

Alibaba Cloud Server Load Balancer

User Guide (New Console)

Issue: 20181016

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.	 Note: 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 <code>cd /d C:/windows</code> command to enter the Windows system folder.
<i>Italics</i>	It is used for parameters and variables.	<code>bae log list --instanceid Instance_ID</code>
[] or [a b]	It indicates that it is a optional value, and only one item can be selected.	<code>ipconfig [-all -t]</code>
{ } or {a b}	It indicates that it is a required value, and only one item can be selected.	<code>swich {stand slave}</code>

Contents

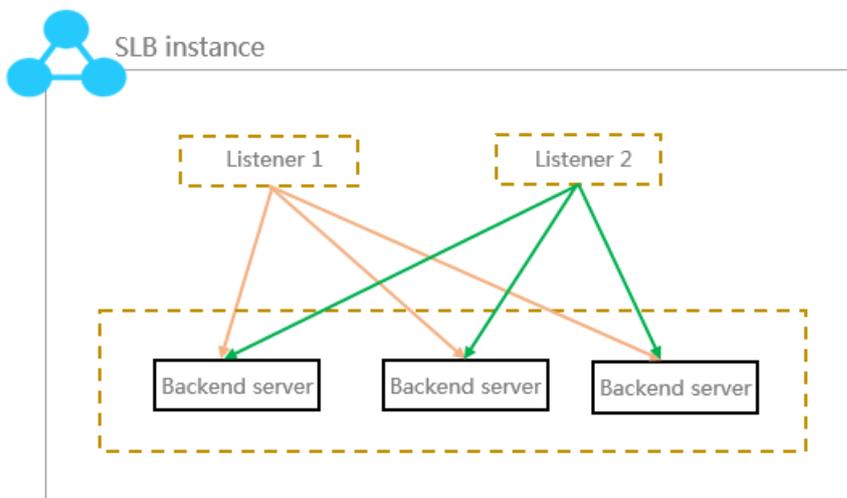
Legal disclaimer	I
Generic conventions	I
1 Server Load Balancer instance	1
1.1 SLB instance overview.....	1
1.2 Guaranteed-performance SLB instances.....	3
1.3 Network traffic flow.....	10
1.4 Create an SLB instance.....	13
1.5 Create an IPv6 instance.....	14
1.6 Start or stop an SLB instance.....	17
1.7 Bind an EIP.....	18
1.8 Release an instance.....	18
1.9 Manage tags.....	20
1.10 Expiring Instances.....	23
1.11 Change the configuration.....	24
2 Listeners	25
2.1 Add a TCP listener.....	25
2.2 Add a UDP listener.....	33
2.3 Add an HTTP listener.....	41
2.4 Add an HTTPS listener.....	50
2.5 Support TLS security policy.....	61
2.6 Manage a domain name extension.....	63
3 Health check	68
3.1 Health check overview.....	68
3.2 Configure health check.....	75
3.3 Close health check.....	79
4 Backend servers	81
4.1 Backend server overview.....	81
4.2 Manage a default server group.....	83
4.3 Manage a VServer group.....	86
4.4 Manage an active/standby server group.....	89
5 Certificate management	92
5.1 Certificate requirements.....	92
5.2 Upload a certificate.....	95
5.3 Generate a CA certificate.....	97
5.4 Convert certificate formats.....	101
5.5 Replace a certificate.....	101
6 Log management	103
6.1 View operation logs.....	103
6.2 Manage health check logs.....	104

6.3 Authorize a RAM user to configure access logs.....	109
6.4 Configure access logs.....	114
7 Access control.....	120
7.1 Configure an access control list.....	120
7.2 Configure access control.....	123
7.3 Migrate to the new access control.....	124
7.4 Configure a whitelist.....	125
8 Monitoring.....	127
8.1 View monitoring data.....	127
8.2 Configure alarm rules.....	129
9 Multiple zone deployment.....	131
10 Achieve cross-region load balancing through Global Traffic Manager.....	135
11 Anti-DDoS Basic.....	141

1 Server Load Balancer instance

1.1 SLB instance overview

An SLB instance is a running entity of the Server Load Balancer service. To use the load balancing service, you must create an SLB instance first, and then add listeners and backend servers to it.



Alibaba Cloud provides Internet SLB service and intranet SLB service. A public or a private IP address is allocated to the SLB instance according to the instance type you select.

Internet SLB instances

An Internet SLB instance distributes client requests over the Internet to backend ECS servers according to configured forwarding rules.

After you create an Internet Server Load Balancer instance, the system will allocate a public IP to the instance. You can resolve a domain name to the public IP to provide public services.

Alibaba Cloud Server Load Balancer

<p>Internet Server Load Balancer Instance</p> <p>Provides a public IP and can be accessed from the Internet.</p>	<p>Intranet Server Load Balancer Instance</p> <p>Provides a private IP and can be accessed from the intranet.</p> <div data-bbox="657 524 1106 698"> <p>Classic network</p> <p>The SLB instance can be accessed from the classic network, and all the ECS instances in the Alibaba Cloud.</p> </div> <div data-bbox="657 745 1106 920"> <p>VPC network</p> <p>The SLB instance can be accessed only from the ECS instances in the same VPC.</p> </div>	<p>Backend Servers</p> <p>The ECS instances of both the classic network and VPC network are supported.</p> <div data-bbox="1144 524 1596 698"> <p>Classic ECS</p> <p>This kind of ECS instances is located in the classic network. Compared with ECS instances in the VPC network, they are not isolated.</p> </div> <div data-bbox="1144 745 1596 920"> <p>VPC ECS</p> <p>This kind of ECS instances is located in a customized VPC. The VPC ECS instances are isolated from the classic ECS instances and other VPC ECS instances.</p> </div>
---	---	--

Intranet SLB instances

Intranet SLB instances can only be used inside Alibaba Cloud and can only forward requests from clients that can access the intranet of SLB.

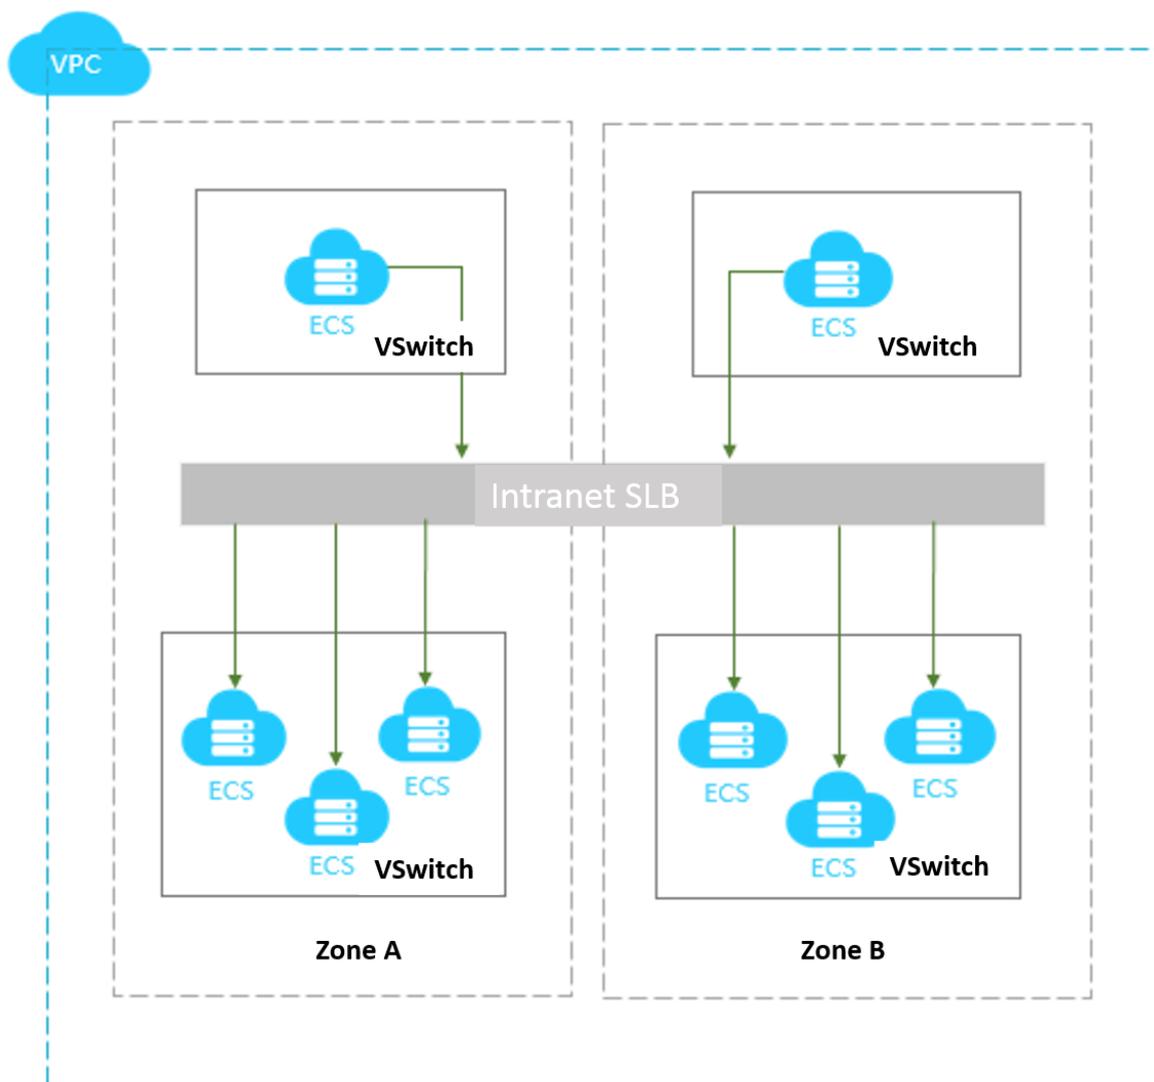
For an intranet SLB instance, you can further select the network type:

- Classic network

If you choose classic network for the intranet SLB instance, the IP of the SLB instance is allocated and maintained by Alibaba Cloud. The classic SLB instance can only be accessed by the classic ECS instances.

- VPC network

If you choose VPC network for the intranet SLB instance, the IP of the SLB instance is allocated from the CIDR of the VSwitch that the instance belongs to. SLB instances of the VPC network can only be accessed by ECS instances in the same VPC.



1.2 Guaranteed-performance SLB instances

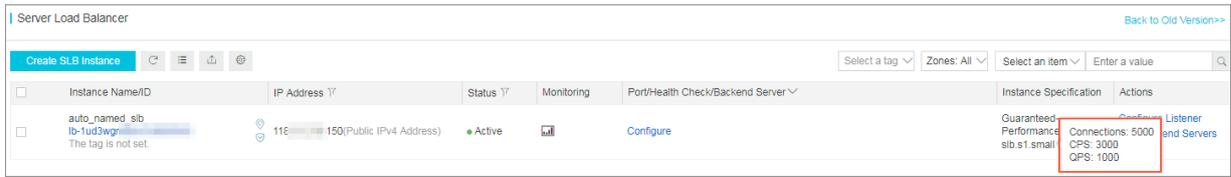
Alibaba Cloud now charges specification fees for guaranteed-performance SLB instances.

1. What are guaranteed-performance instances?

A guaranteed-performance instance provides guaranteed performance metrics (performance SLA) and is opposite to a shared-performance instance. For a shared-performance instance, the performance metrics are not guaranteed and the resources are shared by all instances.

All instances are shared-performance instances before Alibaba launches guaranteed-performance instances. You can view the instance type on the console.

Hover your mouse pointer to the green icon of the target instance to view the performance metrics , as shown in the following figure.



The following are three key performance metrics for guaranteed-performance instances:

- Max Connection

The maximum number of connections to a SLB instance. When the maximum number of connections reaches the limits of the specification, the new connection will be dropped.

- Connection Per Second (CPS)

The rate at which a new connection is established per second. When the CPS reaches the limits of the specification, the new connection will be dropped.

- Query Per Second (QPS)

The number of HTTP/HTTPS requests that can be processed per second. When the QPS reaches the limits of the specification, the new connection will be dropped.

Alibaba Cloud Server Load Balancer provides the following capacities for guaranteed-performance instances:

Type	Specification	Max Connection	CPS	Query Per Second (QPS)
Specification 1	Small I (slb.s1.small)	5,000	3,000	1,000
Specification 2	Standard I (slb.s2.small)	50,000	5,000	5,000
Specification 3	Standard II (slb.s2.medium)	100,000	10,000	10,000
Specification 4	Higher I (slb.s3.small)	200,000	20,000	20,000
Specification 5	Higher II (slb.s3.medium)	500,000	50,000	30,000
Specification 6	Super I (slb.s3.large)	1,000,000	100,000	50,000

If you want to use a larger specification, contact your customer manager.

2. How are guaranteed-performance instances billed?

Guaranteed-performance instances are billed as follows:

Total fee (per instance) = instance fee + traffic fee + specification fee



Note:

The corresponding specification fee is billed for each guaranteed-performance instance regardless of the network type of the instance, and is billed based on the actual usage depending on the specification selected. If the actual performance metrics of an instance occurs between two capacities, the specification fee is charged at the higher specification fee.

The specification fee of a performance-guarantee instance is charged by usage. No matter what kind of specification you choose, the instance specification fee will be charged according to the specification you actually use.

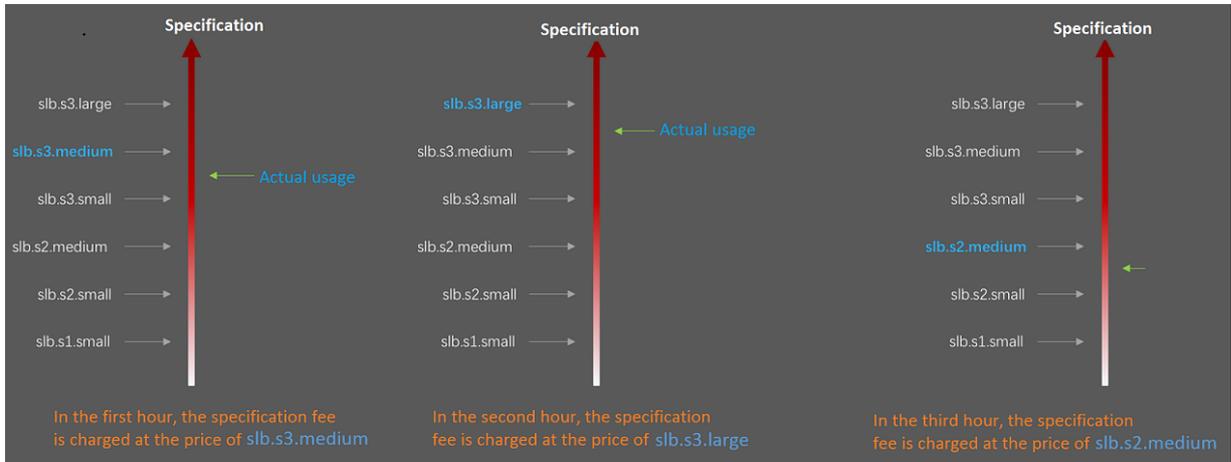
For example, if you purchase the slb.s3.large specification (1,000,000; CPS 500,000; QPS 50,000) and the actual usage of your instance in an hour is as follow:

Max Connection	CPS	QPS
90,000	4,000	11,000

- From the perspective of Max Connection, the actual metrics 90,000 occurs between the limit 50,000 defined in the Standard I (slb.s2.small) specification and the limit 100,000 defined in the Standard II (slb.s2.medium) specification. Therefore, the specification of the Max Connection metrics in this hour is Standard II (slb.s2.medium).
- From the perspective of CPS, the actual metrics 4,000 occurs between the limit 3,000 defined in the Small I (slb.s1.small) specification and the limit 5,000 defined in the Standard I (slb.s2.small) specification. Therefore, the specification of the CPS metrics in this hour is Standard I (slb.s2.small).
- From the perspective of QPS, the actual metrics 11,000 occurs between the limit 10,000 defined in the Standard II (slb.s2.medium) specification and the limit 20,000 defined in the Higher I (slb.s3.small) specification. Therefore, the specification of the QPS metrics in this hour is Higher I (slb.s3.small).

Comparing these three metrics, the specification of the QPS metrics is highest, therefore, the specification fee of the instance in this hour is charged at the price of the Higher I (slb.s3.small) specification.

The following figure is an example showing how the specification fee is billed for an SLB instance in the first three hours:



The billing of the guaranteed-performance instances is flexible. The specification you select when purchasing an instance is the performance limitation of the instance. For example, if slb.s3.medium is selected, the new connections are dropped when the HTTP requests in one second reach 30,000.

3. What is the price of each specification?

The following table lists the price of each specification. In addition to the specification fee, you are also charged for instance fee and traffic fee. For more information, see [Pay-As-You-Go](#).

Region	Type	Max Connection	CPS	QPS	Specification fee (USD/Hour)
China (Hangzhou)	Small I (slb.s1.small)	5,000	3,000	1,000	Free
China (Zhangjiakou)	Standard I (slb.s2.small)	50,000	5,000	5,000	0.05
China (Huhhot)					
China (Qingdao)	Specification 3: Standard II (slb.s2.medium)	100,000	10,000	10,000	0.10
China (Beijing)					
China (Shanghai)	Higher I (slb.s3.small)	200,000	20,000	20,000	0.20
China (Shenzhen)	Higher II (slb.s3.medium)	500,000	50,000	30,000	0.31
	Super I (slb.s3.large)	1,000,000	100,000	50,000	0.51
Singapore	Small I (slb.s1.small)	5,000	3,000	1,000	Free

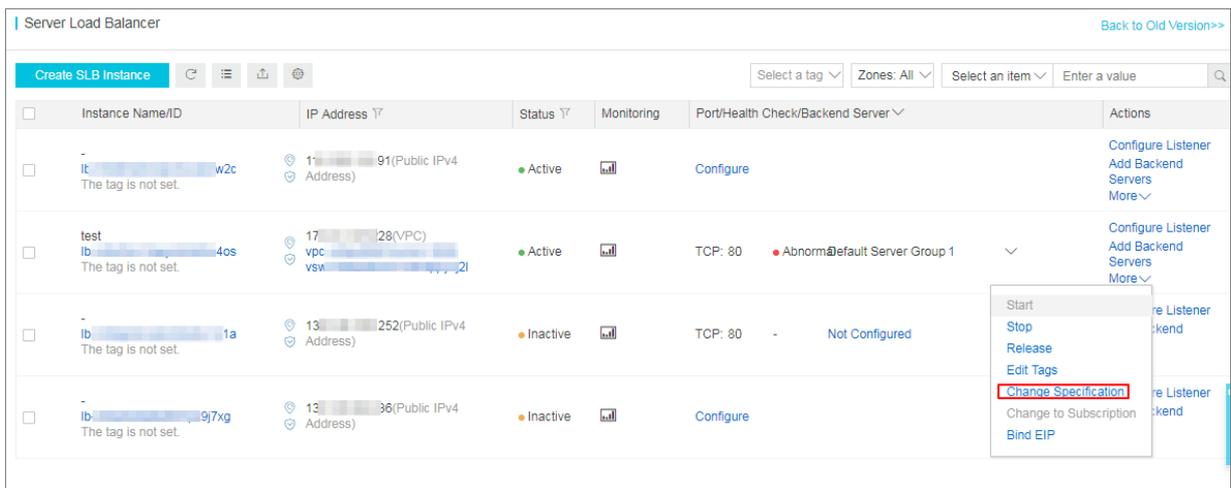
Region	Type	Max Connection	CPS	QPS	Specification fee (USD/Hour)
Malaysia (Kuala Lumpur)	Standard I (slb.s2.small)	50,000	5,000	5,000	0.06
Indonesia (Jakarta)	Standard II (slb.s2.medium)	100,000	10,000	10,000	0.12
India (Mumbai)					
US (Silicon Valley)	Higher I (slb.s3.small)	200,000	20,000	20,000	0.24
US (Virginia)	Higher II (slb.s3.medium)	500,000	50,000	30,000	0.37
China (Hong Kong)					
	Super I (slb.s3.large)	1,000,000	100,000	50,000	0.61

4. How to select a guaranteed-performance instance?

Because the specification fee is billed based on the actual usage, we recommend that you select the largest specification (slb.s3.large). This guarantees the business flexibility (flexibility) and will not cause extra costs. If your traffic does not reach the largest specification, you can select a more reasonable specification, such as slb.s3.medium.

5. Can I modify the specification after the instance is created?

Yes. You can change the specification at any time and the change takes effect immediately.



6. When will the guaranteed-performance instances be charged?

Alibaba Cloud plans to charge specification fee on guaranteed-performance Server Load Balancer instances from April 1st, 2018, and continue to sell shared-performance Server Load Balancer instances.

The fee collection for the guaranteed-performance instances take effect in batches by regions:

- The first batch:

Effective time: From April 1 to April 10

Regions: Singapore, Malaysia (Kuala Lumpur), Indonesia (Jakarta), India (Mumbai), US (Silicon Valley), US (Virginia)

- The second batch:

Effective time: From April 11 to April 20

Effective regions: China (Hangzhou), China (Zhangjiakou), China (Hohhot), China (Hong Kong)

- The third batch:

Effective time: From April 21 to April 30

Effective regions: China (Qingdao), China (Beijing), China (Shanghai), China (Shenzhen)

7. After Alibaba Cloud starts to charge specification fee on guaranteed-performance instances, will extra fees be charged on shared-performance instances?

Not at all.

The billing of the original shared-performance instances is the same if you do not change it to a performance-guaranteed instance. However, if you change the shared-performance instance to the guaranteed-performance one, the specification fee will be charged.

8. Why sometimes guaranteed-performance instances cannot reach the performance limit as defined in the specification?

It applies the cask theory.

Guaranteed-performance instances do not guarantee that the three metrics can reach the specification limits at the same time. The limitation is triggered as long as a metric first reaches the limitation defined in the specification.

For example, you have purchased a guaranteed-performance instance of the Higher I (slb.s3.small) specification. When the QPS of the instance reaches 20,000 but the number of maximum

connections does not reach 200,000, the new connections are still dropped because the QPS has reached the limitation.

9. Can I still buy shared-performance instances?

Yes.

However, the shared-performance instances will be phased out in the future. Please pay attention to the official notifications.

10. Will intranet SLB instances be charged for specification fee?

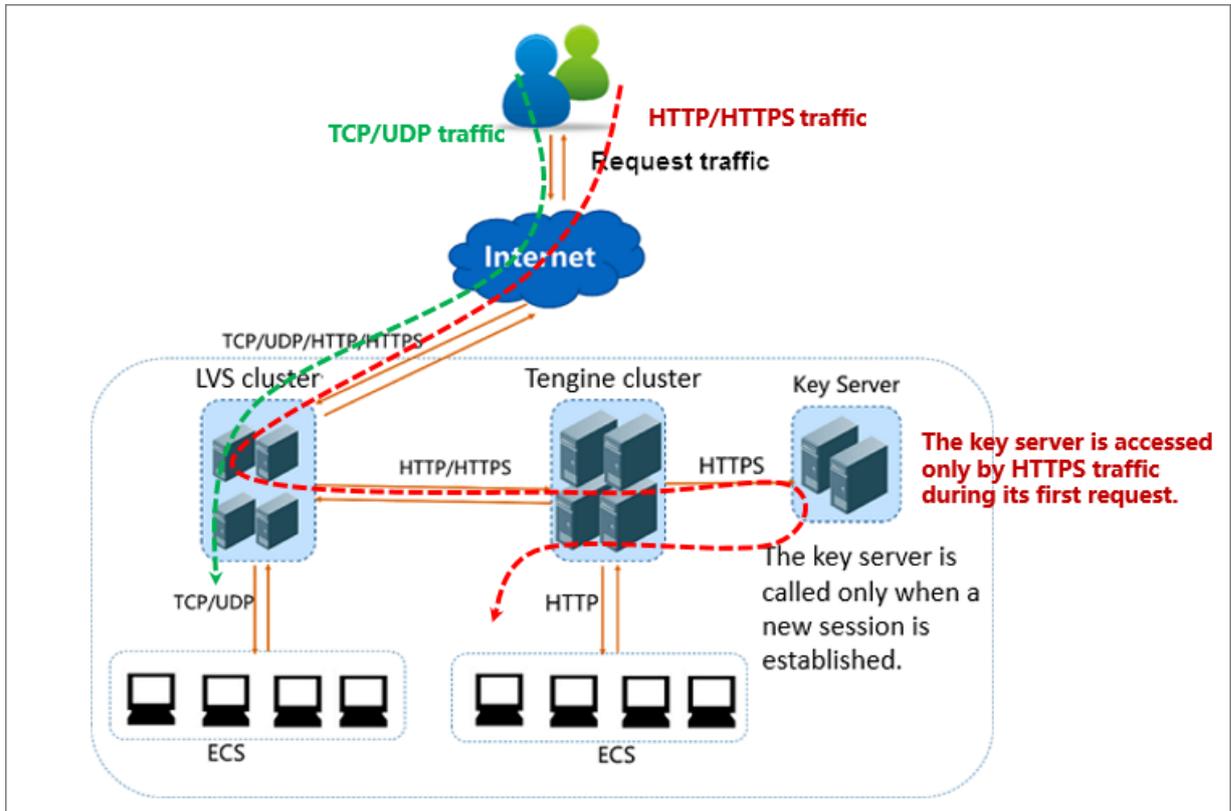
If the intranet SLB instance is a shared-performance instance, no specification fee is charged. If the intranet SLB instance is a guaranteed-performance instance, corresponding specification fee is charged, and no other fees are charged. The specification fees are collected as the same as the Internet guaranteed-performance instances, but intranet guaranteed-performance instances are free from instance fee and traffic fee.

1.3 Network traffic flow

As a traffic forwarding service, SLB forwards requests from clients to backend servers through the SLB cluster and then the backend servers return the responses to SLB through the intranet.

Inbound network traffic

SLB distributes incoming traffic according to forwarding rules configured on the console or API. The inbound network traffic flow is shown as [Figure 1-1: Inbound network traffic](#).

Figure 1-1: Inbound network traffic

1. Regardless of TCP/UDP protocol or HTTP/HTTPS protocol, the incoming traffic must be forwarded through the LVS cluster first.
2. Numerous inbound traffic is distributed evenly among all node servers in the LVS cluster, and the node servers synchronizes session to guarantee high availability.
 - For Layer-4 listeners (the frontend protocol is UDP or TCP), the node servers in the LVS cluster distribute requests directly to backend ECS instances according to the configured forwarding rules.
 - For Layer-7 listeners (the frontend protocol is HTTP), the node servers in the LVS cluster first distribute requests to the Engine cluster. Then, the node servers in the Engine cluster distribute the requests to backend ECS instances according to the configured forwarding rules.
 - For Layer-7 listeners (the frontend protocol is HTTPS), the request distribution is similar to the HTTP protocol. However, before distributing the requests to backend ECS instances, the system will call the Key Server to validate certificates and decrypt data packets.

Outbound network traffic

SLB communicates with backend ECS instances through the intranet.

4. A public IP (configured when you create an ECS instance), EIP, and NAT gateway can all achieve mutual Internet access (access or accessed), but they cannot forward traffic or balance traffic loads.

1.4 Create an SLB instance

Prerequisites

Before creating an SLB instance, make sure that you have properly prepared the environment. For more information, see [Plan and prepare](#).

Procedure

1. Log on to the [SLB console](#).
2. In the left-side navigation pane, click **Instances** > **Server Load Balancer**, and click **Create SLB Instance** in the upper-left corner.
3. Configure the SLB instance according to the following information.

Configuration	Description
Region	<p>Select the region where the SLB instance is located.</p> <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 5px; margin-top: 10px;">  Note: Make sure that the region of the SLB instance is the same as that of backend ECS instances. </div>
Zone Type	<p>Display the zone type of the selected region. The zone of a cloud product refers to a set of independent infrastructure and is usually represented by Internet data centers (IDCs). Different zones have independent infrastructure (network, power supply, air-conditioning and so on). Therefore, an infrastructure fault in one zone will not affect other zones. A zone belongs to a specific region, however, a single region may have one or more zones. SLB has deployed multi-zone in most regions.</p> <ul style="list-style-type: none"> • Single zone: The SLB instance is deployed only in one zone. • Multi-zone: The SLB instance is deployed in two zones. By default, the instance in the primary zone is used to distribute traffic. If the primary zone is faulty, the instance in the backup zone will automatically take over the load balancing service.
Primary Zone	Select the primary zone for the SLB instance. The primary zone carries traffic in normal conditions.

Configuration	Description
Backup Zone	Select the backup zone for the SLB instance. The backup zone only takes over traffic when the primary zone is unavailable.
Instance Spec	Select a performance capacity for the instance. The performance metrics vary by specification. For more information, see Guaranteed-performance instances .
Instance Type	Select the instance type based on your business needs. A public or a private IP address is allocated to the SLB instance based on the instance type. For more information, see SLB instance and network type . <ul style="list-style-type: none"> Internet: An Internet SLB instance only provides an Internet IP and you can access the SLB service from the Internet. Intranet: An intranet SLB instance only provides a private IP and you can only access the SLB service from the intranet.
Network Type	If the selected instance type is Intranet, you have to select a network type for the instance. <ul style="list-style-type: none"> Classic network: The IP of the instance is allocated and managed by Alibaba Cloud in a unified manner. VPC: The IP of the instance is allocated from the VSwitch CIDR block specified by you.
Purchase Quantity	Select the number of instances to create.

4. Click **Buy Now** and complete the payment.

1.5 Create an IPv6 instance

Server Load Balancer supports creating IPv6 instances. After an IPv6 instance is created, the system allocates a public IPv6 address to the instance to forward requests from IPv6 clients.

Context

IPv6 is the abbreviation of Internet Protocol Version 6. IPv6 is the next-generation IP protocol designed by IETF (Internet Engineering Task Force) to replace the current version of IP protocol (IPv4). By extending the length of IPv4 address from 32 bits to 128 bits, it expands the address space by 79,228,162,514,264,337,593,543,950,336 times. After IPv6 is used, each grain of sand on the world can be allocated with an IP address.



Note:

- Currently, only Zone E and Zone F in the China (Hangzhou) region as well as Zone F and Zone G in the China (Beijing) region support creating IPv6 instances and the instances must be guaranteed-performance instances.
- The Internet IPv6 network environment is still in the early stages of construction, and some links may cannot be accessed. If such problem occurs, submit a ticket. Besides, SLA is not provided in the pre-release stage.
- Because IPv6 has a longer IP head than IPv4, when you use a UDP listener on an IPv6 SLB instance, you must ensure that the MTU of the NIC communicating with the SLB on the backend server (ECS instance) is not greater than 1480 (some applications require synchronizing its configuration files based on this MTU value), otherwise the packets may be discarded because they are too large.

If you use a TCP/HTTP/HTTPS listener, no additional configurations are required because the TCP protocol supports MSS auto-negotiation.

SLB IPv6 has the following features:

- Smooth migration, which is not sensed by the service

You can directly bind ECS instances using IPv4 addresses to an IPv6 SLB instance and smoothly migrate the service to IPv6 without transforming the original system.

Adding the IPv6 entry has no impact on the original IPv4 service. If the traffic volume increases , you only need to increase the backend ECS instances.

- IPv6 access control ensures more security and reliable service deployment

Alibaba Cloud SLB supports IPv6 access control. You can configure the access control list according to your business needs.

- A blacklist can effectively block the access of malicious addresses to the SLB service.
- If a whitelist is configured, only addresses in the whitelist can access the SLB service.

Procedure

1. Log on to the [SLB console](#).
2. Select **Instances > Server Load Balancer**.
3. On the **Server Load Balancer** page, click **Create SLB Instance** in the upper-left corner.
4. Configure the SLB instance. For the IP version, select **IPv6**.

Other configurations are the same as configurations of common instances. See [SLB configurations](#).



Note:

Currently, only Zone E and Zone F in the China (Hangzhou) region as well as Zone F and Zone G in the China (Beijing) region support creating IPv6 instances and the instances must be guaranteed-performance instances.

5. Go back to the Server Load Balancer page to view the created IPv6 instance.

Result

Once the IPv6 instance is created, the system allocates an IPv6 address to it.

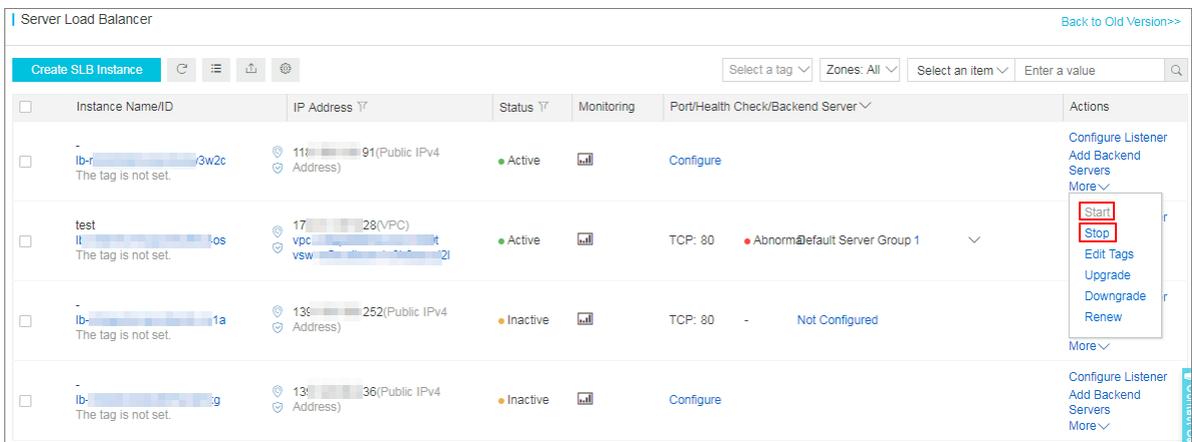
Instance Name/ID	IP Address	Status	Monitoring	Port/Health Check/Backend Server	Actions
auto_named_slb-8u <small>The tag is not set.</small>	24...4(Public IPv6 Address)	Active		Configure	Configure Listener Add Backend Servers More

1.6 Start or stop an SLB instance

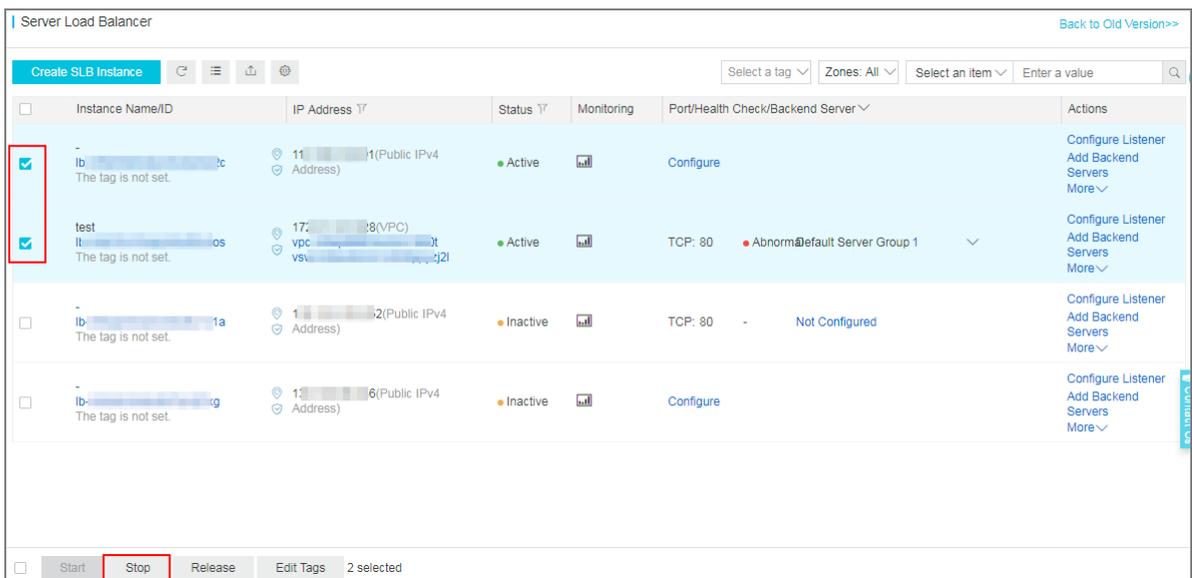
You can start or stop an SLB instance at any time. After being stopped, an SLB instance does not receive or forward requests any more.

Procedure

1. Log on to the [SLB console](#).
2. In the left-side navigation pane, click **Instances > Server Load Balancer**.
3. Select a region and find the target instance.
4. In the **Actions** column, click **More > Start** or **More > Stop**.



5. If you want to start or stop multiple instances at a time, select the target instances and click **Start** or **Stop** at the lower part of the page.



1.7 Bind an EIP

You can bind an EIP to an SLB instance of the VPC network. After being bound to an EIP, the SLB instance can forward requests from the Internet.

Procedure

1. Log on to the [SLB console](#).
2. In the left-side navigation pane, click **Instances** > **Server Load Balancer**.
3. Select a region and find the target instance.



Note:

Ensure that the SLB instance is of the VPC network.

4. Click **More** > **Bind EIP**.

The screenshot shows the 'Server Load Balancer' console interface. At the top, there is a 'Create SLB Instance' button and a search bar. Below is a table of instances:

Instance Name/ID	IP Address	Status	Monitoring	Port/Health Check/Backend Server	Actions
lb-...2c The tag is not set.	1...91(Public IPv4 Address)	Active		Configure	Configure Listener Add Backend Servers More
test lb-...ps The tag is not set.	17...228(VPC)	Active		TCP: 80 Abnormal Default Server Group 1	Configure Listener Add Backend Servers More
lb-...1a The tag is not set.	13...252(Public IPv4 Address)	Inactive		TCP: 80 - Not Configured	Start Stop Release Edit Tags Change Specification Change to Subscription Bind EIP
lb-...xg The tag is not set.	13...36(Public IPv4 Address)	Inactive		Configure	Configure Listener Add Backend Servers More

5. Select an EIP and click **OK**.

1.8 Release an instance

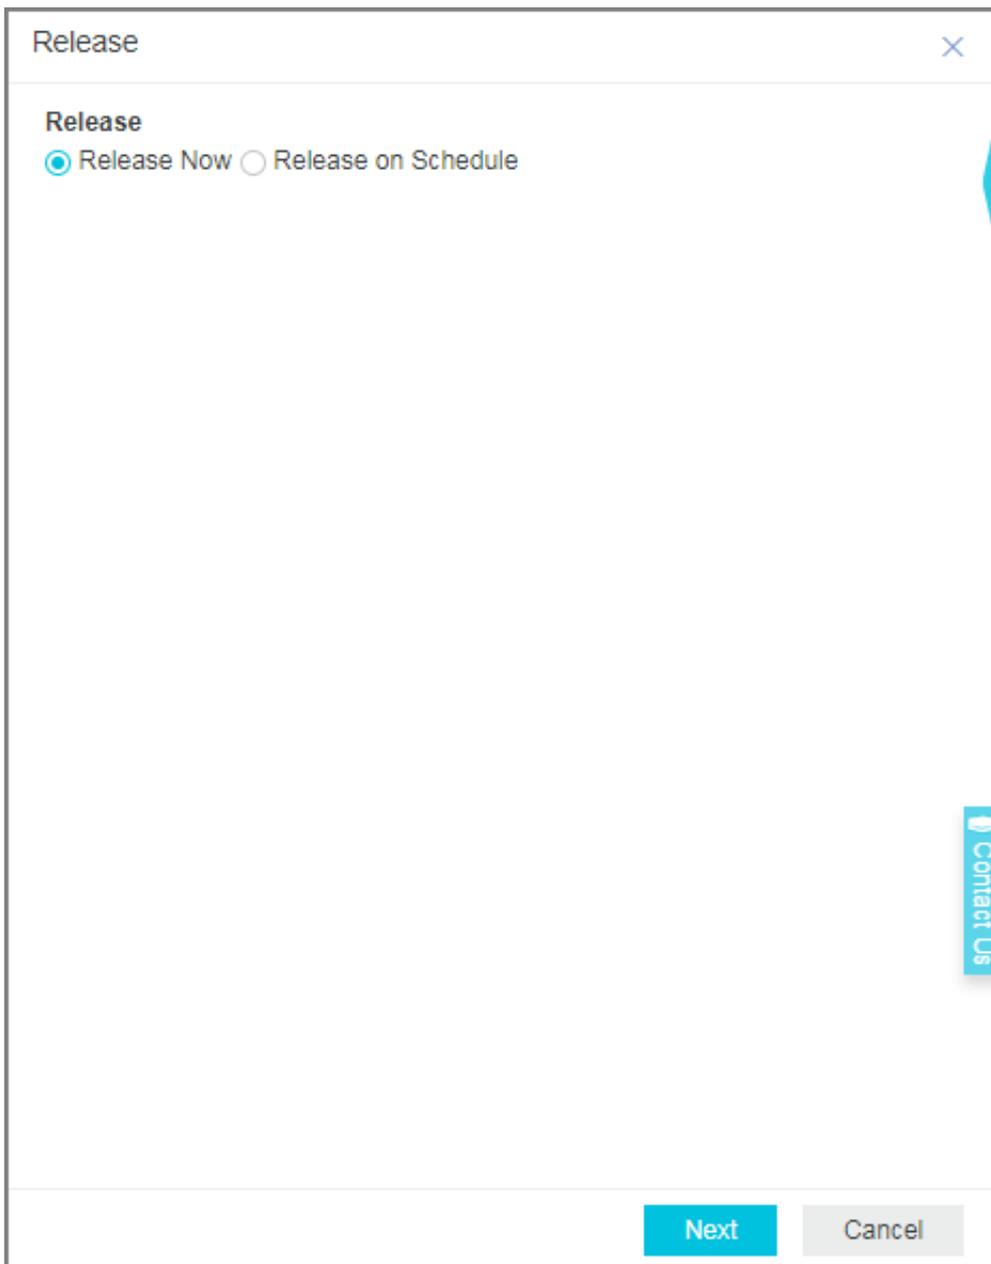
You can release an SLB instance immediately or on schedule.

Context

Procedure

1. Log on to the [SLB console](#).
2. Find the target instance and then click **More** > **Release**.

You can select multiple SLB instances at a time and click **Release** in the lower part of the page to release SLB instance in batches.



Release

Release

Release Now Release on Schedule

Contact Us

Next Cancel

3. On the **Release** page, select to release now or release on schedule.



Note:

The system executes the release operation every one hour or every 30 minutes, but will stop billing at the release time you set.

4. Click **Next**.
5. Confirm the displayed information and click **OK** to release the instance.

1.9 Manage tags

With tagging, you can classify Server Load Balancer instances by tags.

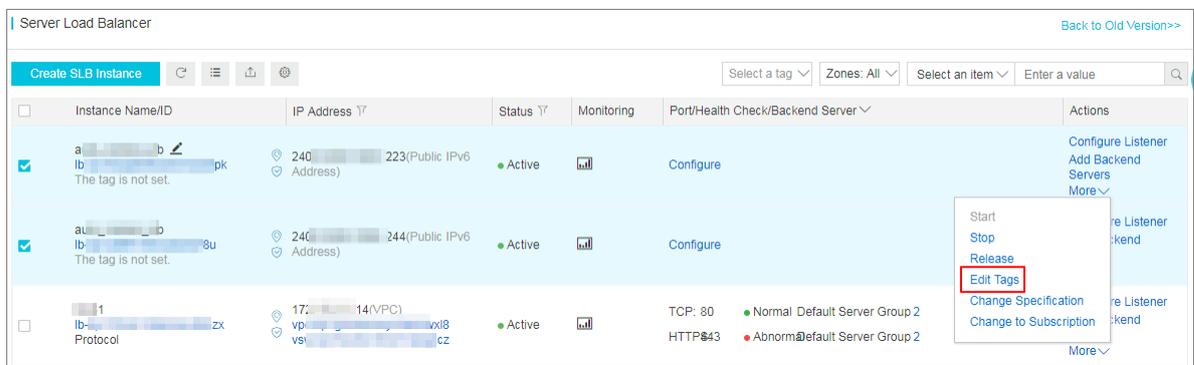
Each tag consists of a key and a value. Note the following limits when using tags:

- A tag cannot exist on its own and must be bound to an SLB instance.
- Up to 10 tags can be bound to an SLB instance.
- The key of each tag added to an instance must be unique. Otherwise, tags of the same key will be overwritten.
- Tags cannot be used across regions and are region-specific resources. For example, tags created in China (Hangzhou) are invisible in China (Shanghai).

Add a tag

To add a tag, complete these steps:

1. Log on to the [SLB console](#).
2. In the left-hand navigation pane, select **Instances > Server Load Balancer**.
3. Select a region and find the target instance.
4. In the **Actions** column, select **More > Edit Tags**.



5. On the **Edit Tags** page, complete these steps:

- a. If there are available tags, click **Saved Tags** and then select the tag to add.
- b. If you want to create a new tag, on the **Edit Tags** page, click **New Tag**, then enter the key and value of the new tag and click **OK**.

Edit Tags

Each resource can have a maximum number of 10 tags. The number of tags that can be added or removed per operation cannot exceed 5.

Add Tags

Key Value **OK** Cancel

Saved Tags

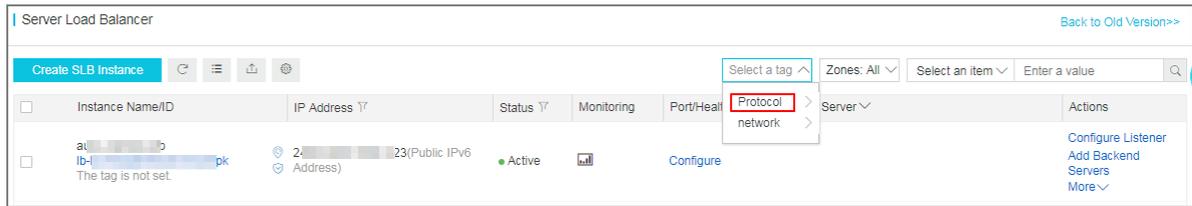
OK Cancel

c. Click **OK**.

Search instances using a tag

To search instances using a tag, complete these steps:

1. Log on to the [SLB console](#).
2. In the left-hand navigation pane, click **Instances** > **Server Load Balancer**.
3. Select a region and find the target instance.
4. Click **Select a tag**, and then select the tag to be used as the search criteria.



5. You can click the delete icon next to the selected tag to clear the filter.

Delete a tag

SLB does not support deleting tags of multiple instances in batches. You can only remove the tags of an instance at a time.

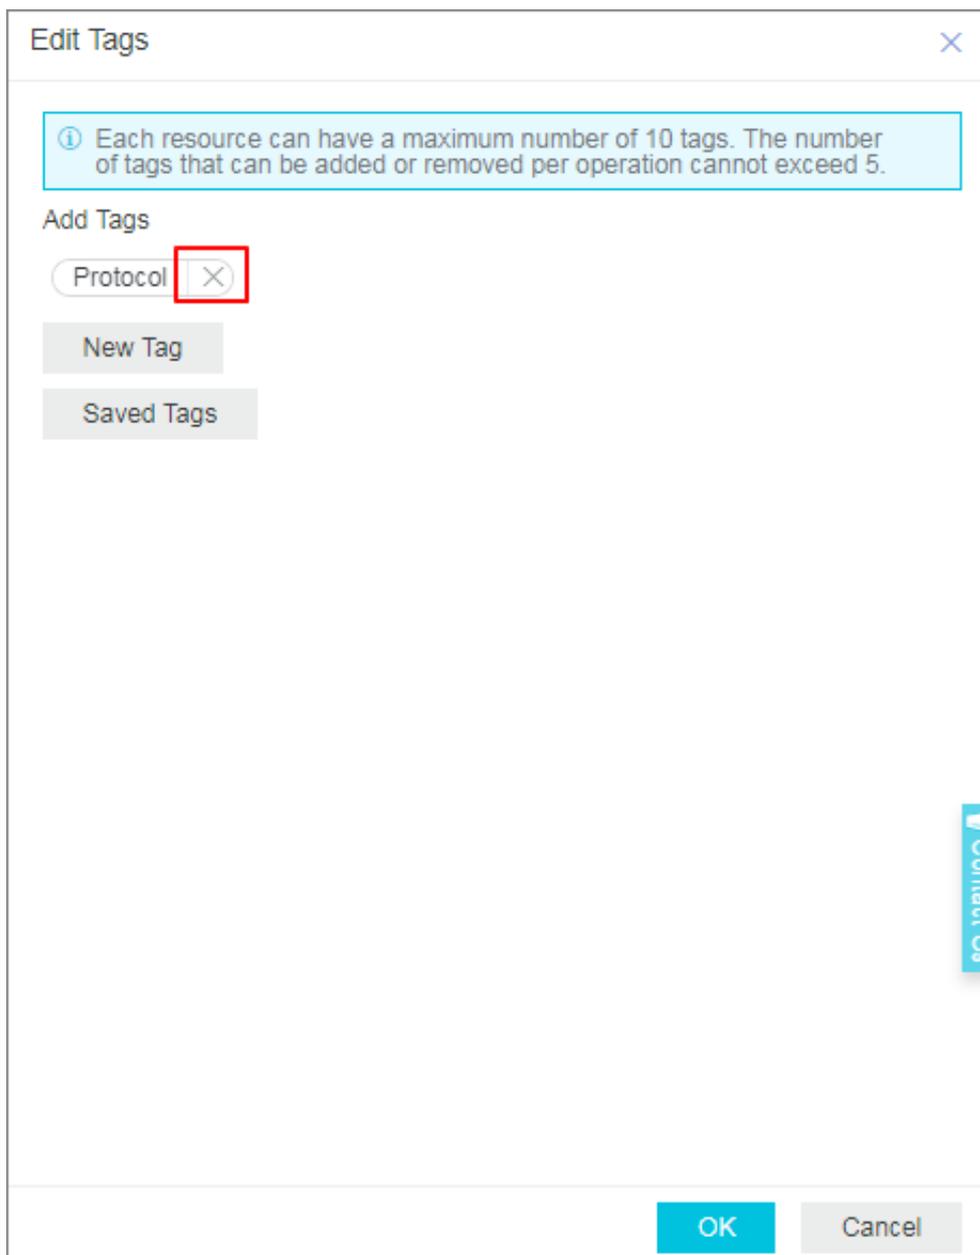
To delete a tag, complete these steps:

1. Log on to the [SLB console](#).
2. In the left-hand navigation pane, click **Instances** > **Server Load Balancer**.
3. Select a region and find the target instance.
4. In the **Actions** column, select **More** > **Edit Tags**.
5. On the **Edit Tags** page, click the delete icon next to the tag to be removed, and then click **OK**.



Note:

If a tag is removed from one instance and is not bound to any other instances, the tag is removed from the system.



1.10 Expiring Instances

You can manage overdue instances.

Context

If you do not renew an expiring instance, the instance will be released automatically.

- After an instance is added to expiring instances, it will only be reserved for one day.

Procedure

1. Log on to the [SLB console](#).
2. Select **Instances > Expiring Instances**.

3. View detailed information of overdue instances.
4. Click **Renew** in the **Actions** column of the target SLB instance, then the instance will added back to the Server Load Balancer list.

1.11 Change the configuration

You can change a shared-performance instance to a guaranteed-performance instance, or modify the capacity of a guaranteed-performance instance.

Context

Before modifying the instance configuration, note the following:

- If you change a shared-performance instance to a guaranteed-performance instance, a brief disconnection of service may occur for 10 to 30 seconds.

We recommend that you change the configuration in the low traffic period, or use DNS to schedule services to other SLB instances first before changing the configuration.

- Once a shared-performance instance is changed to a guaranteed-performance instance, it can no longer be changed back.

You can use the (slb.s1.small) capacity instead after changing to the guaranteed-performance instance.

Procedure

1. Log on to the [SLB console](#).
2. Select a region.
3. Find the target instance, select **More > Change Specification**.
4. In the **Configuration Upgrade** area, select a new specification, and complete the payment.

2 Listeners

2.1 Add a TCP listener

TCP listeners are applicable to scenarios with high requirements on reliability and data accuracy, but with tolerance for low speed, such as file transmission, sending or receiving emails, remote logon, and web applications without special requirements. You can add a TCP listener to forward requests from the TCP protocol.

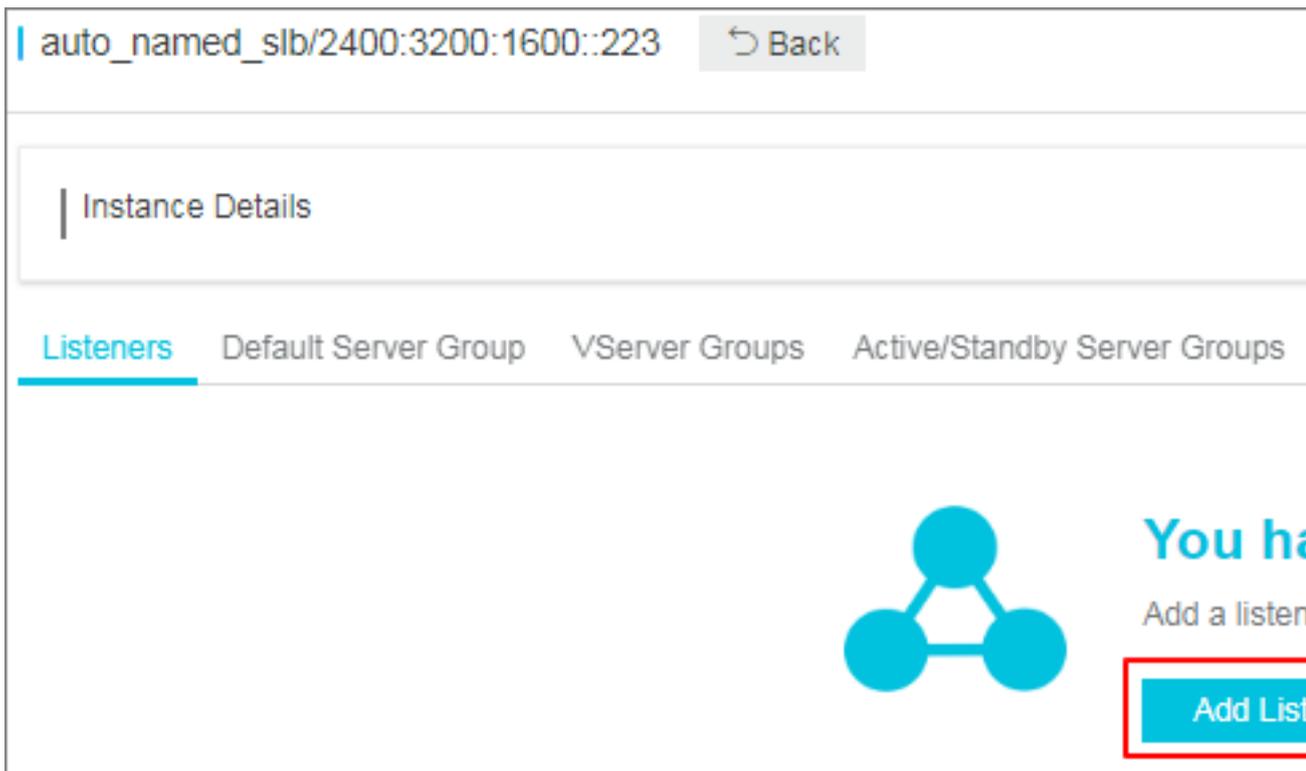
Prerequisites

[Create an SLB instance.](#)

Step 1 Open the listener configuration wizard

To open the listener configuration wizard, complete these steps:

1. Log on to the [SLB console](#).
2. In the left-side navigation pane, select **Instances > Server Load Balancer**.
3. Select a region.
4. Select one of the following methods to open the listener configuration wizard:
 - On the **Server Load Balancer** page, find the target instance and then click **Configure Listener**.
 - On the **Server Load Balancer** page, click the ID of the target SLB instance. On the **Listeners** page, click **Add Listener**.



Step 2 Configure a TCP listener

To configure a TCP listener, complete these steps:

1. On the **Protocol and Listener** page, configure the TCP listener according to the following information.

Configuration	Description
Select Listener Protocol	Select the protocol type of the listener. In this tutorial, select TCP .
Listening Port	The listening port used to receive requests and forward the requests to backend servers. The port number is in the range of 1-65535. <div style="background-color: #f0f0f0; padding: 5px;">  Note: The listening ports must be unique in a Server Load Balancer instance. </div>
Advanced configurations	
Scheduling Algorithm	Server Load Balancer supports three scheduling algorithms: round robin, weighted round robin (WRR), and weighted least connections (WLC).

Configuration	Description
	<ul style="list-style-type: none"> • Weighted Round-Robin (WRR) (default): Backend servers with higher weights receive more requests than those with smaller weights. • Round-Robin (RR): Requests are evenly and sequentially distributed to the backend servers. • Weighted Least Connections (WLC): A server with a higher weight will receive a larger percentage of live connections at any one time. When the weight value is the same, a backend server with a smaller number of connections is more frequently (and probably) accessed. • (Supported in some regions) Consistent Hash (CH): <ul style="list-style-type: none"> — Source IP: The consistent hash based on the source IP address. The same source IP addresses are scheduled to the same backend server. — Tuple : The consistent hash based on the quaternion (source IP + destination IP + source port + destination port). The same steams are scheduled to the same backend server.
Enable Session Persistence	<p>Select whether to enable session persistence.</p> <p>If session persistence is enabled, all session requests from the same client are sent to the same backend server.</p> <p>For TCP listeners, session persistence is based on IP addresses. Requests from the same IP address are forwarded to the same backend server.</p>
Enable Access Control	<p>Select whether to enable the access control function.</p>
Access Control Method	<p>Select an access control method after enabling the access control function:</p> <ul style="list-style-type: none"> • Whitelist: Only requests from IP addresses or CIDR blocks in the selected access control lists are forwarded. It applies to scenarios where the application only allows access from some specific IP addresses. <p>Enabling whitelist poses some business risks. After a whitelist is configured, only the IP addresses in the list can access the listener. If you enable the whitelist without adding any IP entry in the corresponding access control list, all requests are forwarded.</p>

Configuration	Description
	<ul style="list-style-type: none"> • Blacklist: Requests from IP addresses or CIDR blocks in the selected access control lists are not forwarded. It applies to scenarios where the application only denies access from some specific IP addresses. <p>If you enable a blacklist without adding any IP entry in the corresponding access control list, all requests are forwarded.</p>
Access Control List	<p>Select an access control list as the whitelist or the blacklist.</p> <div style="background-color: #f0f0f0; padding: 5px;">  Note: An IPv6 instance can only bind IPv6 access control lists and an IPv4 instance can only bind IPv4 access control lists. For more information, see Configure an access control list. </div>
Enable Peak Bandwidth Limit	<p>Select whether to configure the listening bandwidth. If the SLB instance is billed by bandwidth, you can set different peak bandwidths for different listeners to limit the traffic passing through the listeners. The sum of the peak bandwidths of all listeners under an instance cannot exceed the bandwidth of that instance.</p> <p>By default, all listeners share the bandwidth of the SLB instance.</p> <div style="background-color: #f0f0f0; padding: 5px;">  Note: Instances billed by traffic have no peak bandwidth limit by default. </div>
Idle Timeout	Specify the idle connection timeout in seconds. Valid value: 10-900
Listener Name	Configure the name of the listener.
Get Client Source IP Address	The backend server of a Layer-4 listener can directly obtain the real IP of the client.
Automatically Activate Listener after Creation	Choose whether to enable listener after the listener is configured. The listener is enabled by default.

2. Click **Next**.

Configure Server Load Balancer Back

Protocol and Listener Backend Servers Health Check Submit

Select Listener Protocol

TCP UDP HTTP HTTPS

Listening Port

80

This port is being used by another listener.

Advanced Hide

Scheduling Algorithm

Weighted Round-Robin (WRR) Weighted Least Connections (WLC) Round-Robin (RR)

Step 3 Add backend servers

Add backend servers to process requests. You can use the default server group configured for the instance, or configure a VServer group or an active/standby server group for the listener. For more information, see [Backend server overview](#).

In this tutorial, the default server group is used:

1. Select **Default Server Group** and then click **Add**.

Configure Server Load Balancer Back 监听介绍

Protocol and Listener SSL Certificates **Backend Servers** Health Check Submit

Add Backend Servers

ⓘ Add backend servers to handle the access requests received by the SLB instance.

Forward Requests To

Default Server Group VServer Group Active/Standby Server

Servers Added

You have not added any servers.

2. Select the ECS instances to add and then click **Add to Selected Server List**. Click **OK**.

The screenshot shows the 'Available Servers' dialog box. At the top, there is a search bar for 'ECS Instance Name' and a 'Buy ECS' button. Below is a table with the following columns: (checkbox), ECS Instance ID/Name, Public/Internal IP Address, Status, Zone, SLB Instance, and Actions. The table contains four rows of ECS instances. The first two rows are selected, indicated by checked checkboxes. A red box highlights these two rows. The bottom row of the table has a red box around the 'Add to Selected Server List' button. At the bottom of the dialog, there is a 'Selected Server List' section with a red box around the 'Add to Selected Server List' button, and 'OK' and 'Cancel' buttons.

<input type="checkbox"/>	ECS Instance ID/Name	Public/Internal IP Address	Status	Zone	SLB Instance	Actions
<input checked="" type="checkbox"/>	E-...T i-...5n	4... (Public) 1... 182 (Private)	Running	China East 1 Zone G	Associated SLB Instances 3	Add
<input checked="" type="checkbox"/>	E-...T i-...5n	4... (Public) 1... 183 (Private)	Running	China East 1 Zone G	Associated SLB Instances 3	Add
<input type="checkbox"/>	E-...1 i-...qye	1... 30 (Private)	Running	China East 1 Zone G	Associated SLB Instances 1	Add
<input type="checkbox"/>	E-...2 i-...ba10ghowaye	1... 159 (Private)	Running	China East 1 Zone G	Associated SLB Instances 1	Add

Selected Server List

3. Configure the ports and weights of the added backend servers.

- Port

The port opened on the backend server (ECS instance) to receive requests. The port number is in the range of 1-65535. Ports of backend servers can be the same in an SLB instance.

- Weight

The weight of the backend server (ECS instance). An ECS instance with a higher weight will receive a larger number of connection requests.



Note:

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

Add Backend Servers

① Add backend servers to handle the access requests received by the SLB instance.

Forward Requests To

Default Server Group | VServer Group | Active/Standby Server Group

Servers Added

ECS Instance ID/Name	Public/Internal IP Address	端口
E-4z1	47...0(Public) 17...2(Private) vpc-...wxl8 vsw-...mfi	80
E-m	47...97(Public) ...83(Private) vpc-...3 vsw-...i0mfi	80

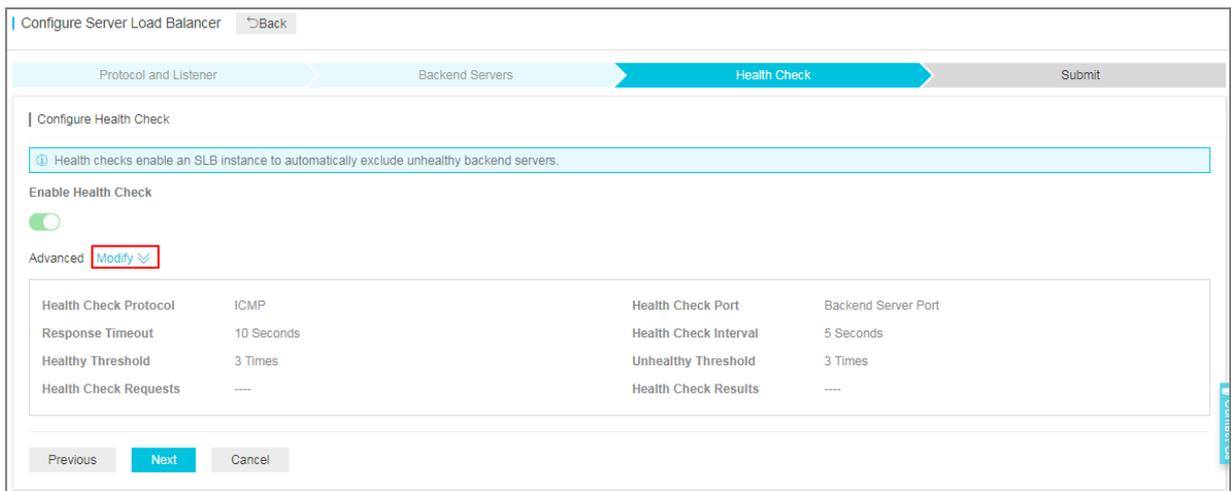
0 servers have been added. 2 servers are to be added, and 0 servers are to be deleted.

Previous | **Next** | Cancel

4. Click **Next**.

Step 4 Configure health check

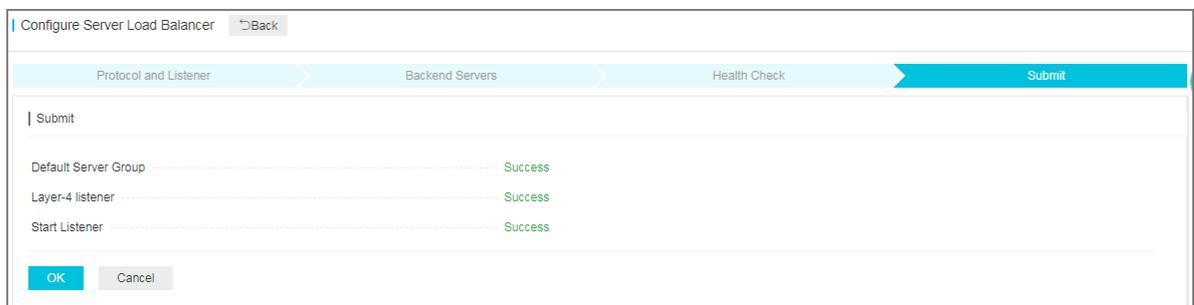
Server Load Balancer checks the service availability of the backend servers (ECS instances) by performing health checks. The health check function improves the overall availability of your services and avoids the impact of backend server failures. Click **Modify** to change health check configurations. For more information, see [Configure health check](#).



Step 5 Submit the configurations

To confirm the listener configurations, complete these steps:

1. On the **Submit** page, check listener configurations. You can click **Modify** to change the configurations.
2. Click **Submit**.
3. On the **Submit** page, click **OK** after the configurations are successful.



After the configurations are successful, you can view the created listener on the Listeners page.

Frontend Protocol/Port	Backend Protocol/Port	Name	Health Status	Monitoring	Forwarding Rule	Session Persistence	Peak Bandwidth	Server Group	Access Control List	Actions
UDP:143	UDP:80	udp_143	Abnormal		Weighted Round-Robin	Disabled	No Limit	Default Server Group	Disabled	Configure Details More
TCP:80	TCP:80	tcp_80	Normal		Weighted Round-Robin	Disabled	No Limit	Default Server Group	Disabled	Configure Details More
HTTPS:443	HTTP:80	-	Abnormal		Weighted Round-Robin	Disabled	No Limit	Default Server Group	Disabled	Configure Details Add Forwarding Rules More

Related operations

- [Configure health check](#)
- [Manage a default server group](#)
- [Manage a VServer group](#)
- [Manage an active/standby server group](#)
- [Configure access control](#)

2.2 Add a UDP listener

Applicable to scenarios with preference for real-time content over reliability, such as video chats and pushes of real-time financial quotations. You can add a UDP listener to forward requests from the UDP protocol.

Limits

Note the following before adding a UDP listener:

- The maximum number of connections per listener: 100,000.
- Currently, fragmented packets are not supported.
- UDP listeners of an SLB instance of the classic network do not support viewing the source IP address.
- In the following two scenarios, UDP listener configurations take effect after five minutes:
 - Remove the backend ECS instances.
 - Set the weight of a backend ECS instance to zero after the instance is declared as unhealthy.
- Because IPv6 has a longer IP head than IPv4, when you use a UDP listener on an IPv6 SLB instance, you must ensure that the MTU of the NIC communicating with the SLB on the backend server (ECS instance) is not greater than 1480 (some applications require synchronizing its configuration files based on this MTU value), otherwise the packets may be discarded because they are too large.

If you use a TCP/HTTP/HTTPS listener, no additional configurations are required because the TCP protocol supports MSS auto-negotiation.

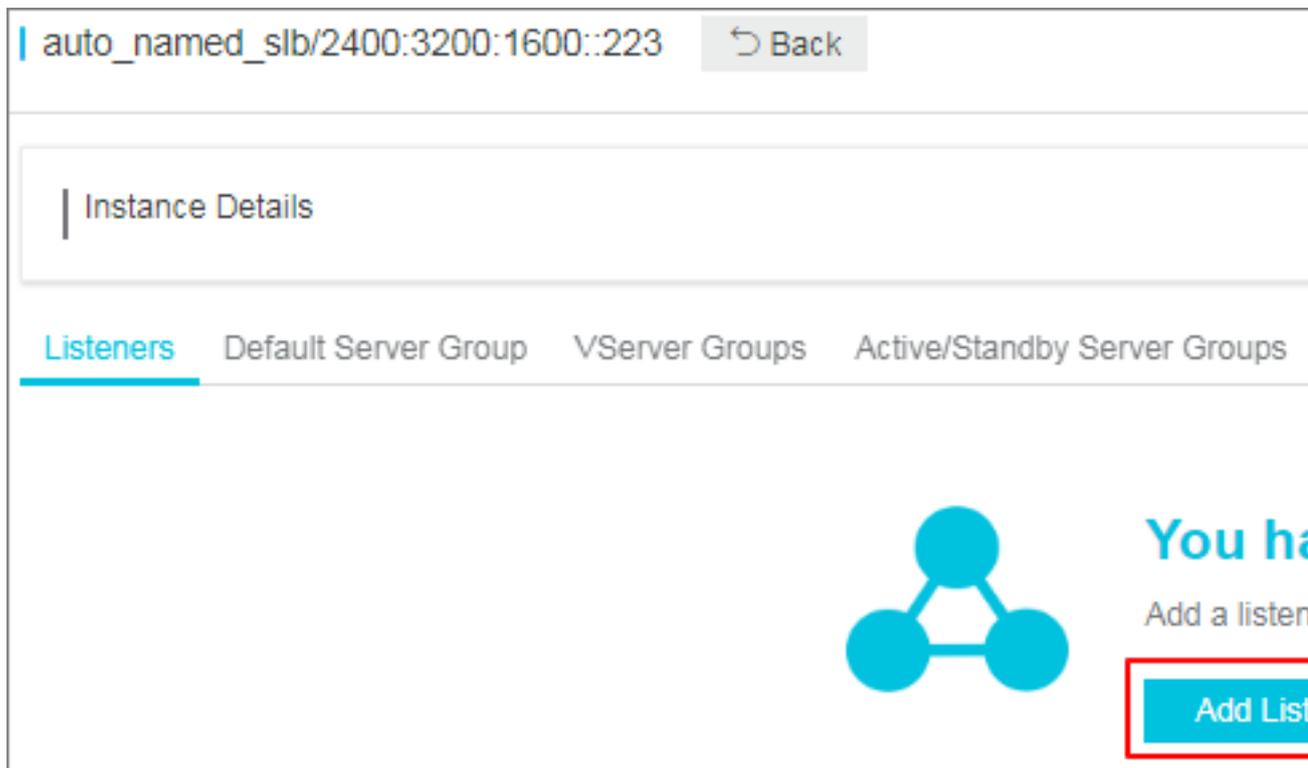
Prerequisites

[Create an SLB instance.](#)

Step 1 Open the listener configuration wizard

To open the listener configuration wizard, complete these steps:

1. Log on to the [SLB console](#).
2. In the left-side navigation pane, select **Instances** > **Server Load Balancer**.
3. Select a region.
4. Select one of the following methods to open the listener configuration wizard:
 - On the **Server Load Balancer** page, find the target instance and then click **Configure Listener**.
 - On the **Server Load Balancer** page, click the ID of the target SLB instance. On the **Listeners** page, click **Add Listener**.



Step 2 Configure a TCP listener

To configure a UDP listener, complete these steps:

1. On the **Protocol and Listener** page, configure the UDP listener according to the following information.

Configuration	Description
Select Listener Protocol	Select the protocol type of the listener.

Configuration	Description
	In this tutorial, select UDP .
Listening Port	<p>The listening port used to receive requests and forward the requests to backend servers. The port number is in the range of 1-65535.</p> <div style="background-color: #f0f0f0; padding: 5px; border: 1px solid #ccc;">  Note: The listening ports must be unique in a Server Load Balancer instance. </div>
Advanced configurations	
Scheduling Algorithm	<p>Server Load Balancer supports three scheduling algorithms: round robin, weighted round robin (WRR), and weighted least connections (WLC).</p> <ul style="list-style-type: none"> • Weighted Round-Robin (WRR) (default): Backend servers with higher weights receive more requests than those with smaller weights. • Round-Robin (RR): Requests are evenly and sequentially distributed to the backend servers. • Weighted Least Connections (WLC): A server with a higher weight will receive a larger percentage of live connections at any one time. When the weight value is the same, a backend server with a smaller number of connections is more frequently (and probably) accessed. • (Supported in some regions)) Consistent Hash (CH): <ul style="list-style-type: none"> — Source IP: The consistent hash based on the source IP address. The same source IP addresses are scheduled to the same backend server. — Tuple : The consistent hash based on the quaternion (source IP + destination IP + source port + destination port). The same steams are scheduled to the same backend server. — QUIC ID: Consistent hash based on the QUIC Connection ID. The same QUIC Connection IDs are scheduled to the same backend server.
Enable Access Control	Select whether to enable the access control function.
Access Control Method	<p>Select an access control method after enabling the access control function:</p> <ul style="list-style-type: none"> • Whitelist: Only requests from IP addresses or CIDR blocks in the selected access control lists are forwarded.

Configuration	Description
	<p>It applies to scenarios where the application only allows access from some specific IP addresses.</p> <p>Enabling whitelist poses some business risks. After a whitelist is configured, only the IP addresses in the list can access the listener. If you enable the whitelist without adding any IP entry in the corresponding access control list, all requests are forwarded.</p> <ul style="list-style-type: none"> Blacklist: Requests from IP addresses or CIDR blocks in the selected access control lists are not forwarded. It applies to scenarios where the application only denies access from some specific IP addresses. <p>If you enable a blacklist without adding any IP entry in the corresponding access control list, all requests are forwarded.</p>
<p>Access Control List</p>	<p>Select an access control list as the whitelist or the blacklist.</p> <div style="background-color: #f0f0f0; padding: 5px;">  <p>Note: An IPv6 instance can only bind IPv6 access control lists and an IPv4 instance can only bind IPv4 access control lists. For more information, see Configure an access control list.</p> </div>
<p>Enable Peak Bandwidth Limit</p>	<p>Select whether to configure the listening bandwidth. If the SLB instance is billed by bandwidth, you can set different peak bandwidths for different listeners to limit the traffic passing through the listeners. The sum of the peak bandwidths of all listeners under an instance cannot exceed the bandwidth of that instance.</p> <p>By default, all listeners share the bandwidth of the SLB instance.</p> <div style="background-color: #f0f0f0; padding: 5px;">  <p>Note: Instances billed by traffic have no peak bandwidth limit by default.</p> </div>
<p>Get Client Source IP Address</p>	<p>The backend server of a UDP listener can directly obtain the real IP of the client.</p> <div style="background-color: #f0f0f0; padding: 5px;">  <p>Note:</p> </div>

Configuration	Description
	UDP listeners of an SLB instance of the classic network do not support viewing the source IP address.
Automatically Activate Listener after Creation	Choose whether to enable listener after the listener is configured. The listener is enabled by default.

2. Click **Next**.

Step 3 Add backend servers

Add backend servers to process requests. You can use the default server group configured for the instance, or configure a VServer group or an active/standby server group for the listener. For more information, see [Backend server overview](#).

In this tutorial, the default server group is used:

1. Select **Default Server Group** and then click **Add**.

2. Select the ECS instances to add and then click **Add to Selected Server List**. Click **OK**.

The screenshot shows the 'Available Servers' dialog box. At the top, there is a search bar for 'ECS Instance Name' and a 'Buy ECS' button. Below is a table with the following columns: (checkbox), ECS Instance ID/Name, Public/Internal IP Address, Status, Zone, SLB Instance, and Actions. The table contains four rows of ECS instances. The first two rows are selected, indicated by checked checkboxes. A red box highlights these two rows. The bottom of the dialog features a 'Selected Server List' section with a red box around the 'Add to Selected Server List' button, along with 'OK' and 'Cancel' buttons.

<input type="checkbox"/>	ECS Instance ID/Name	Public/Internal IP Address	Status	Zone	SLB Instance	Actions
<input checked="" type="checkbox"/>	E-...T i-...5n	4... (Public) 1... 182 (Private)	Running	China East 1 Zone G	Associated SLB Instances 3	Add
<input checked="" type="checkbox"/>	E-...T i-...5n	4... (Public) 1... 183 (Private)	Running	China East 1 Zone G	Associated SLB Instances 3	Add
<input type="checkbox"/>	E-...1 i-...qye	1... 30 (Private)	Running	China East 1 Zone G	Associated SLB Instances 1	Add
<input type="checkbox"/>	E-...2 i-...ba10ghowaye	1... 159 (Private)	Running	China East 1 Zone G	Associated SLB Instances 1	Add

3. Configure the ports and weights of the added backend servers.

- Port

The port opened on the backend server (ECS instance) to receive requests. The port number is in the range of 1-65535. Ports of backend servers can be the same in an SLB instance.

- Weight

The weight of the backend server (ECS instance). An ECS instance with a higher weight will receive a larger number of connection requests.



Note:

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

Add Backend Servers

Add backend servers to handle the access requests received by the SLB instance.

Forward Requests To

Default Server Group
VServer Group
Active/Standby Server Group

Servers Added

ECS Instance ID/Name	Public/Internal IP Address	端口
ECS Instance ID/Name	47...J(Public) 17...2(Private) vpc-...wxl8 vsw-...mfi	80
ECS Instance ID/Name	47...97(Public) ...83(Private) vpc-...3 vsw-...i0mfi	80

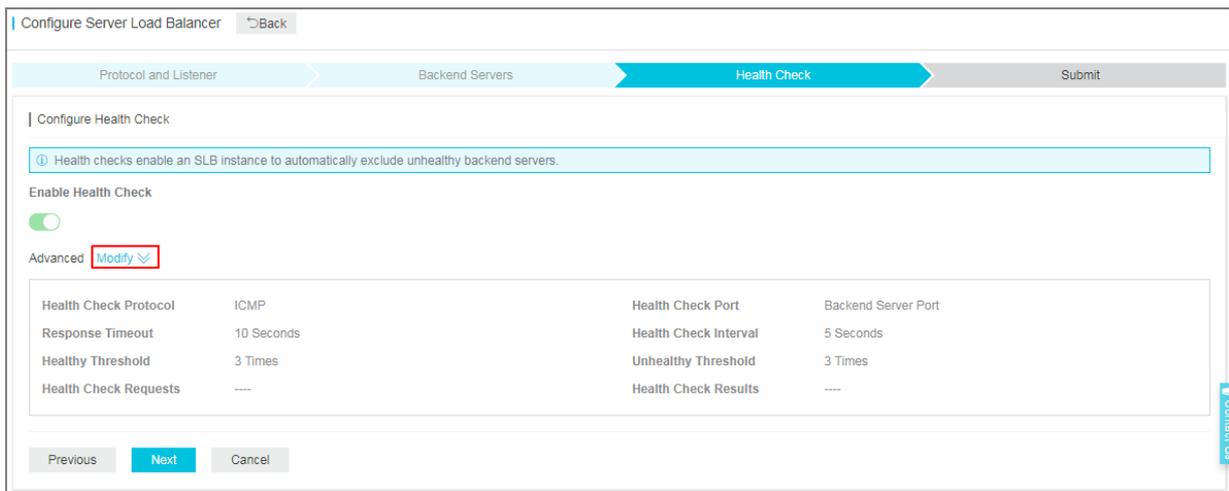
0 servers have been added. 2 servers are to be added, and 0 servers are to be deleted.

Previous
Next
Cancel

4. Click **Next**.

Step 4 Configure health check

Server Load Balancer checks the service availability of the backend servers (ECS instances) by performing health checks. The health check function improves the overall availability of your services and avoids the impact of backend server failures. Click **Modify** to change health check configurations. For more information, see [Configure health check](#).



Step 5 Submit the configurations

To confirm the listener configurations, complete these steps:

1. On the **Submit** page, check listener configurations. You can click **Modify** to change the configurations.
2. Click **Submit**.
3. On the **Submit** page, click **OK** after the configurations are successful.



After the configurations are successful, you can view the created listener on the Listeners page.

Frontend Protocol/Port	Backend Protocol/Port	Name	Health Status	Monitoring	Forwarding Rule	Session Persistence	Peak Bandwidth	Server Group	Access Control List	Actions
UDP:143	UDP:80	udp_143	Abnormal		Weighted Round-Robin	Disabled	No Limit	Default Server Group	Disabled	Configure Details More
TCP:80	TCP:80	tcp_80	Normal		Weighted Round-Robin	Disabled	No Limit	Default Server Group	Disabled	Configure Details More
HTTPS:443	HTTP:80	-	Abnormal		Weighted Round-Robin	Disabled	No Limit	Default Server Group	Disabled	Configure Details Add Forwarding Rules More

Related operations

- [Configure health check](#)
- [Manage a default server group](#)
- [Manage a VServer group](#)
- [Manage an active/standby server group](#)
- [Configure access control](#)

2.3 Add an HTTP listener

It is applicable to applications that need to recognize data contents, such as web applications and small-sized mobile games. You can add an HTTP listener to forward requests from the HTTP protocol.

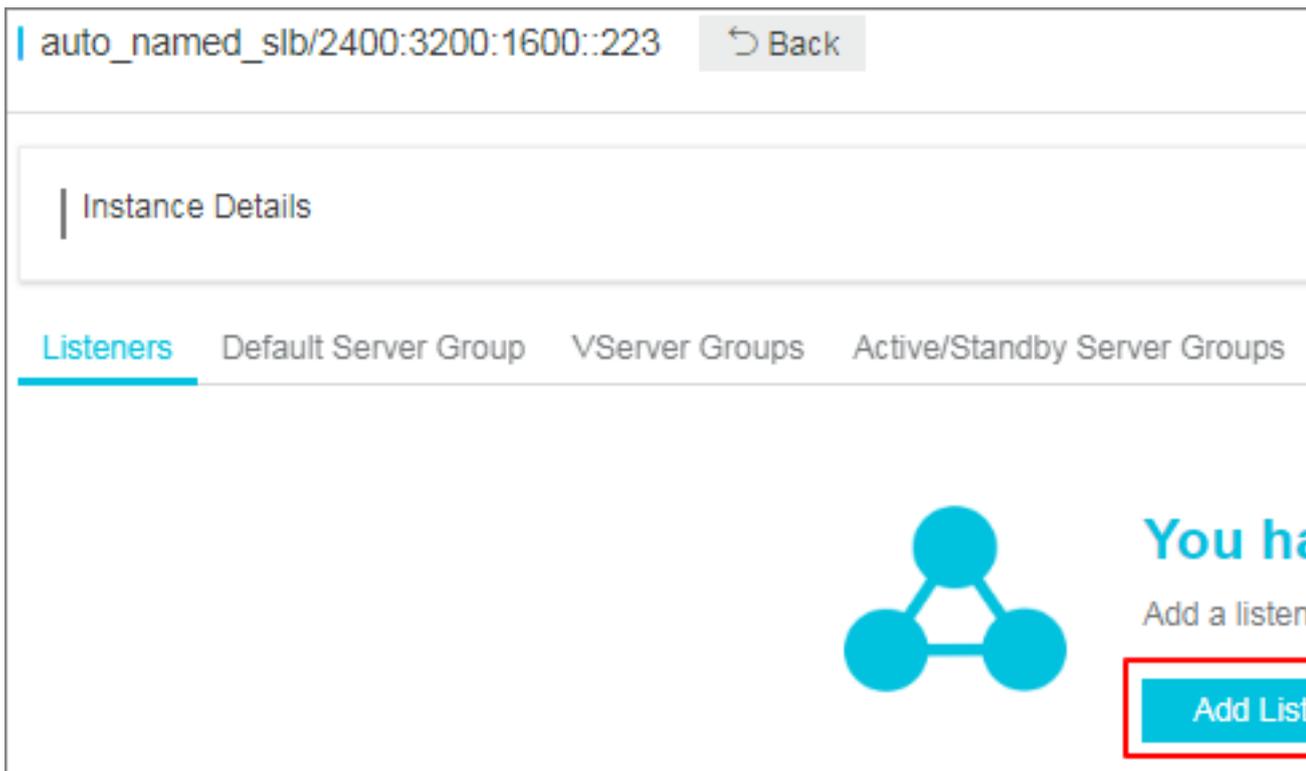
Prerequisites

[Create an SLB instance.](#)

Step 1 Open the listener configuration wizard

To open the listener configuration wizard, complete these steps:

1. Log on to the [SLB console](#).
2. In the left-side navigation pane, select **Instances > Server Load Balancer**.
3. Select a region.
4. Select one of the following methods to open the listener configuration wizard:
 - On the **Server Load Balancer** page, find the target instance and then click **Configure Listener**.
 - On the **Server Load Balancer** page, click the ID of the target SLB instance. On the **Listeners** page, click **Add Listener**.



Step 2 Configure an HTTP listener

To configure a UDP listener, complete these steps:

1. On the **Protocol and Listener** page, configure the HTTP listener according to the following information.

Configuration	Description
Select Listener Protocol	Select the protocol type of the listener. In this tutorial, select HTTP .
Listening Port	The listening port used to receive requests and forward the requests to backend servers. The port number is in the range of 1-65535. <div style="border: 1px solid gray; background-color: #f0f0f0; padding: 5px;">  Note: The listening ports must be unique in a Server Load Balancer instance. </div>
Advanced configurations	
Scheduling Algorithm	Server Load Balancer supports three scheduling algorithms: round robin, weighted round robin (WRR), and weighted least connections (WLC).

Configuration	Description
	<ul style="list-style-type: none"> • Weighted Round-Robin (WRR) (default): Backend servers with higher weights receive more requests than those with smaller weights. • Round-Robin (RR): Requests are evenly and sequentially distributed to the backend servers. • Weighted Least Connections (WLC): A server with a higher weight will receive a larger percentage of live connections at any one time. When the weight value is the same, a backend server with a smaller number of connections is more frequently (and probably) accessed.
<p>Redirection</p>	<p>Select whether to forward traffic of the HTTP listener to an HTTPS listener.</p> <div style="background-color: #f0f0f0; padding: 5px; border: 1px solid #ccc;">  Note: If you enables listener forwarding, make sure that you have created an HTTPS listener. </div>
<p>Session Persistence</p>	<p>Select whether to enable session persistence.</p> <p>If session persistence is enabled, all session requests from the same client are sent to the same backend server.</p> <p>HTTP session persistence is based on cookies. The following two methods are supported:</p> <ul style="list-style-type: none"> • Insert cookie: You only need to specify the cookie timeout period. <p>SLB adds a cookie to the first response from the backend server (insert SERVERID in the HTTP/HTTPS response packet). The next request will contain the cookie and the listener will distribute the request to the same backend server.</p> <ul style="list-style-type: none"> • Rewrite cookie: You can set the cookie to insert to the HTTP/HTTPS response according to your needs. You must maintain the timeout period and lifecycle of the cookie on the backend server. <p>SLB will overwrite the original cookie when it discovers that a new cookie is set. The next time the client carries the new cookie to access SLB, the listener will distribute</p>

Configuration	Description
	<p>the request to the recorded backend server. For more information, see Session persistence.</p>
Enable Access Control	Select whether to enable the access control function.
Access Control Method	<p>Select an access control method after enabling the access control function:</p> <ul style="list-style-type: none"> Whitelist: Only requests from IP addresses or CIDR blocks in the selected access control lists are forwarded. It applies to scenarios where the application only allows access from some specific IP addresses. <p>Enabling whitelist poses some business risks. After a whitelist is configured, only the IP addresses in the list can access the listener. If you enable the whitelist without adding any IP entry in the corresponding access control list, all requests are forwarded.</p> Blacklist: Requests from IP addresses or CIDR blocks in the selected access control lists are not forwarded. It applies to scenarios where the application only denies access from some specific IP addresses. <p>If you enable a blacklist without adding any IP entry in the corresponding access control list, all requests are forwarded.</p>
Access Control List	<p>Select an access control list as the whitelist or the blacklist.</p> <div style="background-color: #f0f0f0; padding: 5px;">  Note: An IPv6 instance can only bind IPv6 access control lists and an IPv4 instance can only bind IPv4 access control lists. For more information, see Configure an access control list. </div>
Enable Peak Bandwidth Limit	<p>Select whether to configure the listening bandwidth. If the SLB instance is billed by bandwidth, you can set different peak bandwidths for different listeners to limit the traffic passing through the listeners. The sum of the peak bandwidths of all listeners under an instance cannot exceed the bandwidth of that instance. By default, all listeners share the bandwidth of the SLB instance.</p>

Configuration	Description
	 Note: Instances billed by traffic have no peak bandwidth limit by default.
Idle Timeout	Specify the idle connection timeout in seconds. Valid value: 1-60 If no request is received during the specified timeout period, Server Load Balancer will close the connection and restart the connection when the next request comes. This function is available in all regions.
Request Timeout	Specify the request timeout in seconds. Valid value: 1-180 If no response is received from the backend server during the specified timeout period, Server Load Balancer will stop waiting and send an HTTP 504 error code to the client. This function is available in all regions.
Enable Gzip Compression	Choose whether to enable Gzip compression to compress files of specific formats. Now Gzip supports the following file types: text/xml, text/plain, text/css, application/javascript, application/x-javascript application/rss+xml, application/atom+xml and application/xml .
Add HTTP Header Fields	Select the custom HTTP headers that you want to add: <ul style="list-style-type: none"> • Use the <code>X-Forwarded-For</code> field to retrieve the client source IP address. • Use the <code>X-Forwarded-Proto</code> field to retrieve the listener protocol used by the SLB instance. • Use the <code>SLB-IP</code> field to retrieve the public IP address of the SLB instance. • Use the <code>SLB-ID</code> field to retrieve the ID of the SLB instance.
Get Client Source IP Address	HTTP listener uses X-Forwarded-For to obtain the real IP of the client.
Automatically Enable Listener After Creation	Choose whether to enable listener after the listener is configured. The listener is enabled by default.

2. Click **Next**.

Configure Server Load Balancer Back

Protocol and Listener Backend Servers Health Check Submit

Select Listener Protocol

TCP UDP **HTTP** HTTPS

Listening Port 22

Advanced Modify

Scheduling Algorithm	Weighted Round-Robin	Session Persistence	Disabled
Access Control	Disabled	Peak Bandwidth	No Limit

Next Cancel

Step 3 Add backend servers

Add backend servers to process requests. You can use the default server group configured for the instance, or configure a VServer group or an active/standby server group for the listener. For more information, see [Backend server overview](#).

In this tutorial, the default server group is used:

1. Select **Default Server Group** and then click **Add**.

Configure Server Load Balancer Back 监听介绍

Protocol and Listener SSL Certificates **Backend Servers** Health Check Submit

Add Backend Servers

① Add backend servers to handle the access requests received by the SLB instance.

Forward Requests To

Default Server Group VServer Group Active/Standby Server

Servers Added

You have not added any servers. Add

Previous Next Cancel

2. Select the ECS instances to add and then click **Add to Selected Server List**. Click **OK**.

The screenshot shows the 'Available Servers' dialog box. At the top, there is a search bar for 'ECS Instance Name' and a 'Buy ECS' button. Below is a table with the following columns: (checkbox), ECS Instance ID/Name, Public/Internal IP Address, Status, Zone, SLB Instance, and Actions. The table contains four rows of ECS instances. The first two rows are selected, indicated by checked checkboxes. A red box highlights these two rows. The bottom row of the table has a red box around the 'Add to Selected Server List' button. At the bottom of the dialog, there are buttons for 'Selected Server List', 'Add to Selected Server List', 'OK', and 'Cancel'.

<input type="checkbox"/>	ECS Instance ID/Name	Public/Internal IP Address	Status	Zone	SLB Instance	Actions
<input checked="" type="checkbox"/>	E-...T i-...5n	4... (Public) 1... 182 (Private)	Running	China East 1 Zone G	Associated SLB Instances 3	Add
<input checked="" type="checkbox"/>	E-...T i-...5n	4... (Public) 1... 183 (Private)	Running	China East 1 Zone G	Associated SLB Instances 3	Add
<input type="checkbox"/>	E-...1 i-...qye	1... 30 (Private)	Running	China East 1 Zone G	Associated SLB Instances 1	Add
<input type="checkbox"/>	E-...2 i-...ba10ghowaye	1... 159 (Private)	Running	China East 1 Zone G	Associated SLB Instances 1	Add

3. Configure the ports and weights of the added backend servers.

- Port

The port opened on the backend server (ECS instance) to receive requests. The port number is in the range of 1-65535. Ports of backend servers can be the same in an SLB instance.

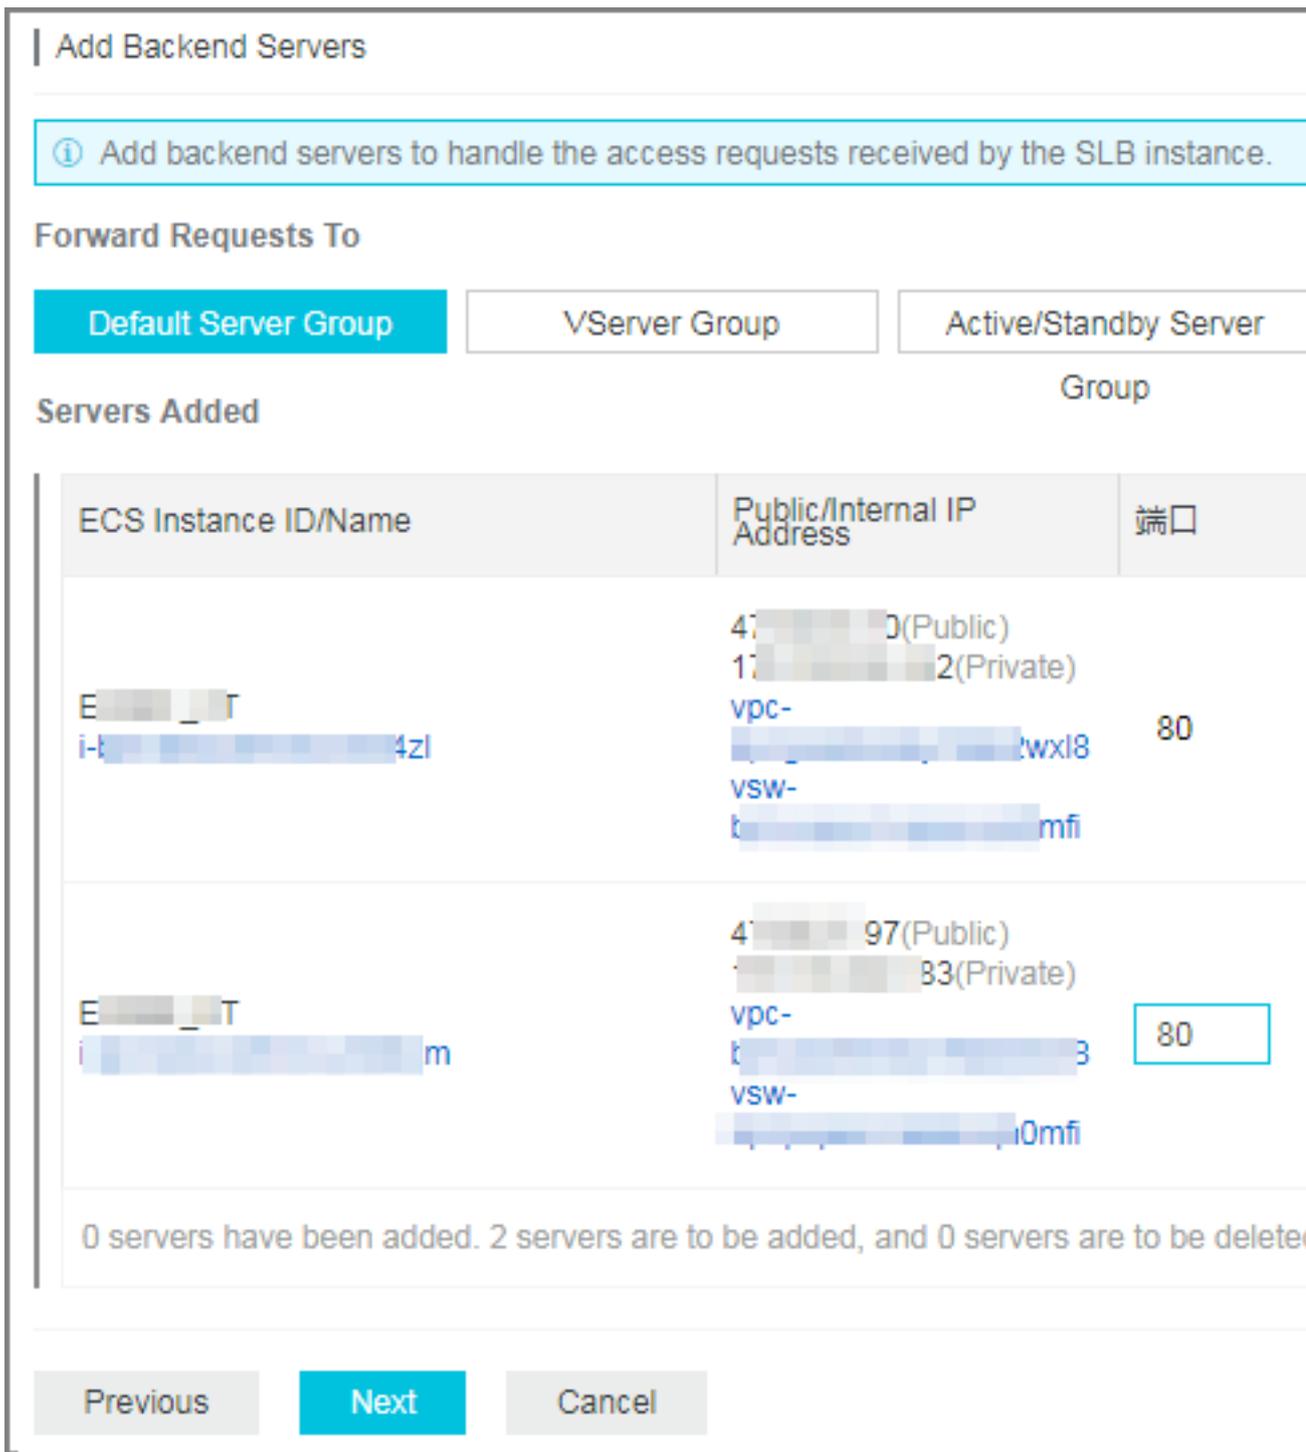
- Weight

The weight of the backend server (ECS instance). An ECS instance with a higher weight will receive a larger number of connection requests.



Note:

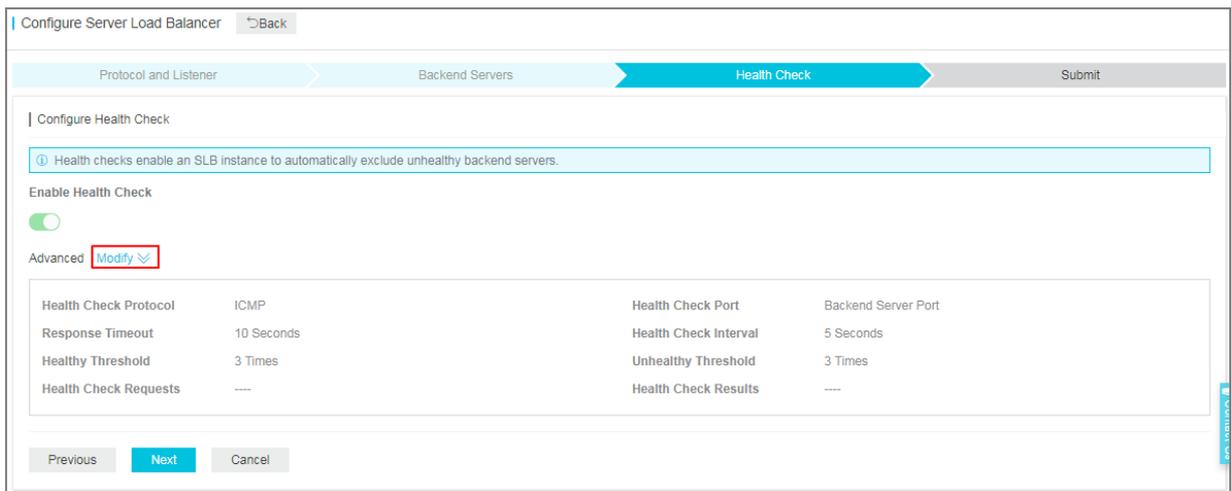
If the weight is set to 0, no requests will be sent to the ECS instance.



4. Click **Next**.

Step 4 Configure health check

Server Load Balancer checks the service availability of the backend servers (ECS instances) by performing health checks. The health check function improves the overall availability of your services and avoids the impact of backend server failures. Click **Modify** to change health check configurations. For more information, see [Configure health check](#).



Step 5 Submit the configurations

To confirm the listener configurations, complete these steps:

1. On the **Submit** page, check listener configurations. You can click **Modify** to change the configurations.
2. Click **Submit**.
3. On the **Submit** page, click **OK** after the configurations are successful.



After the configurations are successful, you can view the created listener on the Listeners page.

Frontend Protocol/Port	Backend Protocol/Port	Name	Health Status	Monitoring	Forwarding Rule	Session Persistence	Peak Bandwidth	Server Group	Access Control List	Actions
UDP:143	UDP:80	udp_143	Abnormal		Weighted Round-Robin	Disabled	No Limit	Default Server Group	Disabled	Configure Details More
TCP:80	TCP:80	tcp_80	Normal		Weighted Round-Robin	Disabled	No Limit	Default Server Group	Disabled	Configure Details More
HTTPS:443	HTTP:80	-	Abnormal		Weighted Round-Robin	Disabled	No Limit	Default Server Group	Disabled	Configure Details Add Forwarding Rules More

Related operations

- [Configure health check](#)
- [Manage a default server group](#)
- [Manage a VServer group](#)
- [Manage an active/standby server group](#)
- [Configure access control](#)
- [Add domain-name based or URL-based forwarding rules](#)
- [Manage a domain name extension](#)

2.4 Add an HTTPS listener

It is applicable to applications requiring encrypted transmission. You can add an HTTPS listener to forward requests from the HTTPS protocol.

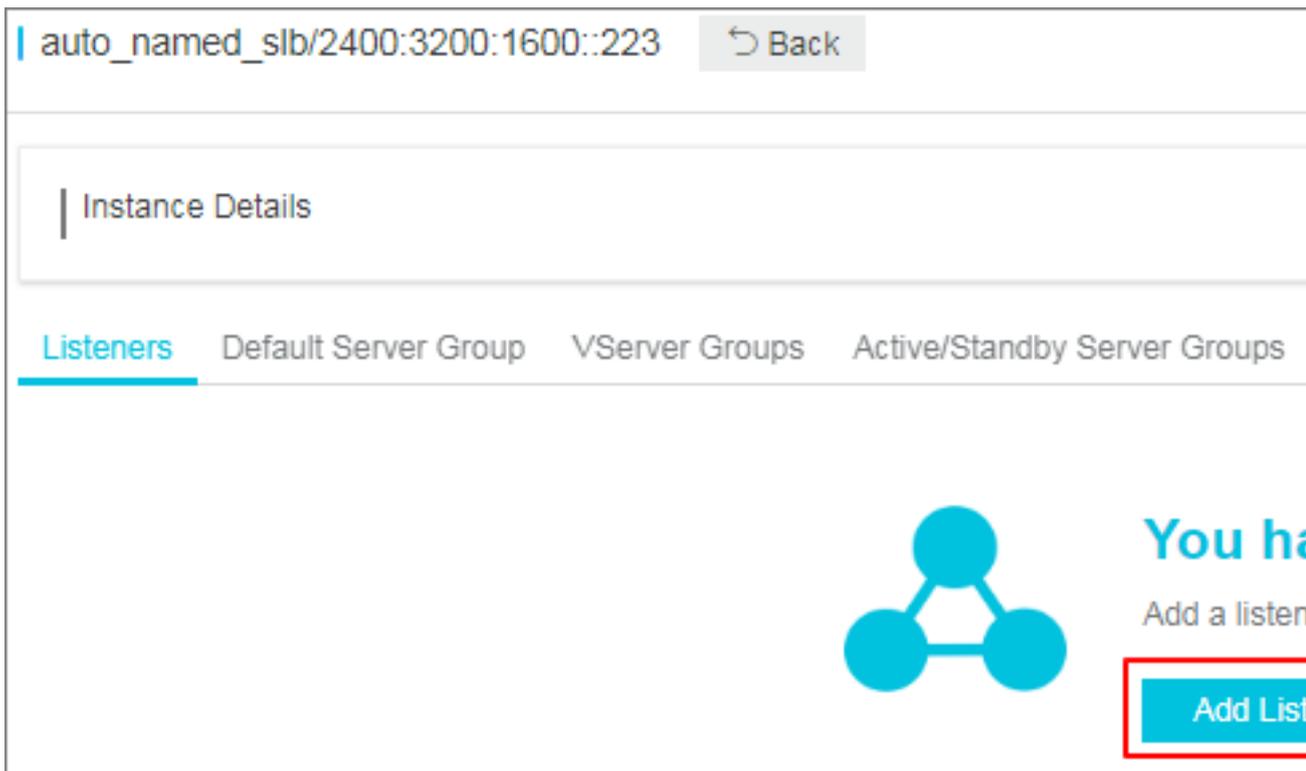
Prerequisites

[Create an SLB instance.](#)

Step 1 Open the listener configuration wizard

To open the listener configuration wizard, complete these steps:

1. Log on to the [SLB console](#).
2. In the left-side navigation pane, select **Instances > Server Load Balancer**.
3. Select a region.
4. Select one of the following methods to open the listener configuration wizard:
 - On the **Server Load Balancer** page, find the target instance and then click **Configure Listener**.
 - On the **Server Load Balancer** page, click the ID of the target SLB instance. On the **Listeners** page, click **Add Listener**.



Step 2 Configure a UDP listener

To configure a TCP listener, complete these steps:

1. On the **Protocol and Listener** page, configure the HTTPS listener according to the following information.

Configuration	Description
Select Listener Protocol	Select the protocol type of the listener. In this tutorial, select HTTPS .
Listening Port	The listening port used to receive requests and forward the requests to backend servers. The port number is in the range of 1-65535. <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 5px;">  Note: The listening ports must be unique in a Server Load Balancer instance. </div>
Advanced configurations	
Scheduling Algorithm	Server Load Balancer supports three scheduling algorithms: round robin, weighted round robin (WRR), and weighted least connections (WLC).

Configuration	Description
	<ul style="list-style-type: none"> • Weighted Round-Robin (WRR) (default): Backend servers with higher weights receive more requests than those with smaller weights. • Round-Robin (RR): Requests are evenly and sequentially distributed to the backend servers. • Weighted Least Connections (WLC): A server with a higher weight will receive a larger percentage of live connections at any one time. When the weight value is the same, a backend server with a smaller number of connections is more frequently (and probably) accessed.
Enable Session Persistence	<p>Select whether to enable session persistence.</p> <p>If session persistence is enabled, all session requests from the same client are sent to the same backend server.</p> <p>HTTP session persistence is based on cookies. The following two methods are supported:</p> <ul style="list-style-type: none"> • Insert cookie: You only need to specify the cookie timeout period. <p>SLB adds a cookie to the first response from the backend server (insert SERVERID in the HTTP/HTTPS response packet). The next request will contain the cookie and the listener will distribute the request to the same backend server.</p> <ul style="list-style-type: none"> • Rewrite cookie: You can set the cookie to insert to the HTTP/HTTPS response according to your needs. You must maintain the timeout period and lifecycle of the cookie on the backend server. <p>SLB will overwrite the original cookie when it discovers that a new cookie is set. The next time the client carries the new cookie to access SLB, the listener will distribute the request to the recorded backend server. For more information, see Session persistence.</p>
Enable HTTP2.0	Select whether to enable HTTP 2.0.
Enable Access Control	Select whether to enable the access control function.
Access Control Method	Select an access control method after enabling the access control function:

Configuration	Description
	<ul style="list-style-type: none"> • Whitelist: Only requests from IP addresses or CIDR blocks in the selected access control lists are forwarded. It applies to scenarios where the application only allows access from some specific IP addresses. Enabling whitelist poses some business risks. After a whitelist is configured, only the IP addresses in the list can access the listener. If you enable the whitelist without adding any IP entry in the corresponding access control list, all requests are forwarded. • Blacklist: Requests from IP addresses or CIDR blocks in the selected access control lists are not forwarded. It applies to scenarios where the application only denies access from some specific IP addresses. If blacklist access is on, but no IP is added to the access policy group, the load balancing listener forwards all requests.
Access Control List	Select an access control list as the whitelist or the blacklist.  Note: An IPv6 instance can only bind IPv6 access control lists and an IPv4 instance can only bind IPv4 access control lists. For more information, see Configure an access control list .
Enable Peak Bandwidth Limit	Select whether to configure the listening bandwidth. 对于按带宽计费的负载均衡实例，您可以针对不同监听设定不同的带宽峰值来限定监听的流量。实例下所有监听的带宽峰值总和不能超过该实例的带宽。 默认不开启，各监听共享实例的总带宽。  Note: Instances billed by traffic have no bandwidth peak limit by default.
Idle Timeout	Specify the idle connection timeout in seconds. Valid value: 1-60 If no request is received during the specified timeout period, Server Load Balancer will close the connection and restart the connection when the next request comes.

Configuration	Description
	This function is available in all regions.
Request Timeout	Specify the request timeout in seconds. Valid value: 1-180 If no response is received from the backend server during the specified timeout period, Server Load Balancer will stop waiting and send an HTTP 504 error code to the client. This function is available in all regions.
TLS Security Policy	Only guaranteed-performance instances support selecting the TLS security policy to use. The TLS security policy contains available TLS protocol versions and supporting encryption algorithm suites. For more information, see Support TLS security policy .
Enable Gzip Compression	Choose whether to enable Gzip compression to compress files of specific formats. Now Gzip supports the following file types: text/xml, text/plain, text/css, application/javascript, application/x-javascript application/rss+xml, application/atom+xml and application/xml .
Add HTTP Header Fields	Select the custom HTTP headers that you want to add: <ul style="list-style-type: none"> • Use the <code>X-Forwarded-For</code> field to retrieve the client source IP address. • Use the <code>X-Forwarded-Proto</code> field to retrieve the listener protocol used by the SLB instance. • Use the <code>SLB-IP</code> field to retrieve the public IP address of the SLB instance. • Use the <code>SLB-ID</code> field to retrieve the ID of the SLB instance.
Get Client Source IP Address	HTTP listener uses X-Forwarded-For to obtain the real IP of the client.
Automatically Enable Listener After Creation	Choose whether to enable listener after the listener is configured. The listener is enabled by default.

Protocol and Listener | SSL Certificates | Backend Servers

Select Listener Protocol

TCP | UDP | HTTP | **HTTPS**

• Listening Port ?

Enter a port number for outbound traffic.

Advanced Hide

• Scheduling Algorithm

Weighted Round-Robin (WRR) | **Weighted Least Connections (WLC)** | Round-Robin (RR)

Enable Session Persistence ?

Enable HTTP/2 ?

Enable Access Control ?

Enable Peak Bandwidth Limit ?

2. Click **Next**.

Configure Server Load Balancer | Back

Protocol and Listener | SSL Certificates | Backend Servers | Health Check | Submit

Select Listener Protocol

TCP | UDP | HTTP | **HTTPS**

• Listening Port ?

443

Advanced Modify

Scheduling Algorithm	Round-Robin	Session Persistence	Disabled
HTTP/2	Enabled	Access Control	Disabled

Next | Cancel

Step 3 Configure the SSL certificate

To add an HTTPS listener, you must upload a server certificate or CA certificate, as shown in the following table.

Certificate	Description	Required for one-way authentication	Required for mutual authentication
Server certificate	Used to identify a server. The client uses it to check whether the certificate sent by the server is issued by a trusted center.	Yes The server certificate must be uploaded to the certificate management system of the Server Load Balancer.	Yes Upload the server certificate to SLB.
Client certificate	Used to identify a client. The client user can prove its true identity when communicating with the server. You can sign a client certificate with a self-signed CA certificate.	No	Yes Install the client certificate on the client.
CA certificate	The server uses the CA certificate to authenticate the signature on the client certificate, as part of the authorization before launching a secure connection. If the authentication fails, the connection will be rejected.	No	Yes The server certificate must be uploaded to the certificate management system of the Server Load Balancer.

Note the following before uploading certificates:

- The uploaded certificate must be in the PEM format. For more information, see [Certificate requirements](#).
- After the certificate is uploaded to SLB, SLB can manage the certificate and you do not need to bind the certificate on backend ECS instances.
- It usually takes one to three minutes to activate the HTTPS listener because the uploading, loading, and validation of certificates take some time. Normally it takes effect in one minute and it will definitely take effect in three minutes.
- The ECDHE algorithm cluster used by HTTPS listeners supports forward secrecy, but does not support uploading security enhancement parameter files required by the DHE algorithm cluster, such as strings containing the `BEGIN DH PARAMETERS` field in the PEM certificate file. For more information, see [Certificate formats](#).
- Currently, Server Load Balancer HTTPS listeners do not support SNI (Server Name Indication). You can use TCP listeners instead, and then configure SNI on the backend ECS instances.

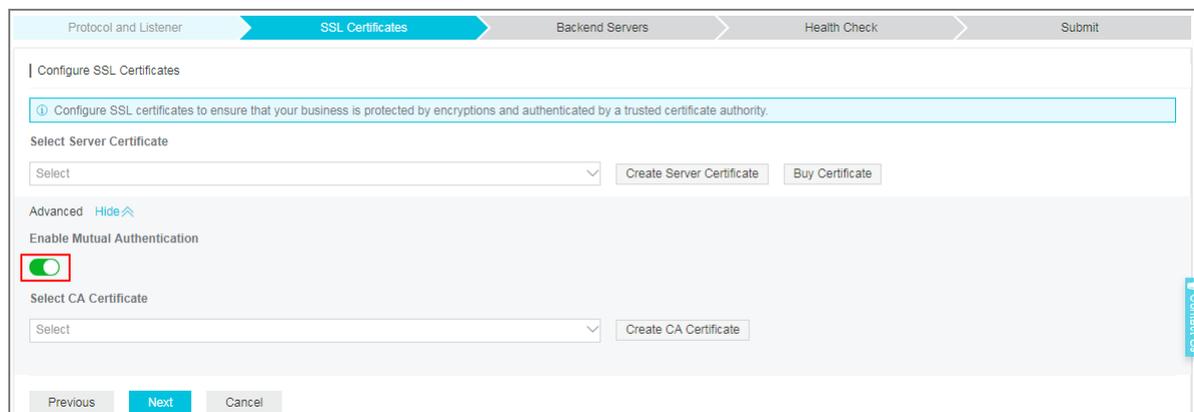
- The session ticket retention time of HTTPS listeners is 300 seconds.
- The actual amount of traffic is larger than the billed traffic amount because some traffic is used for the protocol handshaking.
- In the case of a large number of new connections, HTTPS listeners consume more traffic.

To configure the SSL certificate, complete these steps:

1. Select the server certificate that has been uploaded, or click **Create Server Certificate** to upload a server certificate.

For more information, see [Upload a certificate](#).

2. If you want to enable HTTPS mutual authentication, click **Modify** and enable mutual authentication.



3. Select an uploaded CA certificate, or click **Create CA Certificate** to upload a CA certificate.

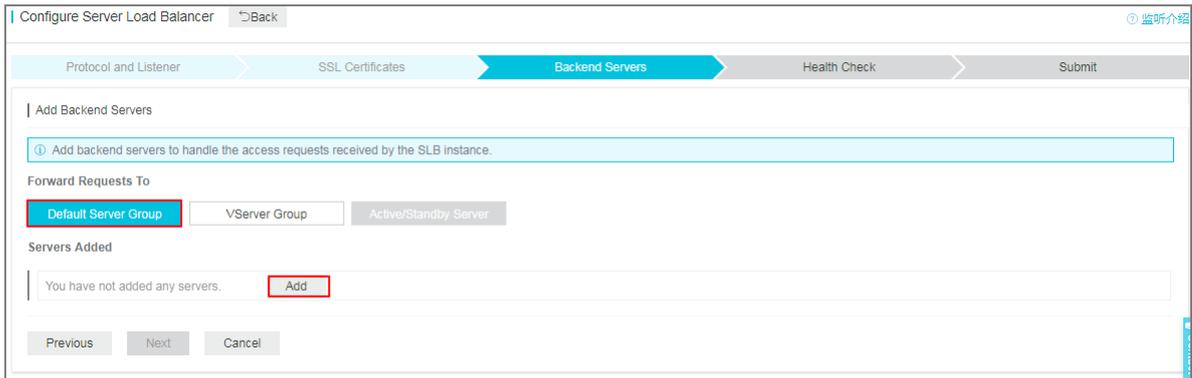
You can use a self-signed CA certificate. For more information, see [Generate a CA certificate](#).

Step 4 Add backend servers

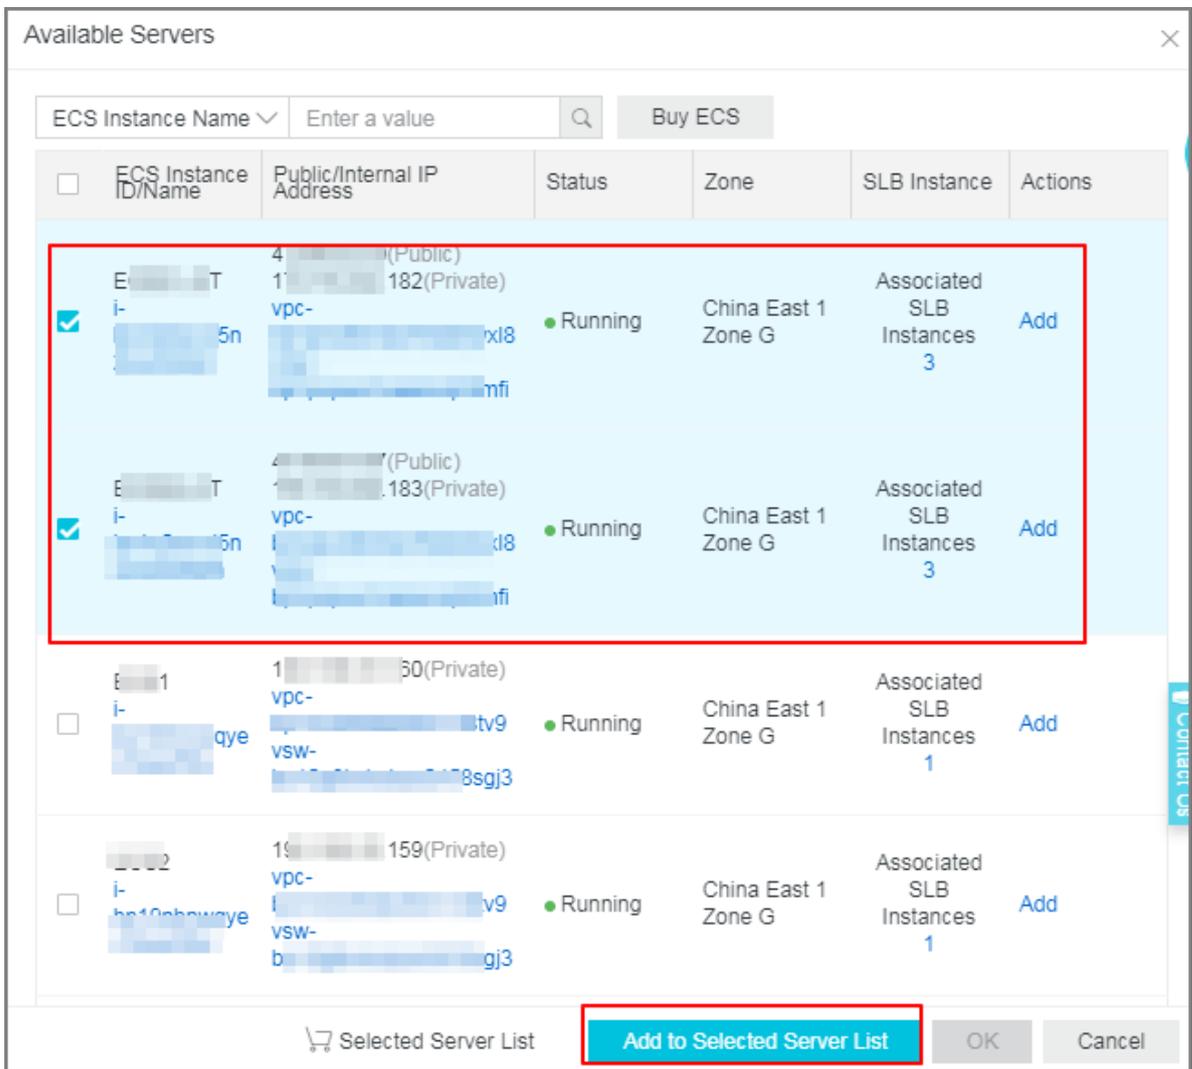
Add backend servers to process requests. You can use the default server group configured for the instance, or configure a VServer group or an active/standby server group for the listener. For more information, see [Backend server overview](#).

In this tutorial, the default server group is used:

1. Select **Default Server Group** and then click **Add**.



2. Select the ECS instances to add and then click **Add to Selected Server List**. Click **OK**.



3. Configure the ports and weights of the added backend servers.

- Port

The port opened on the backend server (ECS instance) to receive requests. The port number is in the range of 1-65535. Ports of backend servers can be the same in an SLB instance.

- Weight

The weight of the backend server (ECS instance). An ECS instance with a higher weight will receive a larger number of connection requests.



Note:

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

Add Backend Servers

① Add backend servers to handle the access requests received by the SLB instance.

Forward Requests To

Default Server Group VServer Group Active/Standby Server

Servers Added

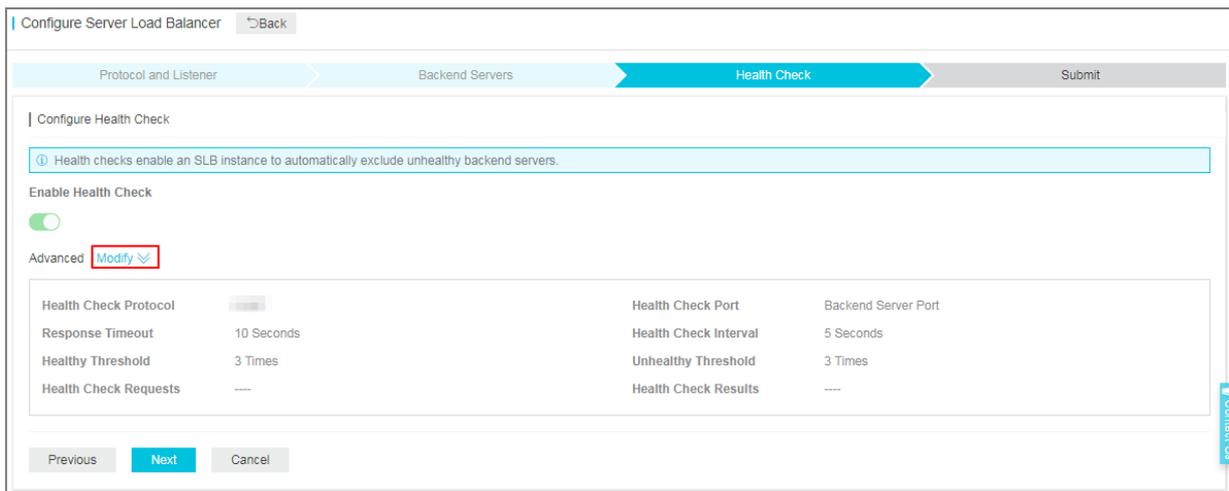
ECS Instance ID/Name	Public/Internal IP Address	端口	Weight	Actions
ECS-4zi	47...0(Public) 12...2(Private) vdc-...wxj8 vsw-...mfi	80	<input type="text" value="100"/>	Delete
ECS-m	47...97(Public) ...33(Private) vdc-...3 vsw-...0mfi	<input type="text" value="80"/>	<input type="text" value="100"/>	Delete

0 servers have been added. 2 servers are to be added, and 0 servers are to be deleted.

4. Click **Next**.

Step 5 Configure health check

Server Load Balancer checks the service availability of the backend servers (ECS instances) by performing health checks. The health check function improves the overall availability of your services and avoids the impact of backend server failures. Click **Modify** to change health check configurations. For more information, see [Configure health check](#).



Step 6 Submit the configurations

To confirm the listener configurations, complete these steps:

1. On the **Submit** page, check listener configurations. You can click **Modify** to change the configurations. Click **Submit**.
2. On the **Submit** page, click **OK** after the configurations are successful.



After the configurations are successful, you can view the created listener on the listeners page.

Frontend Protocol/Port	Backend Protocol/Port	Name	Health Status	Monitoring	Forwarding Rule	Session Persistence	Peak Bandwidth	Server Group	Access Control List	Actions
HTTPS:143	HTTP:80	https_143	● Abnormal		Weighted Round-Robin	Disabled	Share Instance Bandwidth	Default Server Group	Disabled	Configure Details Add Forwarding Rules More

Related operations

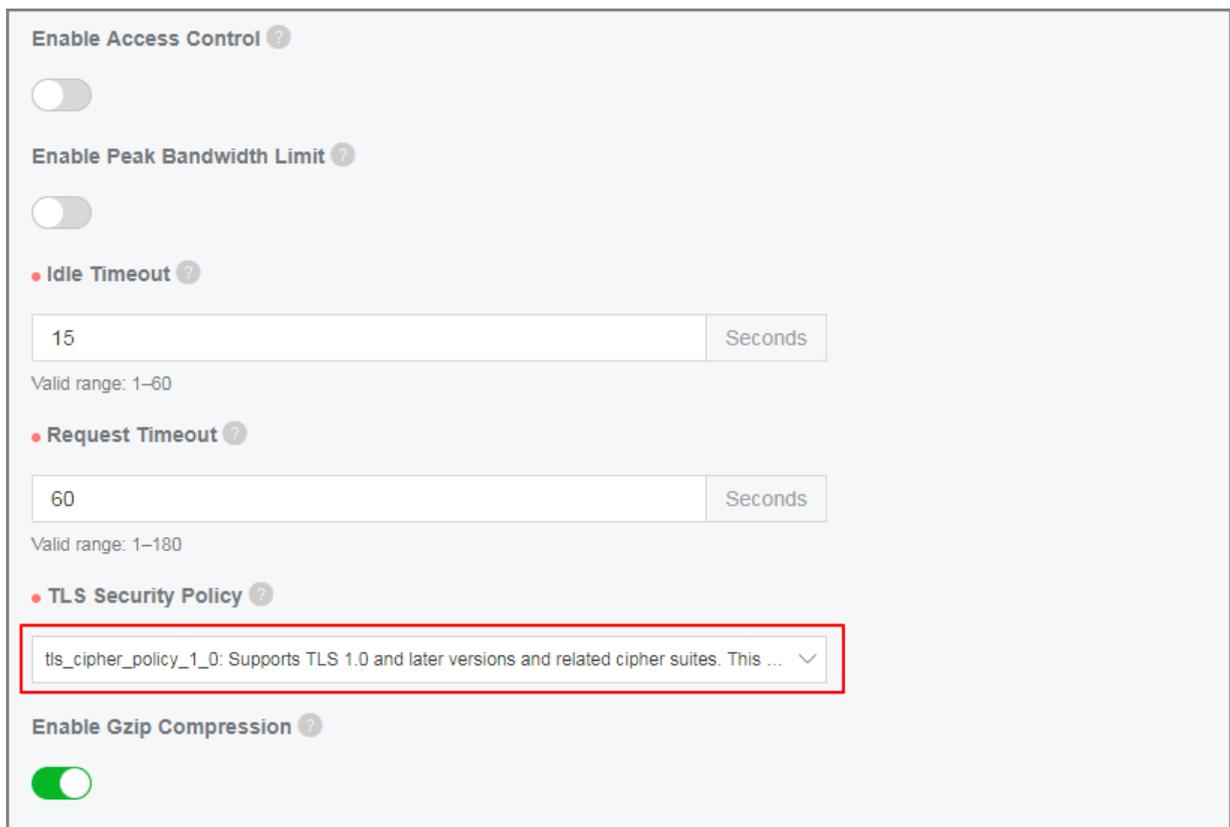
- [Configure health check](#)
- [Manage a default server group](#)
- [Manage a VServer group](#)
- [Manage an active/standby server group](#)

- [Generate a CA certificate](#)
- [Upload a certificate](#)
- [Configure access control](#)
- [Add domain-name based or URL-based forwarding rules](#)
- [Manage a domain name extension](#)

2.5 Support TLS security policy

Guaranteed-performance instances support selecting the TLS security policy to use when you create or configure an HTTP listener.

You can select the TLS security policy when you set advanced configurations of **Protocol and Listener** during adding or configuring an HTTPS listener. For more information, see [Add an HTTPS listener](#).



The screenshot displays a configuration panel for a listener. It includes several sections:

- Enable Access Control**: A toggle switch that is currently turned off.
- Enable Peak Bandwidth Limit**: A toggle switch that is currently turned off.
- Idle Timeout**: A text input field containing the value '15' and a 'Seconds' label. Below the input, it states 'Valid range: 1-60'.
- Request Timeout**: A text input field containing the value '60' and a 'Seconds' label. Below the input, it states 'Valid range: 1-180'.
- TLS Security Policy**: A dropdown menu with a red border. The selected option is 'tls_cipher_policy_1_0: Supports TLS 1.0 and later versions and related cipher suites. This ...'. A question mark icon is visible to the right of the label.
- Enable Gzip Compression**: A toggle switch that is currently turned on.

The TLS security policy contains available TLS protocol versions and supporting encryption algorithm suites.

TLS security policy

Security policy	Features	Supported TLS versions	Supported encryption algorithm suites
tls_cipher_policy_1_	Best compatibility and low security	TLSv1.0, TLSv1.1, and TLSv1.2	Supported encryption algorithm suites: ECDHE-RSA-AES128-GCM-SHA256, ECDHE-RSA-AES256-GCM-SHA384, ECDHE-RSA-AES128-SHA256, ECDHE-RSA-AES256-SHA384, AES128-GCM-SHA256, AES256-GCM-SHA384, AES128-SHA256, AES256-SHA256, ECDHE-RSA-AES128-SHA, ECDHE-RSA-AES256-SHA, AES128-SHA, AES256-SHA and DES-CBC3-SHA.
tls_cipher_policy_1_1	Good compatibility and security	TLSv1.1 and TLSv1.2	ECDHE-RSA-AES128-GCM-SHA256, ECDHE-RSA-AES256-GCM-SHA384, ECDHE-RSA-AES128-SHA256, ECDHE-RSA-AES256-SHA384, AES128-GCM-SHA256, AES256-GCM-SHA384, AES128-SHA256, AES256-SHA256, ECDHE-RSA-AES128-SHA, ECDHE-RSA-AES256-SHA, AES128-SHA, AES256-SHA and DES-CBC3-SHA.
tls_cipher_policy_1_2	Good compatibility and high security	Tlsv1.2	ECDHE-RSA-AES128-GCM-SHA256, ECDHE-RSA-AES256-GCM-SHA384, ECDHE-RSA-AES128-SHA256, ECDHE-RSA-AES256-SHA384, AES128-GCM-SHA256, AES256-GCM-SHA384, AES128-SHA256, AES256-SHA256, ECDHE-RSA-AES128-SHA, ECDHE-RSA-AES256-SHA, AES128-SHA, AES256-SHA and DES-CBC3-SHA.
tls_cipher_policy_1_2_strict	Only support forward security, extremely high security	TLSv1.2	ECDHE-RSA-AES128-GCM-SHA256, ECDHE-RSA-AES256-GCM-SHA384, ECDHE-RSA-AES128-SHA256, ECDHE-RSA-AES256-SHA384, ECDHE-RSA-AES128-SHA and ECDHE-RSA-AES256-SHA.

Differences along TLS security policies

Security policy	tls_cipher_policy_1_0	tls_cipher_policy_1_1	tls_cipher_policy_1_2	tls_cipher_policy_1_2_strict
TLS	1.2/1.1/1.0	1.2/1.1	1.2	1.2

Security policy		tls_cipher _policy_1_ 0	tls_cipher _policy_1_ 1	tls_cipher _policy_1_ 2	tls_cipher _policy_1_ 2_strict
CIPHER	ECDHE-RSA-AES128-GCM-SHA256	✓	✓	✓	✓
	ECDHE-RSA-AES256-GCM-SHA384	✓	✓	✓	✓
	ECDHE-RSA-AES128-SHA256	✓	✓	✓	✓
	ECDHE-RSA-AES256-SHA384	✓	✓	✓	✓
	AES128-GCM-SHA256	✓	✓	✓	
	AES256-GCM-SHA384	✓	✓	✓	
	AES128-SHA256	✓	✓	✓	
	AES256-SHA256	✓	✓	✓	
	ECDHE-RSA-AES128-SHA	✓	✓	✓	✓
	ECDHE-RSA-AES256-SHA	✓	✓	✓	✓
	AES128-SHA	✓	✓	✓	
	AES256-SHA	✓	✓	✓	
	DES-CBC3-SHA	✓	✓	✓	

2.6 Manage a domain name extension

HTTPS listeners of guaranteed-performance SLB support configuring multiple certificates, allowing you to forward requests from different domains to different backend servers.

Introduction to SNI

Server Name Indication (SNI) is an extension to the SSL/TLS protocol, allowing a server to install multiple certificates on the same IP address. Only guaranteed-performance SLB instances support SNI. When a client accesses SLB, the certificate configured for the domain name is used by default. If no certificate is configured for the domain name, the certificate configured for the HTTPS listener is used.

If you want to resolve multiple domain names to the IP address of an SLB instance, and distribute requests from different domains to different backend servers, use the domain name extension function.

The domain name extension function is available in all regions.

Add a domain name extension

1. Log on to the [SLB console](#).
2. Select a region.
3. Click the ID of the SLB instance.
4. In the left-side navigation pane, click **Listeners**.
5. On the **Listeners** page, find the created HTTPS listener, and then click **More > Additional Domains**.

Frontend Protocol/Port	Backend Protocol/Port	Name	Health Status	Monitoring	Forwarding Rule	Session Persistence	Peak Bandwidth	Server Group	Access Control List	Actions
UDP:143	UDP:80	udp_143	Abnormal		Weighted Round-Robin	Disabled	No Limit	Default Server Group	Disabled	Configure Details More
TCP:80	TCP:80	tcp_80	Normal		Weighted Round-Robin	Disabled	No Limit	Default Server Group	Disabled	Configure Details More
HTTPS:443	HTTP:80	-	Abnormal		Weighted Round-Robin	Disabled	No Limit	Default Server Group	Disabled	Configure Details Add Forwarding Rules More Start Stop Remove Set Access Control Additional Domains

6. Click **Add Additional Domain** and configure the domain name:
 - a. Enter the domain name. The domain name can only contain letters, numbers, dashes, or dots.

Domain name forwarding rules support exact match and wildcard match.

- Exact domain name: www.aliyun.com
- Wildcard domain name (generic domain name): *.aliyun.com, *.market.aliyun.com

When a request matches multiple forwarding rules, exact match takes precedence over small-scale wildcard match and small-scale wildcard match takes precedence over large-scale wildcard match, as shown in the following table.

Type	Request URL	Domain name-based forwarding rule		
		www.aliyun.com	*.aliyun.com	*.market.aliyun.com
Exact match	www.aliyun.com	✓	×	×
Wildcard match	Market.aliyun.com	×	✓	×
Wildcard match	info.market.aliyun.com	×	×	✓

b. Select the certificate associated with the domain name.

 **Note:**
The domain name in the certificate must be the same as the added domain name extension.

c. Click **OK**.



Additional Domains

Add Domain Extension

*.example2.com example2

OK Cancel

Domain Extensions

Domain Name	Certificate Name (Domain Name)	Actions
www.example.com()		

7. On the **Listeners** page, find the created HTTPS listener and click **Add Forwarding Rules**.

8. On the **Forwarding Rules** page, click **Add Forwarding Rules**.

9. For more information, see [Add domain-name based or URL-based forwarding rules](#).

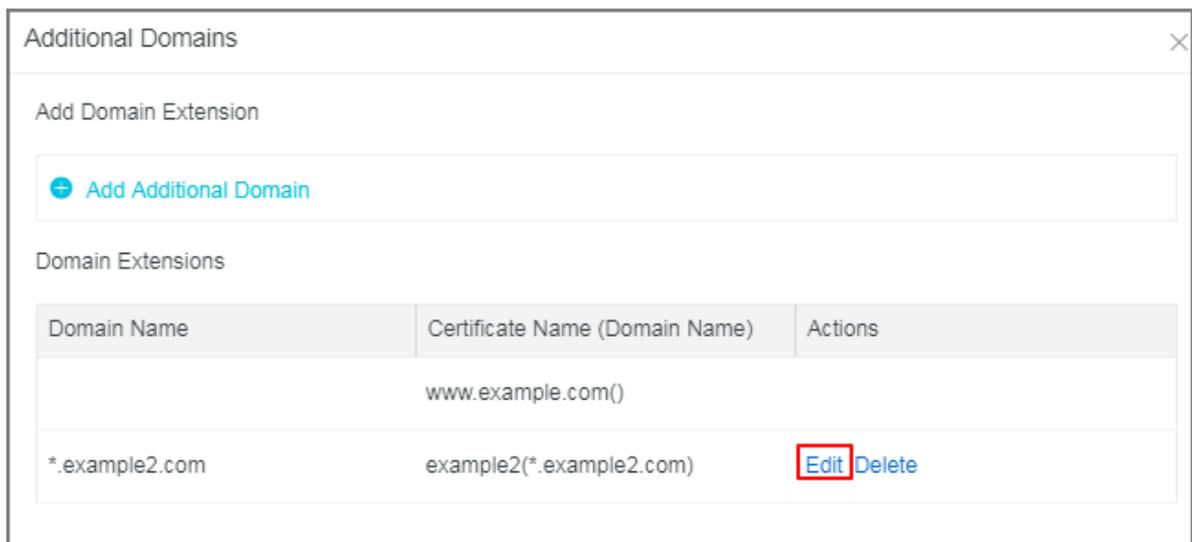
 **Note:**
Make sure that the domain name configured in the forwarding rule is the same as the added domain name extension.

Edit a domain name extension

You can replace the certificate used by an added domain name extension.

To edit a domain name extension, complete these steps:

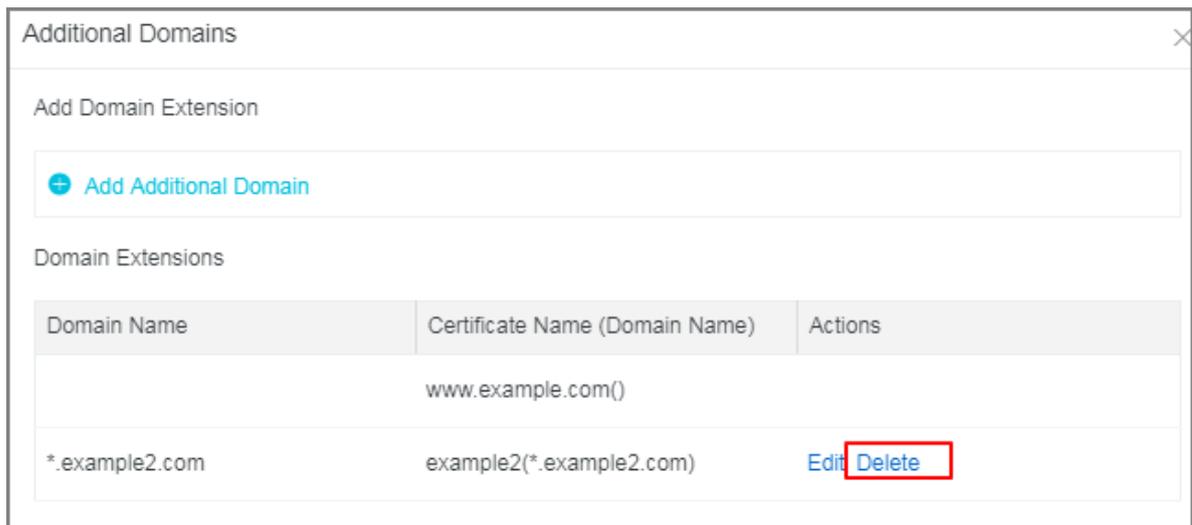
1. Log on to the [SLB console](#).
2. Select a region and all SLB instance in this region are displayed.
3. Click the ID of the SLB instance.
4. In the left-side navigation pane, click **Listeners**.
5. On the **Listeners** page, find the created HTTPS listener, and then click **More > Additional Domains**.
6. Find the target domain name extension and then click **Edit**.
7. In the **Edit Additional Domain** dialog box, select a new certificate and then click **OK**.



Delete a domain name extension

To delete a domain name extension, complete these steps:

1. Log on to the [SLB console](#).
2. Select a region and all SLB instance in this region are displayed.
3. Click the ID of the SLB instance.
4. In the left-side navigation pane, click **Listeners**.
5. On the **Listeners** page, find the created HTTPS listener, and then click **More > Additional Domains**.
6. Find the target domain name extension and then click **Remove**.



7. In the displayed dialog box, click **OK**.

3 Health check

3.1 Health check overview

Server Load Balancer checks the service availability of the backend servers (ECS instances) by performing health checks. The health check function improves the overall availability of your services and avoids the impact of backend server failures.

After enabling the health check function, SLB stops distributing requests to the instance that is discovered as unhealthy and restarts forwarding requests to the instance only when it is declared healthy.

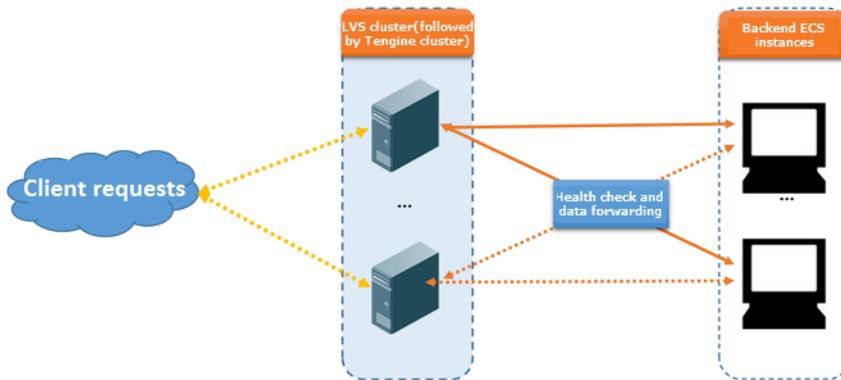
If your business is highly sensitive to traffic load, frequent health checks may impact normal service. You can reduce this impact by reducing the frequency of health checks, increasing the health check interval, or changing the HTTP health check to TCP health check. To guarantee the service availability, we do not recommend removing all health checks.

Health check process

Server Load Balancer is deployed in clusters. Data forwarding and health checks are handled at the same time by the node servers in the LVS cluster and Tengine cluster.

The node servers in the cluster independently perform health checks in parallel, according to the health check configuration. If a node server discovers an ECS instance is unhealthy, the node server will stop distributing requests to the ECS instance. This operation is synchronized through all node servers.

The IP address range used to perform the health check is 100.64.0.0/10. The backend servers cannot block this CIDR block. You do not need to additionally configure a security group rule to allow access from this CIDR block. However, if you have configured security rules such as iptables, allow access from this CIDR block (100.64.0.0/10 is reserved by Alibaba Cloud, and other users cannot use any IP address in this CIDR block, so there is no security risk).



Health check of HTTP/HTTPS listeners

For Layer-7 (HTTP or HTTPS) listeners, SLB detects the status of backend servers by sending HTTP HEAD requests, as shown in the following figure.

For HTTPS listeners, certificates are managed in SLB. Data exchange (including health check data and service interaction data) between SLB and backend ECS instances is not transmitted over HTTPS to improve the system performance.

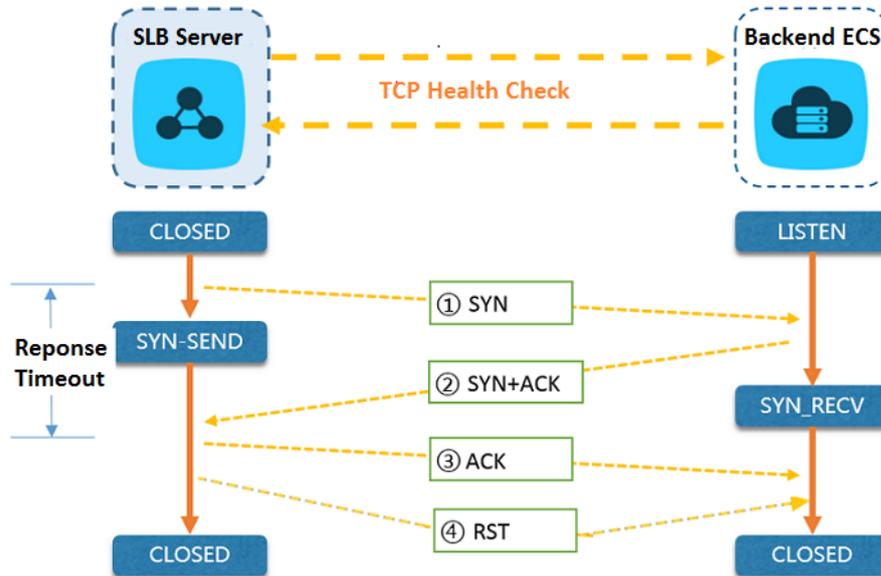


The health check process of a Layer-7 listener is as follows:

1. The Tengine node server sends an HTTP HEAD request to the backend servers at [ECS internal IP:health check port:health check URL] according to the health check settings.
2. After receiving the request, the backend server returns an HTTP status code based on the running status.
3. If the Tengine node server does not receive the response from the backend server within the configured response timeout period, then the ECS instance is declared unhealthy.
4. If the Tengine node server receives the response from the backend ECS instance within the configured response timeout period, then it compares the returned status code with the status code specified in the listener configuration. If the status code is the same, the backend server is declared healthy. Otherwise, the backend server is declared unhealthy.

Health check of TCP listeners

For TCP listeners, SLB detects the status of backend servers by sending TCP detections, as the following figure shows.



The health check process of a TCP listener is as follows:

1. The LVS node server sends a TCP SYN packet to the intranet IP address of the backend ECS instance with the configured health check port.
2. After receiving the request, the backend server returns a TCP SYN and ACK packet if the corresponding port is listening normally.
3. If the LVS node server does not receive the required data packet from the backend server within the configured response timeout period, the ECS instance is declared unhealthy. Then, the LVS node server sends an RST data packet to the backend server to terminate the TCP connection.
4. If the LVS node server receives the data packet from the backend ECS instance within the configured response timeout period, the ECS instance is declared healthy. Then, the LVS node server sends an RST data packet to the backend server to terminate the TCP connection.



Note:

In general, TCP three-way handshakes are conducted to establish a TCP connection. After the LVS node server receives an SYN + ACK data packet from the backend ECS instance, the LVS node server sends an ACK data packet, and then immediately sends an RST data packet to terminate the TCP connection.

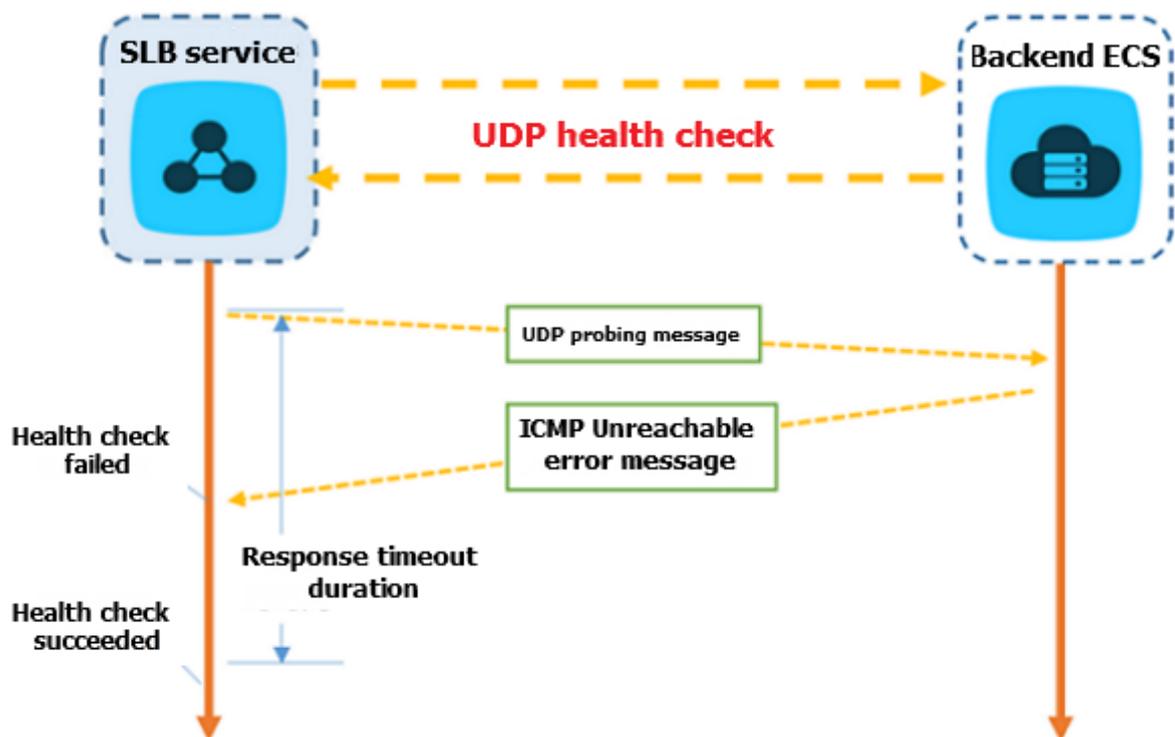
This process may make backend server to think an error occurred in the TCP connection, such as an abnormal exit, and then throw a corresponding error message, such as `Connection reset by peer`.

Resolution:

- Use the HTTP health check.
- If obtaining real IP is enabled, ignore the connection errors caused by the health check.

Health check of UDP listeners

For UDP listeners, Server Load Balancer detects the status of the backend servers through UDP packet detection, as shown in the following figure.



The health check process of a UDP listener is as follows:

1. The LVS node server sends a UDP packet to the intranet IP address of the backend ECS instance with the configured health check port.
2. If the corresponding port of the ECS instance is not listening normally, the system will return an ICMP error message, such as `port XX unreachable`. Otherwise, no message is sent.
3. If the LVS node server receives the ICMP error message within the configured response timeout period, the ECS instance is declared unhealthy.

4. If the LVS node server does not receive any messages within the configured response timeout period, the ECS instance is declared healthy.

**Note:**

For UDP health checks, the real status of the backend server and the health check result may not be the same in the following situation:

If the ECS instance uses a Linux operating system, the speed of sending ICMP messages in high-concurrency scenarios is limited due to the anti-ICMP attack protection in Linux. In this case, even if an exception occurs in the ECS instance, SLB may declare the backend server healthy because the error message `port XX unreachable` is not returned. As a result, the actual service status is different from the health check result.

Resolution:

Set a pair of custom request and response for the UDP health check. If the custom response is returned, the ECS instance is considered healthy. Otherwise, the ECS instance is considered unhealthy. To achieve this, you must add corresponding configurations for the client.

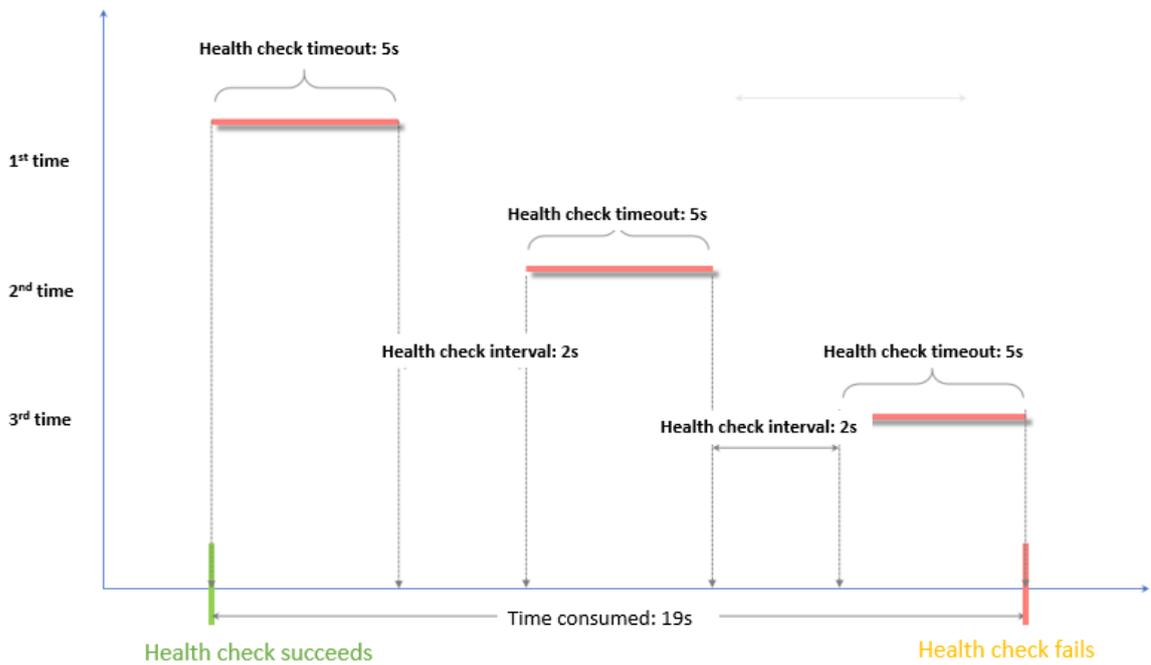
Health check time window

The health check function has effectively improved the availability of your business services . However, to reduce impact on the system availability caused by frequent system switches because of health check failure, SLB declares an ECS instance healthy or unhealthy only after successive successes or failures within a specified timeframe. The health check time window is determined by the following three factors:

- Health check interval (How often the health check is performed.)
- Response timeout (The amount of time to wait for the response.)
- Health check threshold (The number of consecutive successful or failed health checks.)

The health check time window is calculated as follows:

- Health check failure time window = Response Timeout x Unhealthy Threshold + Health Check Interval X (Unhealthy Threshold -1)

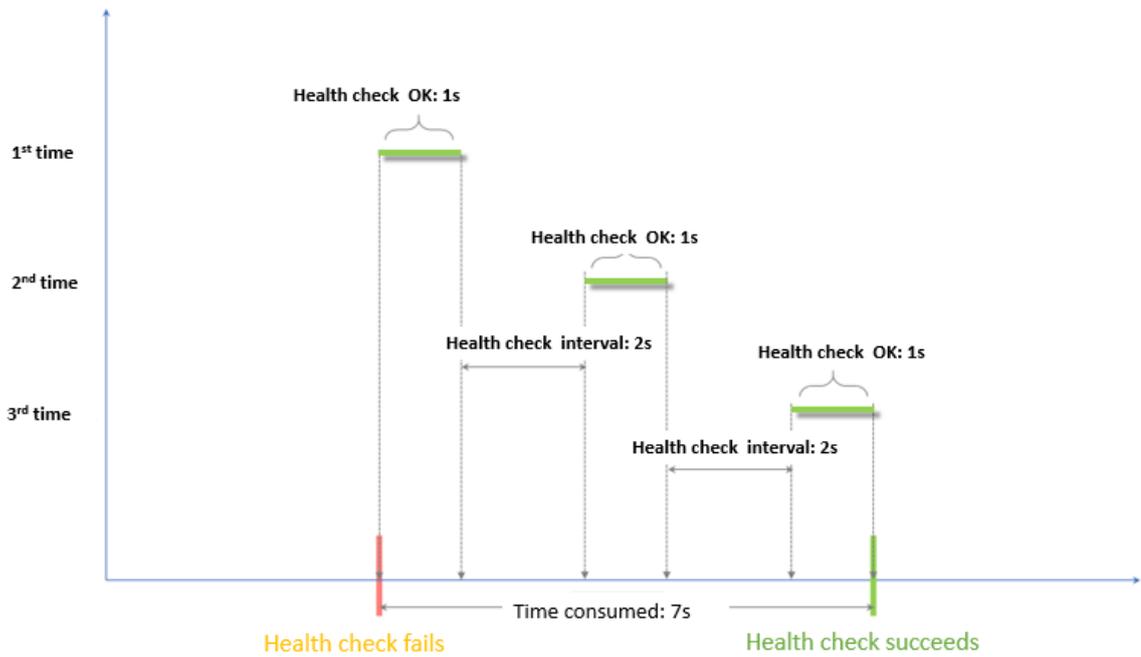


- Health check success time window = (Response Time of a Successful Health Check X Health Threshold) + Health Check Interval X (Health Threshold-1)



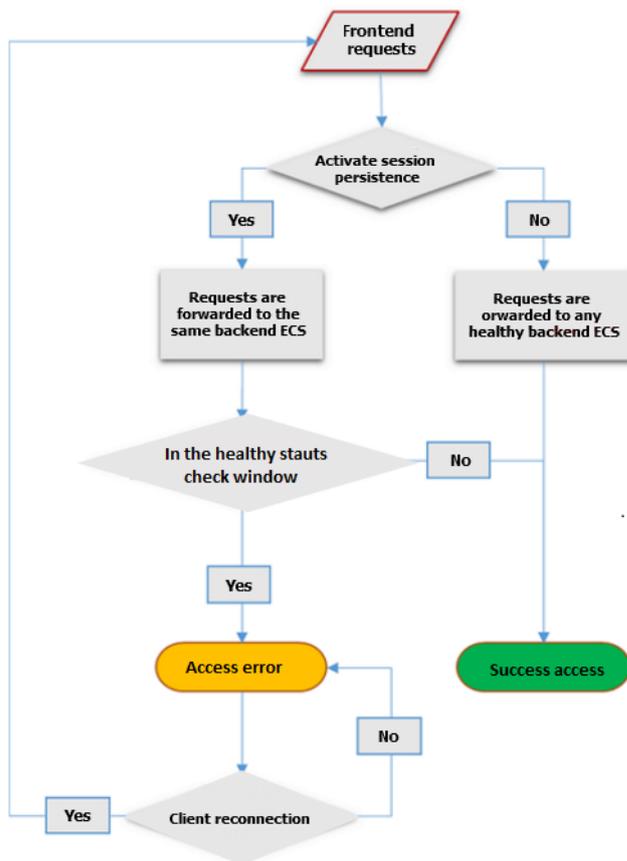
Note:

The success response time of a health check is the time from when the health check request is sent to the time when a response is received. When using TCP for health check, due to the detection of only the port for survival, therefore, the time is very short and can be ignored almost. For TCP health check, the time is very short and almost negligible because TCP health check only the detects whether the port is alive.



The health check result has the following impact on the requests forwarding:

- If the health check of the target ECS instance fails, the request will not be distributed to the ECS instance. Therefore, there is no impact on the client access.
- If the health check of the target ECS instance succeeds, the request will be distributed to it. The client access is normal.
- If a request arrives during a health check failure window, the request is still sent to the ECS instance because the ECS instance is being checked and has not been declared unhealthy. As a result, the client access fails.



3.2 Configure health check

You can configure health check settings when adding a listener. In general, the default settings can meet your requirements.

Configure health check

You can configure the health check of a listener on the console or through API. For more information, see [Health check overview](#) and [Health check FAQ](#).

To configure the health check, complete these steps:

1. Log on to the [SLB console](#).
2. Select a region and all SLB instance in this region are displayed.
3. Click the ID of the target SLB instance.
4. On the **Instance Details** page, click the **Listeners** tab.
5. Click **Add Listener** or the **Configure** option of the target listener.
6. On the **Health Check** page, configure the health check.

We recommend that you use the default values when configuring the health check.

Table 3-1: Health check configurations

Configuration	Description
Health Check Protocol	<p>For TCP listeners, both the TCP health check and HTTP health check are supported.</p> <ul style="list-style-type: none"> TCP health check is based on network layer detection. HTTP health check is performed by sending head requests.
Health Check Path and Domain Name (HTTP health check only)	<p>By default, SLB sends an HTTP head request to the default homepage configured on the application server through the intranet IP address of the backend ECS instance to do the health check. If you do not use the default homepage of the application server to do health check, you must specify the URL for health check. Some application servers verify the host field in the request, therefore, the request header must contain the host field. If so, the health check request will be denied by the server and the health check may fail. Therefore, if your application server verifies the host field in the request, you must configure a domain name to make sure the health check works.</p>
Normal Status Code (HTTP health check only)	<p>Select the HTTP status code indicating that the health check is normal.</p> <p>The default values are http_2xx and http_3xx.</p>
Health Check Port	<p>The detection port used by the health check to access the backend ECS instances.</p> <p>By default, the backend port configured in the listener is used.</p> <div style="background-color: #f0f0f0; padding: 5px; margin-top: 10px;">  Note: If a VServer group or an active/standby server group is configured for the listener, and the ECS instances in the group use different ports, leave this option empty. SLB uses the backend port of each ECS instance to do health check. </div>
Response Timeout	<p>The amount of time in seconds to wait for the response from a health check. If the backend ECS instance does not respond correctly within the specified time, the health check fails.</p> <p>The timeout range is 1-300 seconds. The default value is 10 seconds for UDP listeners and 5 seconds for HTTP/HTTPS/TCP listeners.</p>
Health Check Interval	<p>The time interval between two consecutive health checks.</p> <p>All node servers in the LVS cluster independently and concurrently perform health check on backend ECS instances according to the interval. The statistics from a health check request on a single ECS</p>

Configuration	Description
	instance cannot reflect the health check interval because the health check time of each node server is not synchronized. The time range is 1-50 seconds. The default value is 5 seconds for UDP listeners, and 2 seconds for HTTP/HTTPS/TCP listeners.
Unhealthy Threshold	The number of consecutive failures of health check performed by the same LVS node server on the same ECS instance (from success to failure). Valid value: 2-10. The default value is 3.
Healthy Threshold	The number of consecutive successes of health check performed by the same LVS mode server on the same ECS instance (from failure to success). Valid value: 2-10. The default value is 3.
Health Check Requests and Results	When configuring health check for UDP listeners, you can enter the request contents (such as youraccountID) in Health Check Request and the expected response (such as slb123) in Health Check Response . Add the corresponding health check response logic to the application logic of the backend server. For example, return slb123 when youraccountID is received. If SLB receives the expected response from the backend server, the health check succeeds. Otherwise, the health check fails. This method can maximally guarantee the reliability of health check.

Configure Server Load Balancer Back

Protocol and Listener
Backend Servers
Health Check
Submit

Configure Health Check

i Health checks enable an SLB instance to automatically exclude unhealthy backend servers.

Enable Health Check

Advanced Modify v

Health Check Protocol	TCP	Health Check Port	Backend Server Port
Response Timeout	5 Seconds	Health Check Interval	2 Seconds
Healthy Threshold	3 Times	Unhealthy Threshold	3 Times

Previous
Next
Cancel

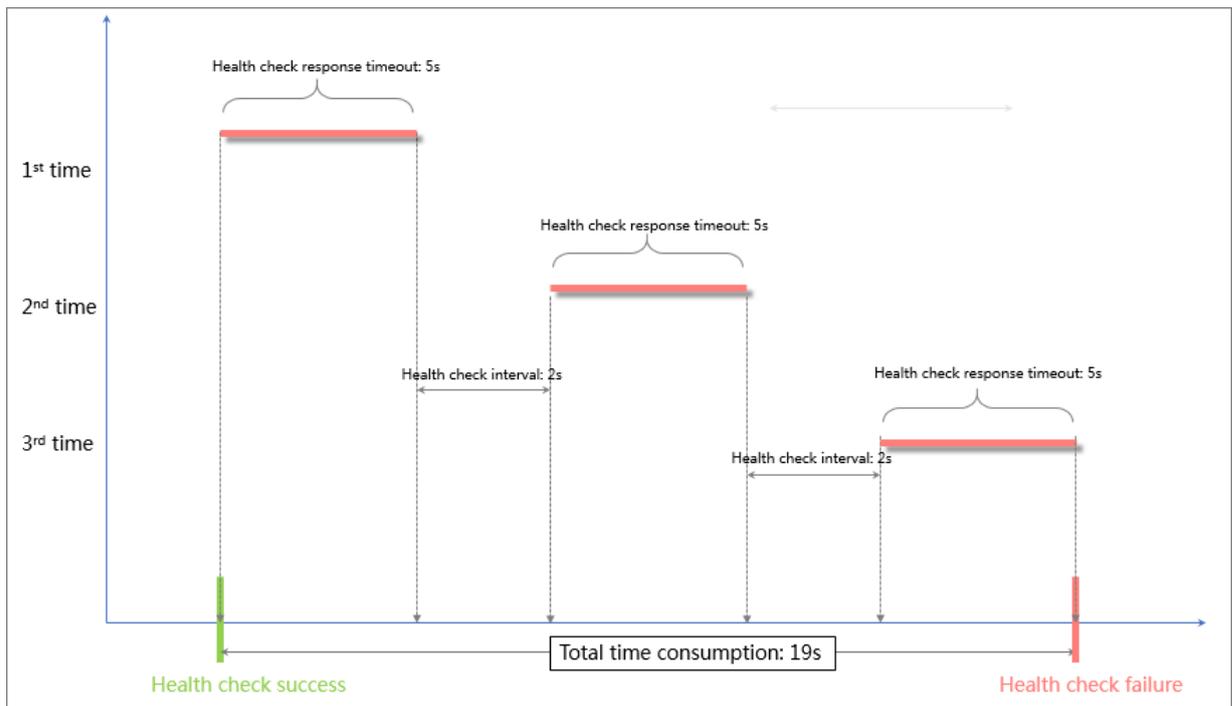
Example of health check response timeout and health check interval

Take the following health check configurations as the example:

- **Response Timeout: 5 seconds**
- **Health Check Interval: 2 seconds**
- **Healthy Threshold: 3 times**
- **Unhealthy Threshold: 3 times**

Health check failure time window = Response Timeout × Unhealthy Threshold + Health Check Interval × (Unhealthy Threshold - 1). That is, $5 \times 3 + 2 \times (3 - 1) = 19s$.

The following figure shows the process to declare an unhealthy backend server:



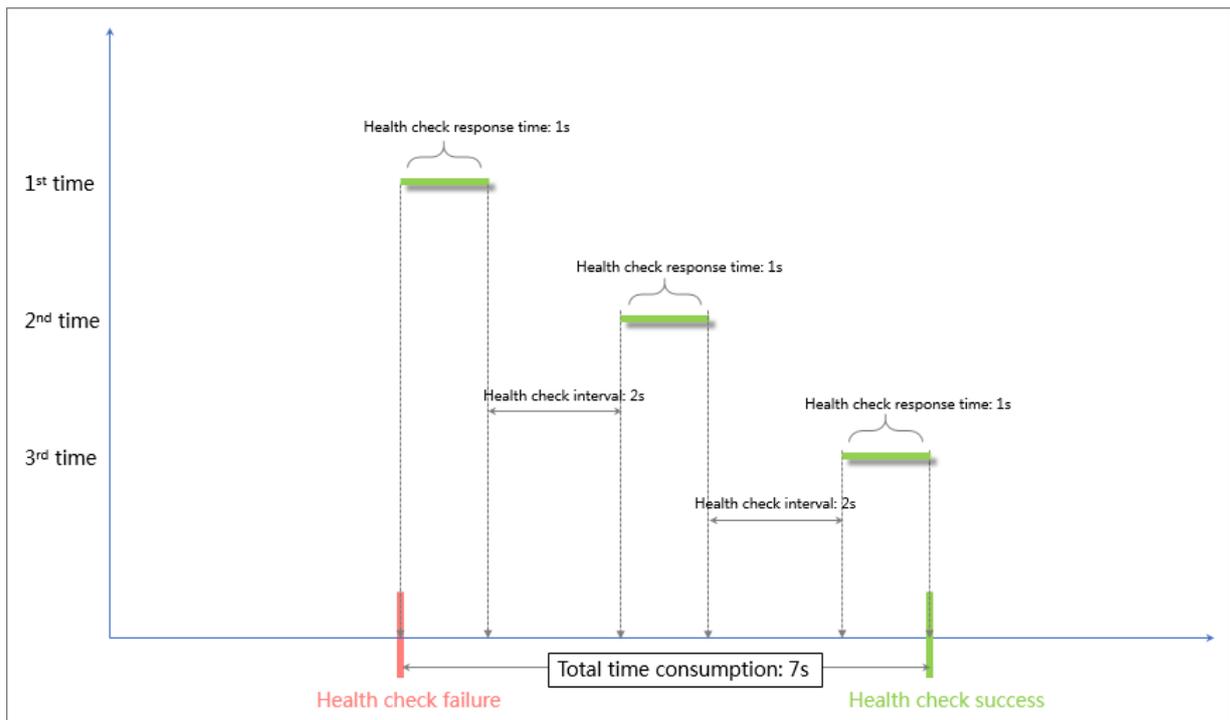
Health check success time window = Health check response time × Healthy Threshold + Health Check Interval × (Healthy Threshold - 1). That is, $(1 \times 3) + 2 \times (3 - 1) = 7s$.



Note:

Health check response time is the time a successful request-response message is sent and received. When using TCP for health check, due to the detection of only the port for survival, therefore, the time is very short and can be ignored almost. When the HTTP health check is used, the response time is usually within seconds, depending on the performance and load of the application server.

The following figure shows the process to declare a healthy backend server (Assume that the backend server takes 1 second to respond to the health check request):



Configure a domain name in HTTP health check

When the HTTP health check is used, you can set a domain name for the health check, but it is not a required option. Some application servers verify the host field in the request, therefore, the request header must contain the host field. If a domain name is configured in the health check, Server Load Balancer adds the domain name to the host field when forwarding the request to the backend server. If so, the health check request will be denied by the server and the health check may fail. Therefore, if your application server verifies the host field in the request, you must configure a domain name to make sure the health check works.

3.3 Close health check

If health check is closed, requests may be distributed to unhealthy ECS instances, which can lead to service interruption. In general, we do not recommend closing health check.

Context



Note:

You can only close health check for HTTP and HTTPS listeners. The health check of UDP and TCP listeners cannot be closed.

Procedure

1. Log on to the [SLB console](#).

2. On the **Instances** page, click the ID of the target instance.
3. Under the **Listeners** tab, find the target listener and then click the **Configure** option.
4. On the **Configure Listener** page, click **Next** until the **Health Check** tab is displayed.
5. Click the switch to disable health check, then click **Next** and **Submit**.

4 Backend servers

4.1 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 application 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 region for the ECS instances and the SLB instance is the same.
- 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 the ECS. The front-end port of the listener is 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 as configured among the backend servers, 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. If session persistence is enabled, you need to strip the access logs for the same IP address. If the weight is configured for SLB, you need to calculate whether the percentage of access traffic recorded in the logs matches the weight ratio.)
- When an ECS instance is undergoing live migration, the persistent connections 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 [Manage a default server group](#) 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 [Manage an active/standby server group](#) 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 [Manage a VServer group](#) to create a VServer group.

4.2 Manage a default server group

Before using the SLB service, you must add at least one default server to receive client requests forwarded by SLB.

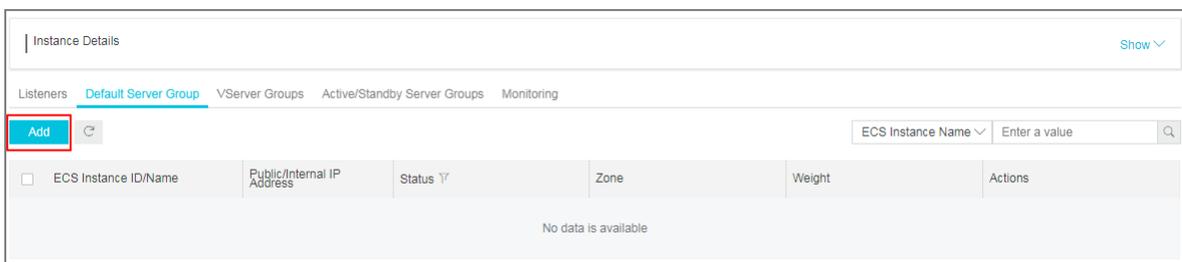
Add default servers

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

- You have [created an SLB instance](#).
- You have created ECS instances and deploy applications to process distributed requests.

To add ECS instances, complete these steps:

1. Log on to the [SLB console](#).
2. On the **Server Load Balancer** page, select a region.
3. Click the ID of the target SLB instance.
4. Click the **Default Server Group** tab.
5. Click **Add**.



6. On the **Available Servers** page, find the target ECS instance and click **Add**, or click multiple target ECS instances and click **Add to Selected Server List** at the bottom of the page.

Available Servers
✕

<input type="checkbox"/>	ECS Instance ID/Name	Public/Internal IP Address	Status	Zone	SLB Instance	Actions
<input checked="" type="checkbox"/>	E-IT i- b-5n 2	47 (Public) 17 (Private) vpc-ay7beb2wxl8 vs- b-azerxqh0mfi	● Running	China East 1 Zone G	Associated SLB Instances 4	Add
<input checked="" type="checkbox"/>	E-IT i- b-5n 2	4 (Public) 1 (Private) vpc-ay7beb2wxl8 vs- b-azerxqh0mfi	● Running	China East 1 Zone G	Associated SLB Instances 4	Add
<input type="checkbox"/>	E- i- b-nwqy e-76tv	1 (Private) vpc-5itm138tv9 vs- b-kpo2458sgj3	● Running	China East 1 Zone G	Associated SLB Instances 1	Add
<input type="checkbox"/>	E- i- b-nwqy e-76tu	1 (Private) vpc-5itm138tv9 vs- b-po2458sgj3	● Running	China East 1 Zone G	Associated SLB Instances 1	Add

7. Click **OK**.

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

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



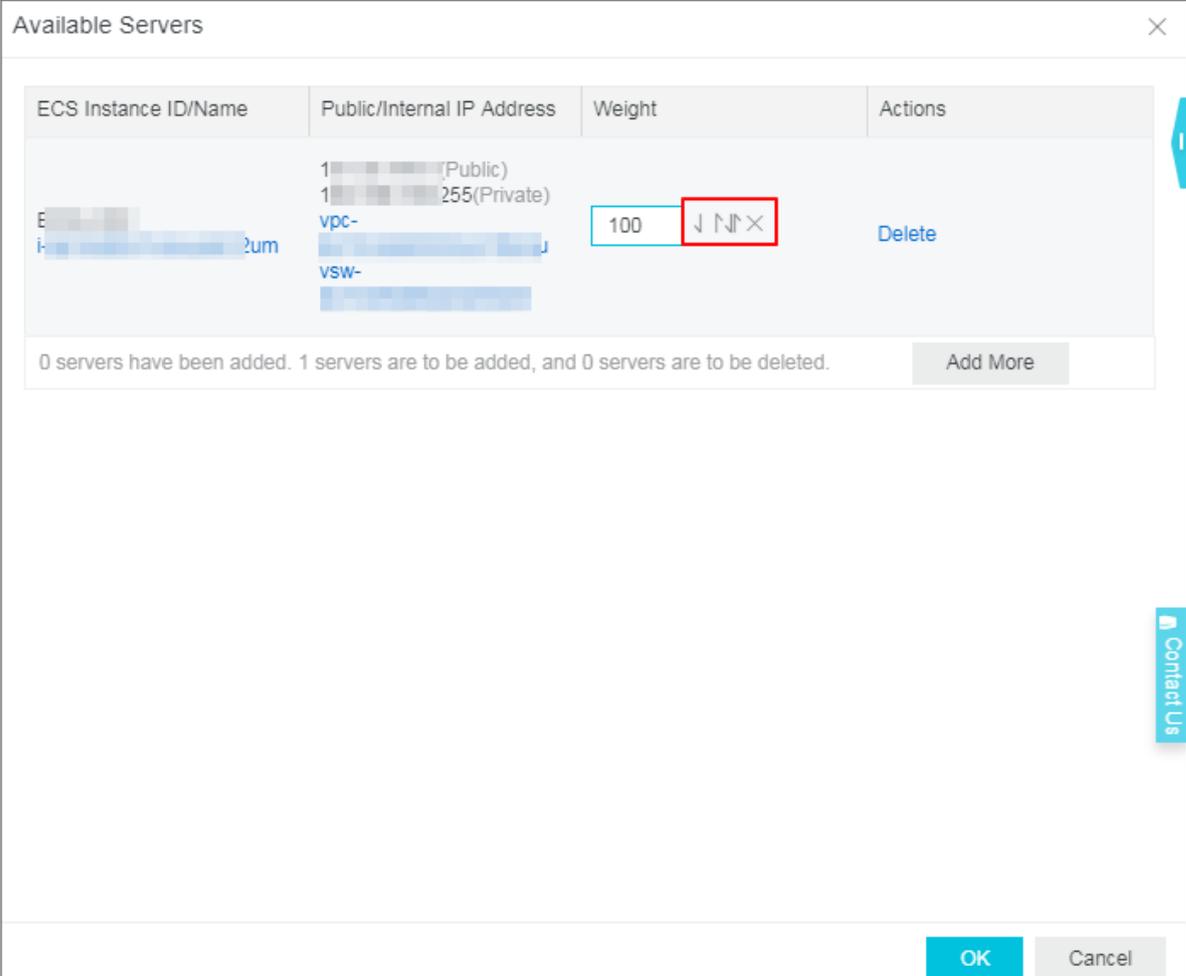
Note:

If the weight is set to 0, no requests will be 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 weight of the current server, the weights of all servers below 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.



ECS Instance ID/Name	Public/Internal IP Address	Weight	Actions
E[redacted]	1[redacted] (Public) 1[redacted] 255 (Private) vpc-[redacted] vsw-[redacted]	100	Delete

0 servers have been added. 1 servers are to be added, and 0 servers are to be deleted. Add More

OK Cancel

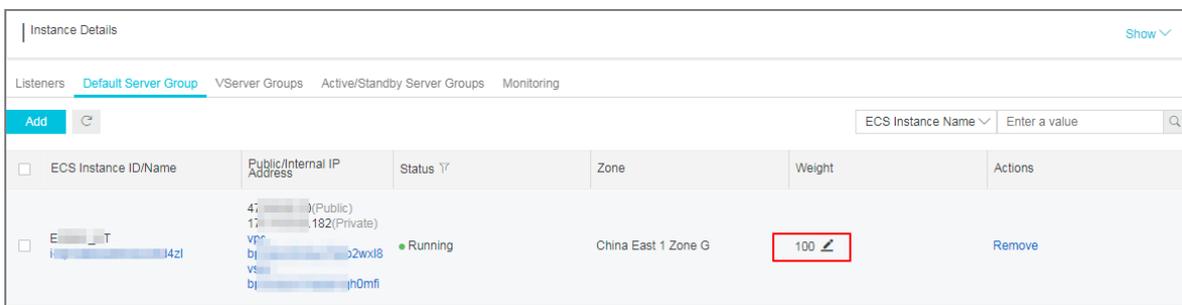
9. Click **OK**.

Edit the weight of a backend server

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

1. Log on to the [SLB console](#).
2. On the **Server Load Balancer** page, select a region.
3. Click the ID of the target SLB instance.
4. Click the **Default Server Group** tab.

5. Hover the mouse to the weight area 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 will receive a larger number of connection requests.



Note:

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

Remove a backend server

To remove a backend server, complete these steps:

1. Log on to the [SLB console](#).
2. On the **Server Load Balancer** page, select a region.
3. Click the ID of the target SLB instance.
4. Click the **Default Server Group** tab.
5. Click **Remove** in the **Actions** column to remove the backend server.

4.3 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.

If you add default backend servers, VServer groups and forwarding rules to the same Layer-7 listener, the sequence of request forwarding is as follows:

- If the requests match a forwarding rule, the requests are distributed to the VServer group associated with the rule.
- If not, 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 creating a VServer group, make sure the following conditions are met:

- You have [created an SLB instance](#).
- You have created ECS instances and deploy applications to process distributed requests.

Note the following when creating a VServer group:

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

To add ECS instances, complete these steps:

1. Log on to the [SLB console](#).
2. On the **Server Load Balancer** page, select a region.
3. Click the ID of the target SLB instance.
4. Click the **VServer Groups** tab.
5. On the **VServer Group** page, click **Create VServer Group**.
6. On the **Create VServer Group** page, complete these steps:
 - a. In the **VServer Group Name** text box, enter the name of the VServer group.
 - b. Click **Add** to select the server to add in the **Available Servers** list.
 - c. Click **Add to Selected Server List** and click **OK**.
 - d. Under the **Servers Added** tab, complete the following configuration and click **OK**.
 - **Port**: The backend port opened on the ECS instance to receive requests.
The backend ports in a Server Load Balancer instance can be the same.
 - **Weight**: An ECS instance with a higher weight receivers more requests.



Note:

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

Create VServer Group ✕

① Note: The network type of the specified SLB is VPC, and the instance type is VPC. You can add either classic ECS instances or VPC ECS instances into the VServer group.

• VServer Group Name

Servers Added

ECS Instance ID/Name	Public/Internal IP Address	端口	Weight	Actions
E(.....)T i-..... b-..... 5d4zl	47.....(Public) 17.....182(Private) vpc-..... vsw-..... b-.....fi	80	100	Delete
E(.....)T i-..... b-..... 5d4zm	47.....(Public) 17.....183(Private) vpc-..... vsw-..... b-.....fi	80	100	Delete

0 servers have been added. 2 servers are to be added, and 0 servers are to be deleted. Add More

OK
Cancel

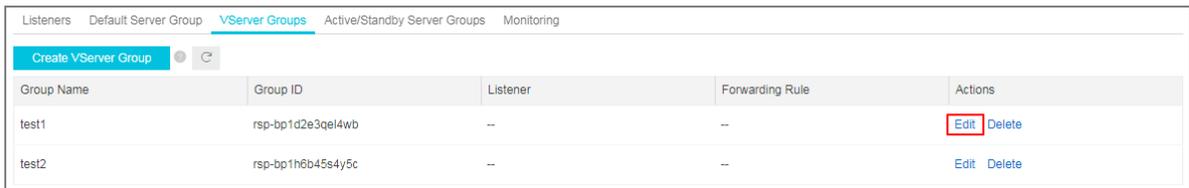
You can modify the ports and weights of added servers in batches.

- Click : Duplicate to below. If you modify the port or weight of the current server, the ports or weights of all servers below 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, complete these steps:

1. Log on to the [SLB console](#).
2. On the **Server Load Balancer** page, select a region.
3. Click the ID of the target SLB instance.
4. Click the **VServer Groups** tab.
5. Find the target VServer group, and then click the **Edit** option.



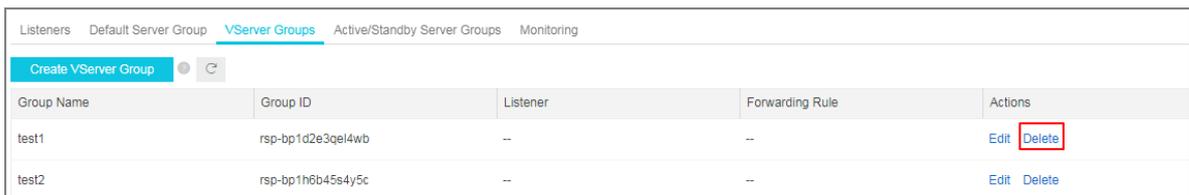
Group Name	Group ID	Listener	Forwarding Rule	Actions
test1	rsp-bp1d2e3qe14wb	--	--	Edit Delete
test2	rsp-bp1h6b45s4y5c	--	--	Edit Delete

6. Modify the port and weight of the ECS instance or click **Delete** to remove the ECS instance from the VServer group, and then click **OK**.

Delete a VServer group

To delete a VServer group, complete these steps:

1. Log on to the [SLB console](#).
2. On the **Server Load Balancer** page, select a region.
3. Click the ID of the target SLB instance.
4. Click the **VServer Groups** tab.
5. Find the target VServer group, and then click the **Delete** option.



Group Name	Group ID	Listener	Forwarding Rule	Actions
test1	rsp-bp1d2e3qe14wb	--	--	Edit Delete
test2	rsp-bp1h6b45s4y5c	--	--	Edit Delete

6. In the displayed dialog box, click **OK**.

4.4 Manage an active/standby server group

If you have traditional active/standby requirement, where one backend server is used as the active server and the other as the standby server, create an active/standby server group. When the active server works normally, requests are distributed to it; when the active server is down, the requests will be distributed to the standby server to avoid service interruption.

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.

**Note:**

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

Create an active/standby server group

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

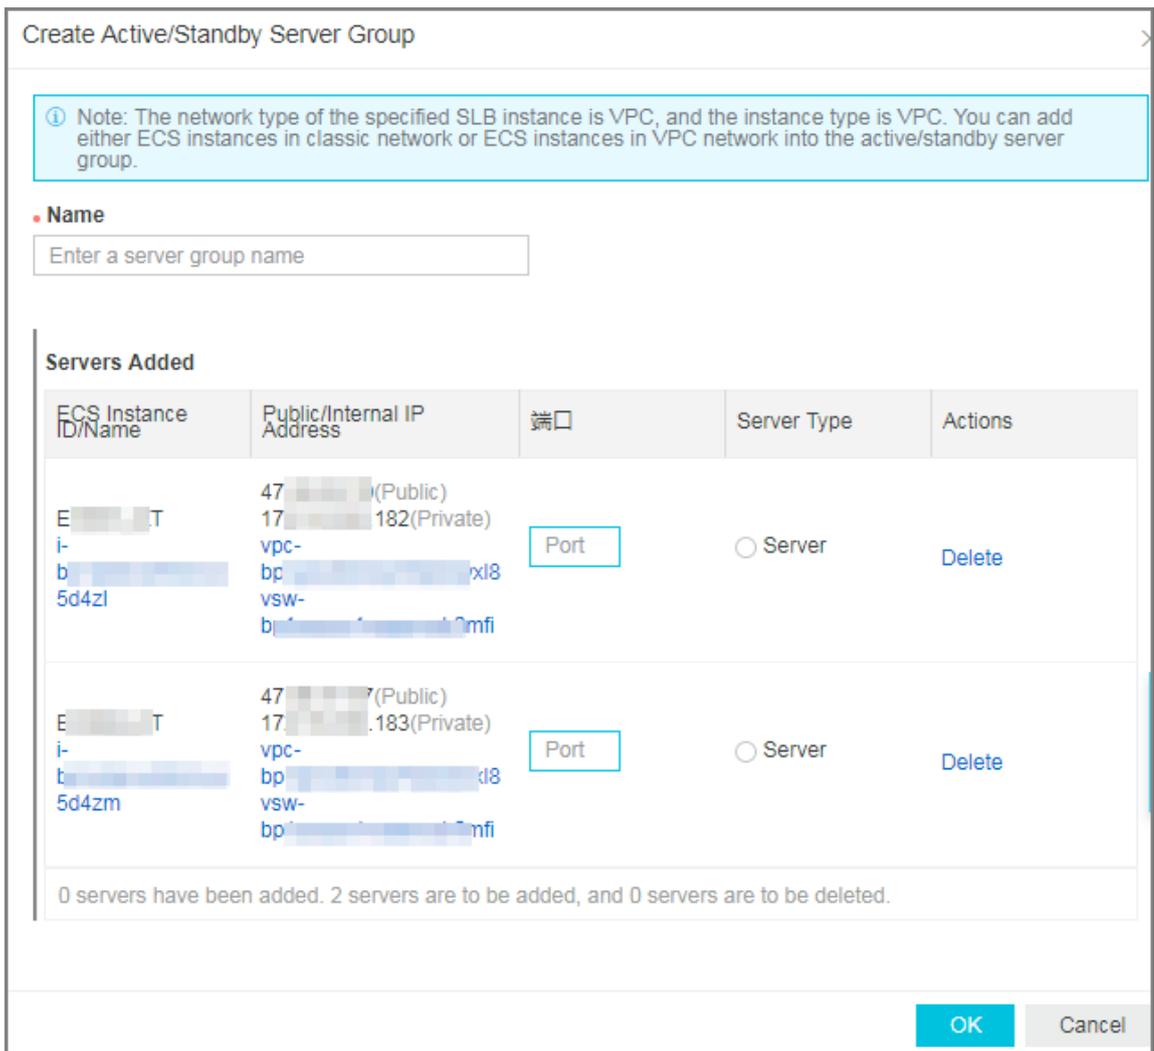
- You have [created an SLB instance](#).
- You have created ECS instances and deploy applications to process distributed requests.

To add ECS instances, complete these steps:

1. Log on to the [SLB console](#).
2. On the **Server Load Balancer** page, select a region.
3. Click the ID of the target SLB instance.
4. Click the **Active/Standby Server Groups** tab.
5. On the **Active/Standby Server Groups** page, click **Create Active/Standby Server Group**.
6. On the **Create Active/Standby Server Group** page, complete these steps:
 - a. In the **Name** text box, enter the name of the active/standby server group.
 - b. Click **Add** to select the server to add in the **Available Servers** list.

You can add up to two ECS instances to an active/standby server group.

- c. Click **Add to Selected Server List** and click **OK**.
- d. In the **Servers Added** tab, complete the following configuration and click **OK**.
 - **Port:** The backend port opened on the ECS instance to receive requests.
The back-end ports in a Server Load Balancer 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, complete these steps:

1. Log on to the [SLB console](#).
2. On the **Server Load Balancer** page, select a region.
3. Click the ID of the target SLB instance.
4. Click the **Active/Standby Server Groups** tab.
5. Click **Delete** next to the target active/standby server group.



6. In the displayed dialog box, click **OK**.

5 Certificate management

5.1 Certificate requirements

Server Load Balancer only supports certificates in the PEM format. Before uploading a certificate, make sure that the certificate, certificate chain, and private key conform to the rules described in this section.

Certificates issued by a root CA

If the certificate is issued by a root CA, the received certificate is the only one required to be uploaded to Server Load Balancer. The website that is configured with the certificate will be trusted by the web browser without configuring additional certificates.

The certificate format must meet the following requirements:

- The certificate content is placed between `-----BEGIN CERTIFICATE-----`, `-----END CERTIFICATE-----`. Include the header and footer when uploading the certificate.
- Each line except the last must contain exactly 64 characters. The last line can contain 64 or fewer characters.
- Space is not allowed in the content.

The following is a sample certificate issued by a root CA.


```
-----BEGIN CERTIFICATE-----  
-----END CERTIFICATE-----  
-----BEGIN CERTIFICATE-----  
-----END CERTIFICATE-----
```

RSA private key

When uploading a server certificate, you also need to upload the private key of the certificate.

The RSA private key format must meet the following requirements:

- The key is placed between `-----BEGIN RSA PRIVATE KEY-----`, `-----END RSA PRIVATE KEY-----`. Include the header and footer when uploading the key.
- Space is not allowed in the content. Each line except the last must contain exactly 64 characters. The last line can contain 64 or fewer characters. For more information, see [RFC1421](#).

If your private key is encrypted. For example, the header and footer are `-----BEGIN PRIVATE KEY-----`, `-----END PRIVATE KEY-----` or `-----BEGIN ENCRYPTED PRIVATE KEY-----`, `-----END ENCRYPTED PRIVATE KEY-----`, or the private key contains `Proc-Type`

: 4, ENCRYPTED, run the following command to convert the private key before uploading it to

Server Load Balancer:

```
openssl rsa -in old_server_key.pem -out new_server_key.pem
```

The following is a sample RSA private key.

```

-----BEGIN RSA PRIVATE KEY-----
MIIEpAIBAAKCAQEAvZiSSSchH67bmT8mFykAxQ1tKCYukwBiWZwk0StFEBTWHy8K
tTHSFD1u9TL6qycrHEG7cjYD4DK+kVIHU/Of/pUWj9LLnrE3W34DaVzQdKA00I3A
Xw9SgrqFJMjclva2khNKA1+tNPSCPJoo9DDrP7wx7cQx7LbMb0dfZ8858KIoluzJ
/fD0XXyuWoqaIePZtK9Qn957ZEPHjtUpVZuhS3409DDM/tJ3Tl8aaNYWhrPBc0
jNcz0Z6XQGf1rZG/Ve520GX6rb5dUYpdcfXzN5WM6xYg8a1L7UHDHPI4AYsatdG
z5TMPnmEfy8yZPUYudTLxgMVAovJr09Dq+5Dm3QIDAQABAoIBAGl68Z/nnFyRHRFi
laF6+Wen8ZvNqkm0hAMQwIjH1Vp1f174//8Qyea/EvUtuJHyB6T/2PZQoNVhxe35
cgQ93Tx424WGPcWUshSfXewfbAYGf3ur8W0xq0uU07BAxaKHNcmNG7dGyo1UowRu
S+yXLrpVzH1YkuH8TT53udd6TeTWi77r8dkGi9KSAZ0pRa19B7t+CHKIzm6ybs/2
06W/zHZ4YAxwkTYLKGHjoiEYs111ah1AJvICVgTc3+LzG2pIpM7I+K0nHC5eswM
i5x9h/OT/ujZsyX9P0PaAyE2bqy0t080tGexM0765Ssv0KVhKFvWjLUnhf6WcqFCD
xqhxkECgYEA+PftNb6eyXl+/Y/U8NM2fg3+rSCms0j9Bg+9+yZzF5GhagHu0edU
ZXIHrJ9u6BlXE1arpijVs/WHmFhYSTm6DbdD7S1tLy0BY4cPTRhziFTKt8AkIXMK
605u0UiWsq0Z8hn1X14lox2cW9ZQa/Hc9udeyQotP4NsMJWgpBV7tC0CgYEAwvNf
0f+/jUjt0HoyxCh4SIAqk4U0o4+hBCQbWcXv5qCz4mRyTawzFEG8/AR3Md2rhmZi
GnJ5fdfe7uY+JsQfX2Q5JjwTad1BW4led0Sa/ukRa04UzVgnYp2aJKxtuWffvVbU
+kf728ZJRA6azSLvGmA8hu/GL6bgfU3fkSkw03ECgYBpYK7TT7JvvnAERMtJf2yS
ICRkQaB3gPSe/LCgzy1nhtaF0UbNxBgeuowLAZR0wrz7X3TZqHEDcYoJ7mK346of
QhGLITyoeHkbYkAUTq038Y04EKh6S/IzMzB0frXiPKg9s8UKQzkU+GSE7ootli+a
R8Xzu835EwxI6BwNN1abpQKBgQC8TialClq1FteXQyGcNdcReLMncUHKIKcP/+xn
R3kVl06MZCfAdqirAjiQWapkh9Bxbp2eHCrB81MFAWLRQSl0k79b/jVmTZMC3upd
EJ/iSWjZKpbw7hCFAeRtPhxyNTJ5idEiu9U8EQid8111giPgn0p3sE0HpDI89qZX
aaiMEQKBgQDK2bsnZE9y0ZWhGTeu94vziKmFrSkJMGH8pLaTiliw1iRhRYWJysZ9
BOIDxnrmwiPa9bCtEpK80zq28dq7axpCs9CavQRcv0Bh5Hx0yy23m9hFRzFDeQ7z
NTKh193HhF1j0NM81LHFyGRFEWWrr0W5gfBudR6USRnR/6iQ11xZXw==
-----END RSA PRIVATE KEY-----

```

5.2 Upload a certificate

Before creating HTTPS listeners, you must upload the required server certificate and CA certificate to SLB. You no longer need to configure certificates on the backend servers after uploading the certificates to SLB.

Prerequisites

- You have purchased a server certificate.
- You have generated a CA certificate and client certificate. For more information, see [Generate a CA certificate](#).

Context

Note the following before uploading certificates:

- If you want to use a certificate in multiple regions, you must upload the certificate to all these regions.
- Up to 100 certificates can be uploaded per account.

Procedure

- Log on to the [SLB console](#).
- In the left-side navigation pane, click **Certificates**.

3. Click **Create Certificate**.



4. On the **Create Certificate** page, upload the certificate content and then click **OK**.

Configuration	Description
Certificate Name	Enter a certificate name. The name must be 1-80 characters long, and can only contain letters , numbers and the following special characters: _/. -
Regions	Select one or more regions where the certificate is uploaded. A certificate cannot be used across regions. If a certificate is to be used in multiple regions, select all these regions.
Certificate Type	Select a certificate type. <ul style="list-style-type: none"> Server Certificate: For HTTPS one-way authentication, only the server certificate and the private key are required. CA Certificate: For HTTPS two-way authentication, both the server certificate and the CA certificate are required.
Certificate Content	Paste the certificate content in the editor. Click Import Sample to view the valid certificate formats. Only certificates in the PEM format are supported. For more information, see Certificate requirements .
Private Key	Paste the private key of the server certificate in the editor. Click Import Sample to view the valid certificate formats. For more information, see Certificate requirements . <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  Note: A private key is only required when uploading a server certificate. </div>

Click **Remove All Expired Certificates** to delete expired certificates in batches.

5.3 Generate a CA certificate

When configuring HTTPS listeners, you can use self-signed CA certificates. Follow the instructions in this document to generate a CA certificate and use the CA certificate to sign a client certificate.

Generate a CA certificate by using Open SSL

1. Run the following commands to create a *ca* folder in the */root* directory and then create four sub folders under the *ca* folder.

```
$ sudo mkdir ca
$ cd ca
$ sudo mkdir newcerts private conf server
```

- *newcerts* is used to store the digit certificate signed by a CA certificate.
- *private* is used to store the private key of the CA certificate.
- *conf* is used to store the configuration files.
- *server* is used to store the server certificate.

2. Create an *openssl.conf* file that contains the following information in the *conf* directory.

```
[ ca ]
default_ca = foo
[ foo ]
dir = /root/ca
database = /root/ca/index.txt
new_certs_dir = /root/ca/newcerts
certificate = /root/ca/private/ca.crt
serial = /root/ca/serial
private_key = /root/ca/private/ca.key
RANDFILE = /root/ca/private/.rand
default_days = 365
default_crl_days= 30
default_md = md5
Unique_subject = No
Policy = policy_any
[ policy_any ]
countryName = match
stateOrProvinceName = match
organizationName = match
organizationalUnitName = match
localityName = optional
commonName = supplied
emailAddress = optional
```

3. Run the following command to generate a private key.

```
$ cd /root/ca
$ sudo openssl genrsa -out private/ca.key
```

The following figure is an example of key generation.

```

root@iZbp1hfvivcqx1jwap31iZ:~/ca/conf# cd /root/ca
root@iZbp1hfvivcqx1jwap31iZ:~/ca# sudo openssl genrsa -out private/ca.key
Generating RSA private key, 2048 bit long modulus
.....++++
....+++
..+++
e is 65537 (0x10001)

```

4. Run the following command and input the required information according to the prompts. Press Enter to generate a *csr* file.

```
$ sudo openssl req -new -key private/ca.key -out private/ca.csr
```



Note:

Common Name is the domain name of the SLB instance.

```

root@iZbp1hfvivcqx1jwap31iZ:~/ca# sudo openssl req -new -key private/ca.key -out private/ca.csr
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [AU]:CN
State or Province Name (full name) [Some-State]:ZheJiang
Locality Name (eg, city) []:HangZhou
Organization Name (eg, company) [Internet Widgits Pty Ltd]:Alibaba
Organizational Unit Name (eg, section) []:Test
Common Name (e.g. server FQDN or YOUR name) []:mydomain
Email Address []:a@alibaba.com

Please enter the following 'extra' attributes
to be sent with your certificate request
A challenge password []:
An optional company name []:
root@iZbp1hfvivcqx1jwap31iZ:~/ca# █

```

5. Run the following command to generate a *crt* file.

```
$ sudo openssl x509 -req -days 365 -in private/ca.csr -signkey private/ca.key -out private/ca.crt
```

6. Run the following command to set the start sequence number for the private key, which can be any four characters.

```
$ sudo echo FACE > serial
```

7. Run the following command to create a CA key library.

```
$ sudo touch index.txt
```

8. Run the following command to create a certificate revocation list for removing the client certificate.

```
$ sudo openssl ca -gencrl -out /root/ca/private/ca.crl -crl days 7 -config "/root/ca/conf/openssl.conf"
```

The response is as follows:

```
Using configuration from /root/ca/conf/openssl.conf
```

Sign the client certificate

1. Run the following command to generate a *users* folder under the *ca* directory to store the client key.

```
$ sudo mkdir users
```

2. Run the following command to create a key for the client certificate.

```
$ sudo openssl genrsa -des3 -out /root/ca/users/client.key 1024
```



Note:

Enter a pass phrase when creating the key. It is the password to protect the private key from unauthorized access. The pass phrase entered is the password for this key.

3. Run the following command to create a *csr* file for requesting certificate sign.

```
$ sudo openssl req -new -key /root/ca/users/client.key -out /root/ca/users/client.csr
```

Enter the pass phrase set in the previous step when prompted.



Note:

A challenge password is the password of the client certificate. Note that it is not the password of the client key.

4. Run the following command to sign the client key.

```
$ sudo openssl ca -in /root/ca/users/client.csr -cert /root/ca/private/ca.crt -keyfile /root/ca/private/ca.key -out /root/ca/users/client.crt -config "/root/ca/conf/openssl.conf"
```

Enter `y` twice when prompted.

```
root@iZbp1hfvivcqx1jwbp31iZ:~/ca# sudo openssl ca -in /root/ca/users/client.csr -cert /root/ca/private/ca.crt -keyfile /root/ca/private/ca.key -out /root/ca/users/client.crt -config "/root/ca/conf/openssl.conf"
Using configuration from /root/ca/conf/openssl.conf
Check that the request matches the signature
Signature ok
The Subject's Distinguished Name is as follows
countryName             :PRINTABLE:'CN'
stateOrProvinceName    :ASN.1 12:'ZheJiang'
localityName           :ASN.1 12:'HangZhou'
organizationName       :ASN.1 12:'Alibaba'
organizationalUnitName :ASN.1 12:'Test'
commonName             :ASN.1 12:'mydomain'
emailAddress           :IA5STRING:'a@alibaba.com'
Certificate is to be certified until Jun  4 15:28:55 2018 GMT (365 days)
Sign the certificate? [y/n]:y

1 out of 1 certificate requests certified, commit? [y/n]y
Write out database with 1 new entries
Data Base Updated
root@iZbp1hfvivcqx1jwbp31iZ:~/ca#
```

5. Run the following command to convert the certificate to a *PKCS12* file.

```
$ sudo openssl pkcs12 -export -clcerts -in /root/ca/users/client.crt -inkey /root/ca/users/client.key -out /root/ca/users/client.p12
```

Enter the password of the client key when prompted. Then, enter the password used for exporting the client certificate. This password is used to protect the client certificate, which is required when installing the client certificate.

6. Run the following command to view the generated client certificate.

```
cd users
```

```
ls
```

5.4 Convert certificate formats

Server Load Balancer supports PEM certificates only. Certificates in other formats must be converted to the PEM format before they can be uploaded to Server Load Balancer. We recommend that you use Open SSL for conversion.

Convert DER to PEM

DER: This format is usually used on a Java platform.

- Run the following command to convert the certificate format.

```
openssl x509 -inform der -in certificate.der -out certificate.pem
```

- Run the following command to convert the private key.

```
openssl rsa -inform DER -outform PEM -in privatekey.der -out privatekey.pem
```

Convert P7B to PEM

P7B: This format is usually used in a Windows server and Tomcat.

Run the following command to convert the certificate format.

```
openssl pkcs7 -print_certs -in incertificate.p7b -out outcertificate.cer
```

Convert PFX to PEM

PFX: This format is usually used in a Windows server.

- Run the following command to extract the certificate format.

```
openssl pkcs12 -in certname.pfx -nokeys -out cert.pem
```

- Run the following command to extract the private key.

```
openssl pkcs12 -in certname.pfx -nocerts -out key.pem -nodes
```

5.5 Replace a certificate

To avoid the impact of certificate expiration on your service, replace the certificate before the certificate expires.

Procedure

1. Create and upload a new certificate.

For more information, see [Upload certificates](#) and [Generate certificates](#).

2. Configure the new certificate in HTTPS listener configuration.

For more information, see [Add an HTTPS listener](#).

3. On the **Certificates** page, find the target certificate, and then click **Delete**.
4. In the displayed dialog box, click **OK**.

6 Log management

6.1 View operation logs

You can view the logs of operations performed on SLB instances, HTTP listeners and server certificates in one month.

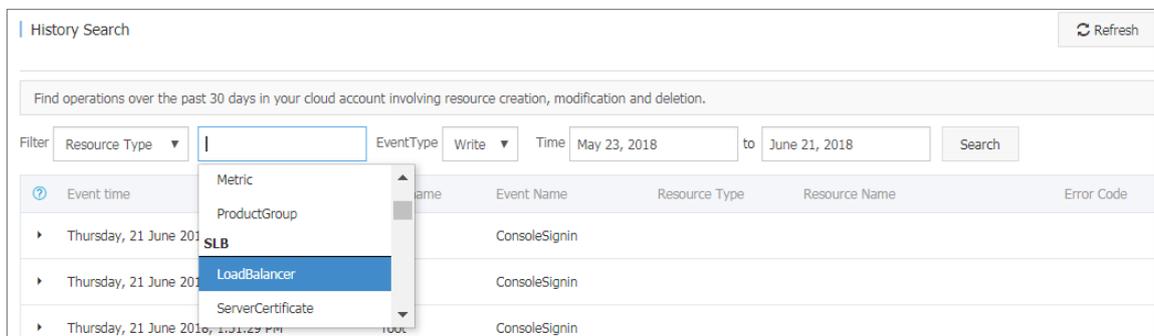
Context

The operation logs are recorded in ActionTrail. ActionTrail records the operations acting upon your Alibaba Cloud resources, you can query operation records and store the records to OSS.

Procedure

1. Log on to the [SLB console](#).
2. In the left-side navigation pane, click **Logs > Operation Log**.
3. Click **View Operation Logs**.
4. On the **History Search** page, complete these steps to view operation logs:
 - a) Select **Resource Type** as a filter.
 - b) Select the SLB resource of which operation logs you want to view.

In this tutorial, **LoadBalancer** is selected.



- c) Select an event type.

In this tutorial, **All** is selected.
- d) Select the time range to search.
- e) Click **Search** to view logs of operations performed on the selected resource.

Expand the resource to view more detailed information.

History Search
Refresh

Find operations over the past 30 days in your cloud account involving resource creation, modification and deletion.

Filter
Resource Type LoadBalancer
EventType Write
Time May 23, 2018 to June 21, 2018
Search

Event time	Username	Event Name	Resource Type	Resource Name	Error Code
Tuesday, 19 June 2018, 12:38:41 AM	root	CreateDomainExtension	LoadBalancer	lb-1ud31bltkhwwjnw3c5k0	

Access key:	Event source: slb-openapi-share.aliyuncs.com
Region: cn-hangzhou	Event Time: Tuesday, 19 June 2018, 12:38:41 AM
Error Code:	Request ID: 115E5A4E-4083-4103-A301-8D62E2F1BF81
Event ID: 115E5A4E-4083-4103-A301-8D62E2F1BF81	Source IP address: 61.149.173.169
Event Name: CreateDomainExtension	Username: root

Related Resources (1)

- LoadBalancer
lb-1ud31bltkhwwjnw3c5k0

View event

6.2 Manage health check logs

You can view the health logs within three days on **Health Check Logs** page. If you want to get the health check logs three days before or longer, you can store the health check logs to OSS. Therefore, you can download complete health check logs.

Store health check logs

You can view the health check logs of the backend servers by using the health check log function . Currently, logs in three days are provided. If you want to view more logs, store the health check logs to OSS.

You can enable and disable the storage function at any time. After the storage function is enabled, SLB will create a folder named *AliyunSLBHealthCheckLogs* in the selected bucket to store the health check logs. The health logs are generated hourly and the system will create a subfolder named after the date to store the log files generated in that day, for example *20170707*.

The log files in a day are named after the time when they are generated. For example, the log file that is generated between 00:00-01:00, the file name will be *01.txt* and the log file that is generated between 01:00-02:00, the file name will be *02.txt*.



Note:

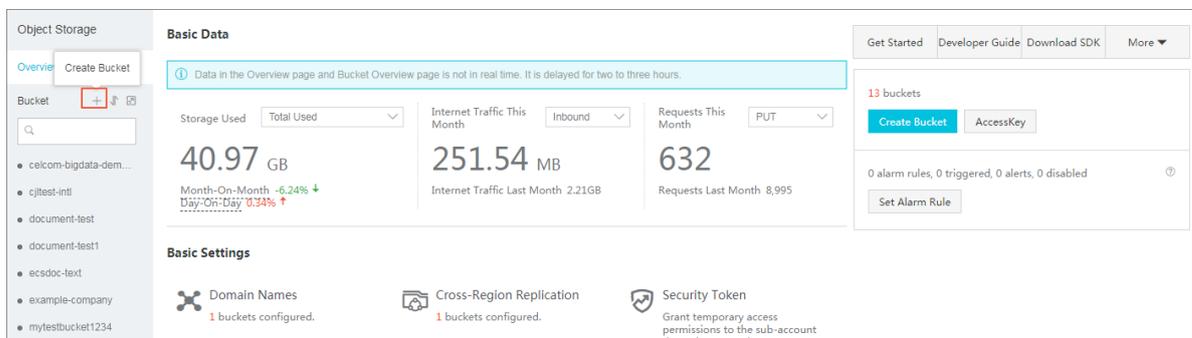
The health check logs are generated only when the backend server is abnormal. Health check logs are generated only when the backend server is abnormal. If no failures occur for all the backend servers in an hour, no health check logs are generated in that hour.

To store health check logs, complete these steps:

1. [Create a bucket](#)
2. [Authorize SLB to access OSS](#)
3. [Configure log storage](#)

Step 1 Create a bucket

1. Open the [OSS product page](#) and click **Buy Now** to activate the OSS service.
2. Log on to the OSS console.
3. Click **Create Bucket**.



4. In the **Create Bucket** dialog box, configure the bucket and click **OK**.



Note:

Make sure that the region of the bucket and the SLB instance are the same.

Step 2 Authorize SLB to access OSS

After creating a bucket, you have to authorize the log role (SLBLogDefaultRole) to access OSS resources.

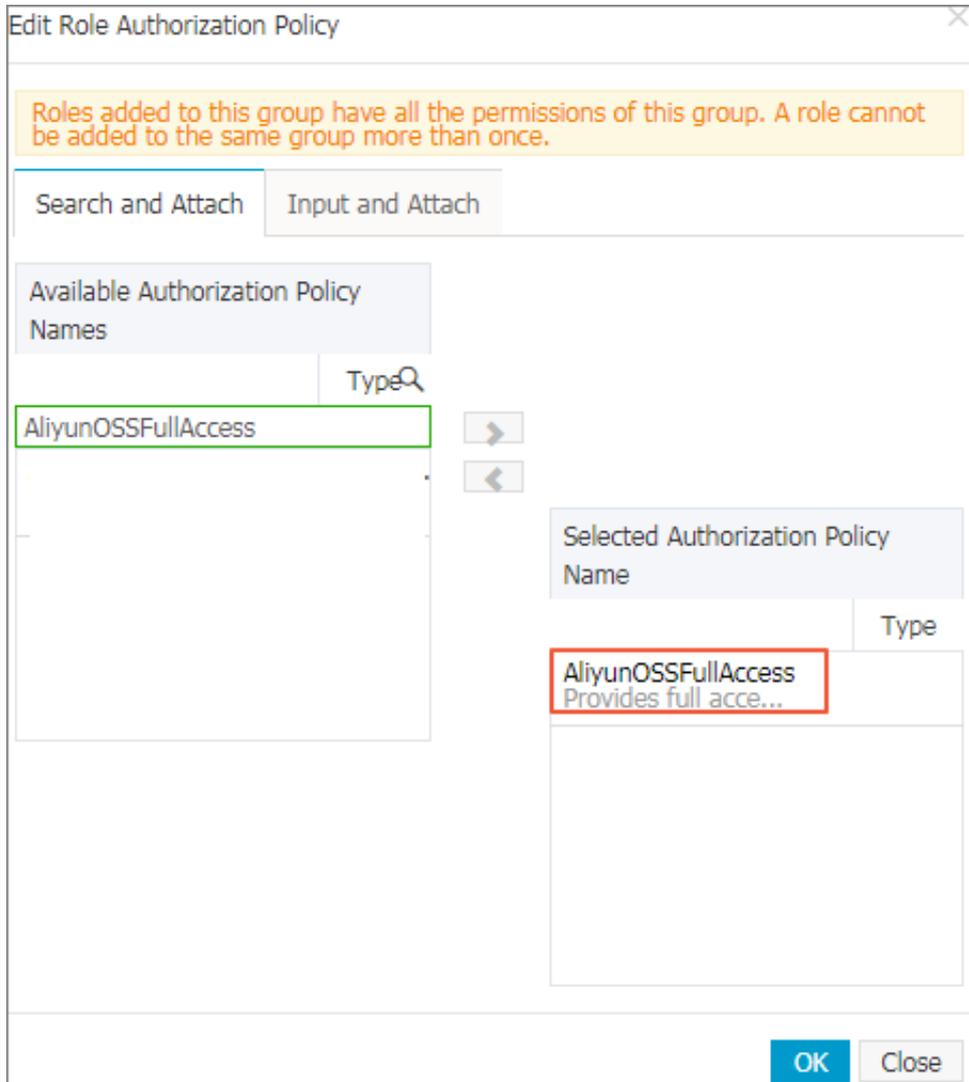


Note:

The authorization is required only for first configuration.

1. On the SLB console, click **Logs > Health Check Logs**.
2. Click **1. Activate OSS**, if OSS has not been activated yet.
3. On the **Health Check Logs** page, click **Add Role Now** in the **2. Add the RAM role to your account** section.
4. Read the authorization description, and then click **Confirm Authorization Policy**.
5. Log on to the RAM console.

6. In the left-side navigation pane, click **Roles** and find the role named SLBLogDefaultRole, and then click **Authorize**.
7. In the **Edit Role Authorization Policy** dialog box, find the **AliyunOSSFullAccess** policy, and then click **OK**.

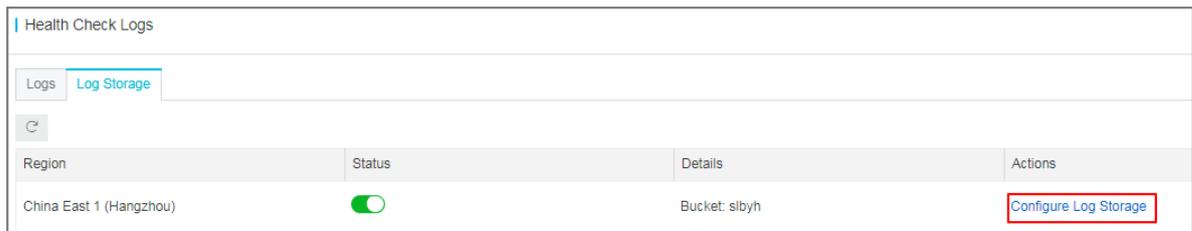


After the authorization, click **SLBLogDefaultRole**, and then click **Role Authorization Policies** to view the attached policy.

SLBLogDefaultRole			Edit Authorization Policy
Authorization Policy Name	Description	Type	Actions
AliyunOSSFullAccess	Provides full access to Object Storage Service(OSS) via Management Console.	System	View Permissions Revoke Authorization

Step 3 Configure log storage

1. Log on to the [SLB console](#).
2. In the left-side navigation pane, click **Logs > Health Check Logs**.
3. On the **Health Check Logs** page, click the **Log Storage** tab.
4. Click **Configure Log Storage** link of the target region.



5. In the **Configure Log Storage** dialog box, select a bucket to store health check logs, and then click **OK**.
6. Click the switch in the status column to enable log storage.

View health check logs

To view the health check logs generated less than three days before, complete these steps:

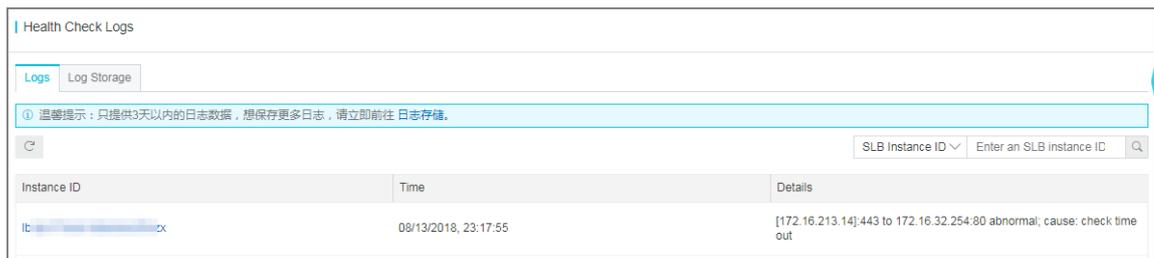
1. Log on to the [SLB console](#).
2. In the left-side navigation pane, click **Logs > Health Check Logs**.
3. On the **Health Check Logs** page, click the **Logs** tab.



Note:

Health check logs are generated only when the health status of a backend server is abnormal. Health check logs are generated every one hour. If no failure occurs to all the backend servers in an hour, no health check logs are generated in that hour.

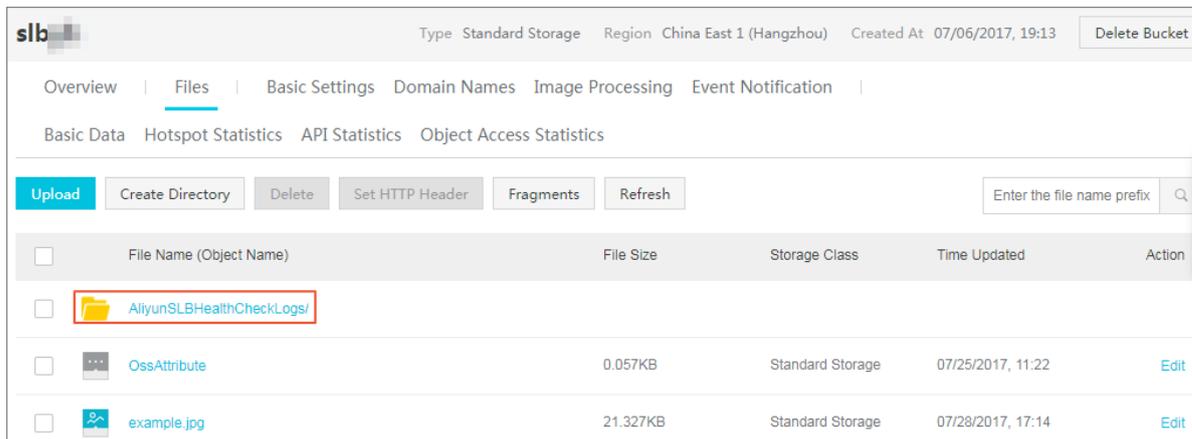
- The `SLB_instance_IP:port to Added_ECS_instance_IP:port abnormal; cause:XXX` log message indicates that the backend server is abnormal. Troubleshoot according to the detailed error message.
- The `SLB_instance_IP:port to Added_ECS_instance_IP:port normal` log message indicates that the backend server becomes normal again.



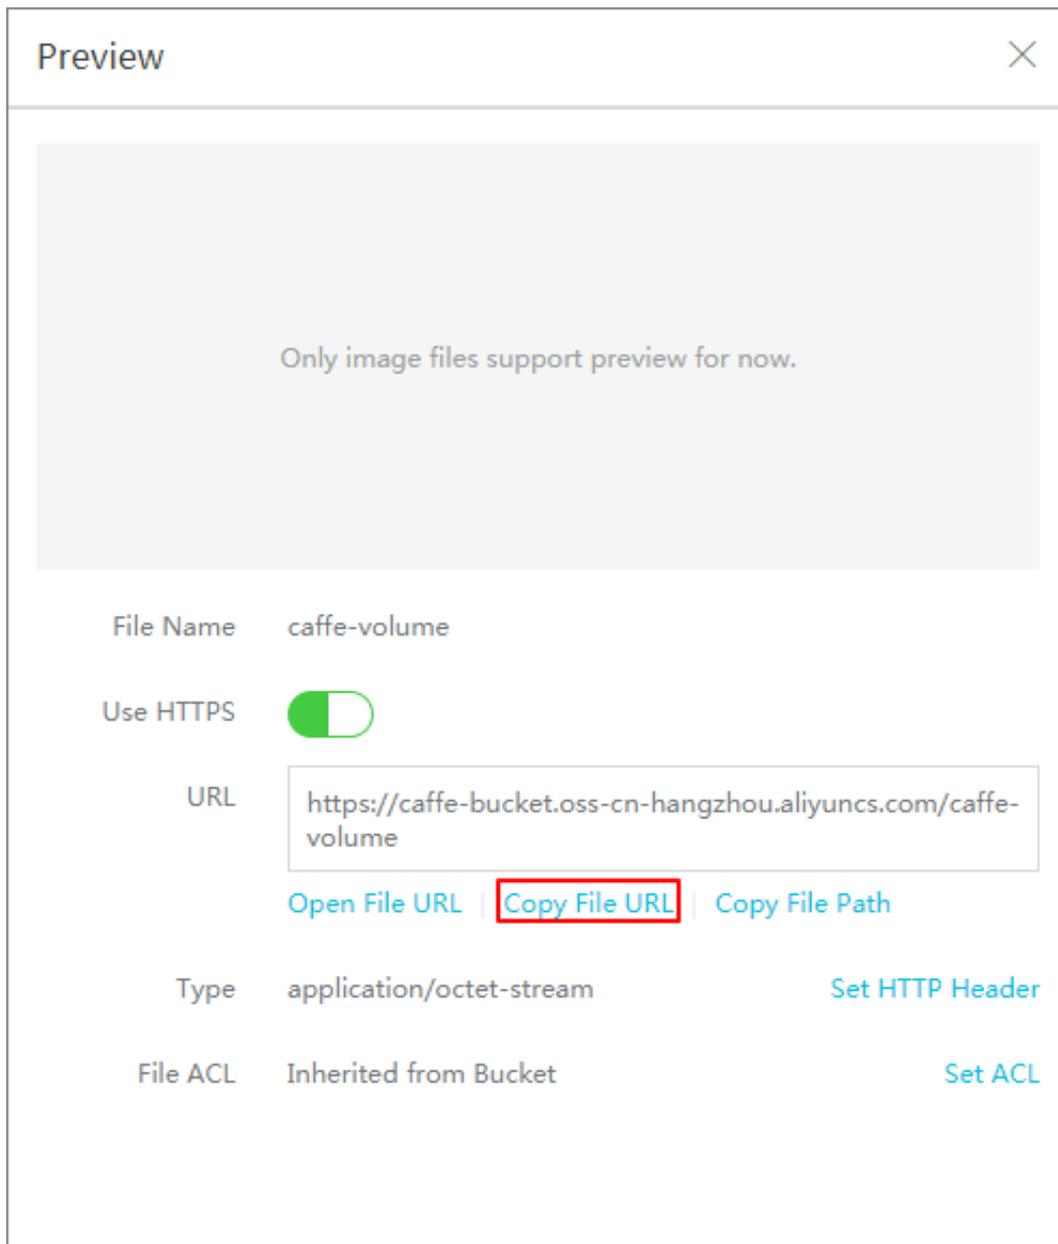
Download health check logs

You can download the completed health check logs stored in OSS.

1. Log on to the OSS console.
2. On the **Overview** page, click the target bucket and then click **Files**.
3. On the **Files** page, click *AliyunSLBHealthCheckLogs/*.



4. Click the folder of the health logs to download.
5. Click **Edit** of the target folder. Then, click **Copy File URL** in the displayed page.



6. Enter the copied URL in the web browser to download the logs.

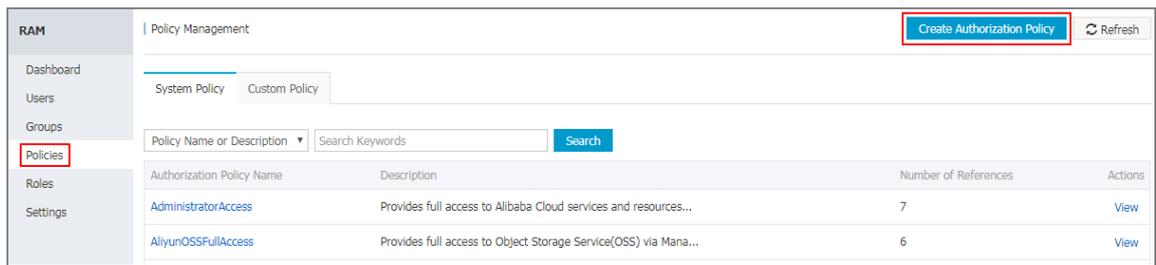
6.3 Authorize a RAM user to configure access logs

Before a RAM user starts to use the access log function, the RAM user must be authorized by the primary account.

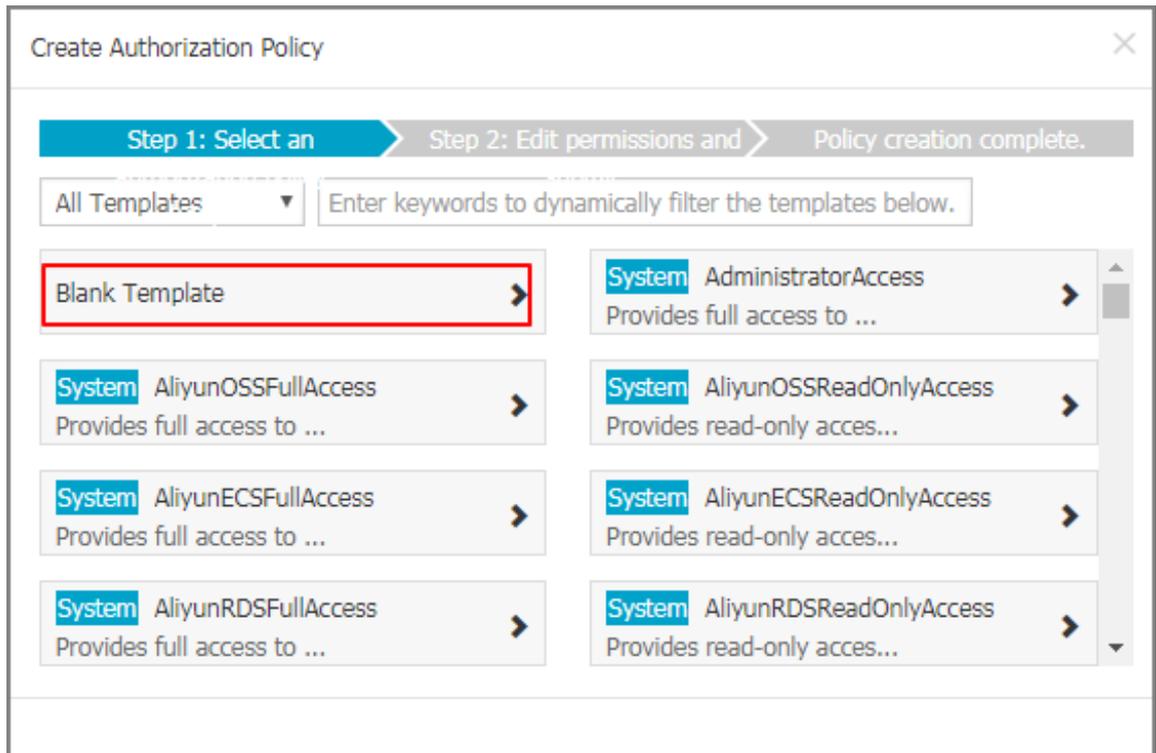
Procedure

1. Create an authorization policy:

- a) Use the primary account to log on to the RAM console.
- b) In the left-side navigation pane, click **Policies**, and then click **Create Authorization Policy**.



c) Click **Blank Template**.



d) Enter a policy name, such as **SlbAccessLogPolicySet**, and then enter the following policy. Click **Create Authorization Policy**.

```

{
  "Statement": [
    {
      "Action": [
        "slb:Create*",
        "slb:List*"
      ],
      "Effect": "Allow",
      "Resource": "acs:log:*:*:project/*"
    },
    {
      "Action": [
        "log:Create*",
        "log:List*"
      ],
      "Effect": "Allow",
      "Resource": "acs:log:*:*:project/*"
    }
  ],
}
```

```

    "Action": [
      "log:Create*",
      "log:List*",
      "log:Get*",
      "log:Update*"
    ],
    "Effect": "Allow",
    "Resource": "acs:log:*:*:project/*/logstore/*"
  },
  {
    "Action": [
      "log:Create*",
      "log:List*",
      "log:Get*",
      "log:Update*"
    ],
    "Effect": "Allow",
    "Resource": "acs:log:*:*:project/*/dashboard/*"
  },
  {
    "Action": "cms:QueryMetric*",
    "Resource": "*",
    "Effect": "Allow"
  },
  {
    "Action": [
      "slb:Describe*",
      "slb>DeleteAccessLogsDownloadAttribute",
      "slb:SetAccessLogsDownloadAttribute",
      "slb:DescribeAccessLogsDownloadAttribute"
    ],
    "Resource": "*",
    "Effect": "Allow"
  },
  {
    "Action": [
      "ram:Get*",
      "ram:ListRoles"
    ],
    "Effect": "Allow",
    "Resource": "*"
  }
],
"Version": "1"
}

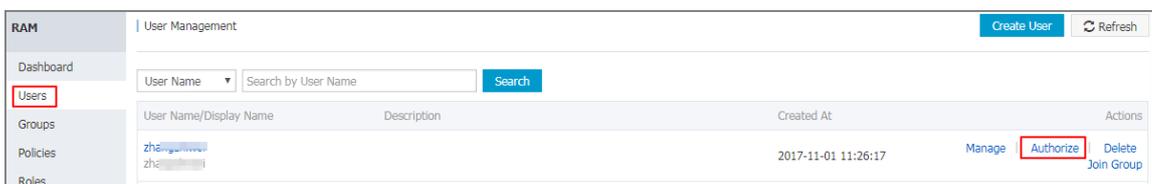
```

a) Click **Close**.

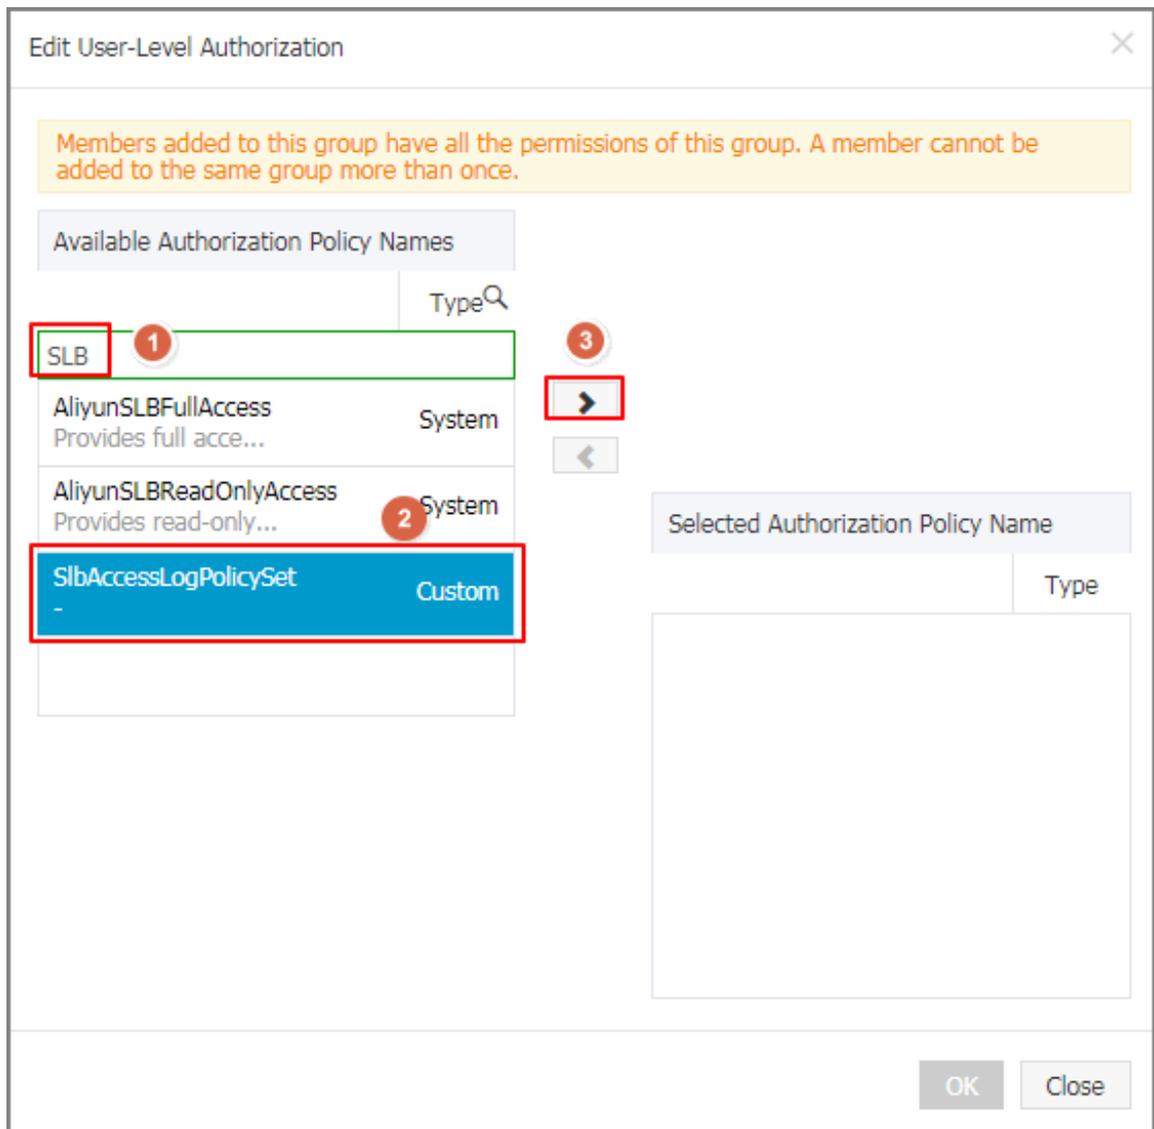
2. Attach the created policy to a RAM user:

a) In the left-side navigation pane, click **Users**.

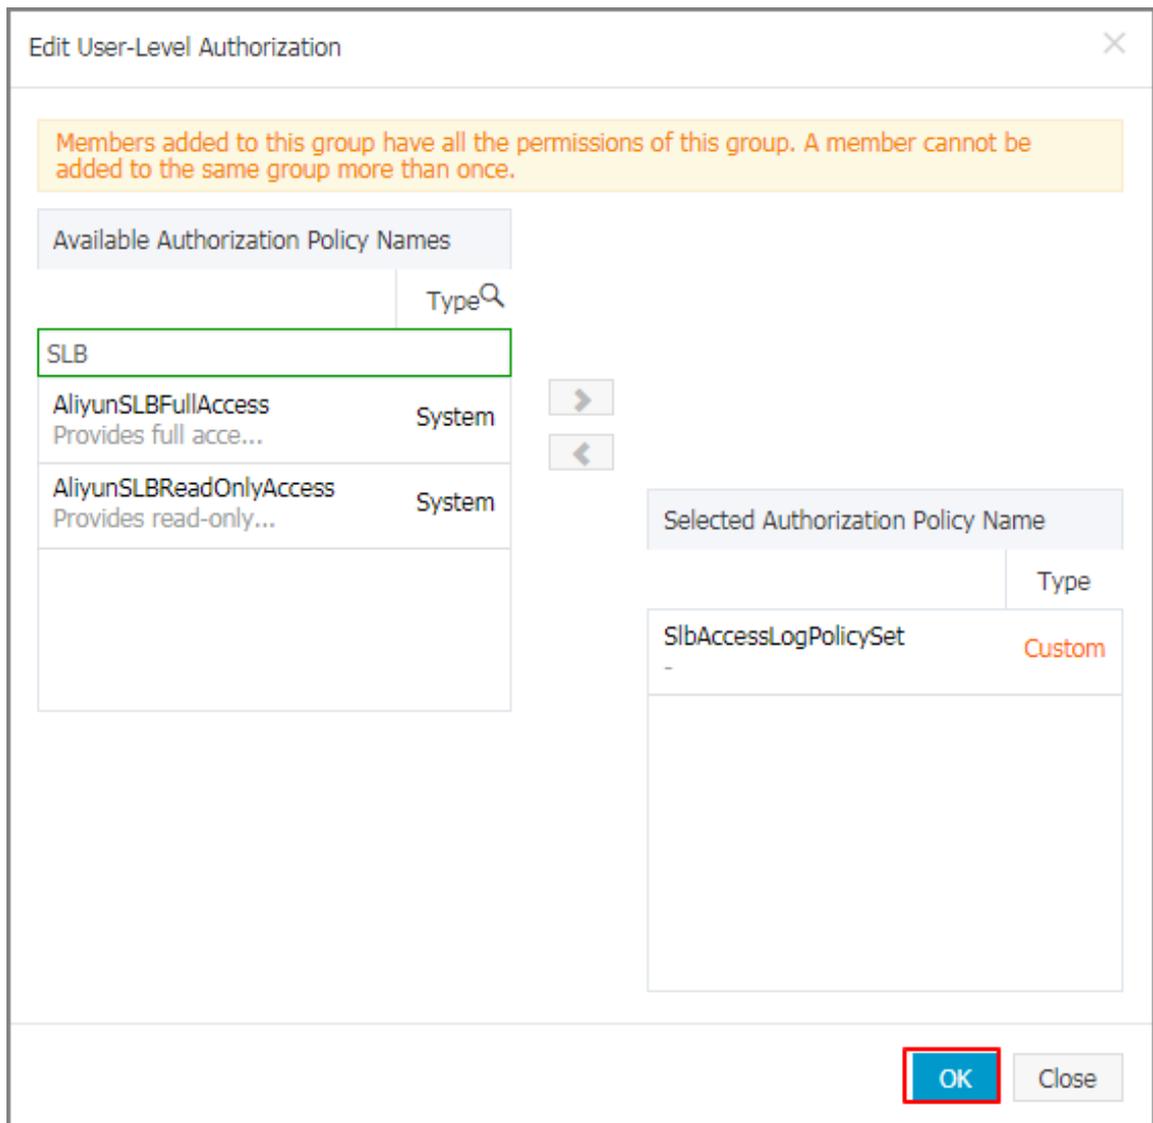
b) Find the target user (the user who uses the SLB Access Log function) and click **Authorize**.



- c) Search the created authorization policy and attach the policy to the RAM user.



- d) Click **OK**.



- e) Go to the User Authorization Policies page to check if the policy has been attached to target RAM user.



6.4 Configure access logs

Integrated with Alibaba Cloud Log Service, you can understand the behavior and geographical distribution of client users, troubleshoot problems by analyzing the access logs of a Server Load Balancer instance.

What are access logs?

The access log collects detailed information of all requests sent to a Server Load Balancer instance, including the request time, client IP address, latency, request URL, server response, and so on. As the entry of Internet access, Server Load Balancer receives massive client requests. You can use access logs to analyze user behavior and geographical distribution, and troubleshoot.

After you enable the SLB access logs, you can store the access logs in the Logstore of SLS to collect and analyze the access logs. You can also disable access logs at any time.

There is no extra fee for Server Load Balancer access logs. You only need to pay for the Log Service.



Note:

- Only Layer-7 Server Load Balancer supports configuring access logs and this function is available in all regions now.
- Make sure that the HTTP header value does not contain | |, otherwise, the exported logs may be misplaced.

Benefits

The following are benefits of Server Load Balancer access logs:

- Easy to use

Free developers and maintenance staff from tedious and time-consuming log processing so that they can concentrate on business development and technical research.

- Cost-effective

Access logs are typically very large. Processing access logs takes a lot of time and consumes a lot of resources. With Log Service, the access log processing is faster and cost-effective than self-build open-source solutions. Log Service can analyze one hundred million logs in one second.

- Real-time

Scenarios such as DevOps, monitoring, and alerting require real-time log data. Traditional data storage and analysis tools cannot meet this requirement. For example, it takes long time to ETL data to Hive where a lot of work is spent on data integration. Powered by its powerful computing capability, Log Service can process and analyze access logs in seconds.

- Flexible

You can enable or disable Server Load Balancer access logs according to the instance capacity. Additionally, you can set the storage period (1 to 365 days) as needed and the Logstore's capacity is scalable to meet increasing business service demands.

Configure access logs

Before configuring access logs, make sure:

1. A Layer-7 listener is added.
2. Log Service is activated.

To configure access logs, complete these steps.

1. Log on to the [SLB console](#).
2. In the left-side navigation pane, click **Logs > Access Logs**.
3. Select a region.
4. Click **Authorize**, and then click **Confirm Authorization Policy** to authorize Server Load Balancer to write logs to Log Service.

If you are using a RAM account, the primary account is required to perform authorization. For more information, see [Authorize a RAM user to use access logging](#).

**Note:**

This operation is required only at the first time.

5. On the **Access Logs** page, find the target Server Load Balancer instance and click **Configure Logging**.
6. Select the LogProject and LogStore of Log Service and then click **OK**.

If there is no available LogStore, click **Log Service console** to create log projects.

**Note:**

Make sure that the name of the LogProject is globally unique and the region of the LogProject is the same as that of the Server Load Balancer instance.

Configure Logging
✕

ⓘ Configure layer-7 access logging.

- **LogProject**

Select
▼
- **Log Store**

Select
▼

OK
Cancel

Search and analyze access logs

After configuring Server Load Balancer access logs, you can search and view logs using the following indexing fields.

Field	Description
body_bytes_sent	The size of HTTP body (in byte) sent to the client.
client_ip	The client IP.
host	The host header in the request.
http_user_agent	The received http_user_agent header in the request.

Field	Description
request_length	The length of the request including startline, HTTP header and HTTP body.
request_method	The request method.
request_time	The interval between the time the Server Load Balancer receives the first request and the time the Server Load Balancer returns a response.
Request_uri	The URL of the received request.
Slbid	The ID of the Server Load Balancer instance.
status	The response status code sent by the SLB.
Upstream_addr	The IP address and port number of the backend server.
upstream_response_time	The interval between the time Server Load Balancer sends a request to the backend server and the time Server Load Balancer sends a response to the client.
Upstream_status	The response status code of the backend server received by SLB.

Search access logs

To search access logs, complete these steps:

1. Go to the log search page. You can navigate to the search page from the Server Load Balancer console or the Log Service Console:

- From the Server Load Balancer console:

On the **Access Logs** page, click **View Logs**.

Instance Name/ID	IP Address	Network Type	Status	Storage Path	Actions
wh2myb8pk	24...223(Public Network)	Classic Internal Network	Active	--	Configure Logging
adrahj78u	24...14(Public Network)	Classic Internal Network	Active	--	Configure Logging
6nzx	1...4(VPC)	VPC	Active	--	Configure Logging
miup	11...2(Public Network)	Classic Internal Network	Inactive	--	Configure Logging
sf	...2(Public Network)	Classic Internal Network	Active	--	Configure Logging
ms	...10(VPC)	VPC	Active	--	Configure Logging
HTTPS lb-q6k	1...34(Public Network)	Classic Internal Network	Active	www1111/www	View Logs Delete

- Log Service Console

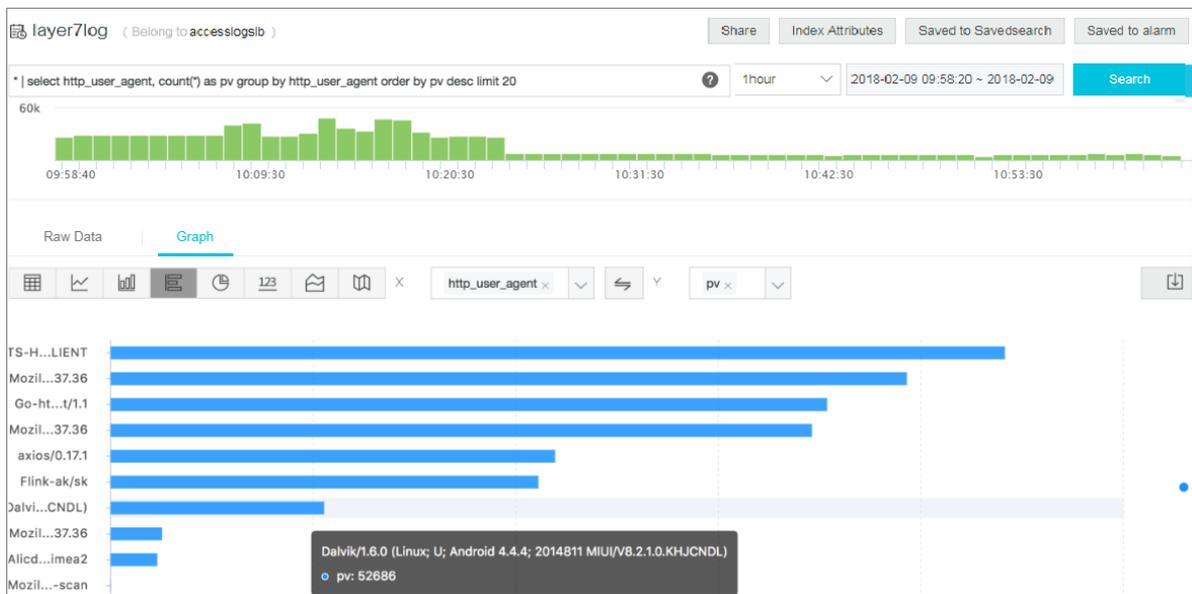
On the **Logstores** page, click **Search** of the target Logstore.

2. Click the target log field to view detailed information.

3. Enter an SQL statement to query access logs.

For example, enter the following SQL statement to query the client of Top20, which is used for analyzing the request resource to assist business decision-making.

```
* | select ip_to_province(client_ip) as client_ip_province, count
(*) as pv group by
client_ip_province order by pv desc limit 50
```



Analyze access logs

You can analyze access logs through the dashboard, which provides various graphic information.

To analyze access logs, complete these steps:

1. On the Log Service console, click the project link of the SLB instance.
2. In the left-side navigation pane, click **Search/Analytics - Query > Dashboard**, and then click the name of the access log.

Disable access logging

To disable access logging, complete these steps:

1. Log on to the *SLB console*.
2. In the left-side navigation pane, click **Logs > Access Logs**.
3. Select a region.
4. On the **Access Logs** page, find the target instance and click **Delete** to disable access logging.

Access Logs						
						SLB Instance ID <input type="text" value="Enter a value"/>
<input type="checkbox"/>	Instance Name/ID	IP Address <input type="text"/>	Network Type <input type="text"/>	Status <input type="text"/>	Storage Path	Actions
<input type="checkbox"/>	lb-2wh2myb8pk	24.23(Public Network)	Classic Internal Network	Active	--	Configure Logging
<input type="checkbox"/>	lb-2drahj78u	24.14(Public Network)	Classic Internal Network	Active	--	Configure Logging
<input type="checkbox"/>	lb-26nzc	1.4(VPC)	VPC	Active	--	Configure Logging
<input type="checkbox"/>	lb-2miup	11.2(Public Network)	Classic Internal Network	Inactive	--	Configure Logging
<input type="checkbox"/>	lb-2sf	1.2(Public Network)	Classic Internal Network	Active	--	Configure Logging
<input type="checkbox"/>	lb-2ms	1.10(VPC)	VPC	Active	--	Configure Logging
<input type="checkbox"/>	HTTPS lb-2q6k	1.34(Public Network)	Classic Internal Network	Active	www1111/www	View Logs Delete

5. In the displayed dialog box, click **OK**.

7 Access control

7.1 Configure an access control list

Server Load Balancer (SLB) provides you with the access control function. You can configure different access control rules (access whitelist or blacklist) for different listeners. Before configuring access control for listeners, you must first configure an access control list.

You can create multiple access control lists. Each list contains multiple IP addresses or CIDR blocks. Limits on access control lists are as follows:

Resource	Limit
The maximum number of access control lists per region.	50
The maximum number of IP addresses added each time.	50
The maximum number of entries per access control list.	300
The maximum number of listeners that an access control list can be added to	50

Create an access control list

To create an access control list, complete these steps:

1. Log on to the [SLB console](#).
2. Select a region.
3. In the left-side navigation pane, click the **Access Control** tab.
4. Click **Create Access Control List**, enter the access control list name, and click **OK**.

Add IP entries

To add IP entries, complete these steps:

1. Log on to the [SLB console](#).
2. Select a region.
3. In the left-side navigation pane, click the **Access Control** tab.
4. Find the target access control list and click **Manage**.
5. Add IP entries:

- Click **Add Multiple Entries**. In the displayed dialog box, add IP addresses or CIDR blocks and click **OK**.

Note the following when adding entries:

- Each line is one entry. Use the Enter key to break lines.
- Use “|” to separate an IP address or CIDR block with the description. For example, “192.168.1.0/24|description”.

Add Multiple IP Entries

Descriptions:

1. One line for each entry. Start a new line by pressing Enter.
2. For each entry, the IP address/IP CIDR block and description should be delimited by a vertical bar (|). For example, 192.168.1.0/24|Description.

Add Multiple Addresses and Descriptions

OK Cancel

- Click **Add Entry**. In the displayed dialog box, add an IP address or CIDR block and the description, and click **OK**.

Add IP Entry

ⓘ Either an IPv4 address or an IPv4 CIDR block. For example, 192.168.1.1 or 192.168.1.1/32.
An IPv4 CIDR block. For example, 192.168.1.0/24.

• **IP Address/IP CIDR Block**

Description

Contact Us

OK Cancel

Delete IP entries

To delete IP entries, complete these steps:

1. Log on to the [SLB console](#).
2. Select a region.
3. In the left-side navigation pane, click the **Access Control** tab.
4. Find the target access control list and click **Manage**.
5. Click **Delete** in the **Actions** column of the target IP entry, or select multiple IP entries and click **Delete** at the bottom of the entry table.
6. In the displayed dialog box, click **OK**.

7.2 Configure access control

Server Load Balancer allows you to configure access control for listeners. You can configure different whitelists or blacklists for different listeners.

You can configure access control when you create a listener or change access control configuration after a listener is created.

This document introduces how to configure access control after a listener is created.

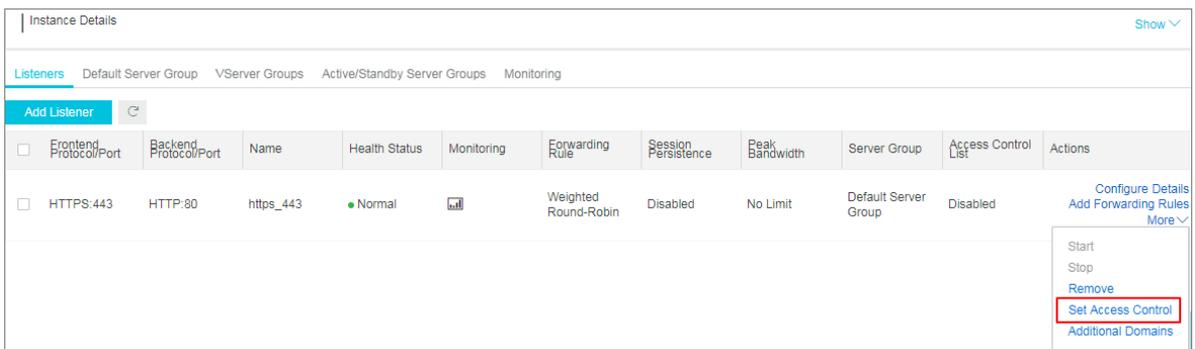
Enable access control

Before enabling access control, make sure:

- You have created an access control list. For more information, see [Configure an access control list](#).
- You have created a listener.

To enable access control, complete these steps:

1. Log on to the [SLB console](#).
2. Select a region.
3. Click the ID of the target SLB instance.
4. On the **Instance Details** page, click the **Listeners** tab.
5. Find the target listener, and then click **More > Set Access Control**.



6. On the **Access Control Settings** page, enable access control, select the access control method and access control list, and click **OK**.

- **Whitelist:** Only requests from IP addresses or CIDR blocks in the selected access control lists are forwarded. It applies to scenarios where the application only allows access from some specific IP addresses.

Enabling whitelist poses some business risks. After a whitelist is configured, only the IP addresses in the list can access the listener. If you enable the whitelist without adding any IP entry in the corresponding access control list, all requests are forwarded.

- **Blacklist:** Requests from IP addresses or CIDR blocks in the selected access control lists are not forwarded. It applies to scenarios where the application only denies access from some specific IP addresses.

If blacklist access is on, but no IP is added to the access policy group, the load balancing listener forwards all requests.

Disable access control

To disable access control, complete these steps:

1. Log on to the [SLB console](#).
2. Select a region.
3. Click the ID of the target SLB instance.
4. On the **Instance Details** page, click the **Listeners** tab.
5. Find the target listener, and then click **More > Set Access Control**.
6. On the **Access Control Settings** page, disable access control and click **OK**.

7.3 Migrate to the new access control

If you have already configured a whitelist for a listener, Server Load Balancer can automatically add the IP addresses or CIDR blocks in the whitelist to an access control list and apply the list to the listener.

Migrate a whitelist to an access control list

To migrate a previously configured whitelist to an access control list, complete these steps:

1. Log on to the [SLB console](#).
2. Select the region where the SLB instance is located, and then click the ID of the target SLB instance.
3. Click the **Listeners** tab.
4. Find the target listener, select **More > Set Access Control**.
5. Click **Use New Access Control Features**.
6. Enter a name of the access control list and click **Create Access Control List**.
7. Click **Apply** to apply the list to the listener as a whitelist.

**Note:**

If you do not apply the list to a listener, the whitelist does not take effect.

View the migrated access control list

To view the migrated access control list, complete these steps:

1. Log on to the [SLB console](#).
2. Select a region.
3. In the left-side navigation pane, click **Access Control**.
4. Find the created access control list and view the associated listener. You can click **Manage** to manage IP entries.

7.4 Configure a whitelist

Whitelist is a method used to control the access of Server Load Balancer. It applies to scenarios where an application only allows access from some specific IP addresses.

Context**Note:**

SLB has released a new version of the access control function, allowing you to configure both whitelists and blacklists. You can migrate the previously configured whitelist to the new version. For more information, see [Migrate to the new access control](#).

Note the following when configuring whitelists:

- Enabling whitelist poses some business risks. After a whitelist is configured, only the IP addresses in the list can access Server Load Balancer.
- If you enable the whitelist function without adding any IP entry in the corresponding access control list, no requests are forwarded.
- When you configure a whitelist, the access to Server Load Balancer may be interrupted for a short time.

Procedure

1. Log on to the [SLB console](#).
2. Select the region where the target SLB instance is located.
3. Click the ID of the target SLB instance.
4. Under the **Listeners** tab, select **More > Set Access Control**.

5. In the displayed dialog box, configure as follows:

a) Click the **Enable Access Control** switch.

b) Enter the IP addresses that are allowed to access the listener.

Separate multiple IP addresses using comma. Up to 40 IP addresses are allowed. You can also enter CIDR blocks.

c) Click **OK**.

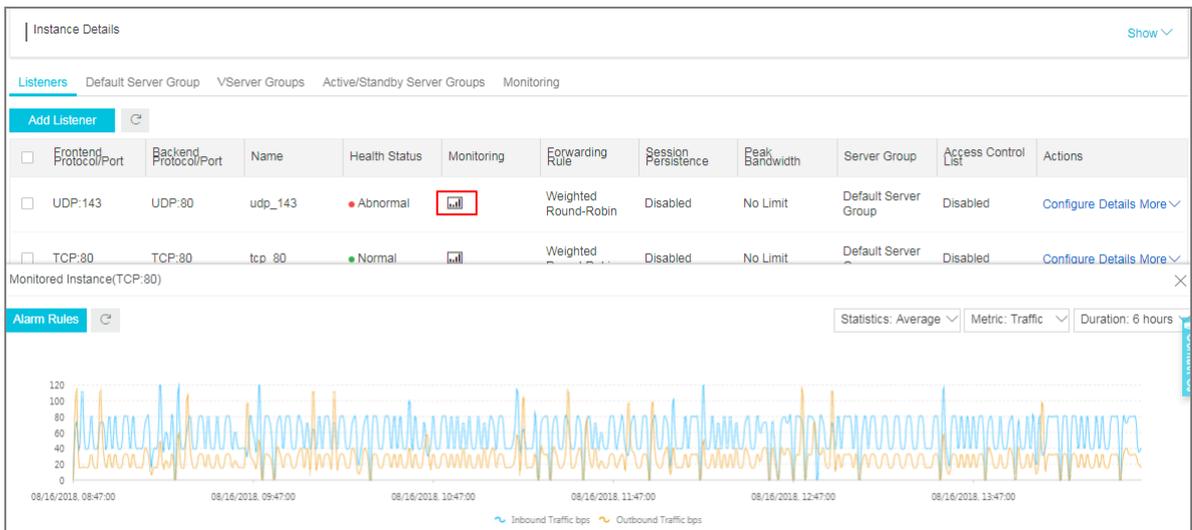
8 Monitoring

8.1 View monitoring data

With the CloudMonitor service, you can view the number of connections, and other traffic information of SLB listeners.

Procedure

1. Log on to the [SLB console](#).
2. Select the region where the SLB instance is located.
3. Click the monitoring icon  next to the target SLB instance.
4. Select the monitor metrics that you want to view.



SLB supports viewing the following monitor metrics.

Monitor metrics	Description
Traffic	<ul style="list-style-type: none"> • Inbound Traffic: The traffic consumed by external access. • Outbound Traffic: The traffic consumed by Server Load Balancer.
Packets	<ul style="list-style-type: none"> • RX Packets Count: The number of request packets received per second. • TX Packets Count: The number of response packets sent per second.

Monitor metrics	Description
Concurrent Connections	<ul style="list-style-type: none"> Active Connections Count: The number of established TCP connections. If persistent connections are used, a connection can transfer multiple file requests at one time. Inactive Connections Count: The number of TCP connections that are not in the established status. You can use <code>netstat -an</code> command to view the active connections. Max Concurrent Connections Count: The total number of TCP connections.
Average Connection Requests Count	The average number of new TCP connections established between clients and the Server Load Balancer in the statistical period.
Dropped Traffic	<ul style="list-style-type: none"> Dropped Inbound Traffic: The amount of inbound traffic dropped per second. Dropped Outbound Traffic: The amount of outbound traffic dropped per second.
Dropped Packets	<ul style="list-style-type: none"> Dropped RX Packets: The number of inbound packets dropped per second. Dropped TX Packets: The number of outbound packets dropped per second.
Dropped Connections	The number of TCP connections dropped per second.
The following metrics are specific to Layer-7 listeners.	
Layer-7 Protocol QPS	The number of HTTP/HTTPS requests that can be handled per second.
Response Time (Listener)	The average response time of Server Load Balancer.
HTTP Status Code 2XX/3XX/4XX/5XX/Others (Listener)	The average number of HTTP response codes generated by the listener.
Response Code 4xx/5xx (Server)	The average number of HTTP response codes generated by the backend server.
Response Time (Server)	The average response time of the backend server.

8.2 Configure alarm rules

After activating the CloudMonitor service, you can configure alarm rules for SLB instances on the CloudMonitor console.

Context



Note:

If a listener or an SLB instance is deleted, its alarm settings are deleted correspondingly.

Procedure

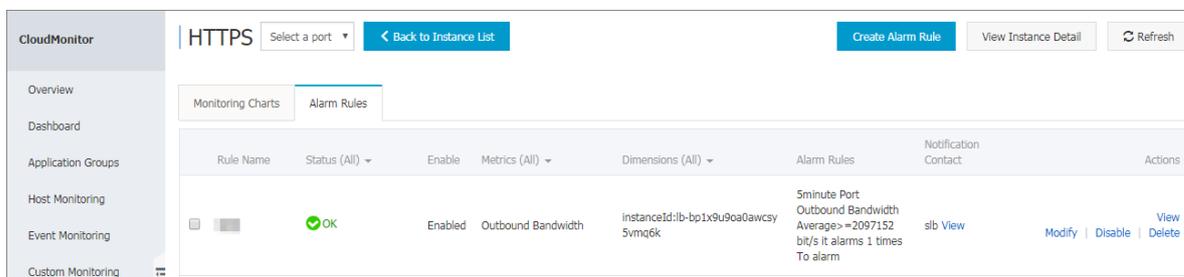
1. Log on to the [SLB console](#).
2. Select the region where the SLB instance is located.
3. Find the target instance and click .



Note:

Make sure the instance has configured with listeners and enabled health check.

4. Click **Alarm Rule**. You are then directed to CloudMonitor console.



5. Click **Create Alarm Rule**.
6. Configure the alarm rule.

1 Related Resource

Products :

Resource Range : When selecting an application group, you can use an alarm template. Click [View alarm template best practices](#).

Region :

Instances :

2 Set Alarm Rules

Alarm Rule :

Rule Describe : unit

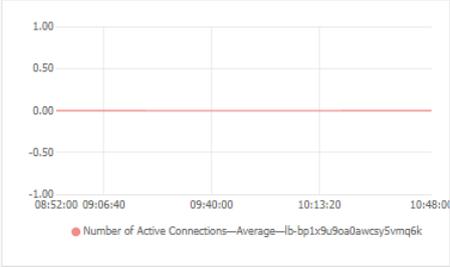
Port :

[+Add Alarm Rule](#)

Mute for :

Triggered when threshold is exceeded for :

Effective Period : To:



Time	Value
08:52:00	0.00
09:06:40	0.00
09:40:00	0.00
10:13:20	0.00
10:48:00	0.00

9 Multiple zone deployment

When creating SLB instances, you can create SLB instances in the region with multiple zones to improve the availability.

What is multiple zone deployment?

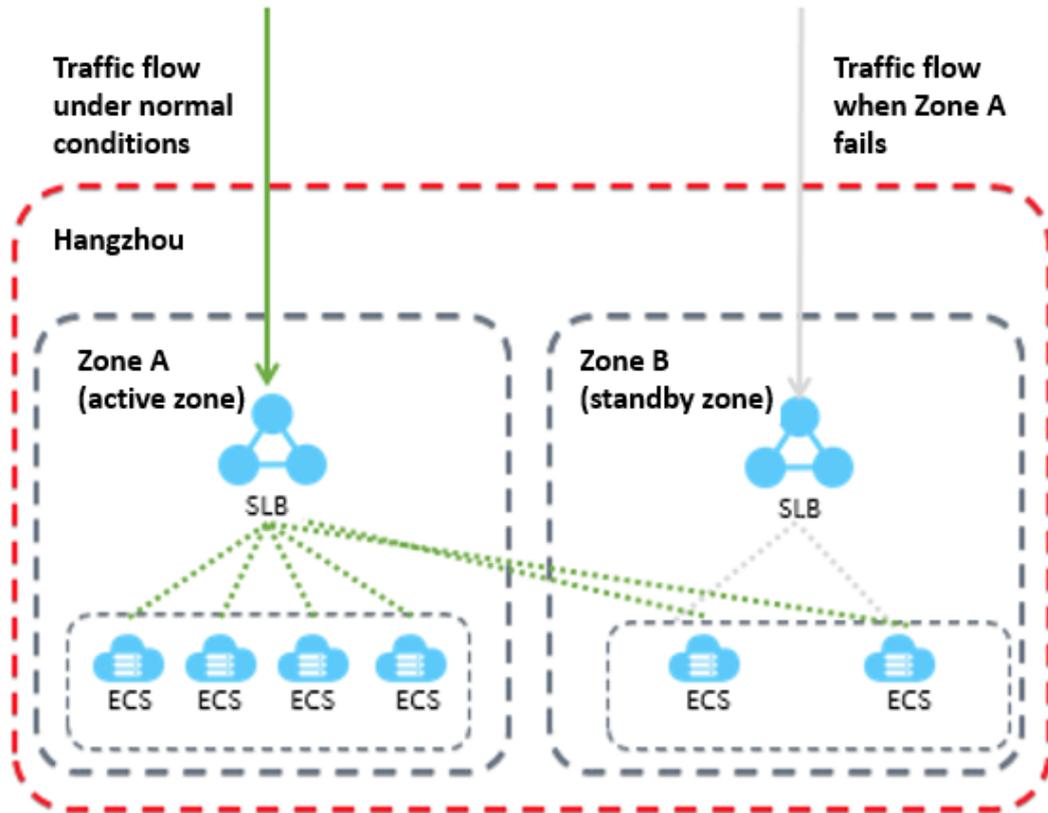
A cloud product zone refers to a set of independent infrastructures. Different zones have independent infrastructures (such as network, power supply and air-conditioning), thus an infrastructure fault in one zone does not affect other zones.

To provide more reliable services, SLB has deployed multiple zones in most regions to achieve disaster recovery across data centers. When the data center in the primary zone is faulty and unavailable, SLB is able to switch to the data center in the backup zone to restore its service capabilities within 30 seconds.

Note the following about SLB primary/backup zones:

- SLB supports attaching ECS instances in different zones as long as the ECS instances and the SLB instance are in the same region. SLB can distribute traffic to the ECS instances in different zones.
- Normally, the SLB instance in the backup zone is in the backup state. You cannot manually switch the primary/backup state of an SLB instance. SLB will switch to the backup zone only when the data center of the primary zone is unavailable such as outage. SLB will not switch to the backup zone if an SLB instance is faulty.
- SLB and ECS are deployed in different clusters. When an SLB instance is unavailable, the ECS instance is still available. Therefore, after SLB switches to the backup zone, the SLB instance in the backup zone still can distribute traffic to the added ECS instances. However, if all clusters in a zone is unavailable or the optical cable is broken, then all the services in the zone including SLB and ECS cannot work anymore.

For more information, see [SLB high availability](#).



Primary/backup zone list

The following table lists the primary/backup zones in different regions. You can call the `DescribeZones` API to obtain available primary/backup zones in a region.

Region	Zone type	Zones	
China (Hangzhou)	Multi-zone	Primary zone	Backup zone
		Zone B	Zone D
		Zone D	Zone E
		Zone E	Zone F
		Zone F	Zone E
China (Shanghai)	Multi-zone	Primary zone	Backup zone
		Zone A	Zone B
		Zone B	Zone A or Zone D
		Zone C	Zone B
		Zone D	Zone B

Region	Zone type	Zones	
China (Shenzhen)	Multi-zone	Primary zone	Backup zone
		Zone A	Zone B
		Zone B	Zone A
		Zone C	Zone B
China (Qingdao)	Multi-zone	Primary zone	Backup zone
		Zone B	Zone C
		Zone C	Zone B
China (Beijing)	Multi-zone	Primary zone	Backup zone
		Zone A	Zone B or Zone D
		Zone B	Zone A or Zone C
		Zone C	Zone B
		Zone D	Zone A
China (Zhangjiakou)	Single-zone	Zone A	Zone A
China (Hohhot)	Single-zone	Zone A	Zone A
Germany (Frankfurt)	Single-zone	Zone A	Zone A
UAE (Dubai)	Single-zone	Zone A	Zone A
Singapore	Multi-zone	Primary zone	Backup zone
		Zone A	Zone B
		Zone B	Zone A
Australia (Sydney)	Single-zone	Zone A	Zone A
Malaysia (Kuala Lumpur)	Single-zone	Zone A	Zone A
Japan (Tokyo)	Single-zone	Zone A	Zone A

Region	Zone type	Zones	
Hong Kong	Multi-zone	Primary zone	Backup zone
		Zone B	Zone C
		Zone C	Zone B
US (Virginia)	Single-zone	Zone A	Zone A
US (Virginia)	Multi-zone	Primary zone	Backup zone
		Zone A	Zone B
		Zone B	Zone A

10 Achieve cross-region load balancing through Global Traffic Manager

Global traffic management

Load balancing is divided into local load balancing and global load balancing according to the geographical structure of its application. Local load balancing balances loads of server groups in the same region. Global load balancing balances server groups that are in different regions and have different network structures.

By using Global Traffic Manager, you can deploy global traffic management above the local traffic balancing to achieve cross-region disaster tolerance, accelerate access from different regions and achieve intelligent resolution.

- Multi-line intelligent resolution service

Using DNS intelligent resolution and health check on the running status of the application, global traffic management directs user accesses to the most appropriate IP addresses so that the users can obtain the fastest and smoothest experience.

- Cross-region disaster tolerance

Global traffic management supports adding IP addresses of different regions to different address pools and configuring health check. In access policy configurations, set the **address pool A** as the default IP address pool and **address pool B** as the failover IP address pool. Then active-standby IP disaster tolerance of the application service can be achieved.

- Accelerate accesses from different regions

By using Global Traffic Manager, you can direct user accesses from different regions to different IP address pools, thus achieving grouped user management and grouped access and helping the application improve user experience.

Deploy Global Traffic Manager

This tutorial takes aliyuntest.club as an example (most users of the website are in Singapore and China) to show how to achieve global load balancing through global traffic management and load balancing.

Step 1 Purchase and configure ECS instances

Purchase and configure at least two ECS instances in each region where the users of the application of the application service are located.

In this tutorial, two ECS instances are purchased in Beijing, Shenzhen and Shanghai separately, and a simple static web page is built on each ECS instance.

- Example of ECS instances in the Beijing region
- Example of ECS instances in the Shenzhen region
- Example of ECS instances in the Singapore region

Step 2 Purchase and configure Server Load Balancing instances

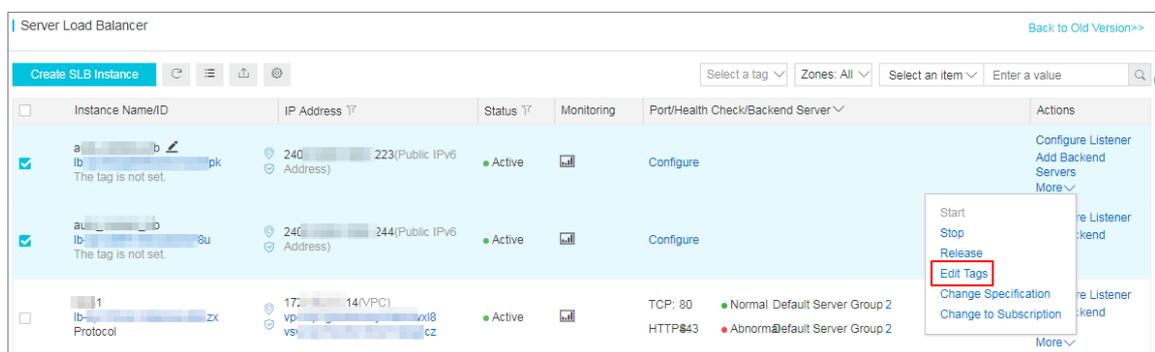
1. Refer to [Create an SLB instance](#) to create three Internet Server Load Balancing instances in Beijing, Shenzhen and Singapore separately.
2. Refer to [Configure an SLB instance](#) to add listeners and add configured ECS instances to backend server pools.

- Example of the SLB instance in the Beijing region
- Example of the SLB instance in the Shenzhen region
- Example of the SLB instance in the Singapore region

Step 3 Configure Global Traffic Manager

1. Purchase a Global Traffic Manager instance.
 - a. Log on to the [Alibaba Cloud DNS console](#).
 - b. In the left-side navigation pane, click **Global Traffic Manager**.
 - c. On the **Global Traffic Manager** page, click **Create Instance**.
 - d. Select the version, purchase quantity and service time.
 - e. Click **Buy now**.

After the instance is successfully purchased, the system automatically allocates a CNAME access domain name.



2. Configure the Global Traffic Manager instance.

- a. On the **Global Traffic Manager** page, click the ID of the Global Traffic Manager instance or click **Configure** in the **Actions** column.
- b. In the left-side navigation pane, click **Configurations**.
- c. In the **Global Settings** tab, click **Edit** to configure the parameters of the Global Traffic Manager instance.

Configure the following parameters and use the default values for the remaining options.

- **Instance Name:** Used for identifying the instance used for a certain application and can be customized.
- **Primary Domain:** The primary domain name is used by you to access the application. In this tutorial, enter aliyuntest.club.
- **Alert Group:** When the global traffic management generates abnormal sound, the message sender is notified and the alert contact groups you create in CloudMonitor are automatically obtained.

- d. Click **Confirm**.

3. Configure the IP address pool.

- a. In the **IP Address Pool Configurations** tab, click **Create Address Pool**.
- b. On the **Create Address Pool** page, configure the IP address pool.

In this tutorial, three IP address pools are to be added and each IP address pool accommodates one of the three SLB addresses in different regions.

- **Address Pool Name:** Custom. For example, China North_Beijing, China East_Hangzhou, Singapore.
- **Address:** The SLB public IP address to be added to this region.

Create Address Pool
✕

*** Address Pool Name:**

*** Address Pool Type** (?)

*** Minimum Available Addresses** (?)

Address	Mode
<input style="width: 100%; height: 20px;" type="text"/>	Smart Return ▼

+ New Row

c. Click **Confirm**.

4. Configure health check

In this operation, you need to configure health check for the three address pools separately.

- a. In the **Address Pool** tab, click **Edit** next to health check.
- b. Configure health check parameters.

Monitoring Node shows the locations of monitoring nodes. Select the monitoring node according to the region of the address pool.

5. Configure the access policy.

In this tutorial, add different access policies for the three different regions.

- a. In the **Access Policy** tab, click **Add Access Policy**.
- b. On the **Add Access Policy** tab, configure the access policy.
 - Configure the corresponding default address pools for different access regions, and set an address pool of another region as the failover address pool.
 - Select the access region. When users in this region access the application, the address pool configured in the access policy is matched.

There must be an access policy with **Global** selected. Otherwise some areas cannot access the application.

6. Configure CNAME access.
 - a. Log on to the Alibaba Cloud DNS console.
 - b. Find the domain name aliyuntest.club and click **Configure** in the **Actions** column.
 - c. On the **DNS Settings** page, click **Add Record**.
 - d. On the **Add Record** page, direct the domain name aliyuntest.club accessed by end users to the alias record of the Global Traffic Manager instance in the form of CNAME.

The screenshot shows the 'Add Record' dialog box with the following fields and values:

- Type: A- IPV4 address
- Host: Enteraliyuntest.club
- ISP Line: Default - Return to the default value when the query is not ...
- * Value: Enter ...
- * TTL: 10 minute(s)
- Synchronize the Default Line

Buttons: Cancel, OK

- e. Click **Confirm**.

Step 4 test

Remove the ECS instances of the SLB instance in the Beijing region so as that the SLB service is unavailable.

Visit the website to see if the access is normal.

**Note:**

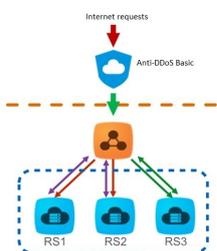
It takes one to two minutes for Global Traffic Manager to make judgement after it monitors that your IP is down. If you set the monitoring frequency to 1 minute, it takes two to three minutes for the link switching caused by exceptions to take effect.

11 Anti-DDoS Basic

You can view Alibaba Cloud Security thresholds of an Internet SLB instance on the SLB console.

Introduction to Anti-DDoS Basic

Alibaba Cloud provides up to 5 Gbps basic anti-DDoS protection for SLB. As shown in the following figure, all traffic from the Internet must first go through Alibaba Cloud Security before arriving at SLB. Anti-DDoS Basic cleans and filters common DDoS attacks and protects your services against attacks such as SYN flood, UDP flood, ACK flood, ICMP flood, and DNS Query flood.



Anti-DDoS Basic sets the cleaning threshold and blackholing threshold according to the bandwidth of the Internet SLB instance. When the inbound traffic reaches the threshold, the cleaning or blackholing is triggered:

- **Cleaning:** When the attack traffic from the Internet exceeds the cleaning threshold or matches certain attack traffic model, Alibaba Cloud Security starts cleaning the attack traffic. The cleaning operation includes packet filtration, traffic speed limitation, packet speed limitation and so on.
- **Blackholing:** When the attack traffic from the Internet exceeds the blackholing threshold, blackholing is triggered and all inbound traffic is dropped.

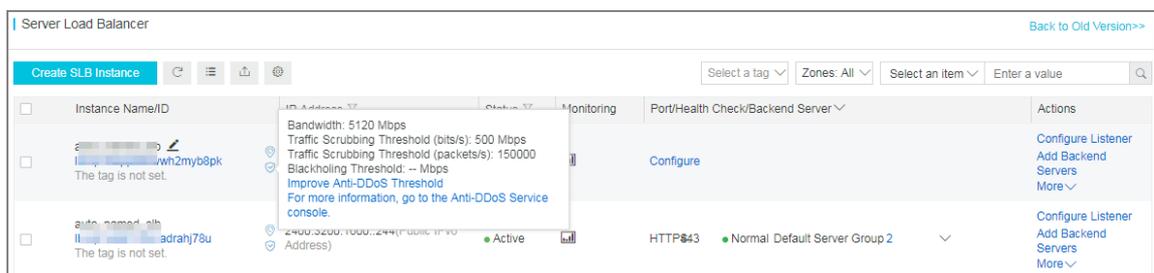
View thresholds

You can view the thresholds of an instance on the SLB console. If you cannot view the thresholds using a RAM account, ask your system administrator to grant the permission for you. For more information, see [Allow read-only access to Anti-DDoS Basic](#).

To view thresholds, complete these steps:

1. Log on to the [SLB console](#).
2. Select a region.

3. Hover the mouse pointer to the DDoS icon next to the target instance to view the following thresholds. You can click the link to go to the DDoS console to view more information.
 - Traffic Scrubbing Threshold (bits/s): When the inbound traffic exceeds the BPS cleaning threshold, cleaning is triggered.
 - Traffic Scrubbing Threshold (packets/s): When the inbound packets exceed the PPS cleaning threshold, cleaning is triggered.
 - Blackholing Threshold: When the inbound traffic exceeds the blackholing threshold, blackholing is triggered.



Allow read-only access to Anti-DDoS Basic

To allow read-only access to Anti-DDoS Basic, complete these steps:



Note:

You have to use the primary account to complete the authorization.

1. Use the primary account to log on to the RAM console.
2. In the left-side navigation pane, click **Users**, find the target RAM account and click **Manage**.



3. Click **User Authorization Policies**, and then click **Edit Authorization Policy**.
4. In the displayed dialog box, search **AliyunYundunDDoSReadOnlyAccess**, and then add it to the Selected Authorization Policy Names list. Click **OK**.

Edit User-Level Authorization

Members added to this group have all the permissions of this group. A member cannot be added to the same group more than once.

Available Authorization Policy Names

	Type
aliyunyundunDDoS	
AliyunYundunDDoSFullAccess Provides full acce...	System

Selected Authorization Policy Name

	Type
AliyunYundunDDoSReadOnlyAccess Provides read-only...	System

OK Close