Alibaba Cloud Server Load Balancer

Backend servers

Issue: 20190816

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 due to product version upgrades , adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall 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 the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults " and "as available" basis. 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 bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, 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 content of the Alibaba Cloud website, including but not limited to works, products, images, archives, information, materials, website architecture, website graphic layout, and webpage design, 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 the Alibaba Cloud website, product programs, or content 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 contact Alibaba Cloud directly if you discover any errors in this document.

Generic conventions

Table -1: Style conventions

Style	Description	Example
	This warning information 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.
	This warning information 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 restore business.
	This indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: Take the necessary precautions to save exported data containing sensitive information.
	This indicates supplemental instructions, best practices, tips, and other content that is good to know for the user.	 Note: You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
Bold	It is used for buttons, menus , page names, and other UI elements.	Click OK.
`Courier font`	It is used for commands.	Run the `cd / d C :/ windows` command to enter the Windows system folder.
Italics	It is used for parameters and variables.	`bae log list --instanceid` *Instance_ID*
[] or [a|b]	It indicates that it is a optional value, and only one item can be selected.	`ipconfig` *[-all|-t]*

Style	Description	Example
{} or {a|b}	It indicates that it is a required value, and only one item can be selected.	`swich {stand | slave}`

Contents

Legal disclaimer..... I
Generic conventions..... I
1 Backend server overview..... 1
2 Manage a VServer group..... 4
3 Manage a default server group..... 8
4 Manage an active/standby server group..... 13
5 Add private IP addresses of ENIs to backend servers..... 16
6 FAQ..... 19
6.1 What can I do if my ECS instance is declared unhealthy after I enable health checks for Server Load Balancer?..... 19
6.2 Backend server FAQs..... 20
7 Default server groups..... 23
7.1 Add default servers..... 23
7.2 Modify the weight of a backend server..... 24
7.3 Remove a backend server..... 25
8 VServer groups..... 26
8.1 Create a VServer group..... 26
8.2 Edit a VServer group..... 27
8.3 Delete a VServer group..... 28
9 Active/standby server groups..... 29
9.1 Create an active/standby server group..... 29
9.2 Delete an active/standby server group..... 30

1 Backend server overview

Before using the load balancing service, you must add one or more ECS instances as the backend servers to an SLB instance to process the distributed client requests.

SLB service virtualizes the added ECS instances in the same region into an applicatio n pool featured with high performance and high availability. You can also manage backend servers through a VServer group. Different listeners can be associated with different server groups so that different listeners of an SLB instance can forward requests to the backend servers with different ports.

 Note:
After a VServer group is configured for a listener, the listener will forward requests to the ECS instances in the associated VServer group instead of the ECS instances in the default server group.

You can increase or decrease the number of the backend ECS instances at any time and specify the ECS instances that receive requests. However, we recommend that you enable the health check function, and there must be at least one normal ECS to maintain service stability.

When adding ECS instances to an SLB instance, note the following:

- SLB does not support cross-region deployment. Make sure that the ECS instances and the SLB instance are in the same region.
- SLB does not limit the operating system used in the ECS instances as long as the applications deployed in the ECS instances are the same, and the data is consistent . However, we recommend that you use the same operating system for better management and maintenance.
- Up to 50 listeners can be added to an SLB instance. Each listener corresponds to an application deployed on backend ECS instances. The listening port of an SLB instance corresponds to the application port opened on the ECS instance.
- You can specify a weight for each ECS instance in the backend server pool. An ECS instance with a higher weight will receive a larger number of connection requests.

- If you have enabled the session persistence function, the requests distributed to the backend ECS instances may be imbalanced. If so, we recommend that you disable the session persistence function to check if the problem persists.

  When the traffic is not distributed evenly, troubleshoot as follows:

  1. Collect the access logs of the web service within a period of time.
  2. Check if the number of logs of multiple ECS instances are different according to SLB configurations. If session persistence is enabled, you need to strip the access logs for the same IP address. If weights are configured for SLB, you need to calculate whether the percentage of access logs matches the weight ratio.)

- When an ECS instance is undergoing live migration, the persistent connection s of the SLB may be interrupted and can be restored by reconnecting them. Be prepared for the reconnection.

Default server group

A default server group contains ECS instances that receive requests. If a listener is not associated with a VServer group or an active/standby server group, requests are forwarded to ECS instances in the default server group by default.

See #unique_4 to create a default server group.

Active/standby server group

An active/standby server group only contains two ECS instances. One acts as the active server and the other acts as the standby server. No health check is performed on the standby server. When the active server is declared as unhealthy, the system forwards traffic to the standby server. When the active server is declared as healthy and restores service, the traffic is forwarded to the active server again.

See #unique_5 to create an active/standby server group.

Note:
Only Layer-4 listeners (TCP and UDP protocols) support configuring active/standby server groups.

VServer group

When you need to distribute different requests to different backend servers, or you want to configure domain name or URL based forwarding rules, you can use VServer groups.

See #unique_6 to create a VServer group.

2 Manage a VServer group

A virtual server group (VServer group) is a group of ECS instances. If you associate a VServer group with a listener, the listener distributes requests to the associated VServer group instead of other backend servers.

For Layer-7 listeners, the following algorithm is used to determine whether requests are forwarded to default backend server groups, or VServer groups, and whether forwarding rules are applied:

- If the requests match a forwarding rule, the requests are distributed to the VServer group associated with the rule.
- If no forwarding rule is matched and a VServer group is configured on the listener, the requests are distributed to the VServer group associated with the listener.
- If no VServer group is configured on the listener, the requests are forwarded to ECS instances in the default server group.

Create a VServer group

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

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

Note the following when you create a VServer group:

- The ECS instances added to a VServer group and the SLB instance must belong to the same region.
- One ECS instance can be added to multiple VServer groups.
- One VServer group can be associated with multiple listeners of an SLB instance.
- A VServer group consists of ECS instances and application ports.

To add ECS instances, follow these steps:

1. Log on to the SLB console.
2. On the Server Load Balancer page, select the region of the target SLB instance.
3. Find the target SLB instance and click the instance ID.
4. Click the VServer Groups tab.

5. On the VServer Groups page, click Create VServer Group.
6. On the Create VServer Group page, complete these steps:
   a. In the VServer Group Name field, enter a name for the VServer group to be created.
   b. Click Add and on the Available Servers page, select the servers to add.
   c. Click Next: Set weight and Port.
   d. Enter the port and weight of each ECS instance, and click OK.
      - Port: The backend port opened on the ECS instance to receive requests.

        The backend ports in an SLB instance can be the same.
      - Weight: An ECS instance with a higher weight receivers more requests.

         Notice:

If the weight is set to 0, no requests are sent to the ECS instance.

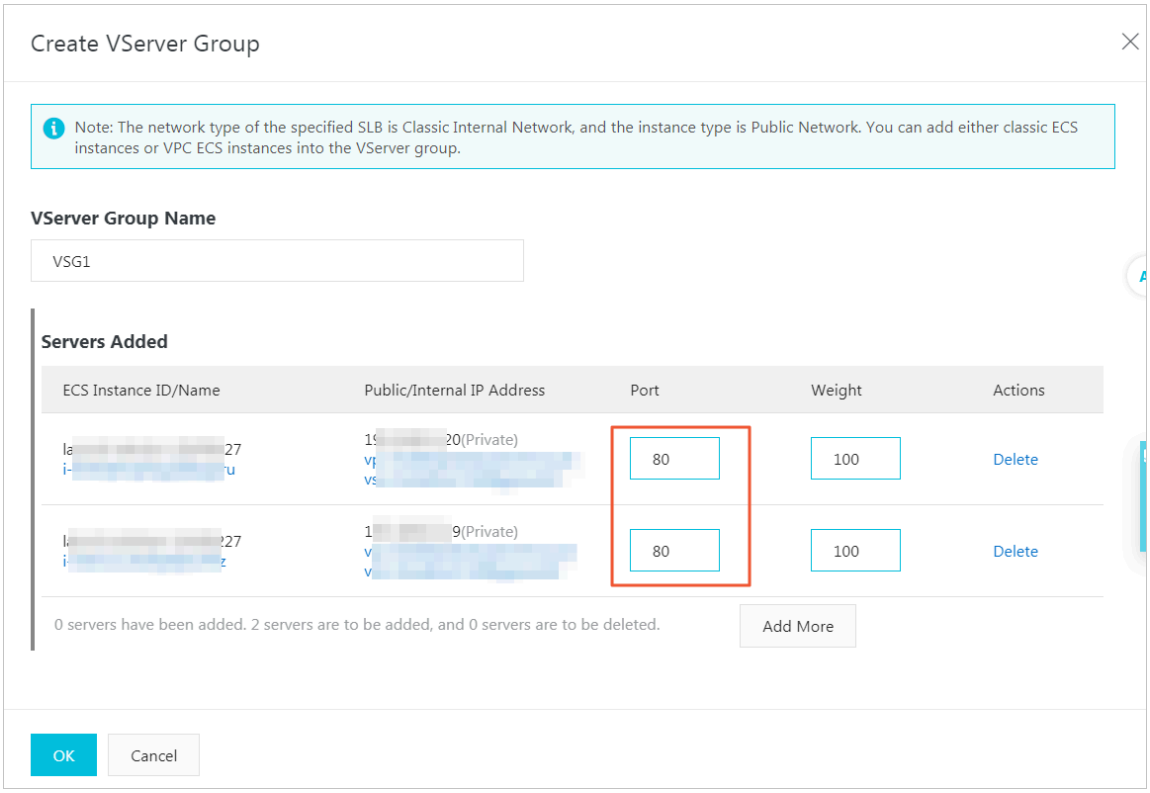

You can modify the ports and weights of added servers in batches.

- Click : Duplicate to below. If you modify the port or weight of the current server, the ports or weights of all servers blow are also changed.
- Click : Duplicate to above. If you modify the port or weight of the current server, the ports or weights of all servers above are also changed.
- Click : Duplicate to all. If you modify the port or weight of the current server, the ports or weights of all servers in the VServer group are also changed.
- Click : Clear all. If you clear the port or weight of the current server, the ports or weights of all servers in the VServer group are also cleared.

Edit a VServer group

To modify the ECS instance configuration in a VServer group, follow these steps:

1. Log on to the SLB console.
2. On the Server Load Balancer page, select the region of the target instance.

3. Find the target SLB instance and click the instance ID.
4. Click the VServer Groups tab.
5. Find the target VServer group, and then click Edit in the Actions column.

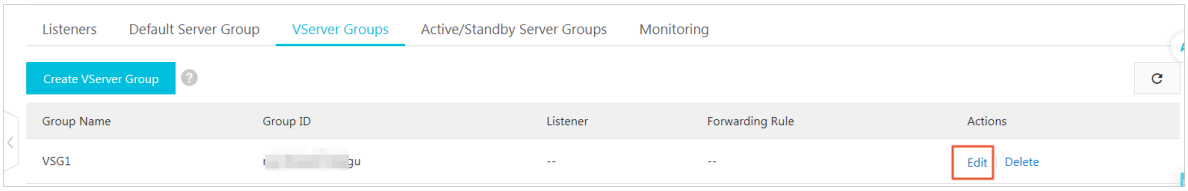

6. Modify the ports and weights of ECS instances or click Delete to remove ECS instances from the VServer group, and then click OK.

Delete a VServer group

To delete a VServer group, follow these steps:

1. Log on to the SLB console.
2. On the Server Load Balancer page, select the region of the target instance.
3. Find the target SLB instance and click the instance ID.
4. Click the VServer Groups tab.
5. Find the target VServer group, and then click Delete in the Actions column.

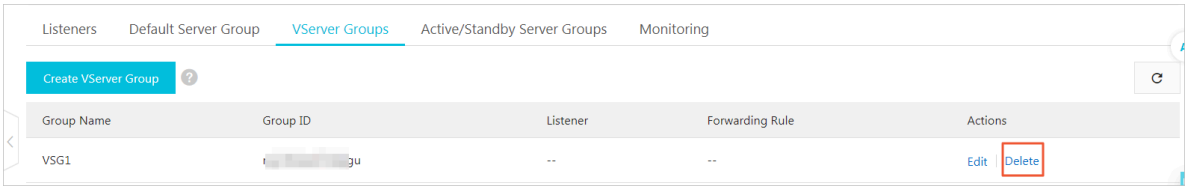

6. In the displayed dialog box, click OK.

3 Manage a default server group

Before you use the Server Load Balancer (SLB) service, you must add at least one default server to the default server group to receive client requests forwarded by SLB.

Add default servers

Before you add ECS instances to the default server group, make sure the following conditions are met:

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

To add backend servers, follow these steps:

1. Log on to the SLB console.
2. On the Server Load Balancer page, select the region of the target SLB instance.
3. Find the target SLB instance and click the instance ID.
4. Click the Default Server Group tab.
5. Click Add.

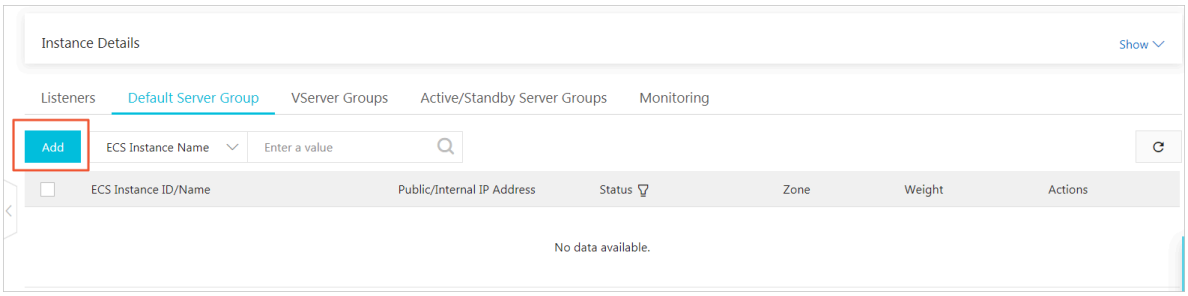

6. On the Available Servers page, select the ECS instances to add to the default server group.

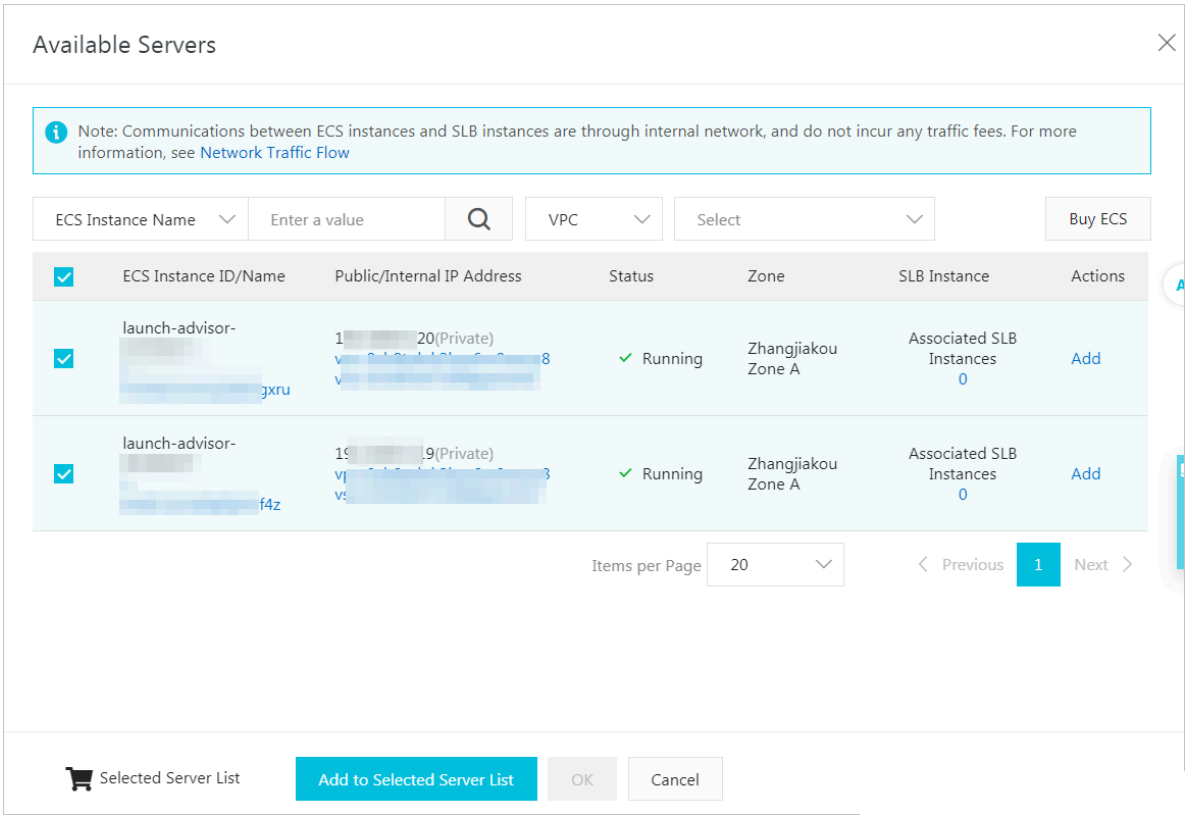

7. Click Next: Set Weight and Port.

8. On the Available Servers page, set the weights and ports of added ECS instances, and click OK.

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

   You can modify server weights in batches:

   - Click  : Duplicate to below. If you modify the weight of the current server, the weights of all servers blow are also changed.
   - Click  : Duplicate to above. If you modify the weight of the current server, the weights of all servers above are also changed.
   - Click  : Duplicate to all. If you modify the weight of the current server, the weights of all servers in the default server group are also changed.
   - Click  : Clear all. If you clear the weight of the current server, the weights of all servers in the default server group are also cleared.

    Notice:

If the weight is set to 0, the server no longer receives new requests.

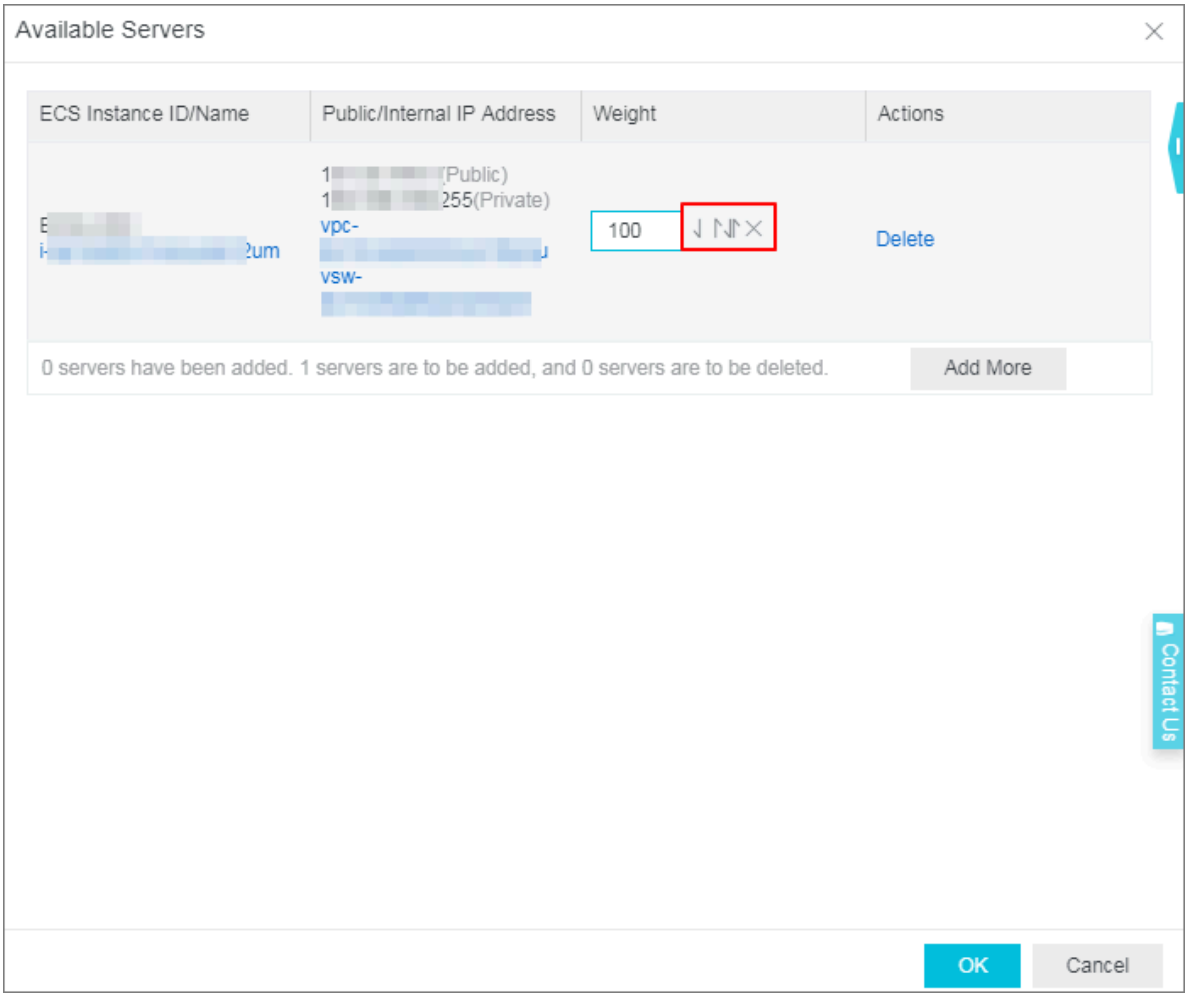

9. Click OK.

Edit the weight of a backend server

To edit the weight of a backend server, follow these steps:

1. Log on to the SLB console.
2. On the Server Load Balancer page, select the region of the target SLB instance.
3. Find the target SLB instance and click the instance ID.
4. Click the Default Server Group tab.
5. Rest the pointer over the weight value of the target backend server, and then click the displayed pencil icon.

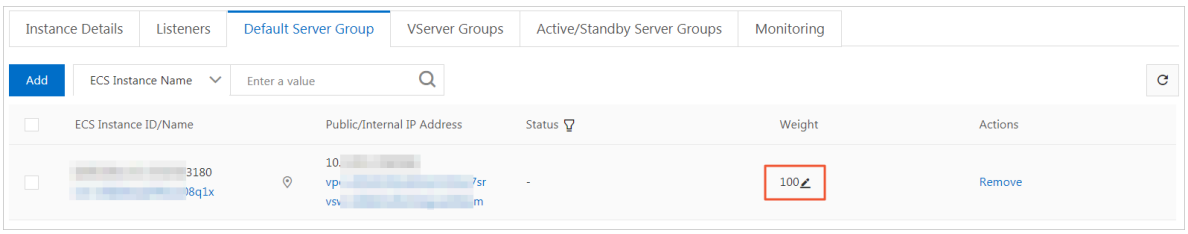

6. Modify the weight and then click OK.

   An ECS instance with a higher weight receives more requests.

   > **Notice:**
   > If the weight is set to 0, no requests are sent to the ECS instance.

Remove a backend server

To remove a backend server, follow these steps:

1. Log on to the SLB console.
2. On the Server Load Balancer page, select the region of the target SLB instance.
3. Find the target SLB instance and click the instance ID.
4. Click the Default Server Group tab.
5. Find the target backend server and click Remove in the Actions column.

4 Manage an active/standby server group

If you need active/standby failover configurations, where one backend server is used as the active server and the other as the standby server, you can create an active/standby server group. When the active server works normally, requests are distributed to the active server. If the active server is down, requests are distributed to the standby server to avoid service interruptions.

An active/standby server group only contains two ECS instances. One acts as the active server and the other acts as the standby server. No health check is performed on the standby server. When the active server is declared as unhealthy, the system forwards traffic to the standby server. When the active server is declared as healthy and restores service, the traffic is forwarded to the active server again.

 Notice:
Only Layer-4 listeners (TCP and UDP protocols) support active/standby server groups.

Create an active/standby server group

Before you create an active/standby server group, make sure the following conditions are met:

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

To create an active/standby server group, follow these steps:

1. Log on to the SLB console.
2. On the Server Load Balancer page, select the region of the target SLB instance.
3. Find the target SLB instance and click the instance ID.
4. Click the Active/Standby Server Groups tab.
5. On the Active/Standby Server Groups tab, click Create Active/Standby Server Group.

6. On the Create Active/Standby Server Group page, complete these steps:
   a. In the Name filed, enter a name for the active/standby server group to be created.
   b. Click Add and on the Available Servers page, select the servers to add.

      You can add up to two ECS instances to an active/standby server group.
   c. Click Next: Set Weight and Port.
   d. In the Servers Added section, set the port, select an active server, and click OK.
      - Port: The backend port opened on the ECS instance to receive requests.

        The backend ports in an SLB instance can be the same.
      - Server: Select a server to act as the active server.

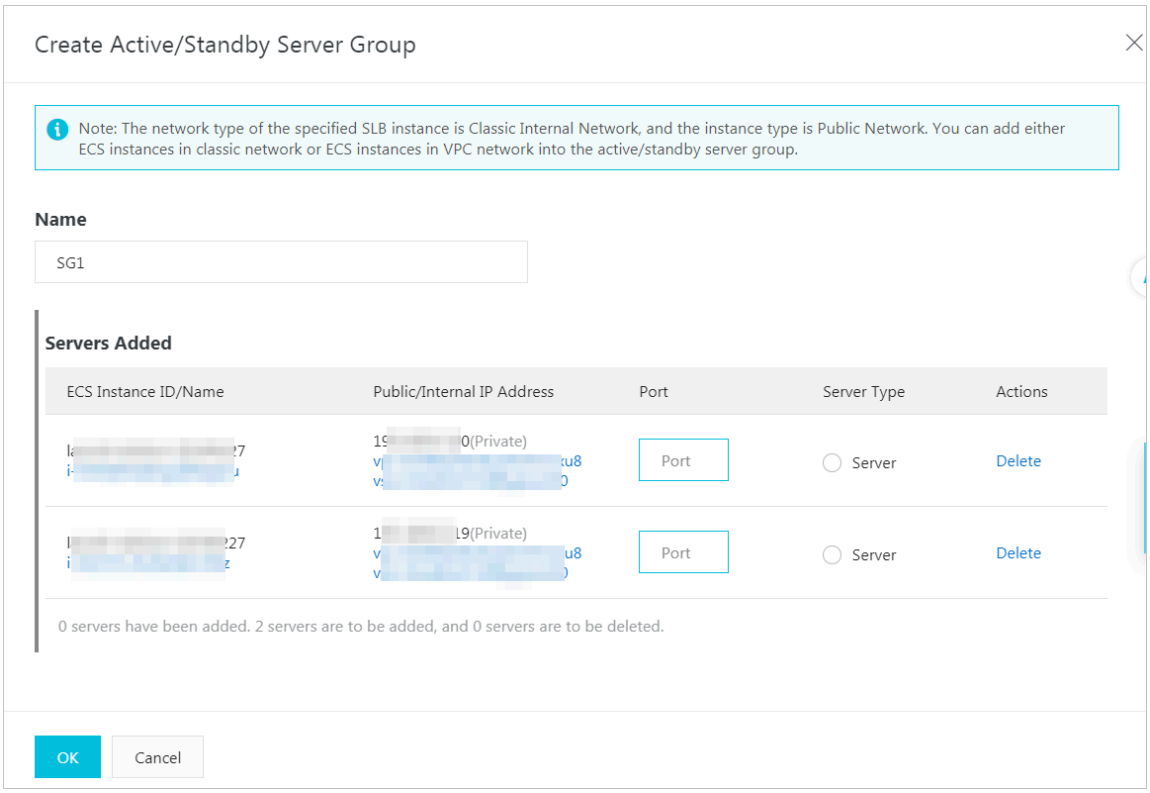

Delete an active/standby server group

To delete an active/standby server group, follow these steps:

1. Log on to the SLB console.
2. On the Server Load Balancer page, select the region of the target SLB instance.
3. Click the ID of the target SLB instance.
4. Click the Active/Standby Server Groups tab.

5. Find the target active/standby server group and click Delete in the Actions column.

6. In the displayed dialog box, click OK.

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 an ECS instance in a VPC. When you add backend servers to a guaranteed-performance Server Load Balancer (SLB) instance, you can choose to add the primary and secondary private IP addresses of ENIs if the ENIs are associated with ECS instances.

Prerequisites

The ECS instances are associated with ENIs.

For more information about how to associate an ENI with an ECS instance, see #unique_12.

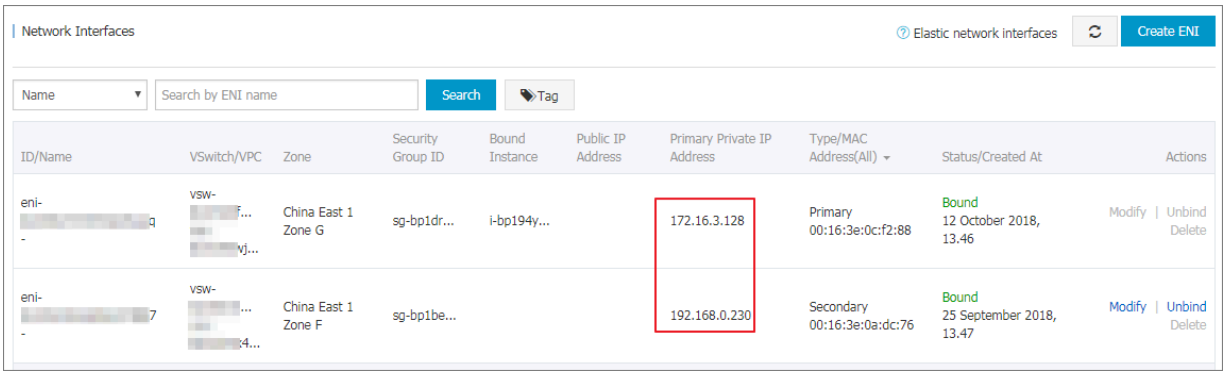

Note:
Only guaranteed-performance SLB instances support adding the primary and secondary private IP addresses of ENIs to backend servers.

Procedure

1. Log on to the SLB console.
2. In the left-side navigation pane, click Server Load Balancer. On the Server Load Balancer page, click the ID of the target SLB instance.
3. Select the backend server group type by clicking the corresponding tab. Default server groups, VServer groups, and active/standby server groups all support adding the primary and secondary private IP addresses of ENIs. In this topic, click the Default Server Group tab and then click Add.

4. On the Available Servers page, turn on Advanced Mode and click to select ENIs and its secondary private IP addresses.

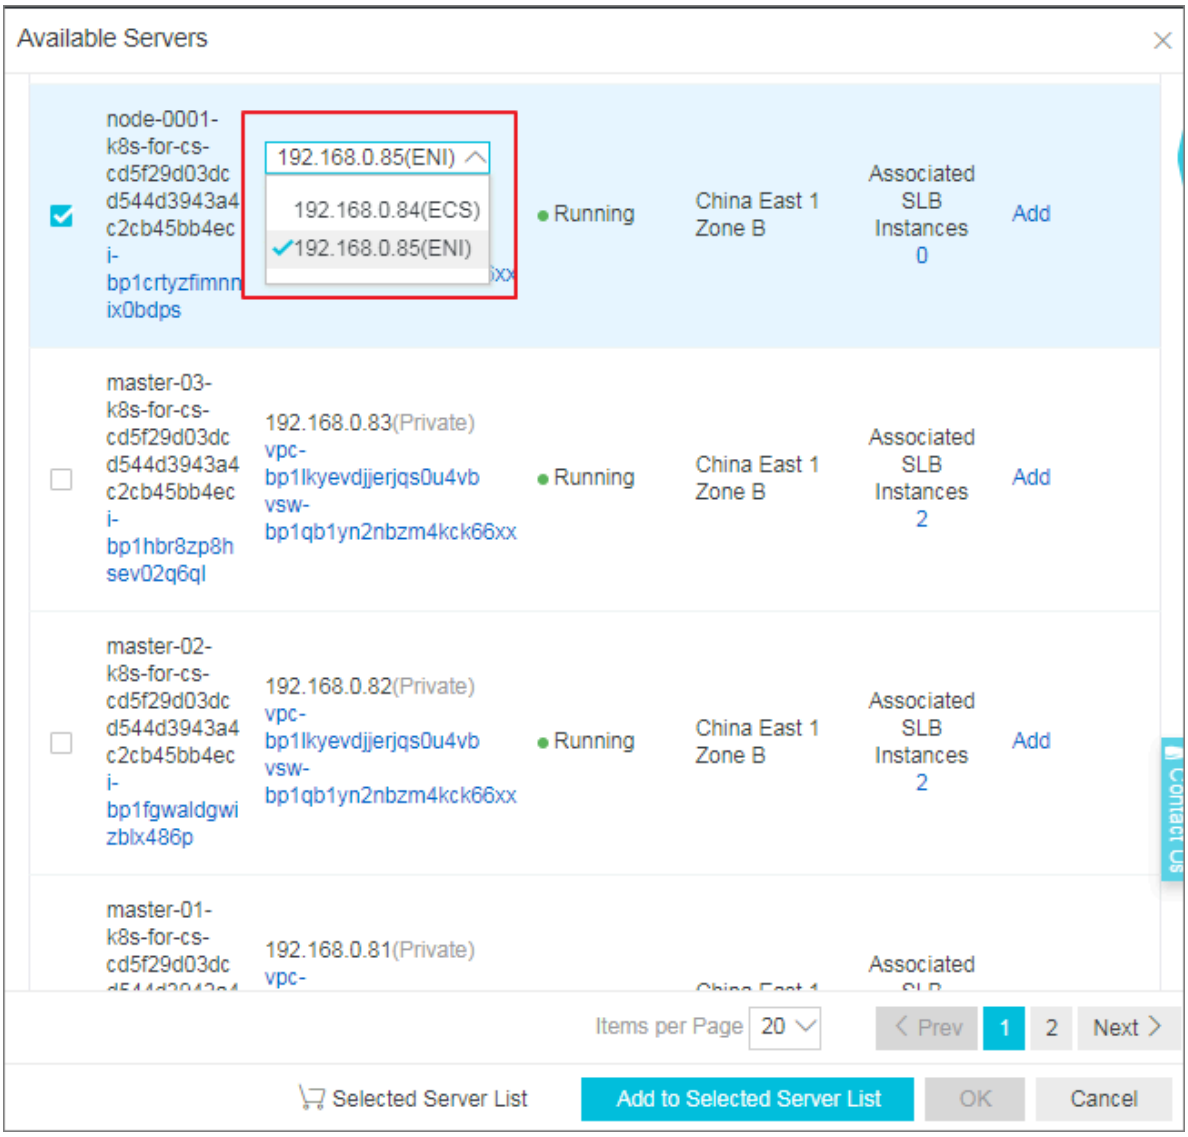

5. Click Next: Set Weight and Port to set the weights and port numbers of the added backend servers.

6. Click OK and you can see the added ENIs and its private IP addresses on the Default Server Group tab.

   If the default server group is added for a listener, you can see the backend servers added with ENIs and secondary private IP addresses on the Server Load Balancer page as follows:

   where,

   -  : Represents an ECS instance.
   -  : Represents an ENI and its secondary private IP address.

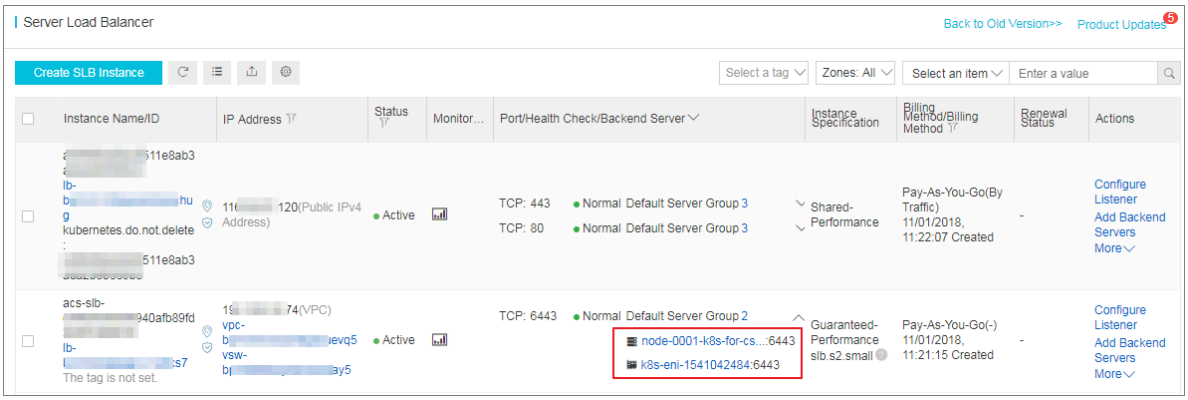

6 FAQ

6.1 What can I do if my ECS instance is declared unhealthy after I enable health checks for Server Load Balancer?

After you enable health checks of Server Load Balancer, when one backend ECS instance is declared as unhealthy, requests are forwarded to other normal ECS instances. When the faulty ECS instance becomes normal, Server Load Balancer forwards requests to the ECS instance again.

For Layer-7 SLB service, when an ECS instance is declared as unhealthy, you can troubleshoot the ECS instance from the following aspects:

- Make sure that you can directly access your service through the ECS instance.
- Make sure that the backend port you configured in the listener is opened on the backend server.
- Check whether the backend ECS instance has installed a firewall or other security protection software. This type of software may easily block the local IP address of the Server Load Balancer service, and thus disable the communication between the Server Load Balancer service and the backend server.
- Check whether the Server Load Balancer health check parameters are correctly set . We recommend that you use default health check parameters.
- We recommend that you use a static page for health checks. If the page you use for health checks isn't the default homepage of the backend ECS instance, you must enter the URL of the health check page in health check configurations. We recommend that you use a simple html page for health checks and the page is only used for checking the response. We do not recommend that you use dynamic scripting languages such as php.
- Check whether the backend ECS instance has high loads, which slow the ECS instance's response in offering services.

Besides, because the Layer-7 SLB service communicates with the backend ECS instance through intranet, the ECS instance must listen to intranet or all-network ports. You can check the ECS instance using the following methods:

1. Check whether the listening function is normal.

   If the frontend port is 80 and the backend port is 80, the intranet IP address of the ECS instance is 10.11.192.1. Run the following command on the server. If you can see the monitoring information of 10.1.1.192.1: 80, or the monitoring information of 0.5.0.0: 80, the listening function of the ports is normal.

   - Run the following command on the Windows server: `netstat -ano | findstr :80`
   - Run the following command on the Linux server: `netstat -anp | grep :80`

2. Check whether the intranet firewall of the server allows port 80. You can disable the firewall temporarily to do the test. Enter the following command to disable the firewall.

   - Windows: `firewall.cpl`
   - Linux: `/etc/init.d/iptables stop`

3. Check whether the backend port is normal.

   - For Layer-4 SLB service, the backend port is normal if you receive response after performing the telnet test. In this topic, use `telnet 10.11.192.1 80` to do the test.
   - For Layer-7 SLB service, the HTTP status code must be a status code that indicates a normal condition, such as 200. The test methods are as follows:
     - Windows: Directly access the intranet IP of the ECS instance. In this topic, access http://10.11.192.1.
     - Linux: Use the `curl -I` command to check if the status is HTTP/1.1 200 OK. In this topic, use `curl -I 10.11.192.1`.

6.2 Backend server FAQs

- #unique_16/unique_16_Connect_42_section_bt3_knx_wdb
- #unique_16/unique_16_Connect_42_section_ct3_knx_wdb
- #unique_16/unique_16_Connect_42_section_dt3_knx_wdb
- #unique_16/unique_16_Connect_42_section_gt3_knx_wdb
- #unique_16/unique_16_Connect_42_section_ot3_knx_wdb
- #unique_16/unique_16_Connect_42_section_pt3_knx_wdb

- #unique_16/unique_16_Connect_42_section_c43_sqx_wdb

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 in an SLB instance at any time and switch between different ECS instances. Before you perform these operations, make sure that health check is enabled and that there is at least one normally running backend ECS instance to avoid service interruption.

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

Yes.

There is no limitation on the operating system used on backend ECS instances as long as applications deployed on the ECS instances are the same and the data is consistent. To facilitate daily management and maintenance, we recommend that you use the same operating system for backend ECS instances.

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

No.

Server Load Balancer does not support cross-region deployment. The ECS instances to be added must belong to the same region as the SLB instance.

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

In addition to forwarding external requests to backend ECS instances by using the intranet IP address of the system server, the SLB system also accesses the ECS instances to perform health checks and monitor service availability.

The IP address range of the SLB system is 100.64.0.0/10（100.64.0.0/10 is reserved by Alibaba Cloud, and will not be used by any user, there is no security risk）, so there are many IP addresses beginning with 100 accessing ECS instances.

To guarantee the service availability, you have to configure appropriate access rules for these IP address ranges.

Why are responses returned by SLB compressed even though my ECS instance is not configured for compression?

The possible reason is that the client web browser supports compression. You can disable Gzip function when creating listeners in the console or use TCP listeners instead.

Is chunked transfer encoding supported if my backend ECS instances use HTTP1.0?

Yes.

Why do my backend ECS instances frequently receive requests where the value of the UA string is KeepAliveClient?

Issue

Backend ECS instances frequently receive GET requests, but there are no visitor IP addresses. Instead, source IP addresses of these requests are intranet IP addresses of Alibaba Cloud, and the value of the User-Agent string is KeepAliveClient.

Cause

TCP listeners are being used, which use the HTTP protocol for health checks. Specifically, when health checks that use 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.

7 Default server groups

7.1 Add default servers

Before you use the Server Load Balancer (SLB) service, you must add at least one default server (ECS instance) to receive client requests forwarded by SLB.

Prerequisites

Before you add ECS instances to the default server group, make sure the following conditions are met:

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

Procedure

1. Log on to the Server Load Balancer console.
2. On the Server Load Balancer page, select the region of the target SLB instance.
3. Find the target SLB instance and click the instance ID.
4. Click the Default Server Group tab.
5. Click Add.
6. On the Available Servers page, select the ECS instances you want to add to the default server group.
7. Click Next: Set Weight and Port.

8. In the Available Servers dialog box, specify the weights of the ECS instances and click OK.

   An ECS instance with a higher weight receives more requests.

   You can modify server weights in batches:

   - Click : Duplicate to below. If you modify the weight of the current server, the weights of all servers blow are also changed.
   - Click : Duplicate to above. If you modify the weight of the current server, the weights of all servers above are also changed.
   - Click : Duplicate to all. If you modify the weight of the current server, the weights of all servers in the default server group are also changed.
   - Click : Clear all. If you clear the weight of the current server, the weights of all servers in the default server group are also cleared.

   > Notice:
   > If the weight is set to 0, the server no longer receives new requests.

9. Click OK.

More information
#unique_20

7.2 Modify the weight of a backend server

After you add a backend server to the default server group, you can modify the weight of the backend server.

Procedure

1. Log on to the Server Load Balancer console.
2. On the Server Load Balancer page, select the region of the target SLB instance.
3. Find the target SLB instance and click the instance ID.
4. Click the Default Server Group tab.

5. Rest the pointer over the weight value of the target backend server, and then click the displayed pencil icon.

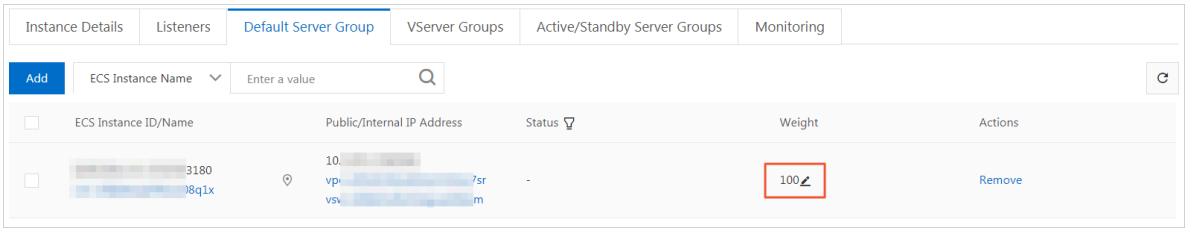

6. Modify the weight and then click OK.

   A backend server with a higher weight receives more requests.

> **Notice:**
> If the weight is set to 0, no request is sent to the backend server.

More information
#unique_22

7.3 Remove a backend server

If you no longer need a backend server in the default server group to forward traffic, you can remove the backend server.

Procedure

1. Log on to the Server Load Balancer console.
2. On the Server Load Balancer page, select the region of the target SLB instance.
3. Find the target SLB instance and click the instance ID.
4. Click the Default Server Group tab.
5. Find the target backend server and click Remove in the Actions column.

More information
#unique_24

8 VServer groups

8.1 Create a VServer group

A virtual server group (VServer group) is a group of ECS instances. If you associate a VServer group with a listener, the listener distributes requests to the associated VServer group instead of other backend servers.

Prerequisites

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

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

Context

Note the following before you create a VServer group:

- The ECS instances added to a VServer group and the SLB instance must belong to the same region.
- One ECS instance can be added to multiple VServer groups.
- One VServer group can be associated with multiple listeners of an SLB instance.
- A VServer group consists of ECS instances and application ports.

Procedure

1. Log on to the Server Load Balancer console.
2. On the Server Load Balancer page, select the region of the target SLB instance.
3. Find the target SLB instance and click the 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, configure the VServer group.
   a) Enter a name for the VServer group to be created in the VServer Group Name field.
   b) Click Add and on the Available Servers page, select the servers to add.
   c) Click Next: Set weight and Port.
   d) Enter the port and weight of each ECS instance, and click OK.

      Set the port and weight according to the following information:

      - Port: The backend port opened on the ECS instance to receive requests.

        The backend ports in an SLB instance can be the same.
      - Weight: An ECS instance with a higher weight receives more requests.

        > Notice:
        > If the weight is set to 0, no request is sent to the ECS instance.

      You can modify the ports and weights of added servers in batches.

      - Click : Duplicate to below. If you modify the port or weight of the current server, the ports or weights of all servers blow are also changed.
      - Click : Duplicate to above. If you modify the port or weight of the current server, the ports or weights of all servers above are also changed.
      - Click : Duplicate to all. If you modify the port or weight of the current server, the ports or weights of all servers in the VServer group are also changed.
      - Click : Clear all. If you clear the port or weight of the current server, the ports or weights of all servers in the VServer group are also cleared.

More information
#unique_27

8.2 Edit a VServer group

After you create a VServer group, you can modify the configurations of the ECS instances in the VServer group.

Procedure

1. Log on to the Server Load Balancer console.
2. On the Server Load Balancer page, select the region of the target SLB instance.

3. Find the target SLB instance and click the instance ID.
4. Click the VServer Groups tab.
5. Find the target VServer group, and then click Edit in the Actions column.

6. Modify the ports and weights of ECS instances or click Delete to remove ECS instances from the VServer group, and then click OK.

#unique_29
#unique_30

8.3 Delete a VServer group

If you no longer need a VServer group to forward traffic, you can delete the VServer group.

Procedure

1. Log on to the Server Load Balancer console.
2. On the Server Load Balancer page, select the region of the target SLB instance.
3. Find the target SLB instance and click the instance ID.
4. Click the VServer Groups tab.
5. Find the target VServer group, and then click Delete in the Actions column.

6. In the displayed dialog box, click OK.

#unique_32
#unique_33

9 Active/standby server groups

9.1 Create an active/standby server group

If you need active/standby failover configurations, where one backend server is used as the active server and the other as the standby server, you can create an active/standby server group. When the active server works normally, requests are distributed to the active server. If the active server fails, requests are distributed to the standby server.

Prerequisites

Before you create an active/standby server group, make sure the following conditions are met:

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

Procedure

1. Log on to the Server Load Balancer console.
2. On the Server Load Balancer page, select the region of the target SLB instance.
3. Find the target SLB instance and click the instance ID.
4. Click the Active/Standby Server Groups tab.
5. On the Active/Standby Server Groups tab, click Create Active/Standby Server Group.

6. On the Create Active/Standby Server Group page, configure the active/standby server group.
   a) In the Name filed, enter a name for the active/standby server group to be created.
   b) Click Add and on the Available Servers page, select the servers you want to add to the active/standby server group.

      You can add only two ECS instances to an active/standby server group.
   c) Click Next: Set Weight and Port.
   d) In the Servers Added section, set the port, select an active server, and click OK.
      - Port: The backend port opened on the ECS instance to receive requests.

        The backend ports in an SLB instance can be the same.
      - Server: Select a server to act as the active server.

More information
#unique_36

9.2 Delete an active/standby server group

If you no longer need an active/standby server group to forward traffic, you can delete the active/standby server group.

Procedure

1. Log on to the Server Load Balancer console.
2. On the Server Load Balancer page, select the region of the target SLB instance.
3. Find the target SLB instance and click the instance ID.
4. Click the Active/Standby Server Groups tab.
5. Find the target active/standby server group and click Delete in the Actions column.
6. In the displayed dialog box, click OK.

More information
#unique_38