the period meets the alarm scaling rule.
    - Minimum Value: If the minimum value of the 25 values exceeds the threshold (50%), the period meets the alarm scaling rule.
    - Average value: If the average value of these 25 values exceeds the threshold (50%), the period meets the alarm scaling rule.
3. If this rule is met for three consecutive periods (15 minutes in total, with the current period judged every 5 minutes), the scaling behavior is triggered.

Running State Examination

AS periodically checks the health of instances in the scaling group and identifies all instances that are in poor health. After the AS marks an instance as unhealthy, it schedules the instance for replacement.

Create a notification
Creating a Scaling Activity Notification

Overview

By creating a notification, AS gives you full control over the scaling group, including:

- You can know all the scaling activities immediately;
- You can then review the details of all scaling activities.
  On the scaling group management page, click Notification and create a notification policy.

Notification content

You can know the following in real-time:

- Scale-out succeeded.
- Scale-out failed.
- Scale-in succeeded.
- Scale-in failed.
- Health replacement succeeded.
- Health replacement failed.

Notification Methods

AS supports the following notification methods:

- Email
- Short Message Service
- Internal Message

Specifying Notification Recipients

If you have created a notification and specified recipients, AS sends notifications to the specified recipients when a scaling activity occurs.

Step 1: Defining Notification Recipients

Enter the user management console, and you can see the contact information of the account registrant. If you want more people to receive notifications, such as other members of the Ops team, you can click Create to create more.

> If you are the only one using your account, you can skip this step because the system has created a Developer account by default.

Step 2: Defining User Groups (Recipient Classification)

AS takes User Group but not Recipients as a unit to send notifications.

Scenario Example

- If you want only yourself to receive notifications, you can create a user group with only your contact information.
- You have defined a number of notification recipients, who may be in the same or different departments. You may want to have a certain type of notifications sent to the recipients of department A, and another type to the recipients of department B. To do so, you can define a user group and add the different notification recipients to the corresponding user groups.

How to Set Up?

1. Enter user group management, click Create User Group, fill in the user name, and click Submit.
2. Then click Add User and add the relevant notification recipients.

Step 3: Using User Groups

When defining the alarm-triggering policy and notification policy of AS, you can see the user groups you have defined in the list of notification recipients. You can set the user groups specified for notification as needed.

Troubleshooting
Scaling Group Failure

Phenomenon

When your account environment cannot create a machine manually or through AS, or when you delete the associated resources (such as load balancing and private networks) after creation. The failure of a scaling group forms a mechanism for detecting risks in advance, which prevents scaling-out failure when scaling-out is really needed and greatly improves the security and availability of your cluster.

Analysis

You can view the scaling group state in the scaling group list. If your scaling group state shows  , it means that the scaling group is in Exception state. At this time, the scaling group is unavailable.

> Note:
> You can hover the mouse pointer over the prompt icon to view the specific cause of the exception.

Exception Causes

- Launch Configuration Exception
- CLB Exception
- VPC Exception
- Insufficient Balance

Troubleshooting

Launch Configuration Exception

Solution: If the launch configuration is invalid, it is because the associated resources (image, snapshot, security group, etc.) of the launch configuration have been deleted.
Processing Method: Click Launch Configuration ID in the row where the scaling group is located, go to the corresponding launch configuration details page, and check whether any resources have been deleted by mistake.

CLB Exception

Processing Method: Click CLB in the row where the scaling group is located, and enter the corresponding CLB details to view the specific situation, including the listener, domain name, and URL path.

VPC Exception

Processing Method: Click Network in the row where the scaling group is located, and enter the corresponding VPC details to view the specific situation, including subnet, routing table, etc.

Impact of Scaling Group Failure

After a scaling group fails, it does not stop working immediately:

- Normal scaling-in activities are not affected.
- The limits on the maximum scaling number, minimum scaling number, and expected number of instances are still working normally.
- The scaling-out activity is no longer performed because your environment no longer meets the conditions for creating a CVM.

Scaling Activity Execution Failed

Phenomenon

When a scaling activity is being executed, there may be some reasons that cause the activity to fail, which require timely attention and processing to ensure that the scaling activity proceeds normally.

You can view the scaling group state in the scaling group list. If your scaling group state shows  , the latest scaling activity in this scaling group has failed.

> Note:
> You can hover the mouse pointer over the prompt icon to view the specific cause of the exception.

Analysis

Viewing Reason Explanation

AS provides an intelligent way to view the reasons for scaling group activity failures. You can view detailed information about scaling activity failures by performing the following steps:

1. Click the ID of the scaling group where the scaling activity failed in the scaling group list.
2. On the scaling group details page, select Scaling Activities to view activity information.

Performing Corresponding Operations According to the Prompts

You can perform corresponding operations based on the reason why the scaling activity failed. Common reasons for scaling activity failure are as follows.

Reasons for the Failure of Scaling Activities

The reasons for the failure of scaling activities can be divided into the following categories:

- CVM Related
- Image Related
- Network Related
- Cloud Block Storage Related
- CLB Related

Troubleshooting

CVM Related

CVM Sold Out

Failure Reason: Insufficient inventory of the specified resource.
Solution: You can configure multi-models in the launch configuration and select other instance specifications.

The CVM Model is Invalid in the Current Availability Zone

Failure Reason: The specified instance specification is offline.
Solution: Select the instance specification on sale in the launch configuration.

The CVM Does Not Work with Cloud Block Storage

Failure Reason: The current system disk type does not support this model instance.
Solution: Check the scaling configuration and change the system disk type. It is recommended that you choose a high-performance cloud disk or SSD cloud disk.

Insufficient CVM Purchase Quota

Failure Reason: Each user has a quota for purchasing CVMs.
If you exceed this quota, AS cannot open the machine.
Solution: Reduce the number of sub-machines for scaling-out.

The CVM Model Does Not Exist

Failure Reason: The model defined in your launch configuration is incorrect or has been removed from the shelves.
Solution: Replace the corresponding launch configuration.

Image Related

The Image Does Not Exist

Failure Reason: The image may no longer exist or is invalid.
Solution: Check the relevant configuration in the launch configuration.

The Image Exceeds the System Disk Capacity

Failure Reason: The system disk capacity is less than the image file size.
Solution: Check the launch configuration, increase the system disk capacity, or use an image that occupies a smaller space.

Network Related

Only Supports VPC

Failure Reason: A model that only supports VPC was selected in the launch configuration.

Solution: Do not use this model in the basic network; deselect it in the launch configuration.

The number of IP addresses in the VPC subnet is less than the number of instances to be scaled-out.

Failure Reason: The number of IP addresses in the VPC subnet is limited.
Solution: Create a new subnet and expand its subnet IP range (CIDR).

Issues Related to Cloud Block Storage

Cloud Disk Sold Out

Failure Reason: Insufficient inventory of the specified resource.
Solution: Configure the default cloud disk feature in the launch configuration.

The CVM Does Not Work with Cloud Block Storage

Failure Reason: The current system disk type does not support this model instance.
Solution: Check the scaling configuration and change the system disk type. It is recommended that you choose a high-performance cloud disk or SSD cloud disk.

Issues Related to CLB

CLB Does Not Exist

Failure Reason: The CLB may no longer exist or is invalid. Check the relevant configuration in the CLB.
Solution: Make sure that the CLB instance associated with the scaling group is working properly.

Listener Does Not Exist

Failure Reason: The listener in the CLB may no longer exist or is invalid. Check the relevant configuration in the CLB.
Solution: Make sure that the CLB instance associated with the scaling group is working properly.

The Forwarding Path Does Not Exist

The listener in the load balancer forwards to the specified domain name/URL path. The domain name/URL path you created may no longer exist or is invalid. Please check the relevant configuration in load balancing.
Solution: Make sure that the CLB associated with the scaling group is working properly.

The Specified CLB Is Busy

Failure Reason: The CLB task is in progress.
Solution: It is recommended that manual operations be reduced when the scaling group is active to avoid operation mutual exclusion.

Best Practice
Setting a Fixed Outbound Access IP for Scaling Groups

This section introduces how to set a fixed outbound access IP if your cluster requires proactive outbound access.

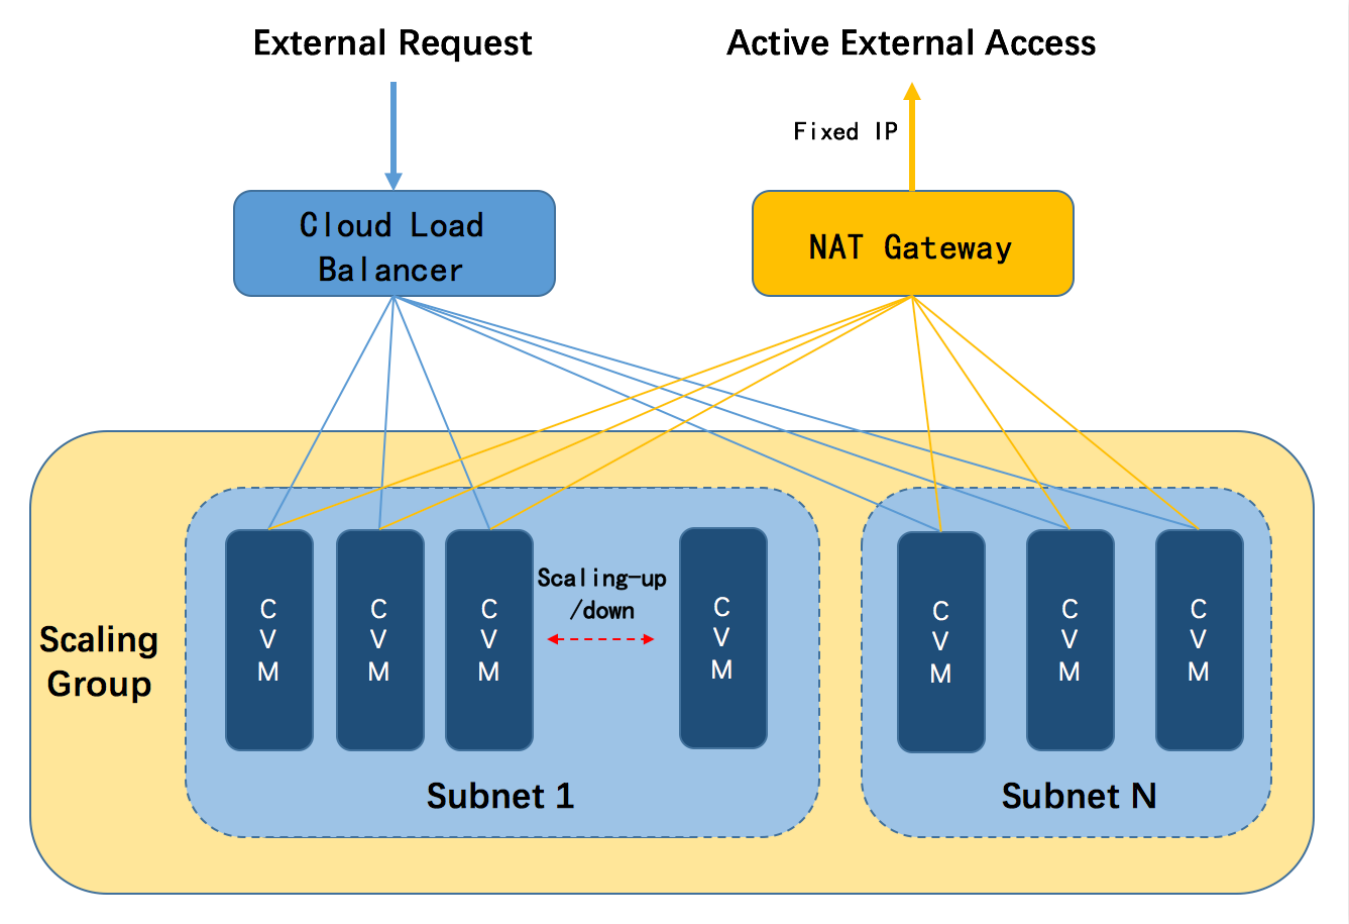
Requirement Scene

This solution is applicable when your clusters in the scaling group have all the three requirements below:

- Accept requests from the Cloud Load Balancer (CLB).
- The cluster machine needs to initiate outbound access.
- For outbound access, you may want to use a fixed public network IP.
  You can set it up through the following scheme.

Solution Overview

1. Receive and respond to external requests by using the CLB.
2. Place the machine into the subnet of VPC and direct the route table to the NAT Gateway, so all proactive outbound access requests are sent through the NAT Gateway's public network IP.
3. The network attribute of the scaling group is set as this subnet so that all scale-out machines will use the NAT Gateway for outbound access.



Methods for Setting

Step 1: Create VPC and Subnet

1. Creating VPC

1. Log in to the cloud console, and click Virtual Private Cloud in the sidebar to enter the VPC console.
2. Select a region from the dropdown list above the list, and click Create to create a VPC.
3. Enter the names of VPC and subnet as well as CIDR, and select the availability zone for the subnet.
4. Click Create.

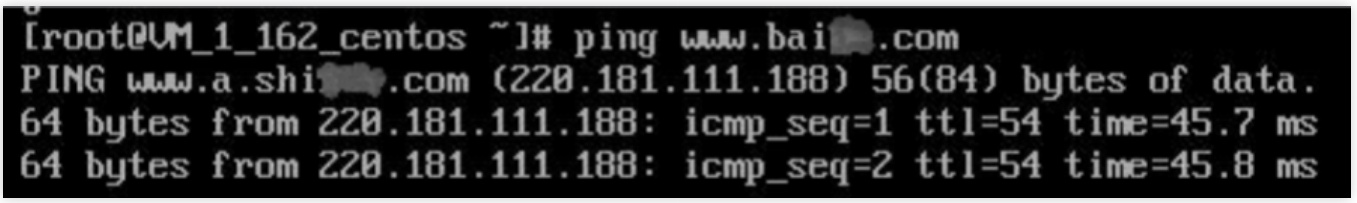
2. Creating subnet

1. Log in to the cloud console, click Virtual Private Cloud in the sidebar, and click Subnet in the left sidebar. Select a region and VPC from the dropdown list.
2. Click Create, fill in the subnet name, CIDR, availability zone, and associated route table.
3. Click Create to confirm.
   After completing the creation, you can purchase machines for this subnet.

Step 2: Create NAT Gateway

1. Purchase
    1. Log in to the cloud console, select the Virtual Private Cloud, and select NAT Gateway.
    2. Click Create in the upper left corner, and enter or confirm the following parameters in the pop-up box:
        - Gateway Name
        - Gateway Type (which can be modified after creation)
        - VPC of the NAT Gateway service (the one you just created)
        - Assign an elastic IP to the NAT Gateway, which will be the static IP for your machine's outbound access in the future.
    3. After selection, click Create button to complete the creation of the NAT Gateway.
    4. After the NAT Gateway is created, you need to configure the routing rules on the route table page in the VPC console to redirect the subnet traffic to the NAT Gateway.

2. Setting the route table (Key step)
    1. Log in to the cloud console, click Virtual Private Cloud in the sidebar to enter the VPC console, and select Route Table.
    2. In the route table list, click the route table ID associated with the subnet that needs to access the Internet to enter the route table details page, then click Add Routing Policy in the routing policy.
    3. Click to add a new line, fill in the destination (for example, in this scene, you can fill 0.0.0.0/130618223492100096), select NAT Gateway for the next hop type, select the created NAT Gateway ID, then click Create.
       Now, machines in this subnet, even without a public network IP, can also actively outbound access through the NAT Gateway, presenting a static IP externally.
       As shown in the figure below, even if one purchases a host without a public network IP and with 0 bandwidth, it can still actively outbound access:



       However, the scaling group needs to recognize this subnet and ensure that all machines are created on this subnet.

Step 3: Set Scaling Group

The purpose of this step is to direct the subnet information to the scaling group, so the new scale-out machines will be launched in this subnet.
This way, the scale-out machines will automatically use the NAT Gateway's IP address for outbound access, achieving the effect of a static exit IP.
In the auto scaling console, click Create:

- Fill in the scaling group name, launch configuration (the launch configuration should be set up in advance),

maximum scaling number, minimum scaling number, starting instance count, and other information.

- Select Network and Subnet, and direct to the VPC and subnet you have just set up (important).
  So far, the setting is complete.

Setting Services to Start at Boot

Set Auto-Scaling Machines to Start Services at Boot

Use Cases

When using auto scaling to scale-out, we aim for no manual intervention throughout the entire process. Therefore, we strongly advocate for setting auto-scaling machines to automatically start services upon boot. For example:

- httpd service
- mysqld service
- php-fpm service
- tomcat service

  It only takes a minute to modify the /etc/rc.d/rc.local file to complete the settings.

Method for Setting (Taking CentOS as an example)

Step 1: Open the rc.local file

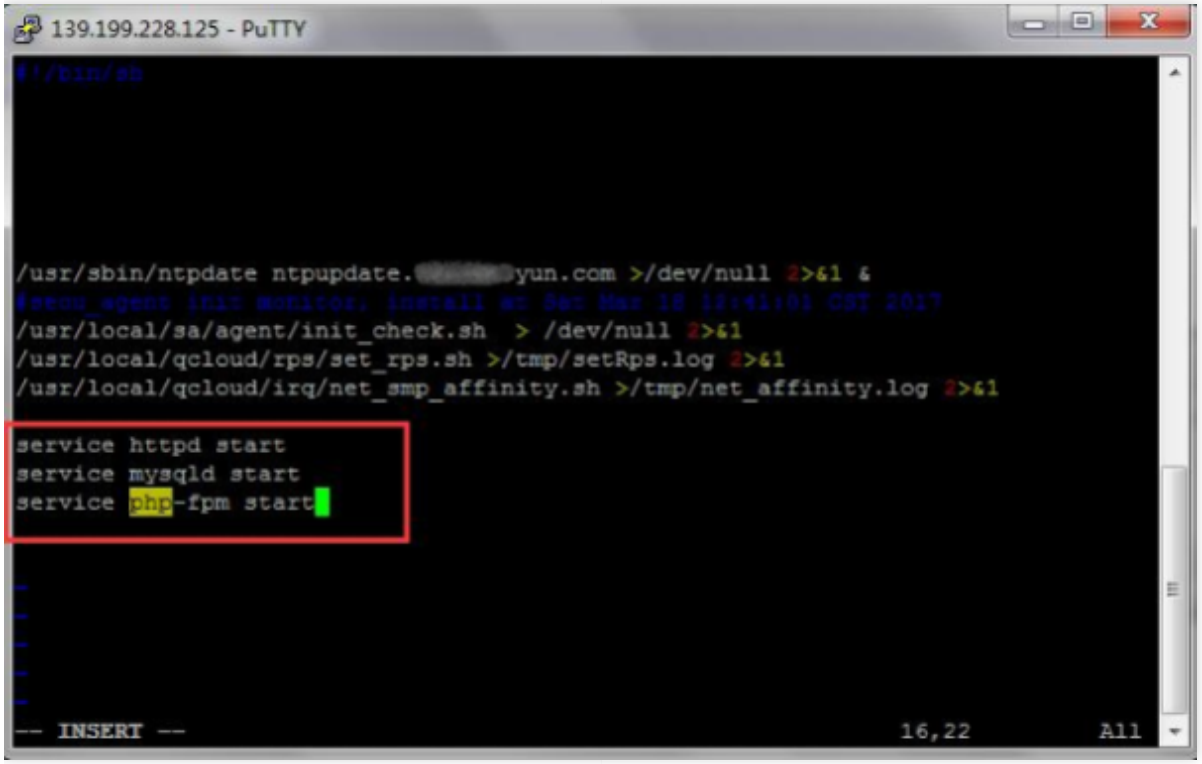
Input vim /etc/rc.d/rc.local
Keep existing content unchanged, and add new content at the end of the file.

> TIPS (Advanced users may skip):
> Enter i to go into vim's insert mode, then you can type the content. At this point, press the ↓ key to move to the end of the file.

Step 2: Specify the services to start

This example is to set up the websites to automatically start httpd, mysqld, php-fpm services upon server boot. Add the following codes at the end of the rc.local file:

```
service httpd start
service mysqld start
service php-fpm start
```



Save and exit. Afterwards, the websites will be accessible automatically once this server boots up. Note that different websites require different services, so customize this step according to your needs.

> TIPS (Advanced users may skip):
> After completing the content entry, press esc, then hold shift + press z twice to exit. That is, enter ZZ.

Step 3: Verification (optional)

Restart the server (you can restart it by entering reboot, or restart it from the console). After the server restarts, do not enter the server; instead, you can directly refresh the website's page to see if there is a response. If there is, then the setting is successful.

Step 4: Create an image based on this machine, and use this image when creating a launch configuration

This step is relatively simple. If you encounter difficulties during the process,see the following tutorials:

Creating custom image

Creating launch configuration

Glossary

Auto Scaling

Auto Scaling is a management service that automatically adjusts elastic computing resources based on user business needs and policies.

Scaling Groups

A Scaling Group is a collection of cloud server instances that follow the same rules and target the same scenario. A scaling group defines the maximum and minimum number of CVM instances in the group, its associated load balancing instances and other attributes.

Launch Configuration

Launch Configuration is a template for automatically creating a CVM, which includes the image ID, CVM instance type, system disk and data disk type and capacity, key pair, security group, etc.

You must specify a launch configuration when you create a scaling group. Once a launch configuration is created, its attributes cannot be edited.

Alarm Scaling

Automatically increase or decrease CVM instances based on cloud monitoring metrics (such as CPU, memory, and network traffic).

Scheduled Action

A Scheduled Scaling Task is a type of auto scaling. When a fixed time point is reached, the CVM instance is automatically increased or decreased, and periodic repetition is supported.

Scaling Policy

Scaling Policy is the condition for executing scaling actions. The trigger condition can be a time or cloud monitoring alarm, and the action can be removing or adding CVM.

Scheduled Scaling Policy

When a fixed time point is reached, the CVM instance is automatically increased or decreased, and periodic repetition is supported.

Cooldown Period

The cooldown period refers to the lock time after a scaling activity is completed in the same scaling group.

API Documentation
Auto Scaling(as) version(2018-04-19)
API overview

API version

V3

Alarm Trigger Policy APIs

Interface name	Interface function
CreateScalingPolicy	CreateScalingPolicy
DeleteScalingPolicy	DeleteScalingPolicy
DescribeScalingPolicies	DescribeScalingPolicies
ExecuteScalingPolicy	ExecuteScalingPolicy
ModifyScalingPolicy	ModifyScalingPolicy

Launch Configuration APIs

Interface name	Interface function
CreateLaunchConfiguration	CreateLaunchConfiguration
DeleteLaunchConfiguration	DeleteLaunchConfiguration
DescribeLaunchConfigurations	DescribeLaunchConfigurations
ModifyLaunchConfigurationAttributes	ModifyLaunchConfigurationAttributes
UpgradeLaunchConfiguration	UpgradeLaunchConfiguration

Lifecycle Hook APIs

Interface name	Interface function
CompleteLifecycleAction	CompleteLifecycleAction
CreateLifecycleHook	CreateLifecycleHook
DeleteLifecycleHook	DeleteLifecycleHook
DescribeLifecycleHooks	DescribeLifecycleHooks
UpgradeLifecycleHook	UpgradeLifecycleHook

Notification APIs

Interface name	Interface function
CreateNotificationConfiguration	CreateNotificationConfiguration
DeleteNotificationConfiguration	DeleteNotificationConfiguration
DescribeNotificationConfigurations	DescribeNotificationConfigurations
ModifyNotificationConfiguration	ModifyNotificationConfiguration

Other APIs

Interface name	Interface function
DescribeAccountLimits	DescribeAccountLimits

Scaling Group APIs

Interface name	Interface function
AttachInstances	AttachInstances
CreateAutoScalingGroup	CreateAutoScalingGroup
DeleteAutoScalingGroup	DeleteAutoScalingGroup

Interface name	Interface function
DescribeAutoScalingActivities	DescribeAutoScalingActivities
DescribeAutoScalingGroupLastActivities	DescribeAutoScalingGroupLastActivities
DescribeAutoScalingGroups	DescribeAutoScalingGroups
DescribeAutoScalingInstances	DescribeAutoScalingInstances
DetachInstances	DetachInstances
DisableAutoScalingGroup	DisableAutoScalingGroup
EnableAutoScalingGroup	EnableAutoScalingGroup
ModifyAutoScalingGroup	ModifyAutoScalingGroup
ModifyDesiredCapacity	ModifyDesiredCapacity
ModifyLoadBalancers	ModifyLoadBalancers
RemoveSelectedInstance	RemoveSelectedInstance
SetInstancesProtection	SetInstancesProtection

Scheduled Task APIs

Interface name	Interface function
CreateScheduledAction	CreateScheduledAction
DeleteScheduledAction	DeleteScheduledAction
DescribeScheduledActions	DescribeScheduledActions
ModifyScheduledAction	ModifyScheduledAction

CallingMethod signaturev1

Converge Cloud authenticates each access request, i.e. each request needs to include authentication information (Signature) in the common parameters to verify the identity of the requester.
The Signature is generated by the security credentials which include SecretId and SecretKey. If you don't have the security credentials yet, go to the API Key page to apply for them; otherwise, you cannot invoke the API.

1. Applying for Security Credentials

Before using the API for the first time, go to the API Key page to apply for security credentials.
Security credentials consist of SecretId and SecretKey:

- SecretId is used to identify the API requester.
- SecretKey is used to encrypt the signature string and verify it on the server.
- You must keep your security credentials private and avoid disclosure.

You can apply for the security credentials through the following steps:

1. Log in to the Converge Cloud Console.
2. Go to the Cloud API Key page.
3. On the API Key Management page, click Create Key to create a SecretId/SecretKey pair.

Note: Each account can have up to two pairs of SecretId/SecretKey.

2. Generating a Signature

With the SecretId and SecretKey, a signature can be generated. The following describes how to generate a signature:

Assume that the SecretId and SecretKey are:

- SecretId: AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE
- SecretKey: Gu5t9xGARNpq86cd98joQYCN3EXAMPLE

Note: This is just an example. For actual operations, please use your own SecretId and SecretKey.

Take the Cloud Virtual Machine's request to view the instance list (DescribeInstances) as an example. When

you invoke this API, the request parameters may be as follows:

Parameter name	Description	Parameter value
Action	Method name	DescribeInstances
SecretId	Key ID	AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE
Timestamp	Current timestamp	1465185768
Nonce	Random positive integer	11886
Region	Region where the instance is located	luzon
InstanceIds.0	ID of the instance to query	ins-09dx96dg
Offset	Offset	0
Limit	Allowed maximum output	20
Version	API version number	2017-03-12

2.1. Sorting Parameters

First, sort all the request parameters in an ascending lexicographical order (ASCII code) by their names. Notes: (1) Parameters are sorted by their names instead of their values; (2) The parameters are sorted based on ASCII code, not in an alphabetical order or by values. For example, InstanceIds.2 should be arranged after InstanceIds.12. You can complete the sorting process using a sorting function in a programming language, such as the ksort function in PHP. The parameters in the example are sorted as follows:

```
{
    'Action' : 'DescribeInstances',
    'InstanceIds.0' : 'ins-09dx96dg',
    'Limit' : 20,
    'Nonce' : 11886,
    'Offset' : 0,
    'Region' : 'luzon',
    'SecretId' : 'AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE',
    'Timestamp' : 1465185768,
    'Version': '2017-03-12',
}
```

When developing in another programming language, you can sort these sample parameters and it will work as long as you obtain the same results.

2.2. Concatenating a Request String

This step generates a request string.
Format the request parameters sorted in the previous step into the form of "parameter name"="parameter value". For example, for the Action parameter, its parameter name is "Action" and its parameter value is "DescribeInstances", so it will become Action=DescribeInstances after formatted.
Note: The "parameter value" is the original value but not the value after URL encoding.

Then, concatenate the formatted parameters with "&". The resulting request string is as follows:

```
Action=DescribeInstances&InstanceIds.0=ins-09dx96dg&Limit=20&Nonce=11886&Offset=0&Region=luzon&SecretId=AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE&Timestamp=1465185768&Version=2017-03-12
```

2.3. Concatenating the Signature Original String

This step generates a signature original string.
The signature original string consists of the following parameters:

1. HTTP method: POST and GET modes are supported, and GET is used here for the request. Please note that the method name should be in all capital letters.
2. Request host: the domain name of the request to view the list of instances (DescribeInstances) is cvm.convergecloud.com. The actual request domain name varies by the module to which the API belongs. For more information, see the instructions of the specific API.
3. Request path: The request path in the current version of Cloud API is fixed to / .
4. Request string: the request string generated in the previous step.

The concatenation rule of the signature original string is: Request method + request host + request path + ? + request string

The concatenation result of the example is:

```
GETcvm.convergecloud.com/?Action=DescribeInstances&InstanceIds.0=ins-09dx96dg&Limit=20&Nonce=11886&Offset=0&Region=luzon&SecretId=AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE&Timestamp=1465185768&Version=2017-03-12
```

2.4. Generating a Signature String

This step generates a signature string.
First, use the HMAC-SHA1 algorithm to sign thesignature original stringobtained in the previous step, and then encode the generated signature using Base64 to obtain the final signature.

The specific code is as follows with the PHP language being used as an example:

```
$secretKey = 'Gu5t9xGARNpq86cd98joQYCN3EXAMPLE';
$srcStr = 'GETcvm.convergecloud.com/?Action=DescribeInstances&InstanceIds.0=ins-09dx96dg&Limit
=20&Nonce=11886&Offset=0&Region=luzon&SecretId=AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPL
E&Timestamp=1465185768&Version=2017-03-12';
$signStr = base64_encode(hash_hmac('sha1', $srcStr, $secretKey, true));
echo $signStr;
```

The final signature is:

```
EliP9YW3pW28FpsEdkXt/+WcGeI=
```

When developing in another programming language, you can sign and verify the original in the example above and it works as long as you get the same results.

3. Encoding a Signature String

The generated signature string cannot be directly used as a request parameter and must be URL encoded.

For example, if the signature string generated in the previous step is EliP9YW3pW28FpsEdkXt/+WcGeI= , the final signature string request parameter (Signature) is EliP9YW3pW28FpsEdkXt%2f%2bWcGeI%3d, which will be used to generate the final request URL.

**Note: If your request method is GET, or the request method is POST and the Content-Type is application/x-www-form-urlencoded, then all the request parameter values need to be URL encoded (except the parameter key and the symbol of =) when sending the request. Non-ASCII characters need to be encoded with UTF-8 before URL encoding. **

Note: The network libraries of some programming languages automatically URL encode all parameters, in which case there is no need to URL encode the signature string; otherwise, two rounds of URL encoding will cause the signature to fail.

Note: Other parameter values also need to be encoded using RFC 3986. Use %XY in percent-encoding for special characters such as Chinese characters, where "X" and "Y" are hexadecimal characters (0-9 and uppercase A-F), and using lowercase will cause an error.

4. Signature Failure

The following situational error codes for signature failure may occur. Please resolve the errors accordingly.

Error code	Error description
AuthFailure.SignatureExpire	The signature is expired

Error code	Error description
AuthFailure.SecretIdNotFound	The key does not exist
AuthFailure.SignatureFailure	Signature error
AuthFailure.TokenFailure	Token error
AuthFailure.InvalidSecretId	Invalid key (not a Cloud API key type)

5. Signature Demo

When calling API 3.0, you are recommended to use the corresponding Converge Cloud SDK 3.0, Swhich encapsulates the signature process, enabling you to focus on only the specific APIs provided by the product when developing. See SDK Center for more information. Currently, the following programming languages are supported:

- Python
- Java
- PHP
- Go
- Node

To further explain the signing process, we will use a programming language to implement the process described above. The request domain name, API and parameter values in the sample are used here. This goal of this example is only to provide additional clarification for the signature process, please see the SDK for actual usage.

The final output URL might be: https://cvm.convergecloud.com/?Action=DescribeInstances&InstanceIds.0=ins-09dx96dg&Limit=20&Nonce=11886&Offset=0&Region=luzon&SecretId=AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE&Signature=EliP9YW3pW28FpsEdkXt%2F%2BWcGeI%3D&Timestamp=1465185768&Version=2017-03-12

Note: The key in the example is fictitious, and the timestamp is not the current time of the system, so if this URL is opened in the browser or called using commands such as curl, an authentication error will be returned: Signature expired. In order to get a URL that can work properly, you need to replace the SecretId and SecretKey in the example with your real credentials and use the current time of the system as the Timestamp.

Note: In the example below, even if you use the same programming language, the order of the parameters in the URL may be different for each execution. However, the order does not matter, as long as all the parameters are included in the URL and the signature is calculated correctly.

Note: The following code is only applicable to API 3.0. It cannot be directly used in other signature processes. Even with an older API, signature calculation errors may occur due to the differences in details. Please refer to the corresponding documentation.

Java

```
import java.io.UnsupportedEncodingException;
import java.net.URLEncoder;
import java.util.Random;
import java.util.TreeMap;
import javax.crypto.Mac;
import javax.crypto.spec.SecretKeySpec;
import javax.xml.bind.DatatypeConverter;

public class CloudAPIDemo {
    private final static String CHARSET = "UTF-8";

    public static String sign(String s, String key, String method) throws Exception {
        Mac mac = Mac.getInstance(method);
        SecretKeySpec secretKeySpec = new SecretKeySpec(key.getBytes(CHARSET), mac.getAlgorithm());
        mac.init(secretKeySpec);
        byte[] hash = mac.doFinal(s.getBytes(CHARSET));
        return DatatypeConverter.printBase64Binary(hash);
    }

    public static String getStringToSign(TreeMap<String, Object> params) {
        StringBuilder s2s = new StringBuilder("GETcvm.convergecloud.com/?");
        // When signing, the parameters need to be sorted in lexicographical order. TreeMap is used here
to guarantee the correct order.
        for (String k : params.keySet()) {
            s2s.append(k).append("=").append(params.get(k).toString()).append("&");
        }
        return s2s.toString().substring(0, s2s.length() - 1);
    }

    public static String getUrl(TreeMap<String, Object> params) throws UnsupportedEncodingException
{
        StringBuilder url = new StringBuilder("https://cvm.convergecloud.com/?");
        // There is no requirement for the order of the parameters in the actual request URL.
        for (String k : params.keySet()) {
            // The request string needs to be URL encoded. As the Key is all in English letters, only the valu
e is URL encoded here.
            url.append(k).append("=").append(URLEncoder.encode(params.get(k).toString(), CHARSET)).app
end("&");
        }
        return url.toString().substring(0, url.length() - 1);
```

```
    }

    public static void main(String[] args) throws Exception {
        TreeMap<String, Object> params = new TreeMap<String, Object>(); // TreeMap enables automati
c sorting
        // A random number should be used when actually calling, for example: params.put("Nonce", new
Random().nextInt(java.lang.Integer.MAX_VALUE));
        params.put("Nonce", 11886); // Common parameter
        // The current time of the system should be used when actually calling, for example:   params.put
("Timestamp", System.currentTimeMillis() / 1000);
        params.put("Timestamp", 1465185768); // Common parameter
        params.put("SecretId", "AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE"); // Common parameter
        params.put("Action", "DescribeInstances"); // Common parameter
        params.put("Version", "2017-03-12"); // Common parameter
        params.put("Region", "luzon"); // Common parameter
        params.put("Limit", 20); // Business parameter
        params.put("Offset", 0); // Business parameter
        params.put("InstanceIds.0", "ins-09dx96dg"); // Business parameter
        params.put("Signature", sign(getStringToSign(params), "Gu5t9xGARNpq86cd98joQYCN3EXAMPLE
", "HmacSHA1")); // Common parameter
        System.out.println(getUrl(params));
    }
}
```

Python

Note: If running in a Python 2 environment, the following requests dependency package must be installed first: pip install requests .

```
# -*- coding: utf8 -*-
import base64
import hashlib
import hmac
import time

import requests

secret_id = "AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE"
secret_key = "Gu5t9xGARNpq86cd98joQYCN3EXAMPLE"

def get_string_to_sign(method, endpoint, params):
    s = method + endpoint + "/?"
    query_str = "&".join("%s=%s" % (k, params[k]) for k in sorted(params))
    return s + query_str

def sign_str(key, s, method):
```

```
    hmac_str = hmac.new(key.encode("utf8"), s.encode("utf8"), method).digest()
    return base64.b64encode(hmac_str)

if __name__ == '__main__':
    endpoint = "cvm.convergecloud.com"
    data = {
        'Action' : 'DescribeInstances',
        'InstanceIds.0' : 'ins-09dx96dg',
        'Limit' : 20,
        'Nonce' : 11886,
        'Offset' : 0,
        'Region' : 'luzon',
        'SecretId' : secret_id,
        'Timestamp' : 1465185768, # int(time.time())
        'Version': '2017-03-12'
    }
    s = get_string_to_sign("GET", endpoint, data)
    data["Signature"] = sign_str(secret_key, s, hashlib.sha1)
    print(data["Signature"])
    # n actual invocation would occur here, which may incur fees after success
    # resp = requests.get("https://" + endpoint, params=data)
    # print(resp.url)
```

signaturev3

Converge Cloud The API will authenticate each access request, meaning that each request must include signature information (Signature) in the public request parameters to verify the requester's identity.
The signature information is generated by the security credentials, which include SecretId and SecretKey . If the user does not have security credentials, please apply for them on the Cloud API Key page, otherwise they cannot call the Cloud API interfaces.

1. Apply for security credentials.

Before using the Cloud API for the first time, please go to the Cloud API Key page to apply for a security credential.
Security credentials include SecretId and SecretKey:

- SecretId The security credentials are used to identify the identity of the API caller.
- SecretKey The secret key is used to encrypt and sign the string for server-side verification of the signature.
- Users must strictly protect their security credentials and avoid disclosure them.

The specific steps to apply for security credentials are as follows:

1. Log in to Converge Cloud Administration Center Console
2. Go to the console page for cloud API keys
3. On the Cloud API Key page, click New to create a pair of SecretId/SecretKey

Note: Developer accounts can have up to two pairs SecretId / SecretKey.

2. TC3-HMAC-SHA256 signature method

Note: For GET method, only Content-Type: application/x-www-form-urlencoded protocol format is supported. For POST method, Content-Type: application/json and Content-Type: multipart/form-data are currently supported. By default, all business interfaces are supported in json format. Only specific business interfaces are supported in multipart format. In this case, the interface cannot be called in json format. Refer to the specific business interface document description.

The following takes Cloud Virtual Machine querying Region-1 instance list as an example to introduce the signature calculation process step by step. We only use two parameters for the query instance list: Limit and Offset, called using the GET method.

Assume that the SecretId and SecretKey of the user are AKIDz 8krbsJ 5yKBZQpn 74WFkmLPx 3EXAMPLE and Gu 5t9xGARNpq 86cd 98joQYCN3EXAMPLE, respectively.

2.1. splice specification request string

Concatenate CanonicalRequest strings in the following format:

```
CanonicalRequest =
  HTTPRequestMethod + '\n' +
  CanonicalURI + '\n' +
  CanonicalQueryString + '\n' +
  CanonicalHeaders + '\n' +
  SignedHeaders + '\n' +
  HashedRequestPayload
```

- HTTPRequestMethod: HTTP request methods (GET, POST), GET in this example;
- CanonicalURI: URI parameter, API 3.0 fixed to forward slash (/);
- CanonicalQueryString: Query string in the URL that initiates the HTTP request. For POST requests, it is fixed as an empty string. For GET requests, it is the question mark (?) in the URL. The following string content, this example value is: Limit=10&Offset=0. Note: CanonicalQueryString needs to be URL encoded.
- CanonicalHeaders: The header information of participating signature, at least including host and content-type, can also be added with customized header to improve the uniqueness and security of its own request. Splicing rules: 1) head key and value are uniformly converted to lowercase, and the first and last spaces are removed, and spliced according to the format of key:value\n;2) multiple heads are spliced according to the dictionary sorting of head key (lowercase). In this example: content-type: application/x-www-form-urlencoded\nhost: cvm.convergecloud.com\n
- SignedHeaders: The header information involved in the signature indicates which headers participate in the signature of this request, and the header content contained in CanonicalHeaders is one-to-one corresponding. content-type and host are required headers. constructed as: 1) The head key is uniformly converted to lowercase; 2) Multiple header keys (lowercase) are concatenated according to dictionary sorting and separated by semicolons (;). In this case: content-type;host
- HashedRequestPayload: The hash value of the request body, calculated as Lowercase(HexEncode(Hash.SHA256(RequestPayload))), SHA256 hash the entire payload of the HTTP request body, then hexadecimal encoding, and finally converting the encoding string to lowercase letters. Note: For GET requests, RequestPayload is fixed as an empty string. For POST requests, RequestPayload is the payload of the HTTP request body.

According to the above rules, the resulting canonical request string in the example is as follows (for clarity, \nnewlines are replaced by a new line printed separately):

```
GET
```

```
/
Limit=10&Offset=0
content-type:application/x-www-form-urlencoded
host:cvm.convergecloud.com

content-type;host
e3b0c44298fc1c149afbf4c8996fb92427ae41e4649b934ca495991b7852b855
```

2.2. Concatenate the string to be signed

Concatenate the string to be signed in the following format:

```
StringToSign =
  Algorithm + \n +
  RequestTimestamp + \n +
  CredentialScope + \n +
  HashedCanonicalRequest
```

- Algorithm: Signature algorithm, currently fixed as TC3-HMAC-SHA256;
- RequestTimestamp: Request timestamp, i.e. the X-TC-Timestamp value of the request header. The request in the example above is 1539084154;
- CredentialScope: Credential range in the format Date/service/tc3_request, containing date, requested service and termination string (tc3_request). **Date is the date of UTC standard time, the value needs to be consistent with UTC standard time converted by common parameter X-TC-Timestamp **;service is the product name, which must be consistent with the product domain name called, such as cvm. The above example request, the value is 2018-10-09/cvm/tc3_request;
- HashedCanonicalRequest: The hash value of the canonical request string obtained by concatenating the above steps is calculated by Lowercase(HexEncode(Hash.SHA256(CanonicalRequest))).

note:

> 1. Date must be calculated from timestamp X-TC-Timestamp and time zone UTC+0. Adding local time zone information to the system, such as East Zone 8, will cause calls to succeed during the day and night, but calls to fail in the early hours of the morning. Assuming the timestamp is 1551113065, the time in the East Eight District is 2019-02-26 00:44:25, but the calculated Date with UTC+0 should be 2019-02-25 instead of 2019-02-26.
> 2. Timestamp must be the current system time and ensure that the system time and standard time are synchronized. If the difference exceeds five minutes, it will fail. If it is not synchronized with the standard time for a long time, it may cause the request to fail after running for a period of time (return signature expiration error).

According to the above rules, the string to be signed in the example is as follows (for clarity,\nnewlines are

replaced by new lines printed):

```
TC3-HMAC-SHA256
1539084154
2018-10-09/cvm/tc3_request
91c9c192c14460df6c1ffc69e34e6c5e90708de2a6d282cccf957dbf1aa7f3a7
```

2.3. calculate the signature

1. Compute the derived signature key, pseudocode as follows

```
SecretKey = "Gu5t9xGARNpq86cd98joQYCN3EXAMPLE"
SecretDate = HMAC_SHA256("TC3" + SecretKey, Date)
SecretService = HMAC_SHA256(SecretDate, Service)
SecretSigning = HMAC_SHA256(SecretService, "tc3_request")
```

- SecretKey: Original SecretKey;
- Date: That is, the Date field information in Credential, as shown in the example above, is 2018-10-09.;
- Service: That is, the Service field information in Credential, as shown in the example above, is cvm;

2. Calculate the signature, pseudocode as follows

```
Signature = HexEncode(HMAC_SHA256(SecretSigning, StringToSign))
```

- SecretSigning: That is, the derivative signature key calculated above;
- StringToSign: That is, the character string to be signed calculated in step 2;

2.4. Splicing Authorization

Splice as follows Authorization:

```
Authorization =
    Algorithm + ' ' +
    'Credential=' + SecretId + '/' + CredentialScope + ', ' +
    'SignedHeaders=' + SignedHeaders + ', '
    'Signature=' + Signature
```

- Algorithm: Signature method, fixed as TC3-HMAC-SHA256;
- SecretId: SecretId in key pair;
- CredentialScope: See above, Scope of Voucher;
- SignedHeaders: See above, header information for participating signatures;
- Signature: signature value

According to the above rules, the values obtained in the example are:

```
TC3-HMAC-SHA256 Credential=AKIDEXAMPLE/Date/service/tc3_request, SignedHeaders=content-type;host, Signature=5da7a33f6993f0614b047e5df4582db9e9bf4672ba50567dba16c6ccf174c474
```

The final complete invocation information is as follows:

```
https://cvm.convergecloud.com/?Limit=10&Offset=0

Authorization: TC3-HMAC-SHA256 Credential=AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE/2018-10-09/cvm/tc3_request, SignedHeaders=content-type;host, Signature=5da7a33f6993f0614b047e5df4582db9e9bf4672ba50567dba16c6ccf174c474
Content-Type: application/x-www-form-urlencoded
Host: cvm.convergecloud.com
X-TC-Action: DescribeInstances
X-TC-Version: 2017-03-12
X-TC-Timestamp: 1539084154
X-TC-Region: luzon
```

3. signing failed

According to the actual situation, there are the following signature failure error codes, please handle according to the actual situation

error code	error description
AuthFailure.SignatureExpire	signature expiration
AuthFailure.SecretIdNotFound	key is not present
AuthFailure.SignatureFailure	signature error
AuthFailure.TokenFailure	token error
AuthFailure.InvalidSecretId	Illegal key (not cloud API key type)

4. Signature Demo

Java

```
import java.io.BufferedReader;
import java.io.InputStream;
```

```
import java.io.InputStreamReader;
import java.net.URL;
import java.text.SimpleDateFormat;
import java.util.Date;
import java.util.Map;
import java.util.TimeZone;
import java.util.TreeMap;
import javax.crypto.Mac;
import javax.crypto.spec.SecretKeySpec;
import javax.net.ssl.HttpsURLConnection;
import javax.xml.bind.DatatypeConverter;

import org.apache.commons.codec.digest.DigestUtils;

public class CloudAPITC3Demo {
    private final static String CHARSET = "UTF-8";
    private final static String ENDPOINT = "cvm.convergecloud.com";
    private final static String PATH = "/";
    private final static String SECRET_ID = "AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE";
    private final static String SECRET_KEY = "Gu5t9xGARNpq86cd98joQYCN3EXAMPLE";
    private final static String CT_X_WWW_FORM_URLENCODED = "application/x-www-form-urlencode
d";
    private final static String CT_JSON = "application/json";
    private final static String CT_FORM_DATA = "multipart/form-data";

    public static byte[] sign256(byte[] key, String msg) throws Exception {
        Mac mac = Mac.getInstance("HmacSHA256");
        SecretKeySpec secretKeySpec = new SecretKeySpec(key, mac.getAlgorithm());
        mac.init(secretKeySpec);
        return mac.doFinal(msg.getBytes(CHARSET));
    }

    public static void main(String[] args) throws Exception {
        String service = "cvm";
        String host = "cvm.convergecloud.com";
        String region = "luzon";
        String action = "DescribeInstances";
        String version = "2017-03-12";
        String algorithm = "TC3-HMAC-SHA256";
        String timestamp = "1539084154";
        //String timestamp = String.valueOf(System.currentTimeMillis() / 1000);
        SimpleDateFormat sdf = new SimpleDateFormat("yyyy-MM-dd");
        // Pay attention to the time zone, otherwise it is easy to make mistakes.
        sdf.setTimeZone(TimeZone.getTimeZone("UTC"));
        String date = sdf.format(new Date(Long.valueOf(timestamp + "000")));

        // ************* Step 1: Assemble the specification request string *************
```

```
        String httpRequestMethod = "GET";
        String canonicalUri = "/";
        String canonicalQueryString = "Limit=10&Offset=0";
        String canonicalHeaders = "content-type:application/x-www-form-urlencoded\n" + "host:" + host
+ "\n";
        String signedHeaders = "content-type;host";
        String hashedRequestPayload = DigestUtils.sha256Hex("");
        String canonicalRequest = httpRequestMethod + "\n" + canonicalUri + "\n" + canonicalQueryStri
ng + "\n"
                + canonicalHeaders + "\n" + signedHeaders + "\n" + hashedRequestPayload;
        System.out.println(canonicalRequest);

        // ************* Step 2: Stitching the string to be signed *************
        String credentialScope = date + "/" + service + "/" + "tc3_request";
        String hashedCanonicalRequest = DigestUtils.sha256Hex(canonicalRequest.getBytes(CHARSET));
        String stringToSign = algorithm + "\n" + timestamp + "\n" + credentialScope + "\n" + hashedCan
onicalRequest;
        System.out.println(stringToSign);

        // ************* Step 3: Calculate the signature *************
        byte[] secretDate = sign256(("TC3" + SECRET_KEY).getBytes(CHARSET), date);
        byte[] secretService = sign256(secretDate, service);
        byte[] secretSigning = sign256(secretService, "tc3_request");
        String signature = DatatypeConverter.printHexBinary(sign256(secretSigning, stringToSign)).toLow
erCase();
        System.out.println(signature);

        // ************* Step 4: Splicing Authorization *************
        String authorization = algorithm + " " + "Credential=" + SECRET_ID + "/" + credentialScope + ", "
                + "SignedHeaders=" + signedHeaders + ", " + "Signature=" + signature;
        System.out.println(authorization);

        TreeMap<String, String> headers = new TreeMap<String, String>();
        headers.put("Authorization", authorization);
        headers.put("Host", host);
        headers.put("Content-Type", CT_X_WWW_FORM_URLENCODED);
        headers.put("X-TC-Action", action);
        headers.put("X-TC-Timestamp", timestamp);
        headers.put("X-TC-Version", version);
        headers.put("X-TC-Region", region);
    }
}
```

Python

```
# -*- coding: utf-8 -*-
```

```
import hashlib, hmac, json, os, sys, time
from datetime import datetime

# key parameters
secret_id = "AKIDz8krbsJ5yKBZQpn74WFkmLPx3EXAMPLE"
secret_key = "Gu5t9xGARNpq86cd98joQYCN3EXAMPLE"

service = "cvm"
host = "cvm.convergecloud.com"
endpoint = "https://" + host
region = "luzon"
action = "DescribeInstances"
version = "2017-03-12"
algorithm = "TC3-HMAC-SHA256"
timestamp = 1539084154
date = datetime.utcfromtimestamp(timestamp).strftime("%Y-%m-%d")
params = {"Limit": 10, "Offset": 0}

# ************* Step 1: Assemble the specification request string *************
http_request_method = "GET"
canonical_uri = "/"
canonical_querystring = "Limit=10&Offset=0"
ct = "x-www-form-urlencoded"
payload = ""
if http_request_method == "POST":
    canonical_querystring = ""
    ct = "json"
    payload = json.dumps(params)
canonical_headers = "content-type:application/%s\nhost:%s\n" % (ct, host)
signed_headers = "content-type;host"
hashed_request_payload = hashlib.sha256(payload.encode("utf-8")).hexdigest()
canonical_request = (http_request_method + "\n" +
                     canonical_uri + "\n" +
                     canonical_querystring + "\n" +
                     canonical_headers + "\n" +
                     signed_headers + "\n" +
                     hashed_request_payload)
print(canonical_request)

# ************* Step 2: Stitching the string to be signed *************
credential_scope = date + "/" + service + "/" + "tc3_request"
hashed_canonical_request = hashlib.sha256(canonical_request.encode("utf-8")).hexdigest()
string_to_sign = (algorithm + "\n" +
                  str(timestamp) + "\n" +
                  credential_scope + "\n" +
                  hashed_canonical_request)
print(string_to_sign)
```

```
# ************* Step 3: Calculate the signature *************
# Compute signature digest function
def sign(key, msg):
    return hmac.new(key, msg.encode("utf-8"), hashlib.sha256).digest()
secret_date = sign(("TC3" + secret_key).encode("utf-8"), date)
secret_service = sign(secret_date, service)
secret_signing = sign(secret_service, "tc3_request")
signature = hmac.new(secret_signing, string_to_sign.encode("utf-8"), hashlib.sha256).hexdigest()
print(signature)

# ************* Step 4: Splicing Authorization *************
authorization = (algorithm + " " +
                 "Credential=" + secret_id + "/" + credential_scope + ", " +
                 "SignedHeaders=" + signed_headers + ", " +
                 "Signature=" + signature)
print(authorization)

# Common parameters added to request header
headers = {
    "Authorization": authorization,
    "Host": host,
    "Content-Type": "application/%s" % ct,
    "X-TC-Action": action,
    "X-TC-Timestamp": str(timestamp),
    "X-TC-Version": version,
    "X-TC-Region": region,
}
```

request

1. Service Address

Region refers to the geographical area of a physical data center. Converge Cloud ensures complete isolation between different regions, providing maximum stability and fault tolerance. To reduce access latency and improve download speed, it is recommended to choose the region closest to your customers.

You can view the complete region list by using the API interface "DescribeRegions".

2. Communications Protocol

All the Converge Cloud APIs communicate via HTTPS, providing highly secure communication tunnels.

3. Request Methods

Supported HTTP request methods:

- POST (recommended)
- GET

PThe Content-Type types supported by POST requests:

- application/json (recommended). The TC3-HMAC-SHA256 signature algorithm must be used.
- application/x-www-form-urlencoded. The HmacSHA1 or HmacSHA256 signature algorithm must be used.
- multipart/form-data (only supported by certain APIs). You must use TC3-HMAC-SHA256 to calculate the signature.

The size of a GET request packet is up to 32 KB. The size of a POST request is up to 1 MB when the HmacSHA1 or HmacSHA256 signature algorithm is used, and up to 10 MB when TC3-HMAC-SHA256 is used.

4. Character Encoding

Only UTF-8 encoding is used.

result

Response for Successful Requests

For example, when calling CAM API (version: 2017-03-12) to view the status of instances (DescribeInstancesStatus), if the request has succeeded, you may see the response as shown below:

```
{
    "Response": {
        "TotalCount": 0,
        "InstanceStatusSet": [],
        "RequestId": "b5b41468-520d-4192-b42f-595cc34b6c1c"
    }
}
```

- The API will return Response, which contains RequestId, as long as it processes the request. It does not matter if the request is successful or not.
- RequestId is the unique ID of an API request. Contact us with this ID when an exception occurs.
- Except for the fixed fields, all fields are action-specified. For the definitions of action-specified fields, see the corresponding API documentation. In this example, TotalCount and InstanceStatusSet are the fields specified by the API DescribeInstancesStatus. 0 TotalCount means that the requester owns 0 CVM instance so the InstanceStatusSet is empty.

Response for Failed Requests

If the request has failed, you may see the response as shown below:

```
{
    "Response": {
        "Error": {
            "Code": "AuthFailure.SignatureFailure",
            "Message": "The provided credentials could not be validated. Please check your signature is cor
rect."
        },
        "RequestId": "ed93f3cb-f35e-473f-b9f3-0d451b8b79c6"
    }
}
```

- The presence of the Error field indicates that the request has failed. A response for a failed request will include Error, Code and Message fields.

- Code is the code of the error that helps you identify the cause and solution. There are two types of error codes so you may find the code in either common error codes or API-specified error codes.
- Message explains the cause of the error. Note that the returned messages are subject to service updates. The information the messages provide may not be up-to-date and should not be the only source of reference.
- RequestId is the unique ID of an API request. Contact us with this ID when an exception occurs.

Common Error Codes

If there is an Error field in the response, it means that the API call failed. The Code field in Error indicates the error code. The following table lists the common error codes that all actions can return.

Error Code	Description
AuthFailure.InvalidSecretId	Invalid key (not a API key type).
AuthFailure.MFAFailure	MFA failed.
AuthFailure.SecretIdNotFound	The key does not exist.
AuthFailure.SignatureExpire	Signature expired.
AuthFailure.SignatureFailure	Signature error.
AuthFailure.TokenFailure	Token error.
AuthFailure.UnauthorizedOperation	The request does not have CAM authorization.
DryRunOperation	DryRun Operation. It means that the request would have succeeded, but the DryRun parameter was used.
FailedOperation	Operation failed.
InternalError	Internal error.
InvalidAction	The API does not exist.
InvalidParameter	Incorrect parameter.
InvalidParameterValue	Invalid parameter value.
LimitExceeded	Quota limit exceeded.
MissingParameter	A parameter is missing.
NoSuchVersion	The API version does not exist.

Error Code	Description
RequestLimitExceeded	The number of requests exceeds the frequency limit.
ResourceInUse	Resource is in use.
ResourceInsufficient	Insufficient resource.
ResourceNotFound	The resource does not exist.
ResourceUnavailable	Resource is unavailable.
UnauthorizedOperation	Unauthorized operation.
UnknownParameter	Unknown parameter.
UnsupportedOperation	Unsupported operation.
UnsupportedProtocol	hHTTPS request method error. Only GET and POST requests are supported.
UnsupportedRegion	API does not support the requested region.

PublicParameter

Common parameters are parameters used to identify users and interface authentication purposes. If not necessary, these parameters are not described in the separate interface document of each interface, but these parameters need to be carried in each request before the request can be normally initiated.

signature method v3

When using TC3-HMAC-SHA256 signature method, the common parameters need to be uniformly placed in the HTTP Header request header, as follows:

parameter name	type	required	described
X-TC-Action	String	Yes	The name of the operation interface. Refer to the description of the public parameter Action in the input parameters of the interface documentation for the value. For example, the value for the query instance list interface of the cloud server is DescribeInstances.
X-TC-Region	String	Yes	region parameters, Data used to identify which region you wish to operate on. The region value accepted by the interface refers to the description of the input parameter common parameter Region in the interface document. Note: Some interfaces do not need to pass this parameter, which is specifically stated in the interface documentation. In this case, even passing this parameter will not take effect.
X-TC-Timestamp	Integer	Yes	The current UNIX timestamp, which records the time when the API request was initiated. for example 1529223702. Note: If the time difference from the server exceeds 5 minutes, it will cause a signature expiration error.
X-TC-Version	String	Yes	Version of API for operation. Refer to the description of the public parameter Version in the interface document for the value. For example, the version of the cloud server. 2017-03-12.
Authorization	String	Yes	HTTP The standard authentication header field is "Authorization". For example: TC3-HMAC-SHA256 Credential=AKIDEXAMPLE/Date/service/tc3_request, SignedHeaders=content-type;host, Signature=fe5f80f77d5fa3beca038a248ff027d0445342fe2855ddc963176630326f1024 Here, - TC3-HMAC-SHA256: Signature method, currently fixed as this value, - Credential: Credential: Signature credential; AKIDEXAMPLE is the SecretId; Date is a date in UTC time, and this value must match the value of X-TC-Timestamp (a common parameter) in UTC time format; service is the name of the product/service, and is generally a domain name prefix. For example,cvm; - SignedHeaders: The headers that contains the authentication information; content-type and host are the required headers; - Signature: Signature digest.
X-TC-Token	String	No	The token used for temporary certificates needs to be used in conjunction with a temporary key. Temporary keys and tokens need to be obtained by calling the access management service interface. Long-term keys do not require tokens.

signature method v1

When using HmacSHA1 and HmacSHA256 signature methods, common parameters need to be placed uniformly in the request string, as follows

parameter name	type	required	described
Action	String	Yes	The name of the interface to be operated. The value refers to the explanation of the public parameter Action in the interface documentation. For example, for the cloud server's query instance list interface, the value is DescribeInstances .
Region	String	Yes	The region parameter is used to identify which region's data you want to operate on. The region value accepted by the interface is specified in the input parameter of the public parameter Region in the interface documentation. Note: Some interfaces do not require this parameter to be passed, which will be specifically mentioned in the interface documentation. Even if this parameter is passed, it will not take effect in this case.
Timestamp	Integer	Yes	The current UNIX timestamp, which can record the time when the API request is initiated. For example, 1529223702. If the difference between the current time is too large, it will cause a signature expiration error.
Nonce	Integer	Yes	A random positive integer combined with Timestamp is used to prevent replay attacks.
SecretId	String	Yes	The SecretId is an identifier applied for the cloud API key to identify the user's identity. A SecretId corresponds to a unique SecretKey , and the SecretKey is used to generate the request signature Signature .
Signature	String	Yes	Request signature is used to verify the legitimacy of this request. Users need to calculate it based on the actual input parameters. The specific calculation method can be found in the interface authentication documentation.
Version	String	Yes	The version of the API being operated on. The value refers to the description of the public parameter Version in the interface documentation. For example, the version of the cloud server. 2017-03-12.
SignatureMethod	String	No	Signature method, currently supports HmacSHA256 and HmacSHA1. HmacSHA256 algorithm is used to verify signatures only if this parameter is specified as HmacSHA 256, otherwise HmacSHA1 is used to verify signatures.
Token	String	No	Token used for temporary certificate needs to be used together with temporary key. The temporary key and Token need to be obtained from the Access Management service invocation interface. Long-term keys do not require tokens.

zone list

A region (Region) refers to the geographical area of a physical data center. Converge Cloud Translation: To deliver verification that completely isolates different regions, ensuring maximum stability and fault tolerance between regions. In order to reduce access latency and improve download speed, it is recommended that you choose the region closest to your customers.

You can view the complete region list by using the API interface Query Region List.

Alarm Trigger Policy APIs
CreateScalingPolicy

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (CreateScalingPolicy) is used to create an alarm trigger policy.

Default interface request rate limit: 20 times per second.

Interface update time: 2022-11-04 09:58:02.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: CreateScalingPolicy
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupId	Yes	No	String	ID of the scaling group. Example: asg-12wjuh0s
ScalingPolicyName	Yes	No	String	Name of the alarm trigger policy. Example: cpu_policy_
AdjustmentType	Yes	No	String	Method that you want to use to change the desired capacity after the alarm is triggered. Valid values: • CHANGE_IN_CAPACITY: Increase or decrease the desired capacity. • EXACT_CAPACITY: Adjust to the specified desired capacity. • PERCENT_CHANGE_IN_CAPACITY: Adjust the desired capacity by percentage. Example: CHANGE_IN_CAPACITY
AdjustmentValue	Yes	No	Int64	Adjusted desired capacity after the alarm is triggered. Valid values: • When the AdjustmentType parameter is set to CHANGE_IN_CAPACITY, some new instances will be added after the alarm is triggered if the AdjustmentValue parameter is set to a positive value, and some existing instances will be removed after the alarm is triggered if the AdjustmentValue parameter is set to a negative value. • When the AdjustmentType parameter is set to EXACT_CAPACITY, the value of the AdjustmentValue parameter specifies the desired number of instances after the alarm is triggered, which must be equal to or greater than 0. • When the AdjustmentType parameter is set to PERCENT_CHANGE_IN_CAPACITY, new instances will be added by percentage after the alarm is triggered if the AdjusmentValue parameter is set to a positive value, and existing instances will be removed by percentage after the alarm is triggered if the AdjusmentValue parameter is set to a negative value. Unit: %. Example: 1
Cooldown	No	No	Uint64	Cooldown period. Unit: seconds. Default value: 300. Example: 60
MetricAlarm	Yes	No	MetricAlarm	Alarm metric. Example:View
NotificationUserGroupIds	No	No	Array of String	Notification group ID, a collection of user group IDs. User group IDs can be queried through DescribeUserGroup. Example: ["1678\r"]

3. Output Parameters

Parameter name	Type	Description
AutoScalingPolicyId	String	ID of the alarm trigger policy. Example: asp-iir70sxv
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
ResourceNotFound.AutoScalingGroupIdNotFound	The scaling group does not exist.
InvalidParameterValue.ScalingPolicyNameDuplicate	The alarm policy name already exists.
InvalidParameterValue.LimitExceeded	The value exceeds the limit.
InvalidParameterValue.UserGroupIdNotFound	The user group does not exist.
InternalError.RequestError	Request error

DeleteScalingPolicy

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DeleteScalingPolicy) is used to delete an alarm trigger policy.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-11-01 11:10:52.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DeleteScalingPolicy
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingPolicyId	Yes	No	String	ID of the alarm policy that you want to delete. Example: asp-lnfqbsoq

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
ResourceNotFound.ScalingPolicyNotFound	The alarm policy does not exist.

DescribeScalingPolicies

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DescribeScalingPolicies) is used to query alarm trigger policies.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 18:33:02.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeScalingPolicies
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingPolicyIds	No	No	Array of String	ID of one or more alarm policies, such as asp-i9vkg894 . You can specify a maximum of 100 instance IDs in each request. AutoScalingPolicyIds and Filters cannot be specified at the same time. Example: ["asp-5zffv598\n"]
Filters	No	No	Array of Filter	Filter conditions. • auto-scaling-policy-id - String - Required: No - (Filter condition) Filter by alarm policy ID. • auto-scaling-group-id - String - Required: No - (Filter condition) Filter by scaling group ID. • scaling-policy-name - String - Required: No - (Filter condition) Filter by alarm policy name. The maximum number of Filters per request is 10. The upper limit for Filter.Values is 5. AutoScalingPolicyIds and Filters cannot be specified at the same time. Example:View
Limit	No	No	Uint64	Number of returned items. Default value: 20. Maximum value: 100. For more information on Limit , see the relevant section in API Introduction. Example: 20
Offset	No	No	Uint64	Offset. Default value: 0. For more information on Offset , see the relevant section in API Introduction. Example: 0

3. Output Parameters

Parameter name	Type	Description
ScalingPolicySet	Array of ScalingPolicy	Details of Auto Scaling alarm trigger policies. Example:View
TotalCount	Uint64	Number of notifications that meet the specified conditions. Example: 1
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameter.Conflict	Multiple parameters specified conflict and cannot co-exist.
LimitExceeded	The quota limit is exceeded.
InvalidFilter	Invalid filter.

ExecuteScalingPolicy

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (ExecuteScalingPolicy) is used to run a scaling policy.

- You can run a scaling policy based on the scaling policy ID.
- When the scaling group to which a scaling policy belongs is performing a scaling activity, the system rejects to run the scaling policy.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 17:05:44.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ExecuteScalingPolicy
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingPolicyId	Yes	No	String	ID of the alarm-based scaling policy. Example: asp-f59pppuh
HonorCooldown	No	No	Bool	Specifies whether to check if the scaling group is in the cooldown period. Default value: false. Example: false

3. Output Parameters

Parameter name	Type	Description
ActivityId	String	ID of the scaling activity. Example: asa-o4v87ae9
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
ResourceUnavailable.AutoScalingGroupInActivity	The auto scaling group is active.
InternalError	Internal error.

ModifyScalingPolicy

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (ModifyScalingPolicy) is used to modify an alarm trigger policy.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 18:33:33.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ModifyScalingPolicy
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingPolicyId	Yes	No	String	ID of the alarm policy. Example: asp-iir70sxv
ScalingPolicyName	No	No	String	Name of the alarm policy. Example: cpu_policy_
AdjustmentType	No	No	String	Method that you want to use to change the desired capacity after the alarm is triggered. Valid values: • CHANGE_IN_CAPACITY: Increase or decrease the desired capacity. • EXACT_CAPACITY: Adjust to the specified desired capacity. • PERCENT_CHANGE_IN_CAPACITY: Adjust the desired capacity by percentage. Example: CHANGE_IN_CAPACITY
AdjustmentValue	No	No	Int64	Adjusted desired capacity after the alarm is triggered. Valid values: • When the AdjustmentType parameter is set to CHANGE_IN_CAPACITY, some new instances will be added after the alarm is triggered if the AdjustmentValue parameter is set to a positive value, and some existing instances will be removed after the alarm is triggered if the AdjustmentValue parameter is set to a negative value. • When the AdjustmentType parameter is set to EXACT_CAPACITY, the value of the AdjustmentValue parameter specifies the desired number of instances after the alarm is triggered, which must be equal to or greater than 0. • When the AdjustmentType parameter is set to PERCENT_CHANGE_IN_CAPACITY, new instances will be added by percentage after the alarm is triggered if the AdjusmentValue parameter is set to a positive value, and existing instances will be removed by percentage after the alarm is triggered if the AdjusmentValue parameter is set to a negative value. Unit: %. Example: 1
Cooldown	No	No	Uint64	Cooldown period. Unit: seconds. Example: 60
MetricAlarm	No	No	MetricAlarm	Alarm metric. Example:View
NotificationUserGroupIds	No	No	Array of String	Notification group ID, a collection of user group IDs. User group IDs can be queried through DescribeUserGroup. If the notification user group needs to be cleared, a specific string NULL should be passed in the list. Example: ["1678Vr"]

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
ResourceNotFound.ScalingPolicyNotFound	The alarm policy does not exist.
InvalidParameterValue.ScalingPolicyNameDuplicate	The alarm policy name already exists.

Launch Configuration APIs
CreateLaunchConfiguration

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (CreateLaunchConfiguration) is used to create a new launch configuration.

*The launch configuration can be partially modified through ModifyLaunchConfigurationAttributes . If a new launch configuration is needed, it is recommended to create a new launch configuration.

*Each project can create up to 20 launch configurations, see Usage Limits for details.

Default interface request rate limit: 20 times per second.

Interface update time: 2021-04-28 10:32:54.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: CreateLaunchConfiguration
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
LaunchConfigurationName	Yes	No	String	Display name of the launch configuration. The name can contain only letters, digits, underscores (_), hyphens (-), and decimal points (.) and must be up to 60 bytes in length. Example: as_test
ProjectId	No	No	Uint64	The project ID to which the instance belongs. This parameter can be obtained by calling the projectId field in the return value from the DescribeProjects. If this parameter is not specified, it defaults to the default project. Example: 0
InstanceType	No	No	String	Instance model. Different instance models specify different resource specifications. The specific values can be obtained by calling the DescribeInstanceTypeConfigs interface to access the latest specifications table or seeing the Instance Type description. InstanceType and InstanceTypes parameters are mutually exclusive. One and only one of them must be filled in. Example: S2.SMALL1
ImageId	Yes	No	String	Image ID, such as img-8toqc6s3 . Image types include: • Public Image • Custom image • Shared image You can use the following methods to obtain available image IDs: • Query image IDs in the console for public images , custom images , and shared images . • Call the DescribeImages API and obtain the image ID from the value of the ImageId response parameter. Example: img-8toqc6s3
SystemDisk	No	No	SystemDisk	System disk configurations of the instance. If this parameter is not specified, the default value is assigned to it. Example:View
DataDisks	No	No	Array of DataDisk	Data disk configurations of the instance. If this parameter is not specified, no data disk will be purchased by default. You can specify a maximum of 11 data disks. Example:View
InternetAccessible	No	No	InternetAccessible	Public network bandwidth settings. If this parameter is not specified, the default public network bandwidth of 0 Mbit/s is used. Example:View
LoginSettings	No	No	LoginSettings	The instance login information. You can use this parameter to specify the login password/key or retain the original login settings of the image. By default, passwords are randomly generated and notified to users via internal messages. Example:View
SecurityGroupIds	No	No	Array of String	The security group to which the instance belongs. This parameter can be obtained by calling the SecurityGroupId field in the return value from the DescribeSecurityGroups. If this parameter is not specified, do not bind the security group by default. Example: []
EnhancedService	No	No	EnhancedService	Enhanced services. You can use this parameter to specify whether to enable Cloud Security and Cloud Monitor services. If this parameter is not specified, the Cloud Monitor and Cloud Security services are enabled by default. Example:View
UserData	No	No	String	Custom Base64-encoded data. The length must be up to 16 KB. Example: TXlVc2VyRGF0YQo=
InstanceChargeType	No	No	String	Billing method of the instance. The default value for a CVM instance is POSTPAID_BY_HOUR. • POSTPAID_BY_HOUR: the pay-as-you-go billing method. You are charged on an hourly basis. • SPOTPAID: pay by bidding. Example: POSTPAID_BY_HOUR
InstanceMarketOptions	No	No	InstanceMarketOptionsRequest	Marketplace-related options of the instance, such as spot instance related parameters. This parameter is required if the InstanceChargeType parameter is set to SPOTPAID. Example:View
InstanceTypes	No	No	Array of String	Instance models. Different instance models specify different resource specifications. You can specify a maximum of 10 instance models. InstanceType and InstanceTypes are mutually exclusive. You must specify one of the parameters and can specify only one of them. Example: ["S2.SMALL4", "S2.SMALL2"]
InstanceTypesCheckPolicy	No	No	String	Instance type verification policy. Valid values include ALL and ANY. Default value: ANY. • ALL: the verification is successful only if all instance types are available. Otherwise, the verification fails. • ANY: the verification is successful if one of the instance types is available. Otherwise, the verification fails. Common reasons for instance type unavailability include: the instance type is sold out and the corresponding cloud disks are sold out. If one of the models specified in the InstanceTypes parameter does not exist or is no longer available, the verification fails regardless of the value of the InstanceTypesCheckPolicy parameter. Example: ALL
InstanceTags	No	No	Array of InstanceTag	List of tags. You can use this parameter to bind tags to newly added instances. You can specify a maximum of 10 tags. Example:View
CamRoleName	No	No	String	Name of the CAM role. You can call the DescribeRoleList API and obtain the role name from the value of the roleName response parameter. Example: roleName001

3. Output Parameters

Parameter name	Type	Description
LaunchConfigurationId	String	ID of the launch configuration. This parameter is returned when you call this API to create a launch configuration. Example: asc-23h37kyn
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
AccountQualificationRestrictions	The requesting account failed to pass the qualification review.
InvalidParameterValue.CvmConfigurationError	Exception with CVM parameter validation.
InvalidPermission	The account is not authorized to perform this operation.
InvalidParameter.Conflict	Multiple parameters specified conflict and cannot co-exist.
InvalidParameterValue.LaunchConfigurationNameDuplicated	The launch configuration name already exists.
InvalidParameter.MustOneParameter	A parameter is missing. One of the two parameters must be specified.
LaunchConfigurationQuotaLimitExceeded	The launch configuration quota is exceeded.
InvalidImageId.NotFound	The image cannot be found.
CallCvmError	CVM API call failed.

DeleteLaunchConfiguration

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DeleteLaunchConfiguration) is used to delete a launch configuration.

- If a launch configuration is effective on a scaling group, the launch configuration cannot be deleted.

Default interface request rate limit: 10 times per second.

Interface update time: 2019-11-01 11:14:07.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DeleteLaunchConfiguration
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
LaunchConfigurationId	Yes	No	String	ID of the launch configuration that you want to delete. Example: asc-aaaabbbb

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidLaunchConfigurationId	Invalid launch configuration ID.
InvalidLaunchConfigurationId.InUse	The launch configuration is being used.
InvalidLaunchConfigurationId.NotFound	The launch configuration was not found.
InvalidLaunchConfiguration	Invalid launch configuration.

DescribeLaunchConfigurations

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DescribeLaunchConfigurations) is used to query the information about launch configurations.

- You can query the launch configuration details based on information such as the launch configuration ID and name. For more information about the supported filters, see Filters .
- If the filtering parameters are empty, a certain number (specified by Limit and 20 by default) of launch configurations that belong to the current user will be returned.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 17:06:10.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeLaunchConfigurations
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
LaunchConfigurationIds	No	No	Array of String	ID of one or more launch configurations to query, such as asc-ouy1ax38 . You can specify a maximum of 100 scaling activity IDs in each request. LaunchConfigurationIds and Filters cannot be specified at the same time. Example: ["asc-fa28v4in"]
Filters	No	No	Array of Filter	Filter conditions. • launch-configuration-id - String - Required: No - (Filter condition) Filter by launch configuration ID. • launch-configuration-name - String - Required: No - (Filter condition) Filter by launch configuration name. The maximum number of Filters per request is 10. The upper limit for Filter.Values is 5. LaunchConfigurationIds and Filters cannot be specified at the same time. Example:View
Limit	No	No	Uint64	Number of returned items. Default value: 20. Maximum value: 100. For more information on Limit , see the relevant section in API Introduction. Example: 20
Offset	No	No	Uint64	Offset. Default value: 0. For more information on Offset , see the relevant section in API Introduction. Example: 0

3. Output Parameters

Parameter name	Type	Description
TotalCount	Uint64	Number of launch configurations that meet the specified conditions. Example: 1
LaunchConfigurationSet	Array of LaunchConfiguration	Information about the launch configurations that meet the specified conditions. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidLaunchConfigurationId	Invalid launch configuration ID.
InvalidPermission	The account is not authorized to perform this operation.
InvalidParameter.Conflict	Multiple parameters specified conflict and cannot co-exist.
InvalidFilter	Invalid filter.
InvalidLaunchConfiguration	Invalid launch configuration.

ModifyLaunchConfigurationAttributes

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (ModifyLaunchConfigurationAttributes) is used to modify specific attributes of a launch configuration.

- The changes in the launch configuration do not affect the existing instances. New instances will be created based on the new launch configuration.
- This API supports modifying certain simple types of attributes.

Default interface request rate limit: 20 times per second.

Interface update time: 2022-11-04 17:56:13.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ModifyLaunchConfigurationAttributes
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
LaunchConfigurationId	Yes	No	String	Launch Configuration ID Example: asc-aaaabbbb
ImageId	No	No	String	Specify a valid image ID, in the format of img-8toqc6s3 . There are four image types: • Public Image • Custom Image • Shared Image • Service Market Image You can obtain available image IDs in the following ways: • Image IDs for the public iImage , custom image , and shared image can be queried by logging in to the console. • By calling the ImageId field in the return information from the DescribeImages interface. Example: img-8toqc6s3
InstanceTypes	No	No	Array of String	List of instance types. Each type specifies different resource specifications. You can specify up to 10 instance types. The launch configuration uses InstanceType to indicate one single instance type and InstanceTypes to indicate multiple instance types. If you specify the InstanceTypes parameter, the original InstanceType parameter automatically becomes invalid. Example: S5.SMALL2
InstanceTypesCheckPolicy	No	No	String	Instance type verification policy, which works when the instance types are actually modified. Valid values: ALL and ANY. Default value: ANY. • ALL: the verification is successful only if all instance types are available. Otherwise, the verification fails. • ANY: the verification is successful if one of the instance types is available. Otherwise, the verification fails. Common reasons for instance type unavailability include: the instance type is sold out and the corresponding cloud disks are sold out. If one of the models specified in the InstanceTypes parameter does not exist or is no longer available, the verification fails regardless of the value of the InstanceTypesCheckPolicy parameter. Example: ALL
LaunchConfigurationName	No	No	String	Display name of the launch configuration. The name can contain only letters, digits, underscores (_), hyphens (-), and decimal points (.) and must be up to 60 bytes in length. Example: myascname
UserData	No	No	String	Custom Base64-encoded data. The length must be up to 16 KB. If you want to clear custom data, set the UserData parameter to an empty string. Example: ""

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
AccountQualificationRestrictions	The requesting account failed to pass the qualification review.
InvalidParameterValue.LaunchConfigurationNameDuplicated	The launch configuration name already exists.
InvalidParameterValue.CvmConfigurationError	Exception with CVM parameter validation.
InvalidImageId.NotFound	The image cannot be found.
CallCvmError	CVM API call failed.
InvalidParameterValue.InstanceTypeNotSupported	Unsupported instance type

UpgradeLaunchConfiguration

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (UpgradeLaunchConfiguration) is used to upgrade a launch configuration.

- This API is used to upgrade a launch configuration in a "completely overriding" manner. It uniformly specifies a new configuration according to the request parameters regardless of the original parameter configurations. If optional parameters are left empty, their default values will be used.
- After the launch configuration is upgraded, the existing instances that have been created based on the original launch configuration are not affected, but new instances will be created based on the new launch configuration.

Default interface request rate limit: 20 times per second.

Interface update time: 2022-11-04 17:58:42.

The interface is both signature verification and authentication.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: UpgradeLaunchConfiguration
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
LaunchConfigurationId	Yes	No	String	ID of the launch configuration. Example: asc-aaaabbbb
DataDisks	No	No	Array of DataDisk	Data disk configurations of the instance. If this parameter is not specified, no data disk will be purchased by default. You can specify a maximum of 11 data disks. Example:View
EnhancedService	No	No	EnhancedService	Enhanced services. You can use this parameter to specify whether to enable Cloud Security and Cloud Monitor services. If this parameter is not specified, the Cloud Monitor and Cloud Security services are enabled by default. Example:View
ImageId	Yes	No	String	Specify a valid image ID, in the format of img-8toqc6s3 . There are four image types: • Public Image • Custom Image • Shared Image • Service Market Image You can obtain available image IDs in the following ways: • Image IDs for the public image , custom image , and shared image can be queried by logging in to the console. • By calling the ImageId field in the return information from the DescribeImages interface. Example: img-aaaabbbb
InstanceChargeType	No	No	String	Billing method of the instance. The default value for a CVM instance is POSTPAID_BY_HOUR. • POSTPAID_BY_HOUR: the pay-as-you-go billing method. You are charged on an hourly basis. • SPOTPAID: pay by bidding. Example: POSTPAID_BY_HOUR
InstanceMarketOptions	No	No	InstanceMarketOptionsRequest	Marketplace-related options of the instance, such as spot instance related parameters. This parameter is required if the InstanceChargeType parameter is set to SPOTPAID. Example:View
InstanceTypes	Yes	No	Array of String	List of instance types. Each instance type specifies different resource specifications. You can specify up to 5 instance types. Example: ["S5.SMALL2"]
InstanceTypesCheckPolicy	No	No	String	Instance type verification policy. Valid values include ALL and ANY. Default value: ANY. • ALL: the verification is successful only if all instance types are available. Otherwise, the verification fails. • ANY: the verification is successful if one of the instance types is available. Otherwise, the verification fails. Common reasons for instance type unavailability include: the instance type is sold out and the corresponding cloud disks are sold out. If one of the models specified in the InstanceTypes parameter does not exist or is no longer available, the verification fails regardless of the value of the InstanceTypesCheckPolicy parameter. Example: ALL
InternetAccessible	No	No	InternetAccessible	Public network bandwidth settings. If this parameter is not specified, the default public network bandwidth of 0 Mbit/s is used. Example:View
LaunchConfigurationName	Yes	No	String	Display name of the launch configuration. The name can contain only letters, digits, underscores (_), hyphens (-), and decimal points (.) and must be up to 60 bytes in length. Example: myascname
LoginSettings	No	No	LoginSettings	The instance login information. You can use this parameter to specify the login password/key or retain the original login settings of the image. By default, passwords are randomly generated and notified to users via internal messages. Example:View
ProjectId	No	No	Int64	The project ID to which the instance belongs. This parameter can be obtained by calling the projectId field in the return value from the DescribeProjects. If this parameter is not specified, it defaults to the default project. Example: 1222
SecurityGroupIds	No	No	Array of String	The security group to which the instance belongs. This parameter can be obtained by calling the SecurityGroupId field in the return value from the DescribeSecurityGroups. If this parameter is not specified, do not bind the security group by default. Example: ["sg-aaaabbbb"]
SystemDisk	Yes	No	SystemDisk	System disk configurations of the instance. If this parameter is not specified, the default value is assigned to it. Example:View
UserData	No	No	String	Custom Base64-encoded data. The length must be up to 16 KB. Example: **
InstanceTags	No	No	Array of InstanceTag	List of tags. You can use this parameter to bind tags to newly added instances. You can specify a maximum of 10 tags. Example:View
CamRoleName	No	No	String	Name of the CAM role. You can call the DescribeRoleList API and obtain the role name from the value of the roleName response parameter. Example: myrolename

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameter.MustOneParameter	A parameter is missing. One of the two parameters must be specified.
InvalidParameter.Conflict	Multiple parameters specified conflict and cannot co-exist.
InvalidParameterValue.LaunchConfigurationNameDuplicated	The launch configuration name already exists.
InvalidParameterValue.CvmError	Exception with CVM parameter validation.
InvalidParameterValue.CvmConfigurationError	Exception with CVM parameter validation.
InvalidImageId.NotFound	The image cannot be found.
CallCvmError	CVM API call failed.
InternalError	Internal error.
InvalidParameterValue	The parameter value is invalid.
InvalidParameterValue.InstanceTypeNotSupported	Unsupported instance type

Lifecycle Hook APIs
CompleteLifecycleAction

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (CompleteLifecycleAction) is used to complete a lifecycle action.

- You can call this API to specify a result ("CONTINUE" or "ABANDON") for a lifecycle hook. If this API is not called at all, the lifecycle hook will be processed based on the "DefaultResult" after the timeout period elapses.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 18:28:51.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: CompleteLifecycleAction
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
LifecycleHookId	Yes	No	String	Lifecycle hook ID. Example: ash-fbjiexz7
InstanceId	No	No	String	Instance ID. You must specify at least one of the InstanceId and LifecycleActionToken parameters. Example: ins-ni8bpmve
LifecycleActionToken	No	No	String	You must specify at least one of the InstanceId and LifecycleActionToken parameters. Example: 4d910016-2590-444d-8f4a-c14940036902
LifecycleActionResult	Yes	No	String	Result of the lifecycle action. Valid values: CONTINUE and ABANDON. Example: CONTINUE

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
ResourceUnavailable.LifecycleActionResultHasSet	The lifecycle action has already been set.
InvalidParameter	Parameter error.
ResourceNotFound.LifecycleHookNotFound	The specified lifecycle hook was not found.

Error Code	Description
ResourceNotFound.LifecycleHookInstanceNotFound	The instance corresponding to the lifecycle hook does not exist.
InternalError	Internal error.
InvalidParameterValue	The parameter value is invalid.

CreateLifecycleHook

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API is used to create a lifecycle hook.

- You can configure message notifications for the lifecycle hook, and Auto Scaling will notify TDMQ for CMQ. Notification content format:


```
{
    "Service": "Cloud Platform Auto Scaling",
    "Time": "2019-03-14T10:15:11Z",
    "AppId": "1251783334",
    "ActivityId": "asa-fznnvrja",
    "AutoScalingGroupId": "asg-rrrrtttt",
    "LifecycleHookId": "ash-xxxxyyyy",
    "LifecycleHookName": "my-hook",
    "LifecycleActionToken": "3080e1c9-0efe-4dd7-ad3b-90cd6618298f",
    "InstanceId": "ins-aaaabbbb",
    "LifecycleTransition": "INSTANCE_LAUNCHING",
    "NotificationMetadata": ""
}
```


Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 18:27:50.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: CreateLifecycleHook
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupId	Yes	No	String	ID of the scaling group. Example: asg-8fbozqja
DefaultResult	No	No	String	The action that is taken when the lifecycle hook of the scaling group times out. Valid values: CONTINUE and ABANDON. Default value: CONTINUE. Example: CONTINUE
HeartbeatTimeout	No	No	Int64	Timeout period of the lifecycle hook. Unit: seconds. Valid values: 30 to 3600. Default value: 300. Example: 300
LifecycleHookName	Yes	No	String	Name of the lifecycle hook. The name can contain only letters, digits, underscores (_), hyphens (-), and decimal points (.) and must be up to 128 bytes in length, Example: one-hook-default
LifecycleTransition	Yes	No	String	Scenario where the lifecycle hook is used. Valid values: INSTANCE_LAUNCHING and INSTANCE_TERMINATING. Example: INSTANCE_LAUNCHING
NotificationMetadata	No	No	String	Additional information that Auto Scaling sends to the notification target. By default, this parameter is left empty. The maximum length cannot exceed 1,024 bytes. Example: ""
NotificationTarget	No	No	NotificationTarget	Notification target Example:View

3. Output Parameters

Parameter name	Type	Description
LifecycleHookId	String	Lifecycle hook ID. Example: ash-8azjzxj9
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue.LifecycleHookNameDuplicated	The lifecycle hook name already exists.
InvalidParameter.Conflict	Multiple parameters specified conflict and cannot co-exist.
InvalidParameter	Parameter error.
MissingParameter	Parameter missing.

Error Code	Description
InvalidParameterValue.Filter	Invalid filter.
InternalError	Internal error.
InvalidParameterValue	The parameter value is invalid.
InternalError.CallTDMQError	TDMQ interface call failed.
ResourceNotFound.ResourceNotFoundTDMQPulsarTopicNotFound	TDMQ-Pulsar theme does not exist.

DeleteLifecycleHook

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DeleteLifeCycleHook) is used to delete a lifecycle hook.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 18:30:37.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DeleteLifecycleHook
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
LifecycleHookId	Yes	No	String	Lifecycle hook ID. Example: ash-5rynk2n5

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameter	Parameter error.
MissingParameter	Parameter missing.
ResourceNotFound.LifecycleHookNotFound	The specified lifecycle hook was not found.
InternalError	Internal error.

DescribeLifecycleHooks

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DescribeLifecycleHooks) is used to query the information about lifecycle hooks.

- You can query the details of lifecycle hooks based on information such as the scaling group ID, lifecycle hook ID, or lifecycle hook name. For more information about the supported filters, see Filters .
- If the filtering parameters are empty, a certain number (specified by Limit and 20 by default) of lifecycle hooks that belong to the current user will be returned.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 18:30:12.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeLifecycleHooks
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
LifecycleHookIds	No	No	Array of String	ID of one or more lifecycle hooks to query, such as ash-8azjzxcl . You can specify a maximum of 100 scaling activity IDs in each request. LifecycleHookIds and Filters cannot be specified at the same time. Example: ["ash-8azjzxcl"]
Filters	No	No	Array of Filter	Filter conditions. • lifecycle-hook-id - String - Required: No - (Filter condition) Filter by lifecycle hook ID. • lifecycle-hook-name - String - Required: No - (Filter condition) Filter by lifecycle hook name. • auto-scaling-group-id - String - Required: No - (Filter condition) Filter by scaling group ID. Filter conditions. • lifecycle-hook-id - String - Required: No - (Filter condition) Filter by lifecycle hook ID. • lifecycle-hook-name - String - Required: No - (Filter condition) Filter by lifecycle hook name. • auto-scaling-group-id - String - Required: No - (Filter condition) Filter by scaling group ID. The maximum number of Filters per request is 10. The upper limit for Filter.Values is 5. LifecycleHookIds and Filters cannot be specified at the same time. Example:View
Limit	No	No	Uint64	Number of returned items. Default value: 20. Maximum value: 100. For more information on Limit , see the relevant section in API Introduction. Example: 1
Offset	No	No	Uint64	Offset. Default value: 0. For more information on Offset , see the relevant section in API Introduction. Example: 0

3. Output Parameters

Parameter name	Type	Description
LifecycleHookSet	Array of LifecycleHook	Array of lifecycle hooks. Example:View
TotalCount	Int64	Total number. Example: 4
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameter.Conflict	Multiple parameters specified conflict and cannot co-exist.
InvalidParameter	Parameter error.
MissingParameter	Parameter missing.
InvalidFilter	Invalid filter.
InternalError	Internal error.

UpgradeLifecycleHook

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (UpgradeLifecycleHook) is used to upgrade a lifecycle hook.

- This API is used to upgrade a lifecycle hook in a "completely overriding" manner. It uniformly specifies a new configuration based on the request parameters regardless of the original parameter configurations. If optional parameters are left empty, their default values will be used.

Default interface request rate limit: 20 times per second.

Interface update time: 2022-10-14 17:29:24.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: UpgradeLifecycleHook
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
LifecycleHookId	Yes	No	String	Lifecycle hook ID. Example: ash-8azjzxj9
DefaultResult	No	No	String	The action that is taken when the lifecycle hook of the scaling group times out. Valid values: CONTINUE and ABANDON. Default value: CONTINUE. Example: CONTINUE
HeartbeatTimeout	No	No	Int64	Timeout period of the lifecycle hook. Unit: seconds. Valid values: 30 to 3600. Default value: 300. Example: 240
LifecycleHookName	Yes	No	String	Name of the lifecycle hook. Example: hook-updated
LifecycleTransition	Yes	No	String	Scenario where the lifecycle hook is used. Valid values: INSTANCE_LAUNCHING and INSTANCE_TERMINATING. Example: INSTANCE_LAUNCHING
NotificationMetadata	No	No	String	Additional information that Auto Scaling sends to the notification target. By default, this parameter is left empty. Example: ""
NotificationTarget	No	No	NotificationTarget	Notification target Example:View

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue.LifecycleHookNameDuplicated	The lifecycle hook name already exists.
InvalidParameter.Conflict	Multiple parameters specified conflict and cannot co-exist.
InvalidParameter	Parameter error.
MissingParameter	Parameter missing.
InvalidParameterValue.Filter	Invalid filter.
ResourceNotFound.LifecycleHookNotFound	The specified lifecycle hook was not found.
InternalError	Internal error.
InvalidParameterValue	The parameter value is invalid.
InternalError.CallTDMQError	TDMQ interface call failed.
ResourceNotFound.ResourceNotFoundTDMQPulsarTopicNotFound	TDMQ-Pulsar theme does not exist.

Notification APIs
CreateNotificationConfiguration

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (CreateNotificationConfiguration) is used to create a notification.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 18:35:51.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: CreateNotificationConfiguration
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupId	Yes	No	String	ID of the scaling group. Example: asg-12wjuh0s
NotificationTypes	Yes	No	Array of String	Notification types you want to subscribe to. Valid values: • SCALE_OUT_SUCCESSFUL: the scale-out operation is successful. • SCALE_OUT_FAILED: the scale-out operation has failed. • SCALE_IN_SUCCESSFUL: the scale-in operation is successful. • SCALE_IN_FAILED: the scale-in operation has failed. • REPLACE_UNHEALTHY_INSTANCE_SUCCESSFUL: the unhealthy CVM instance is replaced successfully. • REPLACE_UNHEALTHY_INSTANCE_FAILED: the unhealthy CVM instance has failed to be replaced. Example: ["SCALE_OUT_SUCCESSFUL\n","SCALE_OUT_FAILED\n"]
NotificationUserGroupIds	Yes	No	Array of String	Notification group ID, a collection of user group IDs. User group IDs can be queried through DescribeUserGroup. Example: ["1678\n"]

3. Output Parameters

Parameter name	Type	Description
AutoScalingNotificationId	String	Notification ID. Example: asn-2sestqbr
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
ResourceNotFound.AutoScalingGroupIdNotFound	The scaling group does not exist.
MissingParameter	Parameter missing.
LimitExceeded	The quota limit is exceeded.
ResourceNotFound.AutoScalingNotificationNotFound	The notification does not exist.
InvalidParameterValue.UserGroupIdNotFound	The user group does not exist.

DeleteNotificationConfiguration

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DeleteNotificationConfiguration) is used to delete the specified notification.

Default interface request rate limit: 20 times per second.

Interface update time: 2022-09-27 17:17:02.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DeleteNotificationConfiguration
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingNotificationId	Yes	No	String	ID of the notification that you want to delete. Example: asn-2sestqbr

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
ResourceNotFound.AutoScalingNotificationNotFound	The notification does not exist.

DescribeNotificationConfigurations

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DescribeNotificationConfigurations) is used to query the details about one or more notifications.

You can query the details of notifications based on information such as the notification ID and scaling group ID. For more information about the supported filters, see Filters .
If the filtering parameters are empty, a certain number (specified by Limit and 20 by default) of notifications of the current user will be returned.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 18:34:26.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeNotificationConfigurations
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingNotificationIds	No	No	Array of String	ID of one or more notifications, such as asn-2sestqbr . You can specify a maximum of 100 instance IDs in each request. AutoScalingNotificationIds and Filters cannot be specified at the same time. Example: ["asn-9bhwvxqh\n"]
Filters	No	No	Array of Filter	Filter conditions. • auto-scaling-notification-id - String - Required: No - (Filter condition) Filter by notification ID. • auto-scaling-group-id - String - Required: No - (Filter condition) Filter by scaling group ID. The maximum number of Filters per request is 10. The upper limit for Filter.Values is 5. AutoScalingNotificationIds and Filters cannot be specified at the same time. Example:View
Limit	No	No	Uint64	Number of returned items. Default value: 20. Maximum value: 100. For more information on Limit , see the relevant section in API Introduction. Example: 20
Offset	No	No	Uint64	Offset. Default value: 0. For more information on Offset , see the relevant section in API Introduction. Example: 0

3. Output Parameters

Parameter name	Type	Description
TotalCount	Uint64	Number of notifications that meet the specified conditions. Example: 1
AutoScalingNotificationSet	Array of AutoScalingNotification	Information about auto scaling event notifications. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameter.Conflict	Multiple parameters specified conflict and cannot co-exist.
InvalidFilter	Invalid filter.

ModifyNotificationConfiguration

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (ModifyNotificationConfiguration) is used to modify a notification policy.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 18:35:27.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ModifyNotificationConfiguration
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingNotificationId	Yes	No	String	ID of the notification to be modified. Example: asn-2sestqbr
NotificationTypes	No	No	Array of String	Notification types you want to subscribe to. Valid values: • SCALE_OUT_SUCCESSFUL: the scale-out operation is successful. • SCALE_OUT_FAILED: the scale-out operation has failed. • SCALE_IN_SUCCESSFUL: the scale-in operation is successful. • SCALE_IN_FAILED: the scale-in operation has failed. • REPLACE_UNHEALTHY_INSTANCE_SUCCESSFUL: the unhealthy CVM instance is replaced successfully. • REPLACE_UNHEALTHY_INSTANCE_FAILED: the unhealthy CVM instance has failed to be replaced. Example: ["SCALE_IN_SUCCESSFUL\n","SCALE_IN_FAILED\n"]
NotificationUserGroupIds	No	No	Array of String	Notification group ID, a collection of user group IDs. User group IDs can be queried through DescribeUserGroup. Example: ["1678\n"]

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
ResourceNotFound.AutoScalingNotificationNotFound	The notification does not exist.
InvalidParameterValue.UserGroupIdNotFound	The user group does not exist.

Other APIs
DescribeAccountLimits

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DescribeAccountLimits) is used to query the limits of user's AS resources.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-11-01 11:12:00.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeAccountLimits
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface

3. Output Parameters

Parameter name	Type	Description
MaxNumberOfLaunchConfigurations	Int64	Maximum number of launch configurations that can be created within the user account. Example: 20
NumberOfLaunchConfigurations	Int64	Current number of launch configurations within the user account. Example: 15
MaxNumberOfAutoScalingGroups	Int64	Maximum number of scaling groups that can be created within the user account. Example: 30
NumberOfAutoScalingGroups	Int64	Current number of scaling groups within the user account. Example: 25
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description

Error Code	Description
InternalError	Internal error.

Scaling Group APIs
AttachInstances

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (AttachInstances) is used to add CVM instances to a scaling group.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 18:26:52.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: AttachInstances
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupId	Yes	No	String	ID of the scaling group. Example: asg-boz1qhnk
InstanceIds	Yes	No	Array of String	IDs of CVM instances. Example: ["ins-cri8d02t\n","ins-osckfnm7\n"]

3. Output Parameters

Parameter name	Type	Description
ActivityId	String	ID of the scaling activity. Example: asa-q59zikez
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
ResourceUnavailable.AutoScalingGroupInActivity	The auto scaling group is active.
ResourceUnavailable.InstancesAlreadyInAutoScalingGroup	The instance already exists in the auto scaling group.
ResourceUnavailable.CvmVpcInconsistent	The instance and the auto scaling group are in different VPCs.
ResourceNotFound.InstancesNotFound	The specified instance does not exist.
ResourceNotFound.AutoScalingGroupIdNotFound	The scaling group does not exist.
ResourceInsufficient.AutoScalingGroupAboveMaxSize	The maximum number of instances in the auto scaling group is exceeded.
InvalidParameterValue.LimitExceeded	The value exceeds the limit.
InternalError	Internal error.

CreateAutoScalingGroup

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (CreateAutoScalingGroup) is used to create an auto scaling group.

Default interface request rate limit: 20 times per second.

Interface update time: 2023-03-03 11:45:31.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: CreateAutoScalingGroup
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupName	Yes	No	String	Name of the scaling group. The name must be unique within your account. The name can contain only letters, digits, underscores (_), hyphens (-), and decimal points (.) and must be up to 55 bytes in length. Example: asgname
DefaultCooldown	No	No	Uint64	Default cooldown period. Unit: seconds. Default value: 300. Example: 300
DesiredCapacity	No	No	Uint64	Desired number of instances. The value must be in the range of the minimum number of instances and the maximum number of instances. Example: 1
LaunchConfigurationId	Yes	No	String	Launch Configuration ID Example: asc-aaaabbbb
LoadBalancerIds	No	No	Array of String	IDs of classic load balancers. Up to 20 IDs are supported. You can only specify one of LoadBalancerIds and ForwardLoadBalancers, or leave both of them empty. Example: ["lb-aaaabbbb"]
MaxSize	Yes	No	Uint64	Maximum number of instances. Valid values: 0 to 2000. Example: 1
MinSize	Yes	No	Uint64	Minimum number of instances. Valid values: 0 to 2000. Example: 1
ProjectId	No	No	Uint64	Project ID. Example: 1
ForwardLoadBalancers	No	No	Array of ForwardLoadBalancer	List of application load balancers. Up to 20 IDs are supported. You can only specify one of LoadBalancerIds and ForwardLoadBalancers, or leave both of them empty. Example:View
SubnetIds	No	No	Array of String	IDs of subnets. You must specify subnets in VPC scenarios. Multiple subnets are prioritized in the order in which they are specified and are tried in sequence until an instance can be successfully created. Example: ["subnet-bh595fp6"]
TerminationPolicies	No	No	Array of String	Termination policy. Currently, you can specify up to one termination policy. Valid values: OLDEST_INSTANCE and NEWEST_INSTANCE. Default value: OLDEST_INSTANCE. • OLDEST_INSTANCE: prioritizes the oldest instance in the scaling group. • NEWEST_INSTANCE: prioritizes the newest instance in the scaling group. Example: OLDEST_INSTANCE
VpcId	Yes	No	String	VPC ID. Specify an empty string for a basic network. Example: vpc-aaaabbbb
Zones	No	No	Array of String	List of availability zones (AZs). You must specify AZs for a basic network. Multiple AZs are prioritized in the order in which they are specified and are tried in sequence until an instance can be successfully created. Example: ["ap-region-1"]
RetryPolicy	No	No	String	Retry policy. Valid values: IMMEDIATE_RETRY, INCREMENTAL_INTERVALS, and NO_RETRY. Valid value: IMMEDIATE_RETRY. • IMMEDIATE_RETRY: retries in a short period and stops retrying after five consecutive failed attempts. • INCREMENTAL_INTERVALS: the retry interval increases as the number of consecutive failures increases. The retry interval ranges from seconds to one day. • NO_RETRY: no retry until a user call or alarm message is received again. Example: IMMEDIATE_RETRY
ZonesCheckPolicy	No	No	String	Availability zone verification policy. Valid values include ALL and ANY. Default value: ANY. • ALL: the verification is successful only if all availability zones and subnets are available. Otherwise, the verification fails. • ANY: the verification is successful if one of the availability zones or subnets is available. Otherwise, the verification fails. Common reasons why an AZ or subnet is unavailable include: the CVM instance types in the AZ are sold out, the CBS disks in the AZ are sold out, the quota in the AZ is insufficient, and the IP addresses in the subnet are insufficient. If the AZs specified in the Zones parameter or the subnets specified in the SubnetIds parameter do not exist, the verification fails regardless of the value of the ZonesCheckPolicy parameter. Example: ALL
Tags	No	No	Array of Tag	Tag list. You can use this parameter to bind tags to the scaling group and bind these tags to the corresponding resource instances. Example:View
ServiceSettings	No	No	ServiceSettings	Service settings, including unhealthy instance replacement. Example:View
Ipv6AddressCount	No	No	Int64	Configuration of the instance with the quantity of IPv6 addresses, values include 0 and 1. The default value is 0. Example: 1

3. Output Parameters

Parameter name	Type	Description
AutoScalingGroupId	String	ID of the scaling group. Example: asg-aaaaqqqq
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue.OnlyVpc	The account only supports VPCs.
InvalidParameter.InScenario	The parameter is invalid in a specific scenario.
InvalidParameterValue.LbProjectInconsistent	The load balancer is in a different project.
InvalidParameterValue.Range	The value is outside the specified range.
InvalidParameterValue.TooLong	Too many values.
ResourceUnavailable.ProjectInconsistent	Project inconsistency.
InvalidParameterValue.GroupNameDuplicated	The auto scaling group name already exists.
LimitExceeded.AutoScalingGroupLimitExceeded	The number of auto scaling groups exceeds the limit.
LimitExceeded.MaxSizeLimitExceeded	The maximum number of instances exceeds the limit.
InvalidParameterValue.Size	The value of maximum, minimum, or desired number of instances in the auto scaling group is invalid.
InvalidParameterValue.LbVpcInconsistent	The load balancer and the scaling group are in different VPCs.
LimitExceeded.MinSizeLimitExceeded	The minimum number of instances is below the limit.

Error Code	Description
LimitExceeded	The quota limit is exceeded.
InvalidParameterValue.CvmError	Exception with CVM parameter validation.
InvalidParameterValue.LaunchConfigurationNotFound	The specified launch configuration was not found.
InvalidParameterValue.SubnetIds	The subnet information is invalid.
MissingParameter.InScenario	A parameter is missing in a specific scenario.
ResourceUnavailable.LaunchConfigurationStatusAbnormal	The launch configuration is exceptional.
InvalidParameterValue.LimitExceeded	The value exceeds the limit.
InternalError	Internal error.
InvalidParameterValue.ForwardLb	The application load balancer is incorrectly specified.
ResourceNotFound.LoadBalancerNotFound	The specified load balancer was not found.

DeleteAutoScalingGroup

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DeleteAutoScalingGroup) is used to delete the specified auto scaling group that has no instances and remains inactive.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-11-01 11:13:41.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DeleteAutoScalingGroup
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupId	Yes	No	String	ID of the scaling group. Example: asg-nvnlpbb8

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request

Parameter name	Type	Description
		needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
ResourceNotFound.AutoScalingGroupIdNotFound	The scaling group does not exist.
ResourceInUse.InstanceInGroup	There are still normal instances in the auto scaling group.
ResourceInUse.ActivityInProgress	The auto scaling group is performing a scaling activity.
InternalError	Internal error.

DescribeAutoScalingActivities

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DescribeAutoScalingActivities) is used to query the activity history of an auto scaling group.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 17:05:09.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeAutoScalingActivities
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
ActivityIds	No	No	Array of String	ID of one or more scaling activities to query, such as asa-5l2ejpfo . You can specify a maximum of 100 scaling activity IDs in each request. ActivityIds and Filters cannot be specified at the same time. Example: ["asa-aaaabbbb"]
Filters	No	No	Array of Filter	Filter conditions. • auto-scaling-group-id - String - Required: No - (Filter condition) Filter by scaling group ID. • activity-status-code - String - Required: No - (Filter condition) Filter by scaling activity status. (INIT: the scaling activity is being initialized.
Limit	No	No	Uint64	Number of returned items. Default value: 20. Maximum value: 100. For more information on Limit , see the relevant section in API Introduction. Example: 100
Offset	No	No	Uint64	Offset. Default value: 0. For more information on Offset , see the relevant section in API Introduction. Example: 0
StartTime	No	No	Datetime_iso	Earliest start time of the scaling activity. If ActivityIds is specified, this parameter is ignored. Set the value to a UTC time in the format of YYYY-MM-DDThh:mm:ssZ according to the ISO8601 standard. Example: 2014-08-03T12:00:00+08:00
EndTime	No	No	Datetime_iso	Latest end time of the scaling activity. If ActivityIds is specified, this parameter is ignored. Set the value to a UTC time in the format of YYYY-MM-DDThh:mm:ssZ according to the ISO8601 standard. Example: 2014-08-03T12:00:00+08:00

3. Output Parameters

Parameter name	Type	Description
TotalCount	Uint64	Number of scaling activities that meet the specified conditions. Example: 1
ActivitySet	Array of Activity	Information set of the scaling activities that meet the specified conditions. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameter.Conflict	Multiple parameters specified conflict and cannot co-exist.
InvalidParameterValue.Filter	Invalid filter.
InvalidParameterValue.LimitExceeded	The value exceeds the limit.
InternalError	Internal error.

DescribeAutoScalingGroupLastActivities

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API is used to query the latest activity history of an auto scaling group.

Default interface request rate limit: 20 times per second.

Interface update time: 2022-07-07 15:18:57.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeAutoScalingGroupLastActivities
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupIds	Yes	No	Array of String	IDs of scaling groups. Example: ["asg-2umy3jbd\n"]

3. Output Parameters

Parameter name	Type	Description
ActivitySet	Array of Activity	Information set of the scaling activities that meet the specified conditions. This parameter is not returned for scaling groups without scaling activities. For example, if you specified 50 scaling group IDs but only 45 records are returned, it indicates that 5 of the scaling groups do not have scaling activities. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue.NoResourcePermission	No resource permission.
InvalidParameterValue.LimitExceeded	The value exceeds the limit.
InternalError	Internal error.

DescribeAutoScalingGroups

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DescribeAutoScalingGroups) is used to query the information about scaling groups.

- You can query the details of scaling groups based on information such as scaling group ID, scaling group name, or launch configuration ID. For more information about the supported filters, see Filters .
- If the filtering parameters are empty, a certain number (specified by Limit and 20 by default) of scaling groups that belong to the current user will be returned.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 18:25:58.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeAutoScalingGroups
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupIds	No	No	Array of String	ID of one or more scaling groups, such as asg-nkdwoui0 . You can specify a maximum of 100 scaling activity IDs in each request. AutoScalingGroupIds and Filters cannot be specified at the same time. Example: ["asg-nkdwoui0\n"]
Filters	No	No	Array of Filter	Filter conditions. • auto-scaling-group-id - String - Required: No - (Filter condition) Filter by scaling group ID. • auto-scaling-group-name - String - Required: No - (Filter condition) Filter by scaling group name. • launch-configuration-id - String - Required: No - (Filter condition) Filter by launch configuration ID. • tag-key - String - Required: No - (Filter condition) Filter by tag key. • tag-value - String - Required: No - (Filter condition) Filter by tag value. • tag:tag-key - String - Required: No - (Filter condition) Filter by tag key-value pair. Replace tag-key with a specific tag key. For more information about its usage, see Example 2. The maximum number of Filters per request is 10. The upper limit for Filter.Values is 5. AutoScalingGroupIds and Filters cannot be specified at the same time. Example:View
Limit	No	No	Uint64	Number of returned items. Default value: 20. Maximum value: 100. For more information on Limit , see the relevant section in API Introduction. Example: 10
Offset	No	No	Uint64	Offset. Default value: 0. For more information on Offset , see the relevant section in API Introduction. Example: 0

3. Output Parameters

Parameter name	Type	Description
AutoScalingGroupSet	Array of AutoScalingGroup	Detailed information about scaling groups. Example:View
TotalCount	Uint64	Number of scaling groups that meet the specified conditions. Example: 1
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidPermission	The account is not authorized to perform this operation.
InvalidParameter.Conflict	Multiple parameters specified conflict and cannot co-exist.
InvalidParameterValue.Filter	Invalid filter.
InvalidFilter	Invalid filter.

DescribeAutoScalingInstances

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DescribeAutoScalingInstances) is used to query the information about instances associated with Auto Scaling.

- You can query the details of instances based on information such as the instance ID and scaling group ID. For more information about the supported filters, see Filters .
- If you specified no filters in the request, the returned number of instances is determined by the Limit parameter in the request. The default value of the Limit parameter is 20.

Default interface request rate limit: 10 times per second.

Interface update time: 2019-10-31 18:28:18.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeAutoScalingInstances
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
InstanceIds	No	No	Array of String	IDs of the CVM instances to query. InstanceIds and Filters cannot be specified at the same time. Example: ["ins-1fswxz1m\n"]
Filters	No	No	Array of Filter	Filter conditions. • instance-id - String - Required: No - (Filter condition) Filter by instance ID. • auto-scaling-group-id - String - Required: No - (Filter condition) Filter by scaling group ID. The maximum number of Filters per request is 10. The upper limit for Filter.Values is 5. InstanceIds and Filters cannot be specified at the same time. Example:View
Offset	No	No	Int64	Offset. Default value: 0. For more information on Offset , see the relevant section in API Introduction. Example: 0
Limit	No	No	Int64	Number of returned items. Default value: 20. Maximum value: 100. For more information on Limit , see the relevant section in API Introduction. Example: 20

3. Output Parameters

Parameter name	Type	Description
AutoScalingInstanceSet	Array of Instance	The information about the instances. Example:View
TotalCount	Uint64	Total number of IP address template groups returned. Example: 1
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameter.Conflict	Multiple parameters specified conflict and cannot co-exist.
InvalidParameterValue.Filter	Invalid filter.
InvalidFilter	Invalid filter.
InternalError	Internal error.

DetachInstances

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DetachInstances) is used to remove CVM instances from a scaling group. This API does not terminate instances.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 18:26:30.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DetachInstances
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupId	Yes	No	String	ID of the scaling group. Example: asg-boz1qhnk
InstanceIds	Yes	No	Array of String	IDs of CVM instances. Example: ["ins-cri8d02t\n","ins-osckfnm7\n"]

3. Output Parameters

Parameter name	Type	Description
ActivityId	String	ID of the scaling activity. Example: asa-bcfxhy55
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
ResourceInsufficient.AutoScalingGroupBelowMinSize	The number of instances in the auto scaling group is below the minimum value.
ResourceUnavailable.AutoScalingGroupInActivity	The auto scaling group is active.
ResourceNotFound.AutoScalingGroupIdNotFound	The scaling group does not exist.
ResourceNotFound.InstancesNotInAutoScalingGroup	The target instance is not in the auto scaling group.
InvalidParameterValue.LimitExceeded	The value exceeds the limit.
InternalError	Internal error.

DisableAutoScalingGroup

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DisableAutoScalingGroup) is used to disable a specified scaling group.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-11-01 11:24:57.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DisableAutoScalingGroup
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupId	Yes	No	String	ID of the scaling group. Example: asg-nvnlpbb8

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
ResourceNotFound.AutoScalingGroupIdNotFound	The scaling group does not exist.
InternalError	Internal error.

EnableAutoScalingGroup

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (EnableAutoScalingGroup) is used to enable the specified auto scaling group.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-11-01 11:12:26.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: EnableAutoScalingGroup
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupId	Yes	No	String	ID of the scaling group. Example: asg-nvnlpbb8

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
ResourceNotFound.AutoScalingGroupIdNotFound	The scaling group does not exist.
InternalError	Internal error.

ModifyAutoScalingGroup

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (ModifyAutoScalingGroup) is used to modify an auto scaling group.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 17:04:15.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ModifyAutoScalingGroup
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupId	Yes	No	String	ID of the scaling group. Example: asg-aaaabbbb
AutoScalingGroupName	No	No	String	Name of the scaling group. The name must be unique within your account. The name can contain only letters, digits, underscores (_), hyphens (-), and decimal points (.) and must be up to 55 bytes in length. Example: asgname
DefaultCooldown	No	No	Uint64	Default cooldown period. Unit: seconds. Default value: 300. Example: 300
DesiredCapacity	No	No	Uint64	Desired number of instances. The value must be in the range of the minimum number of instances and the maximum number of instances. Example: 1
LaunchConfigurationId	No	No	String	Launch Configuration ID Example: asc-aaaabbbb
MaxSize	No	No	Uint64	Maximum number of instances. Valid values: 0 to 2000. Example: 1
MinSize	No	No	Uint64	Minimum number of instances. Valid values: 0 to 2000. Example: 1
ProjectId	No	No	Uint64	Project ID. Example: 1
SubnetIds	No	No	Array of String	IDs of subnets. Example: ["subnet-aaaabbbb"]
TerminationPolicies	No	No	Array of String	Termination policy. Currently, you can specify up to one termination policy. Valid values: OLDEST_INSTANCE and NEWEST_INSTANCE. • OLDEST_INSTANCE: prioritizes the oldest instance in the scaling group. • NEWEST_INSTANCE: prioritizes the newest instance in the scaling group. Example: ["OLDEST_INSTANCE"]
VpcId	No	No	String	VPC ID. Set this parameter to an empty string to specify the basic network. If you specify a VPC ID, you must also specify the SubnetIds parameter. If you specify the basic network, you also must specify the corresponding availability zones. Example: vpc-aaaabbbb
Zones	No	No	Array of String	List of AZ IDs. Example: ["ap-region-1"]
RetryPolicy	No	No	String	Retry policy. Valid values: IMMEDIATE_RETRY, INCREMENTAL_INTERVALS, and NO_RETRY. Valid value: IMMEDIATE_RETRY. • IMMEDIATE_RETRY: retries in a short period and stops retrying after five consecutive failed attempts. • INCREMENTAL_INTERVALS: the retry interval increases as the number of consecutive failures increases. The retry interval ranges from seconds to one day. • NO_RETRY: no retry until a user call or alarm message is received again. Example: IMMEDIATE_RETRY
ZonesCheckPolicy	No	No	String	Availability zone verification policy. Valid values include ALL and ANY. Default value: ANY. This parameter takes effect when you change the resource-related settings for the scaling group (launch configuration, availability zone, and subnet) • ALL: the verification is successful only if all availability zones and subnets are available. Otherwise, the verification fails. • ANY: the verification is successful if one of the availability zones or subnets is available. Otherwise, the verification fails. Common reasons why an AZ or subnet is unavailable include: the CVM instance types in the AZ are sold out, the CBS disks in the AZ are sold out, the quota in the AZ is insufficient, and the IP addresses in the subnet are insufficient. If the AZs specified in the Zones parameter or the subnets specified in the SubnetIds parameter do not exist, the verification fails regardless of the value of the ZonesCheckPolicy parameter. Example: ALL
ServiceSettings	No	No	ServiceSettings	Service settings, including unhealthy instance replacement. Example:View

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue.OnlyVpc	The account only supports VPCs.
InvalidParameter.InScenario	The parameter is invalid in a specific scenario.
InvalidParameterValue.LbProjectInconsistent	The load balancer is in a different project.
InvalidParameterValue.Range	The value is outside the specified range.
InvalidParameterValue.TooLong	Too many values.
ResourceUnavailable.ProjectInconsistent	Project inconsistency.
InvalidParameterValue.GroupNameDuplicated	The auto scaling group name already exists.
LimitExceeded.MaxSizeLimitExceeded	The maximum number of instances exceeds the limit.
ResourceNotFound.AutoScalingGroupIdNotFound	The scaling group does not exist.
InvalidParameterValue.Size	The value of maximum, minimum, or desired number of instances in the auto scaling group is invalid.
MissingParameter	Parameter missing.
InvalidParameterValue.LbVpcInconsistent	The load balancer and the scaling group are in different VPCs.
LimitExceeded.MinSizeLimitExceeded	The minimum number of instances is below the limit.
LimitExceeded	The quota limit is exceeded.
InvalidParameterValue.CvmError	Exception with CVM parameter validation.
InvalidParameterValue.LaunchConfigurationNotFound	The specified launch configuration was not found.
InvalidParameterValue.SubnetIds	The subnet information is invalid.
MissingParameter.InScenario	A parameter is missing in a specific scenario.
ResourceUnavailable.LaunchConfigurationStatusAbnormal	The launch configuration is exceptional.

Error Code	Description
InvalidParameterValue.LimitExceeded	The value exceeds the limit.

ModifyDesiredCapacity

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (ModifyDesiredCapacity) is used to modify the desired number of instances in the specified auto scaling group.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-11-01 11:13:15.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ModifyDesiredCapacity
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupId	Yes	No	String	ID of the scaling group. Example: asg-nvnlpbb8
DesiredCapacity	Yes	No	Uint64	Desired Capacity Example: 2

3. Output Parameters

Parameter	Type	Description

name		
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
LimitExceeded.DesiredCapacityLimitExceeded	The desired number of instances exceeds the limit.
ResourceUnavailable.AutoScalingGroupAbnormalStatus	The auto scaling group is exceptional.
ResourceNotFound.AutoScalingGroupIdNotFound	The scaling group does not exist.
InvalidParameterValue.Size	The value of maximum, minimum, or desired number of instances in the auto scaling group is invalid.
ResourceUnavailable.AutoScalingGroupDisabled	The auto scaling group is disabled.
InternalError	Internal error.

ModifyLoadBalancers

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (ModifyLoadBalancers) is used to modify the load balancers in a scaling group.

- This API can specify a new load balancer configuration for the scaling group. The new configuration overwrites the original load balancer configuration.
- If you want to clear the load balancers of the scaling group, specify only the scaling group ID but not the specific load balancers when you call this API.
- This API instantly modifies the load balancers of the scaling group and generates a scaling activity to asynchronously modify the load balancers of existing instances.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 16:51:25.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ModifyLoadBalancers
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupId	Yes	No	String	ID of the scaling group. Example: asg-12wjuh0s
LoadBalancerIds	No	No	Array of String	IDs of classic load balancers. Up to 20 IDs are supported. You can only specify one of LoadBalancerIds and ForwardLoadBalancers, or leave both of them empty. Example: ["lb-crhgatrf\n"]
ForwardLoadBalancers	No	No	Array of ForwardLoadBalancer	List of application load balancers. Up to 20 IDs are supported. You can only specify one of LoadBalancerIds and ForwardLoadBalancers, or leave both of them empty. Example:View
LoadBalancersCheckPolicy	No	No	String	Load balancer verification policy. Valid values: ALL and DIFF. Default value: ALL. • ALL: the verification is successful only if all load balancers are valid. Otherwise, the verification fails. • DIFF: only the changes in the load balancer parameters are verified. If the changes are valid, the verification is successful. Otherwise, the verification fails. Example: ALL

3. Output Parameters

Parameter name	Type	Description
ActivityId	String	ID of the scaling activity. Example: asa-67izy66g
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue.LbProjectInconsistent	The load balancer is in a different project.
ResourceUnavailable.AutoScalingGroupInActivity	The auto scaling group is active.
ResourceNotFound.AutoScalingGroupIdNotFound	The scaling group does not exist.
InvalidParameterValue.LbVpcInconsistent	The load balancer and the scaling group are in different VPCs.
InvalidParameterValue.ForwardLb	The application load balancer is incorrectly specified.
ResourceNotFound.LoadBalancerNotFound	The specified load balancer was not found.

RemoveSelectedInstance

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (RemoveSelectedInstance) is used to delete and remove CVM instances from a scaling group.

Default interface request rate limit: 20 times per second.

Interface update time: 2020-09-17 17:38:26.

The interface is both signature verification and authentication.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: RemoveSelectedInstance
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupId	Yes	No	String	ID of the scaling group. Example: asg-boz1qhnk
DetachInstanceIds	No	No	Array of String	IDs of CVM instances to be removed. Example: ["ins-61dyt49h"]
RemoveInstanceIds	No	No	Array of String	IDs of CVM instances to be deleted. Example: ["ins-61dyt49h"]

3. Output Parameters

Parameter name	Type	Description
ActivityId	String	ID of the scaling activity. Example: asa-q59zikez
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameterValue.LimitExceeded	The value exceeds the limit.
ResourceInsufficient.AutoScalingGroupBelowMinSize	The number of instances in the auto scaling group is below the minimum value.
ResourceNotFound.AutoScalingGroupIdNotFound	The scaling group does not exist.
InternalError	Internal error.
ResourceNotFound.InstancesNotInAutoScalingGroup	The target instance is not in the auto scaling group.
ResourceUnavailable.AutoScalingGroupInActivity	The auto scaling group is active.

SetInstancesProtection

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (SetInstancesProtection) is used to configure scale-in protection for instances.
When scale-in protection is enabled for an instance, the instance will not be removed in scale-in activities triggered when unhealthy instances are replaced, the alarm threshold is reached, or the desired capacity is changed.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 18:36:16.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: SetInstancesProtection
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupId	Yes	No	String	ID of the scaling group. Example: asg-2umy3jbd
InstanceIds	Yes	No	Array of String	Instance ID. Example: ["ins-b2d33ywt\n"]
ProtectedFromScaleIn	Yes	No	Bool	Specifies whether to enable scale-in protection.

Parameter name	Required	Allow NULL	Type	Description
				Example: true

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
ResourceNotFound.AutoScalingGroupIdNotFound	The scaling group does not exist.
ResourceNotFound.InstancesNotInAutoScalingGroup	The target instance is not in the auto scaling group.

Scheduled Task APIs
CreateScheduledAction

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (CreateScheduledAction) is used to create a scheduled task.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 16:54:41.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: CreateScheduledAction
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
AutoScalingGroupId	Yes	No	String	ID of the scaling group. Example: asg-2nr9xh8h
ScheduledActionName	Yes	No	String	Name of the scheduled task. The name can contain only letters, digits, underscores (_), hyphens (-), and decimal points (.) and must be up to 60 bytes in length. The name must be unique within the same scaling group. Example: scheduled-action-0
MaxSize	Yes	No	Uint64	Maximum number of instances set for the scaling group when the scheduled task is triggered. Example: 10
MinSize	Yes	No	Uint64	Minimum number of instances set for the scaling group when the scheduled task is triggered. Example: 4
DesiredCapacity	Yes	No	Uint64	Desired number of instances set for the scaling group when the scheduled task is triggered. Example: 6
StartTime	Yes	No	Datetime_iso	First triggering time of the scheduled task. Format: YYYY-MM-DDThh:mm:ss+08:00 . Example: 2018-08-28T23:00:00+08:00
EndTime	No	No	Datetime_iso	Scheduled task end time, which is the region time (UTC+8). The value follows the ISO 8601 standard. Format: YYYY-MM-DDThh:MM:ss+08:00 . This parameter should be specified together with Recurrence . The scheduled task will no longer take effect when the end time arrives. Example: 2019-01-01T00:00:00+08:00
Recurrence	No	No	String	Repeating mode of the scheduled task, which must be in standard cron format. This parameter and the EndTime parameter must be specified at the same time. Example: 0 23 /1 *

3. Output Parameters

Parameter name	Type	Description
ScheduledActionId	String	ID of the scheduled task. Example: asst-chwbkq4c
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
LimitExceeded.DesiredCapacityLimitExceeded	The desired number of instances exceeds the limit.
InvalidParameterValue.TooLong	Too many values.
LimitExceeded.ScheduledActionLimitExceeded	The number of scheduled tasks exceeds the limit.
InvalidParameterValue.InvalidScheduledActionNameIncludeIllegalChar	The scheduled task name contains invalid characters.
LimitExceeded.MaxSizeLimitExceeded	The maximum number of instances exceeds the limit.
ResourceNotFound.AutoScalingGroupIdNotFound	The scaling group does not exist.
InvalidParameterValue.CronExpressionIllegal	The cron expression specified for the scheduled task is invalid.
InvalidParameterValue.Size	The value of maximum, minimum, or desired number of instances in the auto scaling group is invalid.
MissingParameter	Parameter missing.
LimitExceeded.MinSizeLimitExceeded	The minimum number of instances is below the limit.
InvalidParameterValue.ScheduledActionNameDuplicate	The scheduled task name already exists.
InvalidParameterValue.EndTimeBeforeStartTime	The end time of the scheduled task is before the start time.
InvalidParameterValue.StartTimeBeforeCurrentTime	The start time of the scheduled task is before the current time.
InvalidParameterValue.TimeFormat	Wrong time format.
InternalError	Internal error.

DeleteScheduledAction

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DeleteScheduledAction) is used to delete the specified scheduled task.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-11-01 11:14:31.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DeleteScheduledAction
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
ScheduledActionId	Yes	No	String	ID of the scheduled task that you want to delete. Example: asst-chwbkq4c

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
ResourceNotFound.ScheduledActionNotFound	The specified scheduled task does not exist.

DescribeScheduledActions

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (DescribeScheduledActions) is used to query the details about one or more scheduled tasks.

- You can query the details of scheduled tasks based on information such as the scheduled task ID, scheduled task name, or scaling group ID. For more information about the supported filters, see Filters .
- If the filtering parameters are empty, a certain number (specified by Limit and 20 by default) of scheduled tasks of the current user will be returned.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-11-01 11:11:22.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: DescribeScheduledActions
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
ScheduledActionIds	No	No	Array of String	ID of one or more scheduled tasks, such as asst-am691zxo . You can specify a maximum of 100 instance IDs in each request. ScheduledActionIds and Filters cannot be specified at the same time. Example: ["asst-caa5ha40"]
Filters	No	No	Array of Filter	Filter conditions. • scheduled-action-id - String - Required: No - (Filter condition) Filter by scheduled task ID. • scheduled-action-name - String - Required: No - (Filter condition) Filter by scheduled task name. • auto-scaling-group-id - String - Required: No - (Filter condition) Filter by scaling group ID. Example:View
Offset	No	No	Uint64	Offset. Default value: 0. For more information on Offset, see the relevant section in API introduction. Example: 0
Limit	No	No	Uint64	Number of returned items. Default value: 20. Maximum value: 100. For more information on Limit, see the relevant section in API Introduction. Example: 20

3. Output Parameters

Parameter name	Type	Description
TotalCount	Uint64	Number of scheduled tasks that meet the specified conditions. Example: 1
ScheduledActionSet	Array of ScheduledAction	Details of scheduled tasks. Example:View
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
InvalidParameter.Conflict	Multiple parameters specified conflict and cannot co-exist.
InvalidParameterValue.Filter	Invalid filter.
InternalError	Internal error.

ModifyScheduledAction

1. Interface Description

Interface request domain name: as.api3.convergecloud.com.

This API (ModifyScheduledAction) is used to modify a scheduled task.

Default interface request rate limit: 20 times per second.

Interface update time: 2019-10-31 16:54:15.

The interface only verifies the signature but does not authenticate.

2. Input Parameters

The following list of request parameters only lists interface request parameters and some public parameters, For a complete list of public parameters see public request parameters.

Parameter name	Required	Allow NULL	Type	Description
Action	Yes	No	String	Public parameter, The value of this interface: ModifyScheduledAction
Version	Yes	No	String	Public parameter, The value of this interface: 2018-04-19
Region	No	No	String	Public parameter,Regional information A list of supported geographies is available through the DescribeRegions interface
ScheduledActionId	Yes	No	String	ID of the scheduled task to be modified. Example: asst-chwbkq4c
ScheduledActionName	No	No	String	Name of the scheduled task. The name can contain only letters, digits, underscores (_), hyphens (-), and decimal points (.) and must be up to 60 bytes in length. The name must be unique within the same scaling group. Example: scheduled-action-0
MaxSize	Yes	No	Uint64	Maximum number of instances set for the scaling group when the scheduled task is triggered. Example: 5
MinSize	Yes	No	Uint64	Minimum number of instances set for the scaling group when the scheduled task is triggered. Example: 0
DesiredCapacity	No	No	Uint64	Desired number of instances set for the scaling group when the scheduled task is triggered. Example: 3
StartTime	No	No	Datetime_iso	First triggering time of the scheduled task. Format: YYYY-MM-DDThh:mm:ss+08:00 . Example: 2018-08-28T23:00:00+08:00
EndTime	No	No	Datetime_iso	Scheduled task end time, which is region time (UTC+8). The value follows the ISO 8601 standard. Format: YYYY-MM-DDThh:MM:ss+08:00 . This parameter should be specified together with Recurrence . The scheduled task will no longer take effect when the end time arrives. Example: 2019-01-01T00:00:00+08:00
Recurrence	No	No	String	Repeating mode of the scheduled task, which must be in standard cron format. This parameter and the EndTime parameter must be specified at the same time. Example: 0 23 */1 *

3. Output Parameters

Parameter name	Type	Description
RequestId	String	Unique request id, Each request will return. The RequestId of the request needs to be provided when locating the problem.

4. Error Code

The following lists only the error codes related to the interface business logic, For other error codes, see public ErrorCode.

Error Code	Description
LimitExceeded.DesiredCapacityLimitExceeded	The desired number of instances exceeds the limit.
InvalidParameterValue.TooLong	Too many values.
LimitExceeded.ScheduledActionLimitExceeded	The number of scheduled tasks exceeds the limit.
InvalidParameterValue.InvalidScheduledActionNameIncludeIllegalChar	The scheduled task name contains invalid characters.
LimitExceeded.MaxSizeLimitExceeded	The maximum number of instances exceeds the limit.
InvalidParameterValue.CronExpressionIllegal	The cron expression specified for the scheduled task is invalid.
InvalidParameterValue.Size	The value of maximum, minimum, or desired number of instances in the auto scaling group is invalid.
LimitExceeded.MinSizeLimitExceeded	The minimum number of instances is below the limit.
InvalidParameterValue.ScheduledActionNameDuplicate	The scheduled task name already exists.
InvalidParameterValue.EndTimeBeforeStartTime	The end time of the scheduled task is before the start time.
InvalidParameterValue.StartTimeBeforeCurrentTime	The start time of the scheduled task is before the current time.
ResourceNotFound.ScheduledActionNotFound	The specified scheduled task does not exist.
InvalidParameterValue.TimeFormat	Wrong time format.
InternalError	Internal error.

Data structure

HostNameSettings

CVM HostName settings

Referenced by the following interfaces: DescribeLaunchConfigurations

name	required	NULL allowed	type	description
HostName	yes	yes	String	Hostname of the CVM. • The name cannot start or end with a period (.) or a hyphen (-), and cannot contain consecutive periods or hyphens. • This parameter is unavailable for Windows-based CVM instances. • Other types of instances (such as Linux-based instances): the name must be 2 to 40 characters in length and can contain multiple period-separated segments. Each segment can contain letters, digits, and hyphens (-). Example: aaa
HostNameStyle	no	yes	String	Style of the CVM hostname. Valid values: ORIGINAL and UNIQUE. Default value: ORIGINAL. • ORIGINAL: Auto Scaling directly transfers the value of the HostName parameter to CVM. CVM may add a serial number to the specified hostname. The hostnames of instances within the scaling group may conflict. • UNIQUE. The value of the HostName parameter is used as the prefix of a hostname. Auto Scaling and CVM expand the hostname, which can ensure that the hostname of each instance in the scaling group is unique. Example: ORIGINAL

LoginSettings

Describes login settings of an instance.

Referenced by the following interfaces: CreateLaunchConfiguration, UpgradeLaunchConfiguration

name	required	NULL allowed	type	description
Password	no	yes	String	Password that is used to log in to the instance. The limit on password complexity varies based on the OS type. The following requirements are provided: • The password for a Linux instance must be 8 to 16 characters in length and contain at least two of the following types: uppercase letters, lowercase letters, digits, and special characters ()`~!@#$%^&*-+=
KeyIds	no	no	Array of String	List of key IDs. After you associate keys, private keys can be used to access the corresponding instances. You can call the DescribeKeyPairs operation to obtain your key ID. Keys and passwords cannot be specified at a time. Keys cannot be specified for Windows OS. You can specify only one key when you purchase an instance. Example: ["skey-k8eypc1"]
KeepImageLogin	no	yes	Bool	Specifies whether to maintain the original login configuration of the image. This parameter cannot be used together with Password or KeyIds.N. You can set this parameter to TRUE only if you use a custom image, a shared image, or an imported external image to create an instance. Valid values: • TRUE: maintains the original login configuration of the image. • FALSE: does not maintain the original login configuration of the image. Default value: FALSE. Example: true

RunAutomationServiceEnabled

Describes the relevant information about Automation Service.

Referenced by the following interfaces: CreateLaunchConfiguration, DescribeLaunchConfigurations, UpgradeLaunchConfiguration

name	required	NULL allowed	type	description
Enabled	no	yes	Bool	Whether to enable Automation Service. Valid values: • TRUE: indicates enabling Automation Service. • FALSE: indicates disabling Automation Service. Example: True

SystemDisk

System disk configuration in a launch configuration. If this parameter is not specified, the default value is assigned to it.

Referenced by the following interfaces: CreateLaunchConfiguration, DescribeLaunchConfigurations, UpgradeLaunchConfiguration

name	required	NULL allowed	type	description
DiskType	yes	yes	String	System disk type. Value range: • LOCAL_BASIC: local disk • LOCAL_SSD: local SSD disk • CLOUD_BASIC: general-purpose cloud disk • CLOUD_PREMIUM: high-performance cloud disk • CLOUD_SSD: SSD cloud disk Default value: LOCAL_BASIC. Example: CLOUD_PREMIUM
DiskSize	no	yes	Uint64	The system disk size. Unit: GB. The default value is 50. Example: 50

InternetAccessible

This describes the internet accessibility of the instance created by a launch configuration and declares the internet usage billing method of the instance and the maximum bandwidth

Referenced by the following interfaces: CreateLaunchConfiguration, DescribeLaunchConfigurations, UpgradeLaunchConfiguration

name	required	NULL allowed	type	description
InternetChargeType	no	yes	String	The billing method of network resources. Valid values: • BANDWIDTH_PREPAID: prepaid based on the bandwidth. • TRAFFIC_POSTPAID_BY_HOUR: postpaid based on the traffic per hour. • BANDWIDTH_POSTPAID_BY_HOUR: postpaid based on the bandwidth per hour. • BANDWIDTH_PACKAGE: bandwidth resource plan. Default value: TRAFFIC_POSTPAID_BY_HOUR. Example: BANDWIDTH_POSTPAID_BY_HOUR
InternetMaxBandwidthOut	no	yes	Uint64	Public network egress bandwidth limit, unit: Mbps. Default value: 0 Mbps. The bandwidth cap varies by model. Example: 1
PublicIpAssigned	no	yes	Bool	Specifies whether to assign the public IP address. Valid values: • TRUE: assigns the public IP address. • FALSE: does not assign the public IP address. If the value of InternetMaxBandwidthOut is greater than 0, you can choose whether to assign the public IP address based on your business requirements. The system assigns the public IP address by default. If the value of InternetMaxBandwidthOut is 0, you cannot assign the public network bandwidth. Example: true

EnhancedService

This describes the conditions of enhancement services for the instance and their settings, such as the Agent of Cloud Security or Cloud Monitor.

Referenced by the following interfaces: CreateLaunchConfiguration, DescribeLaunchConfigurations, UpgradeLaunchConfiguration

name	required	NULL allowed	type	description
SecurityService	no	no	RunSecurityServiceEnabled	Enables the Cloud Security service. If you do not configure this parameter, the cloud security service is enabled by default. Example:View
MonitorService	no	no	RunMonitorServiceEnabled	Enables the Cloud Monitor service. If this parameter is not specified, the system enables the Cloud Monitor service by default. Example:View
AutomationService	no	no	RunAutomationServiceEnabled	Enables Automation Service. If the parameter is not specified, the default logic remains consistent with that of CVMs. Note: This field may return null, indicating that no valid values can be obtained. Example:View

InstanceNameSettings

Settings of CVM instance names.

Referenced by the following interfaces: DescribeLaunchConfigurations

name	required	NULL allowed	type	description
InstanceName	yes	no	String	Name of the CVM instance. The name cannot start or end with a period (.) or a hyphen (-), and cannot contain consecutive periods or hyphens. Other types of instances (such as Linux-based instances): the name must be 2 to 40 characters in length and can contain multiple period-separated segments. Each segment can contain letters, digits, and hyphens (-). This parameter may return null, indicating that no valid values can be obtained. Example: alice-1
InstanceNameStyle	no	no	String	Style of the CVM instance name. Valid values: ORIGINAL and UNIQUE. Default value: ORIGINAL. ORIGINAL: Auto Scaling sends the specified input parameter InstanceName to CVM directly. CVM may append a serial number to the value of the InstanceName parameter. The names of the CVM instances within the scaling group may conflict. UNIQUE: the value of the input parameter InstanceName is used as the prefix of an instance name. Auto Scaling and CVM expand the specified prefix, which can ensure that the name of each CVM instance in the scaling group is unique. This parameter may return null, indicating that no valid values can be obtained. Example: ORIGINAL

AutoScalingGroupAbstract

Brief information of a scaling group.

Referenced by the following interfaces: DescribeLaunchConfigurations

name	required	NULL allowed	type	description
AutoScalingGroupId	yes	no	String	ID of the scaling group. Example: asg-30wif85u

name	required	NULL allowed	type	description
AutoScalingGroupName	yes	no	String	Name of the scaling group. Example: alice-asg

InstanceMarketOptionsRequest

Options related to a CVM bidding request

Referenced by the following interfaces: CreateLaunchConfiguration, DescribeLaunchConfigurations, UpgradeLaunchConfiguration

name	required	NULL allowed	type	description
MarketType	no	yes	String	Marketplace option type. Valid value: spot. Example: spot
SpotOptions	yes	no	SpotMarketOptions	Bidding-related options. Example:View

ScheduledAction

This describes the information of a scheduled task.

Referenced by the following interfaces: DescribeScheduledActions

name	required	NULL allowed	type	description
ScheduledActionId	yes	no	String	ID of the scheduled task. Example: asst-jf898dps
ScheduledActionName	yes	no	String	Name of the scheduled task. Example: RunInstances
AutoScalingGroupId	yes	no	String	ID of the scaling group to which the scheduled task belongs. Example: asg-keqt9eg1
StartTime	yes	no	Datetime_iso	Scheduled task start time, which is region time (UTC+8). The value follows the ISO 8601 standard. Format: YYYY-MM-DDThh:MM:ss+08:00 . Example: 2022-02-22T16:00:00+08:00
Recurrence	yes	no	String	Repeating mode of the scheduled task, Example: 0 23 /2 *
EndTime	yes	no	Datetime_iso	Scheduled task end time, which is region time (UTC+8). The value follows the ISO 8601 standard. Format: YYYY-MM-DDThh:MM:ss+08:00 . Example: 2023-02-23T16:00:00+08:00
MaxSize	yes	no	Uint64	Maximum number of instances set in the scheduled task. Example: 20
DesiredCapacity	yes	no	Uint64	Desired number of instances set in the scheduled task. Example: 5
MinSize	yes	no	Uint64	Minimum number of instances set in the scheduled task. Example: 1
CreatedTime	yes	no	Datetime_iso	Time when the scheduled task was created. Set the value to a UTC time in the format of YYYY-MM-DDThh:mm:ssZ according to the ISO8601 standard. Example: 2022-02-21T02:19:52Z

SpotMarketOptions

Bidding-related options.

Referenced by the following interfaces: CreateLaunchConfiguration, DescribeLaunchConfigurations, UpgradeLaunchConfiguration

name	required	NULL allowed	type	description
MaxPrice	yes	no	String	Bidding price, such as "1.05". Example: 0.99
SpotInstanceType	no	yes	String	Bid request type. Currently, only the "one-time" type is supported. Default value: one-time. Example: one-time

InstanceChargePrepaid

Describes the billing method of an instance.

Referenced by the following interfaces: DescribeLaunchConfigurations

name	required	NULL allowed	type	description
Period	yes	no	Int64	The validity period of the instance. Unit: month. Valid values: 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 24, and 36. Example: 1
RenewFlag	no	no	String	Auto-renewal flag. Valid values: • NOTIFY_AND_AUTO_RENEW: expiration notifications and auto-renewal enabled; • NOTIFY_AND_MANUAL_RENEW: expiration notifications enabled but auto-renewal disabled; • DISABLE_NOTIFY_AND_MANUAL_RENEW: expiration notifications and auto-renewal disabled. Default value: NOTIFY_AND_MANUAL_RENEW. If the value of this parameter is set to NOTIFY_AND_AUTO_RENEW, the system will perform auto-renewal for the instance if the account balance is sufficient. Example: NOTIFY_AND_MANUAL_RENEW

ScalingPolicy

Alarm trigger policy.

Referenced by the following interfaces: DescribeScalingPolicies

name	required	NULL allowed	type	description
AutoScalingGroupId	yes	no	String	ID of the scaling group. Example: asg-1tgr5cf8
AutoScalingPolicyId	yes	no	String	ID of the alarm trigger policy. Example: asg-1tgr5cf8
ScalingPolicyName	yes	no	String	Name of the alarm trigger policy. Example: alice-asp
AdjustmentType	yes	no	String	Method that you want to use to change the desired capacity after the alarm is triggered. Valid values: • CHANGE_IN_CAPACITY: Increase or decrease the desired capacity. • EXACT_CAPACITY: Adjust to the specified desired capacity. • PERCENT_CHANGE_IN_CAPACITY: Adjust the desired capacity by percentage. Example: CHANGE_IN_CAPACITY
AdjustmentValue	yes	no	Int64	Adjusted desired capacity after the alarm is triggered. Example: 2
Cooldown	yes	no	Uint64	Cooldown period. Example: 60
MetricAlarm	yes	no	MetricAlarm	Alarm metric. Example:View
NotificationUserGroupIds	yes	no	Array of String	IDs of notification groups. Example: []

AutoScalingNotification

AS event notification

Referenced by the following interfaces: DescribeNotificationConfigurations

name	required	NULL allowed	type	description
AutoScalingGroupId	yes	no	String	ID of the scaling group. Example: asg-19kw47wk
NotificationUserGroupIds	yes	no	Array of String	IDs of user groups. Example: ["12529"]
NotificationTypes	yes	no	Array of String	List of notification events. Example: ["SCALE_OUT_SUCCESSFUL"]
AutoScalingNotificationId	yes	no	String	ID of the event notification. Example: asn-6r4waxc4

RunSecurityServiceEnabled

The information about the Cloud Security service.

Referenced by the following interfaces: CreateLaunchConfiguration, DescribeLaunchConfigurations, UpgradeLaunchConfiguration

name	required	NULL allowed	type	description
Enabled	no	yes	Bool	Whether to enable the Cloud Security service. Value range: • TRUE: Indicates that the Cloud Security service is enabled • FALSE: Indicates that the Cloud Security service is not enabled Default value: TRUE. Example: true

ServiceSettings

Service settings

Referenced by the following interfaces: CreateAutoScalingGroup, DescribeAutoScalingGroups, ModifyAutoScalingGroup

name	required	NULL allowed	type	description
ReplaceMonitorUnhealthy	no	no	Bool	Whether to enable unhealthy instance replacement. If this feature is enabled, Auto Scaling will replace instances that are flagged as unhealthy by Cloud Monitor. Default value: false. Example: False
ScalingMode	no	no	String	Valid values: CLASSIC_SCALING: this is the typical scaling method, which creates and terminates instances to perform scaling operations. WAKE_UP_STOPPED_SCALING: this scaling method first tries to start stopped instances. If the number of started instances is still insufficient, the system creates instances to perform scale-out operations. For scale-in operations, instances are terminated in the same manner as in the typical method. You can use the StopAutoScalingInstances API to stop instances in the scaling group. Scale-out operations triggered by alarms will still create instances. Default value: CLASSIC_SCALING. Example: CLASSIC_SCALING

MetricAlarm

Alarming metric of AS

Referenced by the following interfaces: CreateScalingPolicy, DescribeScalingPolicies, ModifyScalingPolicy

name	required	NULL allowed	type	description
ComparisonOperator	yes	no	String	Comparison operator. Valid values: • GREATER_THAN • GREATER_THAN_OR_EQUAL_TO • LESS_THAN • LESS_THAN_OR_EQUAL_TO • EQUAL_TO • NOT_EQUAL_TO Example: GREATER_THAN
MetricName	yes	no	String	Metric name. Valid values: • CPU_UTILIZATION: CPU utilization • MEM_UTILIZATION: memory utilization • LAN_TRAFFIC_OUT: outbound bandwidth of the private network • LAN_TRAFFIC_IN: inbound bandwidth of the private network • WAN_TRAFFIC_OUT: outbound bandwidth of the public network • WAN_TRAFFIC_IN: inbound bandwidth of the public network Example: CPU_UTILIZATION
Statistic	no	no	String	Statistics type. Valid values: • AVERAGE: average value • MAXIMUM: maximum value • MINIMUM: minimum value Default value: AVERAGE. Example: AVERAGE
Threshold	yes	no	Uint64	Alarming threshold. Valid values vary based on the metric name. • CPU_UTILIZATION: [1, 100]. Unit: %. • MEM_UTILIZATION: [1, 100]. Unit: %. • LAN_TRAFFIC_OUT: > 0. Unit: Mbit/s. • LAN_TRAFFIC_IN: > 0. Unit: Mbit/s. • WAN_TRAFFIC_OUT: > 0. Unit: Mbit/s. • WAN_TRAFFIC_IN: > 0. Unit: Mbit/s. Example: 10
Period	yes	no	Uint64	Time period in seconds. Enumerated values: 60 and 300. Example: 60
ContinuousTime	yes	no	Uint64	Number of recurrences. Valid values: 1 to 10. Example: 3

InstanceTag

Tags attached to the instance. You can use this parameter to bind tags to newly added instances.

Referenced by the following interfaces: CreateLaunchConfiguration, DescribeLaunchConfigurations, UpgradeLaunchConfiguration

name	required	NULL allowed	type	description
Key	yes	no	String	Tag key. Example: *TYPE*
Value	yes	no	String	Tag value. Example: region

LifecycleActionResultInfo

Result information of the lifecycle hook action

Referenced by the following interfaces: DescribeAutoScalingActivities, DescribeAutoScalingGroupLastActivities

name	required	NULL allowed	type	description
LifecycleHookId	yes	no	String	ID of the lifecycle hook. Example: ash-61dyt49h
InstanceId	yes	no	String	ID of the instance. Example: ins-61dyt49h
NotificationResult	yes	no	String	Notification result, which indicates whether the notification is successfully sent to CMQ. Example: null
LifecycleActionResult	yes	no	String	Execution result of the lifecycle hook action. Values: CONTINUE and ABANDON. Example: CONTINUE
ResultReason	yes	no	String	Reason for the result. Example: HEARTBEAT_TIMEOUT

LaunchConfiguration

Information set of eligible launch configurations.

Referenced by the following interfaces: DescribeLaunchConfigurations

name	required	NULL allowed	type	description
ProjectId	yes	no	Int64	ID of the project to which the instance belongs. Example: 0
LaunchConfigurationId	yes	no	String	ID of the launch configuration. Example: asc-61dyt49h
LaunchConfigurationName	yes	no	String	Name of the launch configuration. Example: TestLauchConfig0_0
InstanceType	yes	no	String	Instance type. Example: SA2.MEDIUM2
SystemDisk	yes	no	SystemDisk	System disk configurations of the instance. Example:View
DataDisks	yes	no	Array of DataDisk	Data disk configurations of the instance. Example:View
LoginSettings	yes	no	LimitedLoginSettings	The instance login information. Example:View
InternetAccessible	yes	no	InternetAccessible	Public network bandwidth settings. Example:View
SecurityGroupIds	yes	no	Array of String	The security group to which the instance belongs. Example: ["sg-brwzhm6d"]
AutoScalingGroupAbstractSet	yes	no	Array of AutoScalingGroupAbstract	Scaling group associated with the launch configuration. Example:View
UserData	yes	yes	String	Custom data. Example: dGVzdA==
CreatedTime	yes	no	Datetime_iso	Time when the launch configuration was created. Example: 2021-05-31T12:38:35Z
EnhancedService	yes	no	EnhancedService	Status of the enhanced services on the instance and their settings. Example:View
ImageId	yes	no	String	Image ID. Example: img-61dyt49h
LaunchConfigurationStatus	yes	no	String	Current status of the launch configuration. Valid values: • NORMAL: healthy. • IMAGE_ABNORMAL: an exception has occurred to the image setting in the launch configuration. • CBS_SNAP_ABNORMAL: an exception has occurred to the data disk snapshot setting in the launch configuration. • SECURITY_GROUP_ABNORMAL: an exception has occurred to the security group setting in the launch configuration. Example: NORMAL

name	required	NULL allowed	type	description
InstanceChargeType	yes	no	String	Billing method of the instance. The default value for a CVM instance is POSTPAID_BY_HOUR. • POSTPAID_BY_HOUR: the pay-as-you-go billing method. You are charged on an hourly basis. • SPOTPAID: pay by bidding. Example: POSTPAID_BY_HOUR
InstanceMarketOptions	yes	yes	InstanceMarketOptionsRequest	Marketplace-related options of the instance, such as spot instance related parameters. This parameter is required if the InstanceChargeType parameter is set to SPOTPAID. Example:View
InstanceTypes	yes	no	Array of String	List of instance models. Example: ["SA2.MEDIUM2"]
InstanceTags	yes	no	Array of InstanceTag	List of tags. Example:View
VersionNumber	yes	no	Int64	Version number. Example: 1
UpdatedTime	yes	no	Datetime_iso	The point in time when the multicast group was last updated. Example: 2023-09-04T11:42:46Z
CamRoleName	yes	no	String	Name of the CAM role. You can call the DescribeRoleList API and obtain the role name from the value of the roleName response parameter. Example: TestRole
LastOperationInstanceTypesCheckPolicy	yes	no	String	Value of InstanceTypesCheckPolicy upon the last operation. Example: ANY
HostNameSettings	yes	no	HostNameSettings	CVM hostname settings. Example:View
InstanceNameSettings	yes	no	InstanceNameSettings	Settings of CVM instance names Example:View
InstanceChargePrepaid	yes	no	InstanceChargePrepaid	Settings related to the prepaid billing method. You can use this parameter to configure attributes such as the subscription duration and whether to enable auto-renewal for subscription instances. This parameter is required for subscription instances. Example:View

Filter

Key-value pair filter, which is used for conditional filtering. For example, ID, name, and status.
*If there are multiple filters , the relationship between filters is logical AND .

- If a filter involves multiple values , the relationship between values of this filter is logical OR .

Take the filter of the [DescribeInstances](/tcloud/api/Compute/CVM/APIs/CVM/Version (2017-03-12)/instance APIs/RunInstances) API as an example. To query instances whose AZ (zone) is zone 1 *and* the billing mode (instance-charge-type) is monthly subscription *or* pay-as-you-go, the filters can be as follows:

```
Filters.0.Name=zone
&Filters.0.Values.1=ap-region-1
&Filters.1.Name=instance-charge-type
&Filters.1.Values.1=PREPAID
&Filters.1.Values.2=POSTPAID_BY_HOUR
```

Referenced by the following interfaces: DescribeAutoScalingActivities, DescribeAutoScalingGroups, DescribeAutoScalingInstances, DescribeLaunchConfigurations, DescribeLifecycleHooks, DescribeNotificationConfigurations, DescribeScalingPolicies, DescribeScheduledActions

name	required	NULL allowed	type	description
Name	yes	no	String	Field that you want to filter. Example: key
Values	yes	no	Array of String	Field value that you want to filter. Example: value

LimitedLoginSettings

This describes the configuration and information related to instance login. For security reasons, sensitive information is not described.

Referenced by the following interfaces: DescribeLaunchConfigurations

name	required	NULL allowed	type	description
KeyIds	yes	no	Array of String	List of key IDs. Example: ["skey-k8eypc1"]

AutoScalingGroup

Scaling group.

Referenced by the following interfaces: DescribeAutoScalingGroups

name	required	NULL allowed	type	description
AutoScalingGroupId	yes	no	String	ID of the scaling group. Example: asg-30wif85u
AutoScalingGroupName	yes	no	String	Scaling group name Example: alice use
AutoScalingGroupStatus	yes	no	String	Current status of the scaling group. Valid values: • NORMAL: healthy. • CVM_ABNORMAL: the launch configuration is abnormal. • LB_ABNORMAL: the load balancers are abnormal. • VPC_ABNORMAL: an error occurs in the VPC. • INSUFFICIENT_BALANCE: the account balance is insufficient. Example: NORMAL
CreatedTime	yes	no	Datetime_iso	Creation time in UTC format. Example: 2021-05-31T12:38:35Z
DefaultCooldown	yes	no	Int64	Default cooldown period. Unit: seconds. Example: 300
DesiredCapacity	yes	no	Int64	Desired Capacity Example: 0
EnabledStatus	yes	no	String	Whether the scaling group is enabled. Valid values: ENABLED and DISABLED . Example: ENABLED
ForwardLoadBalancerSet	yes	no	Array of ForwardLoadBalancer	List of application load balancers. Example:View
InstanceCount	yes	no	Int64	Number of instances. Example: 0
InServiceInstanceCount	yes	no	Int64	Number of instances that are in the IN_SERVICE state. Example: 0
LaunchConfigurationId	yes	no	String	Launch Configuration ID Example: asc-966b7lk2
LaunchConfigurationName	yes	no	String	Name of the launch configuration. Example: alice-lc
LoadBalancerIdSet	yes	no	Array of String	IDs of classic load balancers. Example: []
MaxSize	yes	no	Int64	Max Example: 10
MinSize	yes	no	Int64	Min Example: 0
ProjectId	yes	no	Int64	Project ID. Example: 0
SubnetIdSet	yes	no	Array of String	IDs of subnets. Example: ["subnet-rdm629r2"]
TerminationPolicySet	yes	no	Array of String	Termination policy. Example: OLDEST_INSTANCE
VpcId	yes	no	String	VPC ID. Example: vpc-h3cd5td5
ZoneSet	yes	no	Array of String	List of AZ IDs. Example: []
RetryPolicy	yes	no	String	Retry Policy Example: NO_RETRY
InActivityStatus	yes	no	String	Whether the scaling group is performing a scaling activity. IN_ACTIVITY : yes; NOT_IN_ACTIVITY : no. Example: IN_ACTIVITY
Tags	yes	no	Array of Tag	Tags attached to the scaling group. Example:View
ServiceSettings	yes	no	ServiceSettings	Service settings Example:View
Ipv6AddressCount	yes	no	Int64	Number of IPv6 addresses assigned to the instance. Example: 0
MultiZoneSubnetPolicy	yes	no	String	Policy applied when there are multiple availability zones/subnets. • PRIORITY: when creating instances, the system chooses the availability zone/subnet based on the order in the list from top to bottom. If the first instance is successfully created in the availability zone/subnet of the highest priority, all instances will be created in this availability zone/subnet. • EQUALITY: the system chooses the availability zone/subnet with the least instances for scale-out operations. This gives each availability zone/subnet an opportunity for scale-out and disperses the instances created during multiple scale-out operations across different availability zones/ subnets. Example: PRIORITY

ForwardLoadBalancer

Application load balancer

Referenced by the following interfaces: CreateAutoScalingGroup, DescribeAutoScalingGroups, ModifyLoadBalancers

name	required	NULL allowed	type	description
LoadBalancerId	yes	no	String	Load balancer ID. Example: lb-crhgatrf
ListenerId	yes	no	String	Application load balancer listener ID. Example: lb-61dyt49h
LocationId	no	no	String	ID of the forwarding rule. This parameter is required for Layer-7 listeners. Example: loc-61dyt49h
TargetAttributes	yes	no	Array of TargetAttribute	List of target rule attributes. Example:View
Region	no	no	String	Region of the CLB instance. The AS service region is used by default. The format is the same as that of the common parameter Region. For example: "ap-region". Example: ap-region

LifecycleHook

Lifecycle hook.

Referenced by the following interfaces: DescribeLifecycleHooks

name	required	NULL allowed	type	description
LifecycleHookId	yes	no	String	Lifecycle hook ID. Example: ash-ck4bkcjt
LifecycleHookName	yes	no	String	Name of the lifecycle hook. Example: demo
AutoScalingGroupId	yes	no	String	ID of the scaling group. Example: asg-lk5nn4xi
DefaultResult	yes	no	String	Default result of the lifecycle hook. Example: CONTINUE
HeartbeatTimeout	yes	no	Int64	Wait timeout period of the lifecycle hook. Example: 300
LifecycleTransition	yes	no	String	Applicable scenario of the lifecycle hook. Example: INSTANCE_TERMINATING
NotificationMetadata	yes	no	String	Additional information sent to the notification target. Example: ""
CreatedTime	yes	no	Datetime_iso	Creation time. Example: 2021-05-31T12:38:35Z
NotificationTarget	yes	no	NotificationTarget	Notification target Example:View
LifecycleTransitionType	yes	no	String	Lifecycle Hook Scenario Type Example: NORMAL

TargetAttribute

Load balancer target attribute

Referenced by the following interfaces: CreateAutoScalingGroup, DescribeAutoScalingGroups, ModifyLoadBalancers

name	required	NULL allowed	type	description
Port	yes	no	Uint64	Port. Example: 8080
Weight	yes	no	Uint64	Weight Example: 5

NotificationTarget

Notification target

Referenced by the following interfaces: CreateLifecycleHook, DescribeLifecycleHooks, UpgradeLifecycleHook

name	required	NULL allowed	type	description
TargetType	yes	no	String	Target type. Valid values: CMQ_QUEUE and CMQ_TOPIC. • CMQ_QUEUE: TDMQ for CMQ - queue model. • CMQ_TOPIC: TDMQ for CMQ - topic model. Example: TDMQ_CMQ_QUEUE
QueueName	no	no	String	Queue name. This parameter is required if the TargetType parameter is set to CMQ_QUEUE . Example: alice-queue
TopicName	no	no	String	Topic name. This parameter is required when the TargetType parameter is set to CMQ_TOPIC . Example: alice-topic
EnvironmentId	no	yes	String	Environment ID. Example: 84763
ClusterId	no	yes	String	Cluster ID Example: 29183

RunMonitorServiceEnabled

This describes the information related to the Cloud Monitor service.

Referenced by the following interfaces: CreateLaunchConfiguration, DescribeLaunchConfigurations, UpgradeLaunchConfiguration

name	required	NULL allowed	type	description
Enabled	no	yes	Bool	Whether to enable the Cloud Monitor service. Value range: • TRUE: Indicates that the Cloud Monitor service is enabled • FALSE: Indicates that the Cloud Monitor service is not enabled Default value: TRUE. Example: true

DataDisk

Configuration information of data disks in a launch configuration. If this parameter is not specified, no data disk will be purchased by default. You can specify only one data disk when purchasing it.

Referenced by the following interfaces: CreateLaunchConfiguration, DescribeLaunchConfigurations, UpgradeLaunchConfiguration

name	required	NULL allowed	type	description
DiskType	no	yes	String	Data disk type. Value range: • LOCAL_BASIC: local disk; • LOCAL_SSD: local SSD disk; • CLOUD_BASIC: general-purpose cloud disk • CLOUD_PREMIUM: high-performance cloud disk • CLOUD_SSD: SSD cloud disk Default value: LOCAL_BASIC. Example: CLOUD_PREMIUM
DiskSize	no	yes	Uint64	Data disk size, unit: GB. Minimum adjustment step is 10 GB, and different data disk types have different value ranges. Default value is 0, indicating no purchase of data disk. For more restrictions, see the product document. Example: 0
SnapshotId	no	yes	String	Data disk snapshot ID, such as snap-l8psqwnt . Example: snap-l8psqwnt

ActivtyRelatedInstance

Information of the instances related to the current scaling activity.

Referenced by the following interfaces: DescribeAutoScalingActivities, DescribeAutoScalingGroupLastActivities

name	required	NULL allowed	type	description
InstanceId	yes	no	String	Instance ID. Example: ins-61dyt49h
InstanceStatus	yes	no	String	Status of the instance in the scaling activity. Valid values: • INIT: the scaling activity is being initialized. • RUNNING: the instance is running. • SUCCESSFUL: the scaling activity is successful. • FAILED: the scaling activity has failed. Example: SUCCESSFUL

Instance

Instance information.

Referenced by the following interfaces: DescribeAutoScalingInstances

name	required	NULL allowed	type	description
InstanceId	yes	no	String	Instance ID. Example: ins-089dlszs
AutoScalingGroupId	yes	no	String	ID of the scaling group. Example: asg-02ofd40a
LaunchConfigurationId	yes	no	String	Launch Configuration ID Example: asc-23yf98ie
LaunchConfigurationName	yes	no	String	Name of the launch configuration. Example: Test asc-01_2.3
LifeCycleState	yes	no	String	Lifecycle status. Valid values: IN_SERVICE, CREATING, TERMINATING, ATTACHING, DETACHING, ATTACHING_LB, and DETACHING_LB. Example: IN_SERVICE
HealthStatus	yes	no	String	Health status. Valid values: HEALTHY and UNHEALTHY. Example: HEALTHY
ProtectedFromScaleIn	yes	no	Bool	Whether to enable scale-in protection. Example: false
Zone	yes	no	String	AZ. Example: ap-region-1
CreationType	yes	no	String	Creation type. Valid values: AUTO_CREATION and MANUAL_ATTACHING. Example: AUTO_CREATION
AddTime	yes	no	Datetime_iso	Time when the instance was added to the scaling group. Example: 2023-01-13T03:06:25Z
InstanceType	yes	no	String	The instance type. Example: S3.MEDIUM2
VersionNumber	yes	no	Int64	The version number. Example: 1
AutoScalingGroupName	yes	no	String	Scaling group name Example: Test asg-01_2.3

Activity

Information on eligible scaling activities.

Referenced by the following interfaces: DescribeAutoScalingActivities, DescribeAutoScalingGroupLastActivities

name	required	NULL allowed	type	description
AutoScalingGroupId	yes	no	String	ID of the scaling group. Example: asg-aaaaqqqq
ActivityId	yes	no	String	ID of the scaling activity. Example: asa-fznnvrja
ActivityType	yes	no	String	Type of the scaling activity. Valid values: • SCALE_OUT: scale-out activity • SCALE_IN: scale-in activity • ATTACH_INSTANCES: adding instances • REMOVE_INSTANCES: terminating instances • DETACH_INSTANCES: removing instances • TERMINATE_INSTANCES_UNEXPECTEDLY: terminating instances in the CVM console • REPLACE_UNHEALTHY_INSTANCE: replace unhealthy instances Example: SCALE_OUT
StatusCode	yes	no	String	Status of the scaling activity. Valid values: • INIT: the scaling activity is being initialized. • RUNNING: the scaling activity is running. • SUCCESSFUL: the scaling activity is successful. • PARTIALLY_SUCCESSFUL: the scaling activity is partially successful. • FAILED: the scaling activity has failed. • CANCELLED: the scaling activity is canceled. Example: SUCCESSFUL
StatusMessage	yes	no	String	Description of the scaling activity status. Example: Success
Cause	yes	no	String	Cause of the scaling activity. Example: Activity was launched in response to a difference between desired capacity and actual capacity.
Description	yes	no	String	Description of the scaling activity. Example: alice use
StartTime	yes	no	Datetime_iso	Start time of the scaling activity. Example: 2014-08-03T12:00:00+08:00
EndTime	yes	yes	Datetime_iso	End time of the scaling activity. Example: 2014-08-03T12:00:00+08:00
CreatedTime	yes	no	Datetime_iso	Time when the scaling activity was created. Example: 2014-08-03T12:00:00+08:00
ActivityRelatedInstanceSet	yes	no	Array of ActivtyRelatedInstance	Instance information related to the scaling activity. Example:View
StatusMessageSimplified	yes	no	String	Brief description of the scaling activity status. Example: Success
LifecycleActionResultSet	yes	no	Array of LifecycleActionResultInfo	Execution result of the lifecycle hook in the scaling activity. Example:View

Tag

Resource type and tag key-value pair

Referenced by the following interfaces: CreateAutoScalingGroup, DescribeAutoScalingGroups

name	required	NULL allowed	type	description
ResourceType	no	yes	String	Type of resource to which tags are bound. Valid value: auto-scaling-group. Example: instance
Key	yes	no	String	Tag key. Example: key
Value	yes	no	String	Tag value. Example: val
TagKey	no	yes	String	Tag key. Example: key
TagValue	no	yes	String	Tag value. Example: val

ErrorCode

Feature Description

If there is an Error field in the response, it means that the API call failed. For example:

```
{
    "Response": {
        "Error": {
            "Code": "AuthFailure.SignatureFailure",
            "Message": "The provided credentials could not be validated. Please check your signature is cor
rect."
        },
        "RequestId": "ed93f3cb-f35e-473f-b9f3-0d451b8b79c6"
    }
}
```

Code in Error indicates the error code, and Message indicates the specific information of the error.

Error Code List

Common Error Codes

Error codes	Description
AuthFailure.InvalidSecretId	Invalid key (not a Cloud API key type).
AuthFailure.MFAFailure	MFA failed.
AuthFailure.SecretIdNotFound	Key does not exist. Check if the key has been deleted or disabled in the console, and if not, check if the key is correctly entered. Note that whitespaces should not exist before or after the key.
AuthFailure.SignatureExpire	Signature expired. Timestamp and server time cannot differ by more than five minutes. Please ensure your current local time matches the standard time.
AuthFailure.SignatureFailure	Invalid signature. Signature calculation error. Please ensure you've followed the signature calculation process described in the Signature API documentation.
AuthFailure.TokenFailure	Token error.

Error codes	Description
AuthFailure.UnauthorizedOperation	The request does not have CAM authorization.
DryRunOperation	DryRun Operation. It means that the request would have succeeded, but the DryRun parameter was used.
FailedOperation	Operation failed.
InternalError	Internal error.
InvalidAction	The API does not exist.
InvalidParameter	Incorrect parameter.
InvalidParameterValue	Invalid parameter value.
LimitExceeded	Quota limit exceeded.
MissingParameter	A parameter is missing.
NoSuchVersion	The API version does not exist.
RequestLimitExceeded	The number of requests exceeds the frequency limit.
ResourceInUse	Resource is in use.
ResourceInsufficient	Insufficient resource.
ResourceNotFound	The resource does not exist.
ResourceUnavailable	Resource is unavailable.
UnauthorizedOperation	Unauthorized operation.
UnknownParameter	Unknown parameter.
UnsupportedOperation	Unsupported operation.
UnsupportedProtocol	HTTPS request method error. Only GET and POST requests are supported.
UnsupportedRegion	API does not support the requested region.

Service Error Codes

Error codes	Description
InvalidParameterValue.EndTimeBeforeStartTime	The end time of the scheduled task is before the

Error codes	Description
	start time.
ResourceUnavailable.AutoScalingGroupInActivity	The auto scaling group is active.
ResourceNotFound.LifecycleHookNotFound	The specified lifecycle hook was not found.
InvalidParameterValue.GroupNameDuplicated	The auto scaling group name already exists.
InvalidParameterValue.UserGroupIdNotFound	The user group does not exist.
ResourceInsufficient.AutoScalingGroupBelowMinSize	The number of instances in the auto scaling group is below the minimum value.
ResourceUnavailable.InstancesAlreadyInAutoScalingGroup	The instance already exists in the auto scaling group.
InvalidLaunchConfiguration	Invalid launch configuration.
ResourceInsufficient.AutoScalingGroupAboveMaxSize	The maximum number of instances in the auto scaling group is exceeded.
InvalidParameterValue.CronExpressionIllegal	The cron expression specified for the scheduled task is invalid.
InvalidParameterValue.ForwardLb	The application load balancer is incorrectly specified.
InvalidParameterValue.TimeFormat	Wrong time format.
InvalidFilter	Invalid filter.
InvalidParameter.MustOneParameter	A parameter is missing. One of the two parameters must be specified.
InvalidParameterValue.OnlyVpc	The account only supports VPCs.
ResourceNotFound.InstancesNotFound	The specified instance does not exist.

Error codes	Description
LimitExceeded.DesiredCapacityLimitExceeded	The desired number of instances exceeds the limit.
ResourceUnavailable.ProjectInconsistent	Project inconsistency.
InvalidParameterValue.LaunchConfigurationNameDuplicated	The launch configuration name already exists.
ResourceNotFound.AutoScalingGroupIdNotFound	The scaling group does not exist.
InvalidParameterValue.TooLong	Too many values.
ResourceUnavailable.AutoScalingGroupAbnormalStatus	The auto scaling group is exceptional.
InvalidParameter	Parameter error.
LimitExceeded	The quota limit is exceeded.
InvalidParameterValue.Range	The value is outside the specified range.
ResourceNotFound.ScheduledActionNotFound	The specified scheduled task does not exist.
InvalidLaunchConfigurationId	Invalid launch configuration ID.
MissingParameter	Parameter missing.
InvalidLaunchConfigurationId.NotFound	The launch configuration was not found.
InvalidParameterValue.Filter	Invalid filter.
LimitExceeded.MaxSizeLimitExceeded	The maximum number of instances exceeds the limit.
ResourceUnavailable.LaunchConfigurationStatusAbnormal	The launch configuration is exceptional.
ResourceNotFound.ScalingPolicyNotFound	The alarm policy does not exist.
ResourceUnavailable.LifecycleActionResultHasSet	The lifecycle action has already been set.

Error codes	Description
InvalidParameter.InScenario	The parameter is invalid in a specific scenario.
InternalError.CallTDMQError	TDMQ interface call failed.
InvalidParameterValue.LbVpcInconsistent	The load balancer and the scaling group are in different VPCs.
InvalidParameterValue.InstanceTypeNotSupported	Unsupported instance type
LimitExceeded.AutoScalingGroupLimitExceeded	The number of auto scaling groups exceeds the limit.
InvalidParameterValue.InvalidScheduledActionNameIncludeIllegalChar	The scheduled task name contains invalid characters.
ResourceNotFound.InstancesNotInAutoScalingGroup	The target instance is not in the auto scaling group.
AccountQualificationRestrictions	The requesting account failed to pass the qualification review.
InvalidParameterValue.LbProjectInconsistent	The load balancer is in a different project.
MissingParameter.InScenario	A parameter is missing in a specific scenario.
InvalidImageId.NotFound	The image cannot be found.
ResourceNotFound.LoadBalancerNotFound	The specified load balancer was not found.
ResourceNotFound.LifecycleHookInstanceNotFound	The instance corresponding to the lifecycle hook does not exist.
ResourceUnavailable.AutoScalingGroupDisabled	The auto scaling group is disabled.
LimitExceeded.MinSizeLimitExceeded	The minimum number of instances is below the limit.
ResourceInUse.ActivityInProgress	The auto scaling group is performing a scaling activity.

Error codes	Description
InvalidParameterValue.LimitExceeded	The value exceeds the limit.
InvalidParameterValue.CvmConfigurationError	Exception with CVM parameter validation.
ResourceInUse.InstanceInGroup	There are still normal instances in the auto scaling group.
InvalidParameterValue.ScalingPolicyNameDuplicate	The alarm policy name already exists.
ResourceNotFound.ResourceNotFoundTDMQPulsarTopicNotFound	TDMQ-Pulsar theme does not exist.
LaunchConfigurationQuotaLimitExceeded	The launch configuration quota is exceeded.
InvalidParameterValue.StartTimeBeforeCurrentTime	The start time of the scheduled task is before the current time.
InvalidParameterValue.CvmError	Exception with CVM parameter validation.
InvalidParameter.Conflict	Multiple parameters specified conflict and cannot co-exist.
InvalidParameterValue.LifecycleHookNameDuplicated	The lifecycle hook name already exists.
InvalidParameterValue.Size	The value of maximum, minimum, or desired number of instances in the auto scaling group is invalid.
InvalidParameterValue.NoResourcePermission	No resource permission.
InvalidParameterValue.SubnetIds	The subnet information is invalid.
InvalidParameterValue	The parameter value is invalid.
InvalidParameterValue.ScheduledActionNameDuplicate	The scheduled task name already exists.

Error codes	Description
ResourceUnavailable.CvmVpcInconsistent	The instance and the auto scaling group are in different VPCs.
InvalidLaunchConfigurationId.InUse	The launch configuration is being used.
InternalError.RequestError	Request error
LimitExceeded.ScheduledActionLimitExceeded	The number of scheduled tasks exceeds the limit.
InvalidParameterValue.LaunchConfigurationNotFound	The specified launch configuration was not found.
InternalError	Internal error.
CallCvmError	CVM API call failed.
ResourceNotFound.AutoScalingNotificationNotFound	The notification does not exist.
InvalidPermission	The account is not authorized to perform this operation.

Purchase Guide

Auto Scaling Service – Pricing Overview

The Auto Scaling service itself is free of charge. You do not incur any additional costs for using Auto Scaling to automatically manage the number of computing instances based on demand. This means that setting up scaling groups, configuring scaling policies, and using the Auto Scaling interface or APIs will not result in extra charges.

However, it is important to note that while Auto Scaling as a feature or management tool is free, the underlying resources it controls are not. You will be billed for any services that Auto Scaling creates, removes, or manages on your behalf. These associated services typically include:

1. CVM Instances (Cloud Virtual Machines)
    - When Auto Scaling automatically launches new CVM instances to handle increased traffic or workloads, you will be charged for those instances at the standard rates.
    - This includes charges based on the instance type, CPU, memory, storage (system and data disks), bandwidth, and any other attached resources.
    - If you manually add CVM instances to an Auto Scaling group, those too will continue to accrue regular charges.

2. Storage Services
    - If your CVM instances use cloud storage (e.g., Cloud Block Storage, SSDs, HDDs), these will incur separate storage costs.
    - Auto Scaling does not manage storage directly, but instances it provisions may have associated disks, which are billable.

3. Load Balancers
    - Auto Scaling is often used in conjunction with Load Balancers (such as CLB or ELB) to distribute traffic across scaled instances.
    - The use of load balancers incurs charges based on usage, data transfer, and number of connections.

4. Network and Bandwidth
    - Internet or internal traffic generated by scaled instances will be billed based on your configured

bandwidth package or pay-as-you-go model.

- Auto Scaling doesn't add extra fees for network configuration, but scaled instances using bandwidth will still be charged.