

Auto Scaling

FAQs

FAQs

Environment configuration

How can I ensure the configuration environments are consistent with Auto Scaling ECS instances? Alternatively, how can I directly add a previous system environment to the Auto Scaling instances?

When creating scaling configuration for ECS instances, you can use an ECS custom image template to create instances. If you need to sync internal system data (such as a previous system environment) when the ECS instances are running, we recommend that you install a custom rsync program.

In the created instances, after a restart, why is 127.0.0.1 added after /etc/hosts are cleared?

When your content is added after /etc/hosts in an image, and the custom image generated by this image is used to create an instance, its configuration will be restored to the system default settings. This means the added content will be cleared. If these settings need to be retained, add the script code in rc.local. Then, check if the information is in /etc/hosts. If not, you must add the script code again.

My Auto Scaling service is set up to automatically create instances, however, there is no fixed quantity. How can I ensure that my instances are scaled normally when using images in the image market?

If you need to scale to N instances that use the same image, you must buy N images from the image market in advance.

Can I buy images from the image market in batches?

Batch purchasing of images is not supported.

If a previously used image from the image market no longer exists,

how can I ensure that the scaling group instances set up can be scaled normally?

We recommend that you select a suitable replacement image from the image market to ensure your scaling group can scale normally.

Can I use a single product code to get images from different regions?

Yes. However, you must ensure that your desired regions support the image.

I have bought 100 images with the same product code. Can I use these images in all regions?

Images in the image market have region attributes. You must ensure the purchased images are supported by your desired regions.

Password and login troubleshooting

When Auto Scaling automatically creates instances, how do I view their passwords and subsequently log in to these instances?

Instances automatically created by Auto Scaling do not have the same password. In a Linux environment, we recommend that you set a public/private key certificate for SSH login without password.

If you do not want to set a public/private key certificate for SSH login without password, you must reset the password on the console, and then restart the instance to apply the new password, before you can log in.

Why are the passwords of instances created by Auto Scaling different from the password for my custom image?

Created ECS instances do not have the same password as the custom image. To ensure password security, we recommend that you set a public/private key certificate for SSH password-free logon.

If you do not want to set a public/private key certificate for SSH login without password, you must reset the password on the console, and then restart the instance to apply the new password, before you can log in.

When using a custom image to generate Linux system instances, can I manage the instances through SSH login without password?

You can set a public/private key certificate for SSH login without password as follows:

Establish a public key and private key in the custom image's ECS server-end instance.

Copy the ECS instance `idc.pub` to the client.

Delete the public key from the server-end.

Modify the SSH configuration file on the server-end.

Configure the client software.

Using SecureCRT configuration as an example:

Select the corresponding remote connection information.

Right-click on the **Attribute** option and select **SSH2**.

Clear the **Password** option and select **PublicKey**.

Click the **Attribute** button on the right side and select **Use Session Public Key Settings**

Select **Use ID or Certificate File**, and `idc.pub` (the public and private key files you previously copied from the server).

Scaling configuration and rules

What information should I provide for Auto Scaling troubleshooting?

We recommend that you provide your Auto Scaling activity ID and relevant logs to facilitate troubleshooting.

When creating a scaling configuration for Auto Scaling, why can I not select ECS instances that I have already bought?

Auto Scaling allows you to add existing ECS instances under the following conditions:

- The ECS instance you want to add and the scaling group are in the same region.
- The ECS instance you want to add and the instance in the scaling configuration are of the same type.
- The ECS instance you want to add is in **Running** status and cannot be found in another scaling group.
- The ECS instance you want to add is a Subscription or Pay-As-You-Go instance. Note that when removed from the scaling group, the ECS instance will not be released.

Can I add existing subscription instances to Auto Scaling?

Auto Scaling automatically creates Pay-As-You-Go instances by default. However, you can add your existing Monthly Subscription or Pay-As-You-Go instances to the Auto Scaling service.

Can I add a single ECS instance to different scaling groups?

No. This feature is not currently supported.

Can I add ECS instances with different rules to Auto Scaling groups?

You can only set one scaling configuration type (CPU or Memory) for each Auto Scaling group. However, each scaling group can be set with a different configuration type.

Can I increase the maximum number of instances in the Auto Scaling service?

If you want to apply for a higher instance quota (more than 100 instances), submit a ticket to Alibaba Cloud Technical Support.

In the Auto Scaling service, can the ECS instances in scaling groups have 8-core and 16-core configurations?

Yes. To request high-configurations for ECS instances, submit a ticket for **Request Higher Configuration** to Alibaba Cloud Technical Support.

When I remove an ECS instance from a scaling group and release the instance, can I save the ECS instance data?

Data from released ECS instances cannot be saved.

Because Auto Scaling will automatically release ECS instances, Auto Scaling ECS instances cannot be used to save application status information (for example, session) or related data (for example, databases and logs). Status information can be saved to an independent state server, database (for example, RDS) or centralized log storage (for example, Log Service).

If I use API to call DisableScalingGroup, will this automatically release added ECS instances?

Pay-As-You-Go instances created by Auto Scaling will not be automatically released when you execute DisableScalingGroup.

Will ECS instances that are automatically added into a scaling group be automatically included into the RDS and OCS IP address whitelists?

When ECS instances are added or removed, the instance IP addresses will be automatically added to or removed from the RDS access whitelist. However, OCS whitelists are not supported.

Is vertical scaling, and CPU and memory upgrades, supported by ECS instances added to scaling groups?

Auto Scaling does not support vertical scaling, which means it cannot automatically upgrade or downgrade an ECS instance's CPU, memory, or bandwidth.

Does Auto Scaling support periodic task repetition to repeat tasks within a specified time period?

On the Auto Scaling console, go to **Auto-trigger Task Management > Scheduled Tasks** to set scheduled tasks. If you do not set a repetition period, the task is only executed on the specified date and time. Otherwise, the designated execution time defaults as the time point of executing this task periodically.

How can I ensure that manually added ECS instances are not

removed from the scaling group?

If you need to add N subscription ECS instances to the scaling group, and do not want these instances to be automatically removed, you must set the following configuration:

- Set the minimum number of instances (MinSize) to **N** or greater.
- Set the RemovalPolicy's first selection rule to **OldestScalingConfiguration**.

Based on Auto Scaling rules, manually added ECS instances will not match any scaling configuration (since they were not created by a scaling configuration). Therefore, Auto Scaling will remove the automatically created ECS instances from the scaling group. Manually added ECS instances will only be removed after all the automatically created ECS instances have been removed.

Note that automatically created ECS instances will be automatically released when they are removed from a scaling group. Manually created and added ECS instances will not be released if they are removed from a scaling group.

NOTE: The preceding logic applies for healthy instances. If you stop a manually added ECS instance, Auto Scaling will view it as **unhealthy** and remove it. This is because Auto Scaling must ensure that the ECS instances in the scaling group are **healthy**.

Server Load Balancer and RDS

After Auto Scaling creates an ECS instance, will the new instance be automatically added to a Sever Load Balancer instance?

If a Server Load Balancer instance is specified in a scaling group, the scaling group will automatically add the ECS instances in the group to the specified Server Load Balancer instance.

When a scaling group is added in Auto Scaling, can I bind multiple Sever Load Balancer instances to the group?

By default, you can only bind one Server Load Balancer instance to each scaling group. To bind multiple Server Load Balancer instances, submit a ticket to apply for a higher quota to Alibaba Cloud Technical Support.

When Auto Scaling creates an ECS instance, can the instance be added to multiple Server Load Balancer instances?

Yes. However, you must submit a ticket to Alibaba Cloud Technical Support to apply for multiple Server Load Balancer instances to bind to a specified ECS instance.

Can I modify the weights of ECS instances added to an Auto Scaling group Server Load Balancer Instance?

Yes. By default, ECS instances under Auto Scaling Server Load Balancer instances have a weight of 50. This is suitable for most scenarios, as backend ECS instances of Auto Scaling groups normally carry the same services and are the same type. The Server Load Balancer also distributes traffic based on the weight ratio, not the actual number. This means that, if you have two backend ECS instances weighted 50 and 50 (with a ratio of 1:1), this is the same as if they were weighted 100 and 100.

I have a public network Server Load Balancer instance. If I create a scaling configuration, will its ECS instances need public bandwidth?

When a scaling configuration is created, you do not have to allocate public bandwidth to its ECS instances. However, we recommend that you buy at least 1 Mbps of ECS bandwidth, for easier ECS instance management.

Do I need to use Server Load Balancer, CloudMonitor, and RDS in combination with Auto Scaling?

Auto Scaling is an open elastic scaling platform, and can independently scale up or down ECS instances. It can be deployed either separately or in combination with the Server Load Balancer and RDS. Auto Scaling allows CloudMonitor to trigger up or down scaling actions for ECS instances. You can also use Auto Scaling's OpenAPI to connect to your desired monitoring system in order to trigger scaling actions.

Monitoring and automation

How does Auto Scaling determine if its ECS instances are available?

If the Server Load Balancer is available in the desired Auto Scaling group, it will check that the ports of the backend ECS instances are functional before forwarding requests to the ECS instances.

What are the triggering conditions for Auto Scaling alarms?

Monitoring alarms in Auto Scaling are triggered based on the CPU load, memory usage, average system load, and Internet and intranet inbound and outbound traffic. These are used to automatically increase or decrease the number of ECS instances.

Can Auto Scaling support dynamic scaling based on custom alarms

in CloudMonitor?

No. Dynamic scaling based on custom monitoring settings is not supported.

How can I ensure that application operations are completed before ECS instances are automatically released?

To ensure that Auto Scaling releases an instance after a task is completed on the instance, store an execution script in a custom image, and set up a command to automatically execute the script when the operating system is turned off.

Using CentOS as an example, you can create the following shell test script:

```
#!/bin/sh
# chkconfig: 0 10 90
# description: Test Service
echo "hello world!" <!-- unsure why the preview displays this text in heading format, please publish it as normal
body text -->
```

Here, 0 is the default start level. There are a total of 7 levels ranging from 0-6.

- Level 0: Shutdown
- Level 1: Single user mode
- Level 2: Multiuser command line mode with no network connection
- Level 3: Multiuser command line mode with network connection
- Level 4: Unavailable
- Level 5: Multiuser mode with graphic interface
- Level 6: Restart
- 10 is the start priority and 90 is the stop priority. The priority range is 0-100. The higher the number, the lower the priority.

Put the test file in the `/etc/rc.d/init.d/` directory and execute `chkconfig —level 0 test on`. This test script will then be executed each time the system shuts down.

How can I automate the deployment of the ECS applications created in a scaling group?

To automatically install or update a program, or automatically load code after an ECS instance is automatically created in a scaling group, you must store an execution script in a custom image and set up a command to automatically execute this script upon operating system startup.

Using CentOS as an example, you can create the following shell test script:

```
#!/bin/sh
# chkconfig: 6 10 90
```

```
# description: Test Service
echo "hello world!"
```

In the preceding output, 6 is the default start level. There are a total of 7 levels ranging from 0-6.

- Level 0: Shutdown
 - Level 1: Single user mode
 - Level 2: Multiuser command line mode with no network connection
 - Level 3: Multiuser command line mode with network connection
 - Level 4: Unavailable
 - Level 5: Multiuser mode with graphic interface
 - Level 6: Restart
- 10 is the start priority and 90 is the stop priority. The priority range is 0-100. The higher the number, the lower the priority.

Put the test file in the `/etc/rc.d/init.d/` directory and execute `chkconfig --level 6 test on`. This test script will then be executed each time the system starts up.

The following example shows how to use a script to install Phpwind. Put the Phpwind installer in the script for execution (you will need to enter the database password). An example output is as follows:

```
cd /tmp

echo "phpwind"

yum install -y \
unzip \
wget \
httpd \
php \
php-fpm \
php-mysql \
php-mbstring \
php-xml \
php-gd \
php-pear \
php-devel
chkconfig php-fpm on \
&& chkconfig httpd on
wget http://pwwfiles.oss-cn-hangzhou.aliyuncs.com/com/soft/phpwind_v9.0_utf8.zip \
&& unzip -d pw phpwind_v9.0_utf8.zip \
&& mv pw/phpwind_v9.0_utf8/upload/* /var/www/html \
&& wget http://ess.oss-cn-hangzhou.aliyuncs.com/ossupload_utf8.zip -O ossupload_utf8.zip \
&& unzip -d ossupload ossupload_utf8.zip \
&& /bin/cp -rf ossupload/ossupload_utf8/* /var/www/html/src/extensions/ \
&& chown -R apache:apache /var/www/html
service httpd start && service php-fpm start
echo "Install CloudMonitor"
wget http://update2.aegis.aliyun.com/download/quartz_install.sh
chmod +x quartz_install.sh
bash quartz_install.sh
```

```
echo "Installation complete"
```