Alibaba Cloud

Server Load Balancer Backend servers

Document Version: 20220705

C-J Alibaba Cloud

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.

- You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloudauthorized 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 and/or its affiliates 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 |
|-----------------|--|--|
| <u>↑</u> 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. |
| O 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. |
| C) 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. |
| Courier font | Courier font is used for commands | Run the cd /d C:/window command to enter the Windows system folder. |
| Italic | Italic formatting is used for parameters and variables. | bae log listinstanceid Instance_ID |
| [] or [a b] | This format is used for an optional value, where only one item can be selected. | ipconfig [-all -t] |
| {} or {alb} | This format is used for a required value, where only one item can be selected. | switch {active stand} |

Table of Contents

| 1.Backend server overview | 05 |
|---|----|
| 2.Default server groups | 07 |
| 2.1. Add a default backend server | 07 |
| 2.2. Change the weight of a backend server | 09 |
| 2.3. Remove a backend server | 10 |
| 3.Create a vServer group | 11 |
| 4.Active/standby server groups | 14 |
| 4.1. Create a primary/secondary server group | 14 |
| 4.2. Delete a primary/secondary server group | 15 |
| 5.Add private IP addresses of ENIs to backend servers | 17 |
| 6.Backend server FAQ | 19 |

1.Backend server overview

Before you use a Classic Load Balancer (CLB) instance, you must specify Elastic Compute Service (ECS) instances as the backend servers of the CLB instance to receive client requests.

Introduction

CLB creates a server group for multiple ECS instances in the same region and sets a virtual IP address for the server group to ensure high performance and high availability. You can also use vServer groups to manage backend servers. You can associate listeners with different server groups. This way, CLB can distribute requests to backend servers that use different ports.

Note If you associate a listener with a vServer group, the listener distributes requests to the ECS instances in the vServer group instead of the ECS instances in the default server group.

Limits

You can add ECS instances to or remove ECS instances from a CLB instance anytime. CLB distributes network traffic across groups of ECS instances. Before you use CLB, make sure that health checks are enabled and at least one ECS instance works as expected to ensure service stability.

When you add backend servers, take note of the following items:

- CLB does not support cross-region deployment. The ECS instances that you add to a CLB instance and the CLB instance must belong to the same region.
- You can add two ECS instances that run different operating systems to a CLB instance. However, the applications deployed on the ECS instances must be the same and have consistent data. To simplify management and maintenance, we recommend that you add ECS instances that use the same operating system to a CLB instance.
- You must configure a listener for each application deployed on an ECS instance. Each CLB instance supports up to 50 listeners. CLB uses listening ports to receive client requests and forward the requests to backend ports that are used by the applications on ECS instances.
- You can specify a weight for each ECS instance in a server group. An ECS instance with a higher weight receives more requests.
- If session persistence is enabled, requests may not be evenly distributed to each backend server. We recommend that you disable session persistence and check whether the issue persists.

If requests are not evenly distributed, troubleshoot the issue by performing the following steps:

- i. Count the numbers of access log entries generated on different ECS instances within a specified time period.
- ii. Check whether the numbers of access log entries generated on different ECS instances have deviations based on the CLB configurations. If session persistence is enabled, you must exclude access log entries that contain the same IP address. If the ECS instances have different weights, you must check whether the numbers of access log entries are also weighted.
- When an ECS instance performs hot migration, persistent connections to CLB may be closed. Make sure that your applications are configured with the automatic reconnection mechanism.

Primary/secondary mode

In primary/secondary mode, a listener is associated with a primary server group and a secondary server group. If the number of unhealthy ECS instances in the primary server group reaches the failover threshold, requests are distributed to the secondary server group. Each server group contains one or more ECS instances. A listener associated with a primary server group and a secondary server group provides higher reliability than a listener associated with a primary server and a secondary server.

Before you enable the primary/secondary mode for a listener, make sure that the listener is associated with at least two vServer groups. For more information about how to create a vServer group, see Create a vServer group. You can enable the primary/secondary mode for a listener when you add backend servers to the listener.

Note Only TCP and UDP list eners support the primary/secondary mode.

Default server group

You can add ECS instances to the default server group of a listener to receive requests. If a listener is not associated with a vServer group or a primary/secondary server group, requests are distributed to the ECS instances in the default server group.

Before a CLB instance can process requests, you must add at least one backend server to the default server group to receive requests. For more information, see Add a default backend server.

vServer groups

You can create vServer groups for CLB to distribute different requests to different backend servers. To allow CLB to distribute requests based on domain names and URLs, you can specify vServer groups in domain name-based forwarding rules and URL-based forwarding rules. For more information, see Create a vServer group.

Primary/secondary server groups

A primary/secondary server group contains only two ECS instances. One ECS instance serves as the primary server and the other ECS instance serves as the secondary server. CLB does not perform health checks on the secondary server in a primary/secondary server group. If the primary server is down, network traffic is automatically distributed to the secondary server. When the primary server recovers, traffic is switched back to the primary server. For more information, see Create a primary/secondary server group.

? Note You can add primary/secondary server groups only to TCP and UDP listeners.

Related information

- Add a default backend server
- Create a vServer group
- Create a primary/secondary server group

2.Default server groups 2.1. Add a default backend server

This topic describes how to add a default backend server. Before you use the Classic Load Balancer (CLB) service, you must add at least one default backend server to receive client requests.

Prerequisites

Before you add an Elastic Compute Service (ECS) instance to the default server group, make sure that the following requirements are met:

- A CLB instance is created. For more information, see Create a CLB instance.
- ECS instances are created and applications are deployed on the ECS instances to receive requests.

Procedure

1.

2.

- 3. On the **Instances** page, select the region where the CLB instance that you want to manage is deployed.
- 4. Find the CLB instance that you want to manage and click its ID.
- 5. Click the **Default Server Group** tab.
- 6. Click Add.

| nstance Details | Listeners | VServer Groups | Default Server Group | Active/Standby S | Server Groups | | | |
|-----------------|---------------|----------------|----------------------|-------------------|---------------|--------|---------|---|
| dd ECS Inst | ance Name 🛛 🗸 | Enter a value | Q | | | | | C |
| ECS Instand | ce ID/Name | VPC | Public/In | ternal IP Address | Status 🙄 | Weight | Actions | |
| | | | | No dat | ta available. | | | |
| | | | | | | | | |

7. In the **My Servers** panel, for **Select Servers**, select one or more ECS instances that you want to add to the default server group.

| Available Servers | | | | | | | | |
|---|--------------------|-----------------------|-----------------------|----------------|-----------------------------|--|--|--|
| Search by server name, ID, or IP addres Display Available Instances Advance | ect Server | C V Sele | 2 | Configure Port | and Weight Buy ECS 🗗 | | | |
| ✓ ECS Instance ID/Name | Zone | Private IP Address | Public IP Address/VPC | Status | Associated SLB Instances | | | |
| | Hangzhou Zone B | 10.001.00 | vp vs bp | ✓ Running | 1 | | | |
| ■ | Hangzhou Zone B | 10248.230 | vi vi | ✓ Running | 1 | | | |
| | | | | | | | | |
| | | | | | | | | |
| | | | | | | | | |
| | | | | | | | | |
| | _ | | | | | | | |
| You have selected 2 servers. Next | Cancel | | | | | | | |

- 8. Click Next.
- 9. For **Configure Ports and Weights**, specify the weight of each ECS instance.

An ECS instance that has a higher weight receives more requests.

You can change the weight of a server and then move the pointer over 🔲 to change the weights

of other servers:

- If you click **Replicate to Below**, the weights of all servers below the current server are set to the weight of the current server.
- If you click **Replicate to Above**, the weights of all servers above the current server are set to the weight of the current server.
- If you click **Replicate to All**, the weights of all servers in the default server group are set to the weight of the current server.
- If you click **Reset**, when you clear the weight of the current server, the weights of all servers in the default server group are cleared.

♥ Notice

- Valid values of weights: 0 to 100. If you set the weight of a server to 0, the server does not receive requests.
- If two servers have the same weight, only one server receives requests.

| ailable Servers | | | | |
|---|-----------------|--------------------|--|-----------|
| \bigcirc | Select Server | | — 2 Configure Port ar | nd Weight |
| ECS Instance ID/Name | Region | Private IP Address | Weight 👔 Reset | Actions |
| N. S. | Hangzhou Zone B | 15 | Duplicate to Below Duplicate to Above Duplicate to All | Delete |
| Maged. | Hangzhou Zone B | 19 | 100 | Delete |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |

- 10. Click Add.
- 11. Click **OK**.

Related information

AddBackendServers

2.2. Change the weight of a backend server

This topic describes how to change the weight of a backend server to adjust the proportion of requests sent to the backend server.

Procedure

- 1.
- 2.
- 3. On the Server Load Balancers page, select the region of the target SLB instance.
- 4. Find the target SLB instance and click its instance ID.
- 5. Click the Default Server Group tab.

icon.

6. Move the pointer over the weight value of the target backend server and click the

| Server Loa | ad Balancer / Server Load Balancers | 8.159 | | | Start ● Stop | Edit Tags | ∠ Change Specifica | ation |
|------------|---|--|-------------------------------|--------------------------------------|-----------------|-----------|--------------------|-------|
| Instar | nce Details Listener VSe | rver Groups Default Server Group | Primary/Secondary | Server Groups | | | | |
| 🕜 Ar | n SLB instance has a default server gro | up to which you can directly add backend serve | ers. However, all listeners u | nder an SLB share the default server | group. | | | × |
| Add | Add IDC Server ECS Insta | nce Name 🗸 Enter a value | Q | | | | | C |
| | ECS Instance ID/Name | VPC | | Public/Internal IP Address | Status 🙄 | Weight | Actions | |
| | | enderse and the second | | 11.4.10 (mail) | ✓ Running | 100 🗹 | Remove | |
| | | | | | | | | |

7. Change the weight and then click **OK**.

A backend server (ECS instance or IDC server) with a higher weight receives more requests.

Notice The weight value ranges from 0 to 100. If the weight of a backend server is set to 0, no requests are sent to the backend server.

Related information

• Set BackendServers

2.3. Remove a backend server

This topic describes how to remove a backend server that is no longer needed.

Procedure

1.

- 2. On the Server Load Balancers page, select the region of the target SLB instance.
- 3. Find the target SLB instance and click its instance ID.
- 4. Click the Default Server Group tab.
- 5. Find the target backend server and click **Remove** in the **Actions** column.
- 6. In the dialog box that appears, click **OK**.

Related information

RemoveBackendServers

3.Create a vServer group

This topic describes how to create a vServer group for a Server Load Balancer (SLB) instance. A vServer group contains Elastic Compute Service (ECS) instances that function as backend servers. If you associate a vServer group with a listener, the listener distributes requests only to backend servers in the vServer group.

Prerequisites

Before you create a vServer group, make sure that the following conditions are met:

- An SLB instance is created. For more information, see Create a CLB instance.
- ECS instances are created and applications are deployed on the ECS instances to process requests.

Context

Take note of the following items before you create a vServer group for an SLB instance:

- ECS instances are added to a vServer group and the corresponding SLB instance must be deployed in the same region.
- An ECS instance can be added to multiple vServer groups.
- A vServer group can be associated with multiple list eners of an SLB instance.
- A vServer group consists of ECS instances and application ports.

Create a vServer group

- 1.
- 2. On the **Instances** page, select the region of the SLB instance for which you want to create a vServer group.
- 3. Find the SLB instance and click its instance ID.
- 4. Click the VServer Groups tab.
- 5. On the VServer Groups tab, click Create VServer Group.
- 6. On the Create VServer Group page, set the parameters.
 - i. In the VServer Group Name field, enter a name for the vServer group.
 - ii. Click Add. On the My Servers wizard page, select the ECS instances that you want to add.
 - iii. Click Next .

iv. Set the Port and Weight parameters for each ECS instance, and then click Add.

Set the Port and Weight parameters based on the following information:

• **Port** : The backend port opened on an ECS instance to receive requests.

You can set the same port number for multiple backend servers of an SLB instance.

• Weight : An ECS instance with a higher weight receives more requests.

 \bigcirc Notice If the weight of an ECS instance is set to 0, the ECS instance no longer receives new requests.

| Available Servers | | | | | |
|---|-----------------|--------------------|------------|--|-----------------|
| | | | | | |
| (| Select Server | | | 2 Configure P | ort and Weight |
| ECS Instance ID/Name | Region | Private IP Address | Port Reset | Weight 🕜 Reset | Actions |
| 10440 Gott 101 Hog 19 to 44d Date Shiugp | Hangzhou Zone B | 10.101/101/101 | | 80 | Add Port Delete |
| | | | | Duplicate to Below Duplicate to Above Duplicate to All | Delete |
| | | | ۵ | 20 | Delete |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| Previous Add | Cancel | | | | |

You can click () to specify the ports and weights of the added ECS instances in batches.

- **Replicate to Below**: The ports or weights of all servers below the current server are set to the port or weight of the current server.
- Replicate to Above: The ports or weights of all servers above the current server are set to the port or weight of the current server.
- Replicate to All: The ports or weights of all servers in the vServer group are set to the port or weight of the current server.
- Reset: If the port or weight of the current server is cleared, the ports or weights of all servers in the vServer group are also cleared.
- 7. Click Create.

Modify a VServer group

1.

- 2. On the Server Load Balancers page, select the region of the target SLB instance.
- 3. Find the target SLB instance and click its instance ID.
- 4. Click the VServer Groups tab.

5. Find the target VServer group and then click Edit in the Actions column.

| nstance Details | Listener | VServer Groups | Default Server Group | Primary/Secondary Server G | roups Monitoring |
|-----------------------------------|-------------------|--------------------------|----------------------------------|---|----------------------------|
| | | | | | |
| By default, SLB i | nstance maintain | is backend server groups | in the instance dimension, an | d all listeners of the instance use the | same backend server |
| group. VServer | group mode allow | ws you to set the custom | server groups in the listener of | limension. In this case, you can use v | arious backend servers for |
| different listene | rs to meet the re | quirements of domain fo | rwarding and URL redirections | | |
| | | 1 | | | |
| | | | | | |
| | | | | | |
| Create VServer Grou | IP I | | | | |
| Create VServer Grou | р | | | | |
| Create VServer Grou | p Group | | Listener | | Asting |
| Create VServer Grou Group Name | P Group I | ID | Listener | Forwarding Rule | Actions |
| Create VServer Grou Group Name | p Group I | ID | Listener | Forwarding Rule | Actions |

6. Modify the ports and weights of ECS instances or IDC servers, and then click Save.

Delete a VServer group

You can delete a VServer group that is no longer needed for traffic distribution.

- 1.
- 2. On the Server Load Balancers page, select the region of the target SLB instance.
- 3. Find the target SLB instance and click its instance ID.
- 4. Click the VServer Groups tab.
- 5. Find the target VServer group, and then click **Delete** in the Actions column.

| Instance Details | Listener | VServer Groups | Default Server Group | Primary/Secondary Server Gro | oups Monitoring | |
|------------------------------------|--------------------|--------------------------|-----------------------------------|---|-----------------------------|---|
| By default SLR in | stance maintair | s backend server group | in the instance dimension an | d all listeners of the instance use the s | ame backend server group | /Server group mode allows you to set the custom service |
| listener dimensio | n. In this case, y | /ou can use various back | end servers for different listene | ers to meet the requirements of domai | n forwarding and URL redire | ections. |
| | | | | | | |
| Create VServer Group | > | | | | | |
| Create VServer Group Group Name | | Group ID | | Listener | Forwarding Rule | e Actions |

6. In the dialog box that appears, click **OK**.

Related information

- AddVServerGroupBackendServers
- SetVServerGroupAttribute
- ModifyVServerGroupBackendServers
- DeleteVServerGroup
- DescribeVServerGroups

4.Active/standby server groups 4.1. Create a primary/secondary server group

This topic describes how to create a primary/secondary server group and then add Elastic Compute Service (ECS) instances to the primary/secondary server group. A primary/secondary server group contains a primary server and a secondary server that can fail over to prevent service interruption. By default, the primary server handles all requests that are distributed by the SLB instance. When the primary server fails, requests are redirected to the secondary server.

Prerequisites

Before you create a primary/secondary server group, make sure that the following requirements are met:

- A Server Load Balancer (SLB) instance is created. For more information, see Create a CLB instance.
- ECS instances are created and applications are deployed on the ECS instances to process requests.

② **Note** Only TCP and UDP listeners support primary/secondary server groups.

Procedure

- 1.
- 2.
- 3. On the **Instances** page, select the region of the SLB instance that you want to manage.
- 4. Find the SLB instance and click its instance ID.
- 5. Click the Primary/Secondary Server Groups tab.
- 6. On the Primary/Secondary Server Groups tab, click Create Primary/Secondary Server Group.
- 7. On the **Create Primary/Secondary Server Group** page, configure the primary/secondary server group.
 - i. In the **Primary/Secondary Server Group Name** field, enter a name for the primary/secondary server group.
 - ii. Click Add. In the My Servers panel, select the ECS instances that you want to add.You can add only two ECS instances to a primary/secondary server group.
 - iii. Click Next.

iv. On the **Configure Ports and Weights** wizard page, specify the backend ports that you want to open on the ECS instances to receive requests. If you want to open more than one port on an ECS instance, click **Add Port** in the **Actions** column.

You can open the same port on different backend servers that are connected to the same SLB instance.

| Available Servers | | | | | | |
|---|-----------------|------------------------|------------|-------------------------|--|--|
| (| Select Server | | 2 Cor | nfigure Port and Weight | | |
| ECS Instance ID/Name | Region | Private IP Address | Port Reset | Actions | | |
| 104-00-001111 Hog 191-4440-0055hiugp | Hangzhou Zone B | 192,768.2 1440/https:/ | | Add Port Delete | | |
| SEARCH Hot Phyliothigh Logo | Hangzhou Zone B | 192,768.2.1450Privated | | Add Port Delete | | |
| | | | | | | |
| | | | | | | |
| | | | | | | |
| | | | | | | |
| Previous Add Ca | ancel | | | | | |

- v. Click Add.
- 8. On the **Create Primary/Secondary Server Group** page, select an ECS instance in the **Type** column as the primary server.
- 9. Click Create.

Related information

• CreateMasterSlaveServerGroup

4.2. Delete a primary/secondary server group

This topic describes how to delete a primary/secondary server group of a Server Load Balancer (SLB) instance. If a primary/secondary server group is no longer needed to forward traffic, you can delete the primary/secondary server group.

Procedure

1.

2.

- 3. On the **Instances** page, select the region of the SLB instance that you want to manage.
- 4. Find the SLB instance and click its instance ID.

- 5. Click the Primary/Secondary Server Groups tab.
- 6. On the **Primary/Secondary Server Groups** tab, find the primary/secondary server group that you want to delete and click **Delete** in the **Actions** column.

| nstance Details | Listener | VServer Groups | Default Server Group | Primary/Secondary Server Groups | Monitoring | |
|--|---|--|---|---|--|--------------------------------|
| By default, SLB in groups to forwar primary/seconda | istance maintair d traffic in the l ry server group | ns backend server group: istener dimension. If you mode is only available fe | s in the instance dimension, an Ir business strongly depends or or TCP/UDP listeners. | d all listeners of the same instance use the san n the traditional primary/secondary server mo | me backend server group. Primary, de, you can configure primary/sec | /Secondary se ondary server |
| Create Primary/Secor | ndary Server Gro | pup | | | | |
| lame | | 10 | D | Listener | A | ctions |
| and the second sec | | rs | 5 | | View | Delete |

7. In the message that appears, click **OK**.

Related information

• DeleteMasterSlaveServerGroup

5.Add private IP addresses of ENIs to backend servers

An Elastic Network Interface (ENI) is a virtual network interface that can be attached to a VPC-type ECS instance. You can use ENIs to deploy high-availability clusters and perform low-cost failovers and finegrained network management. This topic describes how to add the primary and secondary private IP addresses of ENIs to backend servers when you add backend servers to a guaranteed-performance Server Load Balancer (SLB) instance.

Context

The ECS instances to be added as backend servers are attached with multiple ENIs before you add a backend server group to an SLB instance.

For more information about how to attach an EIP to an ECS instance, see Bind an ENI.

| Network Interfaces | | | | | | | ⑦ Ela | astic network interfaces | Create ENI |
|--------------------|------------------|------------------------|----------------------|-------------------|----------------------|-------------------------------|--------------------------------|--------------------------------------|---------------------------|
| Name 🔻 S | earch by ENI nam | e | Searc | h 📎 Tag | | | | | |
| ID/Name | VSwitch/VPC | Zone | Security Group ID | Bound Instance | Public IP Address | Primary Private IP Address | Type/MAC Address(All) 👻 | Status/Created At | Actions |
| eni- 9 | vsw- F vj | China East 1 Zone G | sg-bp1dr | i-bp194y | | 172.16.3.128 | Primary 00:16:3e:0c:f2:88 | Bound 12 October 2018, 13.46 | Modify Unbind Delete |
| eni- 7 | VSW | China East 1 Zone F | sg-bp1be | | | 192.168.0.230 | Secondary 00:16:3e:0a:dc:76 | Bound 25 September 2018, 13.47 | Modify Unbind Delete |

Onte You can add the primary and secondary private IP addresses of ENIs to backend servers of only guaranteed-performance SLB instances.

Procedure

1.

2.

- 3. In the left-side navigation pane, choose **Instances > Instances**.
- 4. On the **Instances** page, click the ID of the SLB instance for which you want to create a backend server group.
- 5. Click the VServer Groups, Default Server Group, or Primary/Secondary Server Groups tab.

(?) Note You can add the primary and secondary private IP addresses of ENIs to default server groups, VServer groups, and primary/secondary server groups. A VServer group is created in the example.

- 6. On the VServer Groups tab, click Create VServer Group.
- 7. On the Create VServer Group page, click Add.
- In the My Servers panel, turn on Advanced Mode and select Elastic Compute Service (ECS)/Elastic Network Interface (ENI) from the Select Backend Server Type drop-down list.

- 9. Select the ECS instances that you want to add to the server group and click **Next** to configure the ports and weights for the instances.
- 10. Click Add.

If the default server group is added for a listener, you can view the backend server group that is added with ENIs and secondary private IP addresses on the **Instances** page.

- 📑 indicates an ECS instance.
- 📷 indicates an ENI and its secondary private IP address.

| Instance Name/ID | IP Address 77 | Status | Monitor | Port/Health | Check/Backend Server >> | Instance Specification | Billing Method/Billing Method 가 | Renewal Status | Actions |
|---|--|----------------------------|---------|---------------------|---|--|---|-------------------|---|
| sille8ab3 b- b- kubernetes.do.not.delete 511e8ab3 | 11 120(Public IPv4 Address) | Active | I | TCP: 443 TCP: 80 | Normal Default Server Group 3 Normal Default Server Group 3 | [∨] Shared- _∨ Performance | Pay-As-You-Go(By Traffic) 11/01/2018, 11:22:07 Created | - | Configure Listener Add Backend Servers More V |
| acs-slb- 340afb89fd b- 557 The tag is not set. | 15 74(VPC) vpc- b ievq5 vsw- bj jsy5 | Active | al | TCP: 6443 | Normal Default Server Group 2 mode-0001-k8s-for-cs:6443 k8s-eni-1541042484:6443 | Guaranteed- Performance slb.s2.small | Pay-As-You-Go(-) 11/01/2018, 11:21:15 Created | - | Configure Listener Add Backend Servers More V |

6.Backend server FAQ

The following questions about backend ECS instances are frequently asked:

- Can I adjust the number of backend ECS instances while my SLB instance is running?
- Can I use different operating systems for different backend ECS instances?
- Can Lattach ECS instances from different regions to the same SLB instance?
- Why do my records show frequent access to my backend ECS instances from IP addresses that start with 100?
- Why are responses from SLB compressed even though my ECS instances are not configured for compression?
- Is chunked transfer encoding supported if my backend ECS instances use HTTP/1.0?
- Why do my backend ECS instances frequently receive requests whose UA string value is KeepAliveClient?

Can I adjust the number of backend ECS instances while my SLB instance is running?

Yes.

You can increase or decrease the number of backend ECS instances for an SLB instance at any time and switch between ECS instances. Before you perform these operations, make sure that health check is enabled and at least one backend ECS instance is running. This ensures service stability.

Can I use different operating systems for different backend ECS instances?

Yes.

You can use different operating systems for the backend ECS instances of an SLB instance. However, the applications deployed on the ECS instances must be the same and have the same data. To facilitate maintenance, we recommend that you use the same operating system for backend ECS instances.

Can I attach ECS instances from different regions to the same SLB instance?

Yes.

SLB supports cross-region deployment. You can add the IP addresses of ECS instances in different regions to the SLB whitelist to use the instances as backend servers. This feature is available for public preview. Submit a ticket or contact your business manager to use this feature.

Why do my records show frequent access to my backend ECS instances from IP addresses that start with 100?

SLB forwards external requests to backend ECS instances by using the internal IP address of the system server. SLB also accesses the ECS instances to perform health checks and monitor service availability.

The IP address range of the SLB system is . Therefore, many IP addresses that start with 100 are accessing the ECS instances.

To ensure service availability, you must configure appropriate access rules for these IP addresses.

Why are responses from SLB compressed even though my ECS instances are not configured for compression?

A possible cause is that the client web browser supports compression. You can disable the Gzip feature when you create listeners in the console, or use TCP listeners.

Is chunked transfer encoding supported if my backend ECS instances use HTTP/1.0?

Yes.

Why do my backend ECS instances frequently receive requests whose UA string value is KeepAliveClient?

Problem description:

Backend ECS instances frequently receive GET requests, but no users are sending access requests to the ECS instances. Source IP addresses of these requests are internal IP addresses of Alibaba Cloud, and the value of the User-Agent string is KeepAliveClient.

Cause:

TCP listeners are used, but the HTTP protocol is used to perform health checks. When health checks that use the HTTP protocol are performed in TCP listeners, GET requests are used by default.

Solution:

We recommend that you use the same protocol for both listeners and health checks.