

Alibaba Cloud

Auto Scaling FAQ

Document Version: 20220309

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.
3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).
5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.
6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

Style	Description	Example
 Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	 Danger: Resetting will result in the loss of user configuration data.
 Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	 Warning: Restarting will cause business interruption. About 10 minutes are required to restart an instance.
 Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: If the weight is set to 0, the server no longer receives new requests.
 Note	A note indicates supplemental instructions, best practices, tips, and other content.	 Note: You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click Settings > Network > Set network type .
Bold	Bold formatting is used for buttons, menus, page names, and other UI elements.	Click OK .
<code>Courier font</code>	Courier font is used for commands	Run the <code>cd /d C:/window</code> command to enter the Windows system folder.
<i>Italic</i>	Italic formatting is used for parameters and variables.	<code>bae log list --instanceid</code> <i>Instance_ID</i>
[] or [a b]	This format is used for an optional value, where only one item can be selected.	<code>ipconfig [-all -t]</code>
{ } or {a b}	This format is used for a required value, where only one item can be selected.	<code>switch {active stand}</code>

Table of Contents

1.FAQ about Auto Scaling billing -----	05
2.FAQ about scaling groups -----	06
3.FAQ about event-triggered tasks and scheduled tasks -----	16

1. FAQ about Auto Scaling billing

This topic provides answers to commonly asked questions about Auto Scaling billing.

- [Do I need to pay for Auto Scaling?](#)
- [Why are all my ECS instances released immediately after a payment in my account becomes overdue?](#)
- [Am I still charged for ECS instances in a scaling group after the scaling group is disabled?](#)

Do I need to pay for Auto Scaling?

Auto Scaling is free of charge. For example, you are not charged for activating Auto Scaling or creating scaling groups. However, you must pay for the resources of other Alibaba Cloud services used in Auto Scaling, such as Elastic Compute Service (ECS), Server Load Balancer (SLB), and ApsaraDB for RDS instances.

Why are all my ECS instances released immediately after a payment in my account becomes overdue?

Auto Scaling performs health checks on ECS instances. If an ECS instance is not in the running state, Auto Scaling considers the instance to be unhealthy. Auto Scaling automatically removes the unhealthy instance from the scaling group, or even releases it. For more information, see [ECS instance lifecycle](#).

If you have overdue payments in your account, pay-as-you-go and preemptible instances will be stopped. Auto Scaling considers stopped ECS instances to be unhealthy, and will remove and release these instances.



Warning Make sure that you have sufficient balance in your account. For information about status changes of ECS instances with overdue payments, see [欠费说明](#).

Am I still charged for ECS instances in a scaling group after the scaling group is disabled?

Yes, you are still charged for the ECS instances in a scaling group after the scaling group is disabled. You are charged not for the scaling group but for the ECS instances in the scaling group. If ECS instances still exist in the disabled scaling group, you must continue to pay for these instances.

2.FAQ about scaling groups

This topic provides answers to commonly asked questions about scaling groups.

- FAQ about scaling group configurations
 - Must Auto Scaling be used with Alibaba Cloud services such as SLB, Cloud Monitor, and RDS?
 - How many ECS instances can I add to a scaling group?
 - Can I add existing ECS instances to a scaling group?
 - Can I add existing subscription ECS instances to a scaling group?
 - Can I add an ECS instance to multiple scaling groups?
 - Can I control how Auto Scaling removes ECS instances from a scaling group?
 - After I disable a scaling group, does Auto Scaling release ECS instances that are automatically added to the scaling group?
- FAQ about scaling configurations
 - Can I specify multiple ECS instance types for a scaling configuration?
 - Can I configure 8-core or 16-core ECS instances for a scaling group?
 - How do I set the capacity of data disks for ECS instances that are automatically created by Auto Scaling?
 - If Auto Scaling uses images from Alibaba Cloud Marketplace to create ECS instances, how do I make sure that the ECS instances can be created?
 - Can I buy multiple images from Alibaba Cloud Marketplace at a time?
 - If Alibaba Cloud Marketplace no longer provides an image that is used to create ECS instances in Auto Scaling, how do I make sure that ECS instances can be created with the image?
 - Why is an error of unsupported Alibaba Cloud Marketplace images returned when Auto Scaling creates an ECS instance?
 - Can images of different regions use the same product code?
 - I purchased 100 images with the same product code. Can I use them in any region?
- FAQ about scaling activities
 - What information do I need to provide when I submit a ticket about Auto Scaling?
 - Why is a resource error returned when Auto Scaling creates an ECS instance?
 - How do I avoid scale-out failures due to insufficient inventory of an instance type?
 - Why is an ECS instance for which I enabled release protection automatically released from a scaling group?
 - How do I prevent Auto Scaling from automatically releasing an ECS instance that I manually added to a scaling group?
 - When Auto Scaling automatically adds or removes ECS instances to or from a scaling group, does Auto Scaling automatically add or remove the IP addresses of the ECS instances to or from the whitelists of the RDS instances or ApsaraDB for Memcache instances associated with the scaling group?
 - How do I prevent Auto Scaling from automatically removing an ECS instance that I manually added to a scaling group?
 - After ECS instances are removed from scaling groups and released, is the data of the ECS instances retained?

- [How do I delete ECS instances that are automatically created by Auto Scaling in a scaling group?](#)
- [When Auto Scaling automatically creates multiple ECS instances, will Auto Scaling attempt to recreate instances that fail to be created?](#)
- **FAQ about ECS instances in scaling groups**
 - [For an ECS instance that is automatically created by Auto Scaling, how do I obtain its password and log on to it?](#)
 - [Why am I unable to use the password set in the custom image to log on to an ECS instance that is automatically created by Auto Scaling?](#)
 - [How do I synchronize data to ECS instances in a scaling group?](#)
 - [Why is IP Address 127.0.0.1 deleted from the /etc/hosts file of ECS instances that are automatically created by Auto Scaling?](#)
- **FAQ about SLB instances associated with scaling groups**
 - [What does an SLB instance do after it is associated with a scaling group?](#)
 - [How do SLB instances work with scaling groups?](#)
 - [After an ECS instance is added to a scaling group, does Auto Scaling automatically add the ECS instance to an SLB instance as a backend server?](#)
 - [Can Auto Scaling add an automatically created ECS instance to multiple SLB instances as a backend server?](#)
 - [After an ECS instance is added to an SLB instance as a backend server, can I change the weight of the ECS instance?](#)
 - [After I associate a public SLB instance with a scaling group, do I need to set the public bandwidth for ECS instances when I create a scaling configuration for the scaling group?](#)
 - [Why is a health check error about the SLB instance returned when I create a scaling group?](#)
 - [How does an SLB instance determine whether a newly created ECS instance can process requests?](#)
 - [Why does the timeout period of a Layer-7 HTTP listener of an SLB instance exceed 60 seconds?](#)
- **FAQ about RDS instances associated with scaling groups**
 - [What does an RDS instance do after it is associated with a scaling group?](#)
 - [How do RDS instances work with scaling groups?](#)

Must Auto Scaling be used with Alibaba Cloud services such as SLB, Cloud Monitor, and RDS?

No, Auto Scaling is an open and flexible service and can be used alone without the need to combine Server Load Balancer (SLB), ApsaraDB RDS (RDS), or Cloud Monitor. We recommend that you use Auto Scaling with these Alibaba Cloud services to improve performance. You can deploy Auto Scaling with SLB or RDS and use Cloud Monitor to trigger scaling activities in Auto Scaling. For information about these Alibaba Cloud services, see the following topics:

- SLB: [What is CLB?](#)
- RDS: [What is ApsaraDB RDS?](#)
- Cloud Monitor: [Event-triggered task overview](#)

How many ECS instances can I add to a scaling group?

For information about the limits related to Auto Scaling, see [使用限制](#).

Can I add existing ECS instances to a scaling group?

Yes, you can add existing ECS instances to a scaling group. Make sure that the ECS instances meet the following conditions:

- The ECS instances and the scaling group must be in the same region. For more information, see [Regions and zones](#).
- The ECS instances are in the **Running** state. For more information, see [Instance lifecycle](#).
- The ECS instances are not added to other scaling groups.

Can I add existing subscription ECS instances to a scaling group?

Yes, you can add existing subscription or pay-as-you-go instances to a scaling group. You can add existing ECS instances to a scaling group, including subscription and pay-as-you-go instances. Additionally, Auto Scaling can automatically create pay-as-you-go and preemptible instances.

Can I add an ECS instance to multiple scaling groups?

No, you can add an ECS instance only to a single scaling group.

Can I control how Auto Scaling removes ECS instances from a scaling group?

Yes, you can control how Auto Scaling removes ECS instances from a scaling group. You can set a removal policy for a scaling group to preferentially remove the ECS instances created based on the Earliest Instance Created Using Scaling Configuration, Earliest Created Instance, or Most Recent Created Instance option. For more information, see [Create a scaling group](#).

After I disable a scaling group, does Auto Scaling release ECS instances that are automatically added to the scaling group?

No, Auto Scaling does not automatically release ECS instances after you disable a scaling group by using the Auto Scaling console or calling an API operation. For more information, see [Disable a scaling group](#).

Can I specify multiple ECS instance types for a scaling configuration?

Yes, you can specify multiple ECS instance types for a scaling configuration. This increases the success rate of creating ECS instances. However, make sure that the number of ECS instance types specified for a scaling configuration does not exceed the upper limit. For more information, see [使用限制](#).

Can I configure 8-core or 16-core ECS instances for a scaling group?

Yes, you can configure 8-core or 16-core ECS instances for a scaling group. If the current available instance types cannot meet your requirements, [submit a ticket](#) to request the use of more ECS instance types.

How do I set the capacity of data disks for ECS instances that are automatically created by Auto Scaling?

You can specify the capacity of data disks in the Storage section when you create a scaling configuration. For more information, see [Create a scaling configuration \(Elastic Compute Service\)](#).

If Auto Scaling uses images from Alibaba Cloud Marketplace to create ECS instances, how do I make sure that the ECS instances can be created?

If you want Auto Scaling to create N instances that use the same Alibaba Cloud Marketplace image, you must purchase the image N times from Alibaba Cloud Marketplace in advance.

Can I buy multiple images from Alibaba Cloud Marketplace at a time?

No, you can buy only a single image from Alibaba Cloud Marketplace at a time.

If Alibaba Cloud Marketplace no longer provides an image that is used to create ECS instances in Auto Scaling, how do I make sure that ECS instances can be created with the image?

We recommend that you use other images available in Alibaba Cloud Marketplace.

Why is an error of unsupported Alibaba Cloud Marketplace images returned when Auto Scaling creates an ECS instance?

If you specify an Alibaba Cloud Marketplace image that you have not purchased in the scaling configuration, the following error is returned when an ECS instance is created based on the scaling configuration:

```
Fail to create Instance into scaling group("The specified image is from the image market. You have not bought it or your quota has been exceeded.").
```

Auto Scaling cannot create ECS instances by using Alibaba Cloud Marketplace images that you have not purchased. To use Alibaba Cloud Marketplace images to create ECS instances in Auto Scaling, you must purchase the images from [Alibaba Cloud Marketplace](#) in advance.

Can images of different regions use the same product code?

Yes, images of different regions can use the same product code as long as the images are supported in the regions.

I purchased 100 images with the same product code. Can I use them in any region?

No, you cannot use an image in all regions. Alibaba Cloud Marketplace images are region-specific. If you want to use an image in a specific region, purchase the image for that region.

What information do I need to provide when I submit a ticket about Auto Scaling?

When you [submit a ticket](#), provide the ID of the scaling activity or relevant logs to facilitate troubleshooting. The ID of a scaling activity is specified by the `ScalingActivityId` parameter.

For information about how to query scaling activities, see [View the details of a scaling activity](#).

Why is a resource error returned when Auto Scaling creates an ECS instance?

If the following errors are returned, ECS resources may be insufficient. We recommend that you change the zone for the scaling group and try again.

```
Fail to create Instance into scaling group("The resource is out of usage").
```

```
Fail to create Instance into scaling group(The specified region is in resource control, please try later.).
```

How do I avoid scale-out failures due to insufficient inventory of an instance type?

When you create a scaling group, specify multiple zones by selecting vSwitches in different zones. When you create a scaling configuration for the scaling group, select multiple ECS instance types. Then, if an instance type is unavailable in one zone, Auto Scaling automatically selects another instance type with sufficient inventory in that zone. If all instance types are unavailable in that zone, Auto Scaling switches to another zone. For more information, see [Create a scaling group](#) and [Create a scaling configuration \(Elastic Compute Service\)](#).

Why is an ECS instance for which I enabled release protection automatically released from a scaling group?

Auto Scaling can automatically release an ECS instance created during a scale-out event even if you enabled release protection by using the ECS console or calling the [ModifyInstanceAttribute](#) operation.

To prevent an ECS instance from being automatically released, you must put the ECS instance into the protected state. For more information, see [Put an ECS instance into the Protected state](#).

How do I prevent Auto Scaling from automatically releasing an ECS instance that I manually added to a scaling group?

To prevent an ECS instance from being automatically released, you must put the ECS instance into the protected state. For more information, see [Put an ECS instance into the Protected state](#).

When Auto Scaling automatically adds or removes ECS instances to or from a scaling group, does Auto Scaling automatically add or remove the IP addresses of the ECS instances to or from the whitelists of the RDS instances or ApsaraDB for Memcache instances associated with the scaling group?


Auto Scaling can automatically add or remove the IP addresses of ECS instances to or from the whitelists of RDS instances, but not the ApsaraDB for Memcache instances.

How do I prevent Auto Scaling from automatically removing an ECS instance that I manually added to a scaling group?

Assume that you want to prevent Auto Scaling from automatically removing the 100 ECS instances that you manually added to a scaling group. Perform the following operations to configure the scaling group:

- Set the minimum number of ECS instances in the scaling group to a value greater than or equal to 100.
- Set the removal policy to **Earliest Instance Created Using Scaling Configuration** for the first step.

Manually added ECS instances are not created based on scaling configurations. Therefore, the preceding settings ensure that Auto Scaling preferentially selects, removes, and releases automatically created ECS instances. Auto Scaling selects and removes manually added ECS instances only after all automatically created ECS instances are removed. When Auto Scaling removes a manually created ECS instance, Auto Scaling does not release the ECS instance.

 **Note** To prevent Auto Scaling from automatically removing a manually added ECS instance, do not stop the ECS instance. If the ECS instance is stopped, Auto Scaling considers that the ECS instance is unhealthy and automatically removes it.

After ECS instances are removed from scaling groups and released, is the data of the ECS instances retained?

No, Auto Scaling does not retain any data of ECS instances after the instances are removed and released. Therefore, do not store application status information or important data, such as sessions, databases, and logs, on ECS instances that belong to a scaling group. If the status information of your applications must be stored, we recommend that you store the status information on an independent state server or in a database or service. For example, you can store the status information on an independent ECS instance, in RDS, or in Log Service.

How do I delete ECS instances that are automatically created by Auto Scaling in a scaling group?

Log on to the Auto Scaling console, find the ECS instances that you want to delete, and then delete them. For more information, see [Manually remove or delete an ECS instance](#).

When Auto Scaling automatically creates multiple ECS instances, will Auto Scaling attempt to recreate instances that fail to be created?

No, Auto Scaling will not attempt to recreate instances that fail to be created. Auto Scaling does not guarantee that all ECS instances are created in a scaling activity. If some instances fail to be created during a scaling activity, Auto Scaling considers that the scaling activity is complete without attempting to recreate the failed instances. You can view the status of a scaling activity in the Auto Scaling console. For more information, see [View the details of a scaling activity](#).

Assume that Auto Scaling attempts to automatically create 20 ECS instances. 19 instances are created, and one instance fails to be created. Then, Auto Scaling adds the 19 created instances to the scaling group but does not attempt to recreate the failed one. The scaling activity is complete but is in the **Warning** state.

Basic Information	<div><input type="text" value="Enter a scaling activity ID"/></div> <div><input type="button" value="Search"/></div>					<div><input type="button" value="Refresh"/></div> <div><input type="button" value="List"/></div> <div><input type="button" value="Chart"/></div>
ECS Instances						
Rolling Update						
Scaling Activities	Scaling Activity ID	Status (All)	Instances After Scaling	Start Time	End Time	Description
	asa-1-	Successful	-	Aug 17, 2020 9:36 AM	Aug 17, 2020 9:37 AM	Remove "198" ECS instances
Instance Configurati...	asa-1-	Warning	198	Aug 17, 2020 9:35 AM	Aug 17, 2020 9:35 AM	Add "5" ECS instances
Scaling Rules	asa-1-	Successful	196	Aug 17, 2020 9:34 AM	Aug 17, 2020 9:35 AM	Remove "2" ECS instances
Event Notifications	asa-1-	Failed	-	Aug 17, 2020 9:34 AM	Aug 17, 2020 9:34 AM	Add "5" ECS instances
Lifecycle Hooks						
Scaling Activity ID: asa-1-						
Status	Warning		Instances After Scaling	198		
Start Time	Aug 17, 2020 9:35 AM		End Time	Aug 17, 2020 9:35 AM		
Cause	A user requests to execute scaling rule "asr-1-", changing the Total Capacity from "196" to "201".		Status Description	"2" ECS instances are added		
Details	Ignore to create "3" instances("Backend server quota exceeded in load balancer "lb-bj-";) new ECS instances "i-bj-", "i-bj- are created.		Description	Add "5" ECS instances		

For an ECS instance that is automatically created by Auto Scaling, how do I obtain its password and log on to it?

If the ECS instance is a Linux instance and a Secure Shell (SSH) key pair is set in the scaling configuration for creating the ECS instance, you can use the SSH key pair to log on to the ECS instance. Auto Scaling does not allow you to set a unified password for automatically created ECS instances. For Linux ECS instances, we recommend that you set an SSH key pair in the scaling configuration.

If you do not set an SSH key pair in the scaling configuration, you must reset the password for the ECS instance in the ECS console. The new password takes effect after the ECS instance is restarted. Then, you can use the new password to log on to the ECS instance.

Why am I unable to use the password set in the custom image to log on to an ECS instance that is automatically created by Auto Scaling?

For security reasons, Auto Scaling does not use the password set in the custom image as the logon password of automatically created ECS instances. We recommend that you set an SSH key pair in the scaling configuration for creating an ECS instance. Then, you can use the SSH key pair to log on to the ECS instance.

If you do not set an SSH key pair in the scaling configuration, you must reset the password for the ECS instance in the ECS console. The new password takes effect after the ECS instance is restarted. Then, you can use the new password to log on to the ECS instance.

How do I synchronize data to ECS instances in a scaling group?

You can configure a custom image in the scaling configuration for creating ECS instances. In this way, you can pass custom data to the created ECS instances. To synchronize data to or from a running ECS instance, we recommend that you install rsync on the ECS instance and use rsync to transfer data.

Why is IP Address 127.0.0.1 deleted from the `/etc/hosts` file of ECS instances that are automatically created by Auto Scaling?

If you use a custom image where the `/etc/hosts` file is modified to create ECS instances, the file is restored to default settings when Auto Scaling automatically creates ECS instances from the custom image. Therefore, your modifications to the `/etc/hosts` file are lost. To retain the modifications to the `/etc/hosts` file, add a script to the `rc.local` file to check whether the corresponding information exists in the `/etc/hosts` file and add the information if the information does not exist.

What does an SLB instance do after it is associated with a scaling group?

The SLB instance forwards traffic to ECS instances in the scaling group based on the routing method. By associating an SLB instance with a scaling group, you can improve the performance and availability of your applications. For more information about SLB, see [What is CLB?](#).

How do SLB instances work with scaling groups?

You can specify an SLB instance for one or more scaling groups. Then, Auto Scaling automatically adds ECS instances in the scaling groups to the SLB instance as backend servers. By default, the weight of an ECS instance added to an SLB instance as a backend server is 50. For more information about SLB backend servers, see [Add a default backend server](#).

After an ECS instance is added to a scaling group, does Auto Scaling automatically add the ECS instance to an SLB instance as a backend server?

Yes, Auto Scaling automatically adds the ECS instance as a backend server to the SLB instance associated with the scaling group. You must associate the SLB instance with the scaling group in advance.

Can Auto Scaling add an automatically created ECS instance to multiple SLB instances as a backend server?

Yes, Auto Scaling can add an automatically created ECS instance to multiple SLB instances associated with the scaling group. Only a limited number of SLB instances can be associated with a scaling group. For more information, see [使用限制](#).

If you want to associate more SLB instances with a scaling group, [submit a ticket](#).

After an ECS instance is added to an SLB instance as a backend server, can I change the weight of the ECS instance?

Yes, you can change the weight of an ECS instance after it is added to an SLB instance as a backend server. For more information, see [Change the weight of a backend server](#).

An SLB instance balances loads among its backend servers based on the ratio of their weights, instead of the weight values. Assume that you have two ECS instances and their weights are both 50. If you change both the weights to 100, the SLB instance balances loads among the two ECS instances in the same way. This is because the weight ratio remains unchanged and is still 1:1. Typically, ECS instances in a scaling group provide the same service and are of the same instance type. Therefore, their weights are the same. By default, the weight of an ECS instance added to an SLB instance as a backend server is 50.

After I associate a public SLB instance with a scaling group, do I need to set the public bandwidth for ECS instances when I create a scaling configuration for the scaling group?

No, public bandwidth is an optional setting for ECS instances. However, we recommend that you set the public bandwidth to at least 1 Mbit/s when you create a scaling configuration. This facilitates the management of ECS instances.

Why is a health check error about the SLB instance returned when I create a scaling group?

The following error is returned if health check is disabled for the SLB instance:

```
The current health check type of load balancer "xxxx" does not support this action.
```

Make sure that health check is enabled for SLB instances that are associated with scaling groups. For more information, see [Configure health checks](#).

How does an SLB instance determine whether a newly created ECS instance can process requests?

After a newly created ECS instance is added to a scaling group, the SLB instance associated with the scaling group checks whether the corresponding ports on the ECS instance can respond to requests. The SLB instance forwards requests to the new ECS instance only when the ports on the ECS instance can respond to requests.

Why does the timeout period of a Layer-7 HTTP listener of an SLB instance exceed 60 seconds?

Problem: The timeout period of an HTTP listener is about 60 seconds on an SLB instance. However, the timeout period exceeds 60 seconds for all HTTP requests, or the SLB instance directly returns a 504 error code for an HTTP request.

Cause: You can set the timeout period of an HTTP listener to make sure that an HTTP request is responded within the specified period. However, the total timeout period depends on the number of ECS instances configured for the SLB instance.

If an HTTP request times out on the first ECS instance, the SLB instance forwards the HTTP request to the second ECS instance. The same rule applies to the remaining ECS instances until all ECS instances are polled or until the request is processed by an ECS instance. Assume that three ECS instances are added as backend servers to an SLB instance. The total HTTP timeout period reaches about 180 seconds.

Other services may also affect the HTTP timeout period of an SLB instance. We recommend that you do not rely on the HTTP timeout period configured on the SLB instance to monitor the processing of HTTP requests. Instead, we recommend that you configure the timeout period in applications deployed on ECS instances.

What does an RDS instance do after it is associated with a scaling group?

RDS is a stable, reliable, and scalable online database service. By associating an RDS instance with a scaling group, you can back up data of the scaling group to the RDS instance based on custom backup policies. This improves the security and reliability of the data and allows you to restore the data when it is deleted by mistake. For more information about RDS, see [What is ApsaraDB RDS?](#).

How do RDS instances work with scaling groups?

After an RDS instance is associated with a scaling group, Auto Scaling automatically adds the internal IP addresses of newly created ECS instances in the scaling group to the whitelist of the RDS instance. This allows the ECS instances to access the RDS instance over the internal network. You can associate multiple RDS instances with a scaling group. For more information about whitelists of RDS instances, see [Use a database client or the CLI to connect to an ApsaraDB RDS for MySQL instance](#).

3. FAQ about event-triggered tasks and scheduled tasks

This topic provides answers to commonly asked questions about event-triggered tasks and scheduled tasks.

- [Can I create tasks that are executed periodically?](#)
- [What conditions are used by event-triggered tasks to trigger scaling activities?](#)
- [How do I configure conditions for event-triggered tasks?](#)
- [How do I use event-triggered tasks to delete instances created by Auto Scaling?](#)
- [Does Auto Scaling support automatic scaling based on custom Cloud Monitor metrics?](#)
- [Which is prioritized to be executed: an event-triggered task or a scheduled task?](#)
- [Does Auto Scaling support automatic scaling of data disks?](#)

Can I create tasks that are executed periodically?

Yes, you can create scheduled tasks that are executed periodically. For more information, see [Create a scheduled task](#).

What conditions are used by event-triggered tasks to trigger scaling activities?

Event-triggered tasks can trigger scaling activities based on Cloud Monitor metrics such as the CPU utilization, memory usage, average system load, and inbound or outbound traffic.

How do I configure conditions for event-triggered tasks?

Before you configure conditions for event-triggered tasks, you must install the latest version of Cloud Monitor on your Elastic Compute Service (ECS) instances. For more information, see [Install the Cloud Monitor Java agent](#).

Then, you can select required conditions when you create event-triggered tasks. For more information, see [Create an event-triggered task](#).

How do I use event-triggered tasks to delete instances created by Auto Scaling?

To use an event-triggered task to delete instances created by Auto Scaling, set the trigger rule of the event-triggered task to a rule that deletes instances created by Auto Scaling. For more information, see [Create a scaling rule](#) and [Create an event-triggered task](#).

Does Auto Scaling support automatic scaling based on custom Cloud Monitor metrics?

Yes, Auto Scaling can scale ECS instances based on custom Cloud Monitor metrics. For more information, see [Custom monitoring event-triggered tasks](#).

Which is prioritized to be executed: an event-triggered task or a scheduled task?

Event-triggered tasks and scheduled tasks cannot be triggered at the same time. They are independent of each other and are not prioritized over each other.

If an event-triggered task is rejected and the trigger conditions are still met, the event-triggered task is executed after the current scaling activity is complete.

You can set the `Retry Interval` parameter for a scheduled task. This ensures that the scheduled task can be triggered again after it is rejected. For more information, see [Create a scheduled task](#).

Does Auto Scaling support automatic scaling of data disks?

No, Auto Scaling does not support automatic scaling of data disks. Auto Scaling can automatically increase or decrease the number of ECS instances in a scaling group. Auto Scaling cannot automatically increase or decrease the number or sizes of data disks on an ECS instance.