

## Elastic Load Balance

# Getting Started

Issue	04
Date	2022-02-11



**Copyright © Huawei Technologies Co., Ltd. 2024. All rights reserved.**

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

## **Trademarks and Permissions**



HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

## **Notice**

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

# Security Declaration

## Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process*. For details about this process, visit the following web page:

<https://www.huawei.com/en/psirt/vul-response-process>

For vulnerability information, enterprise customers can visit the following web page:

<https://securitybulletin.huawei.com/enterprise/en/security-advisory>

---

# Contents

---

1 Overview..... 1

2 Process Flowchart..... 3

3 Preparations for Using ELB..... 5

4 Using Shared Load Balancers (Entry Level)..... 8

5 Using Shared Load Balancers (Advanced Level)..... 15

6 Change History..... 24

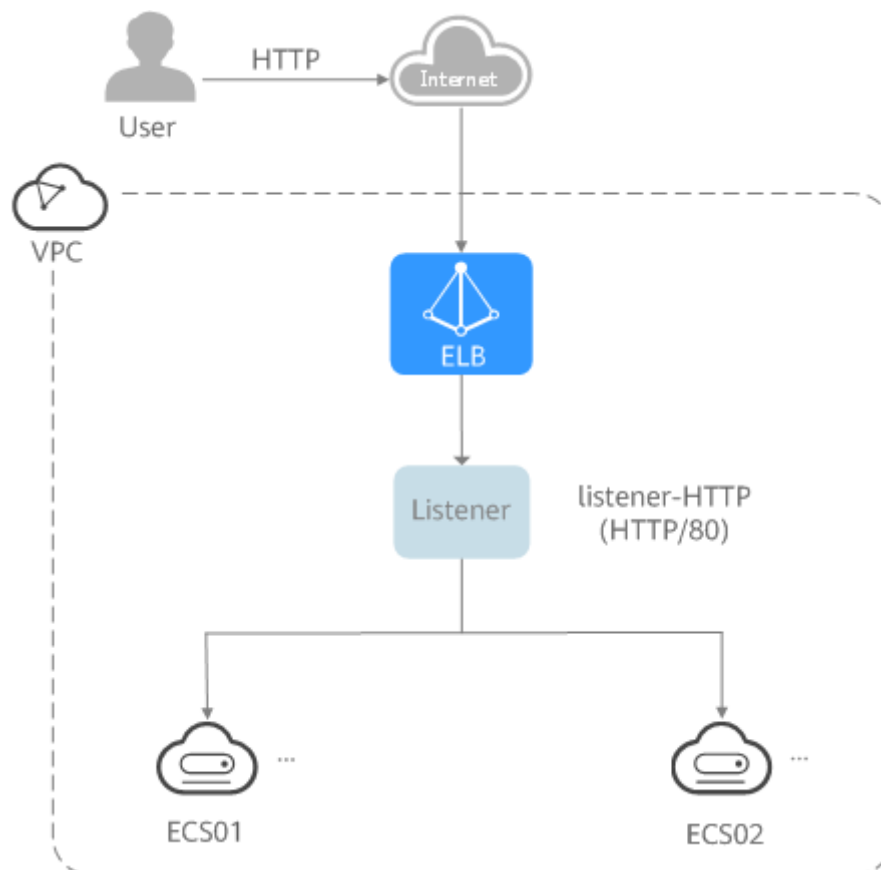
# 1 Overview

To use ELB to distribute traffic across backend servers, you need to create a dedicated load balancer or a shared load balancer.

Two examples are given to show how you can quickly create a shared load balancer to distribute incoming traffic across backend servers.

- **Entry level:** An application deployed on separated ECSs needs to handle a large number of requests. Health checks are required to monitor the health of the servers to ensure that incoming traffic is routed only to healthy servers to eliminate SPOFs and improve service availability.

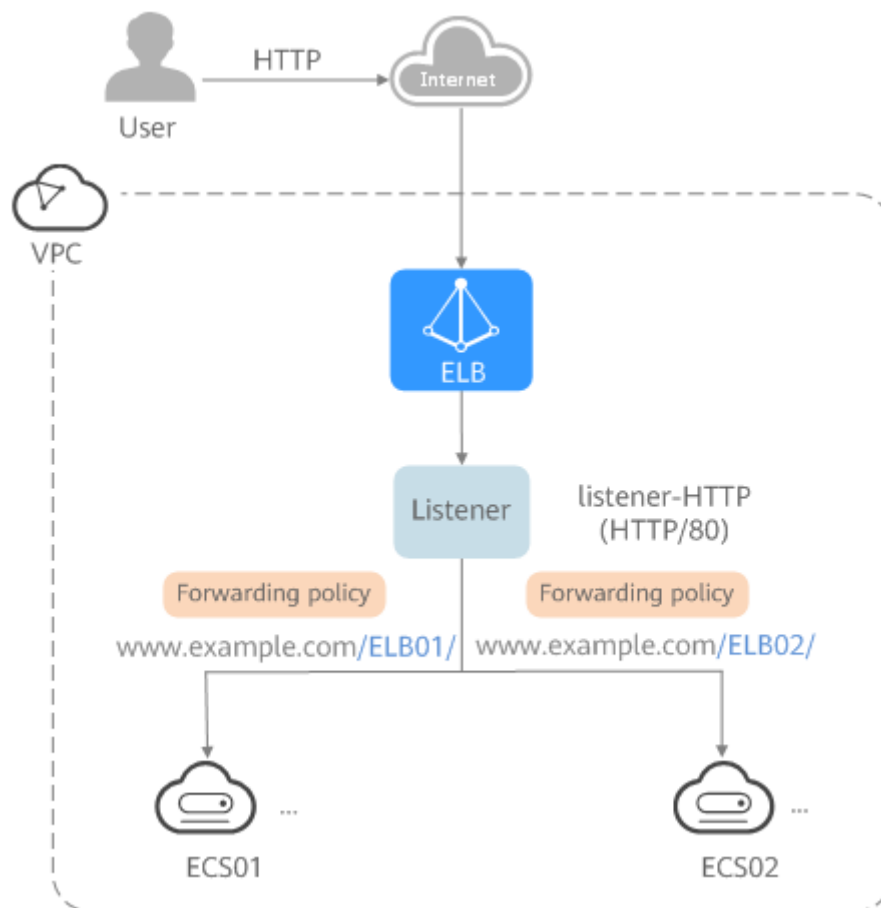
**Figure 1-1** Entry level



As the incoming traffic increases, you can add more backend servers to balance the load.

- **Advanced level:** An application deployed on separated servers uses one domain name but different URLs to provide services, and requests are routed to different servers based on the URLs. Forwarding policies are required to forward requests from different URLs to the servers in the corresponding backend server groups.

**Figure 1-2** Advanced level

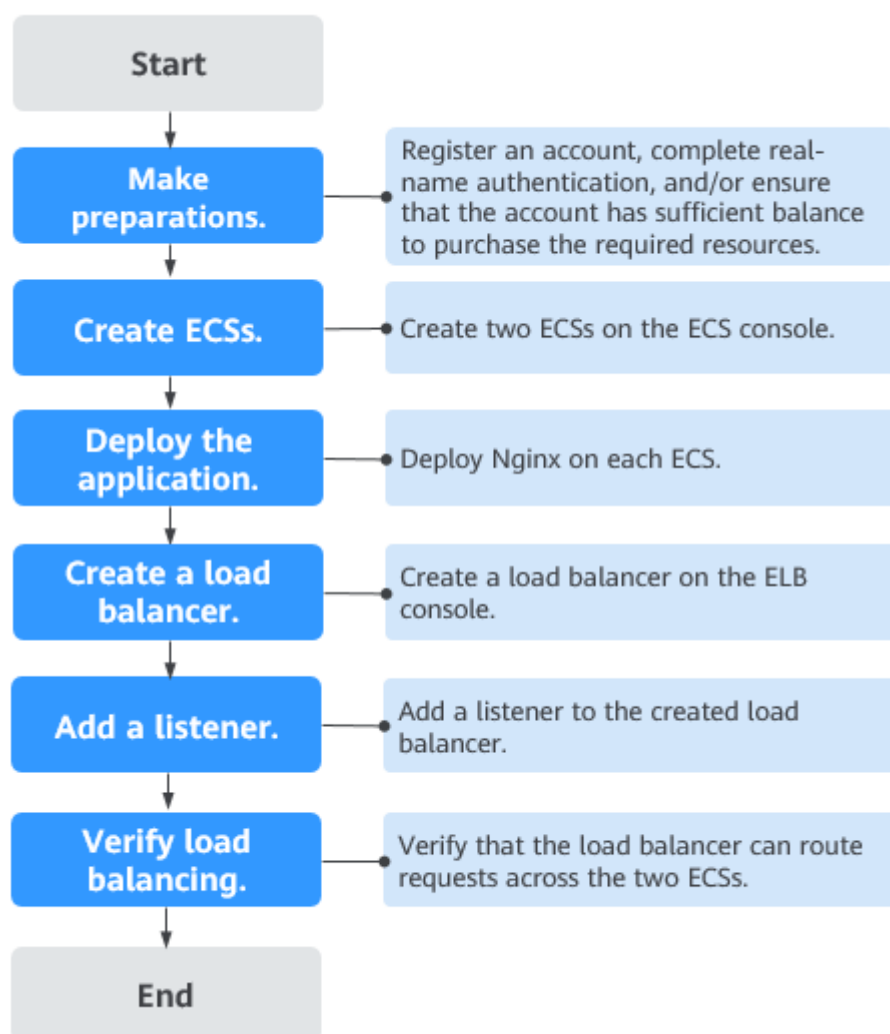


As the incoming traffic increases, you can add more backend servers to the two backend server groups. You can also configure health checks to monitor the health of backend servers to ensure that incoming traffic is routed only to healthy backend servers.

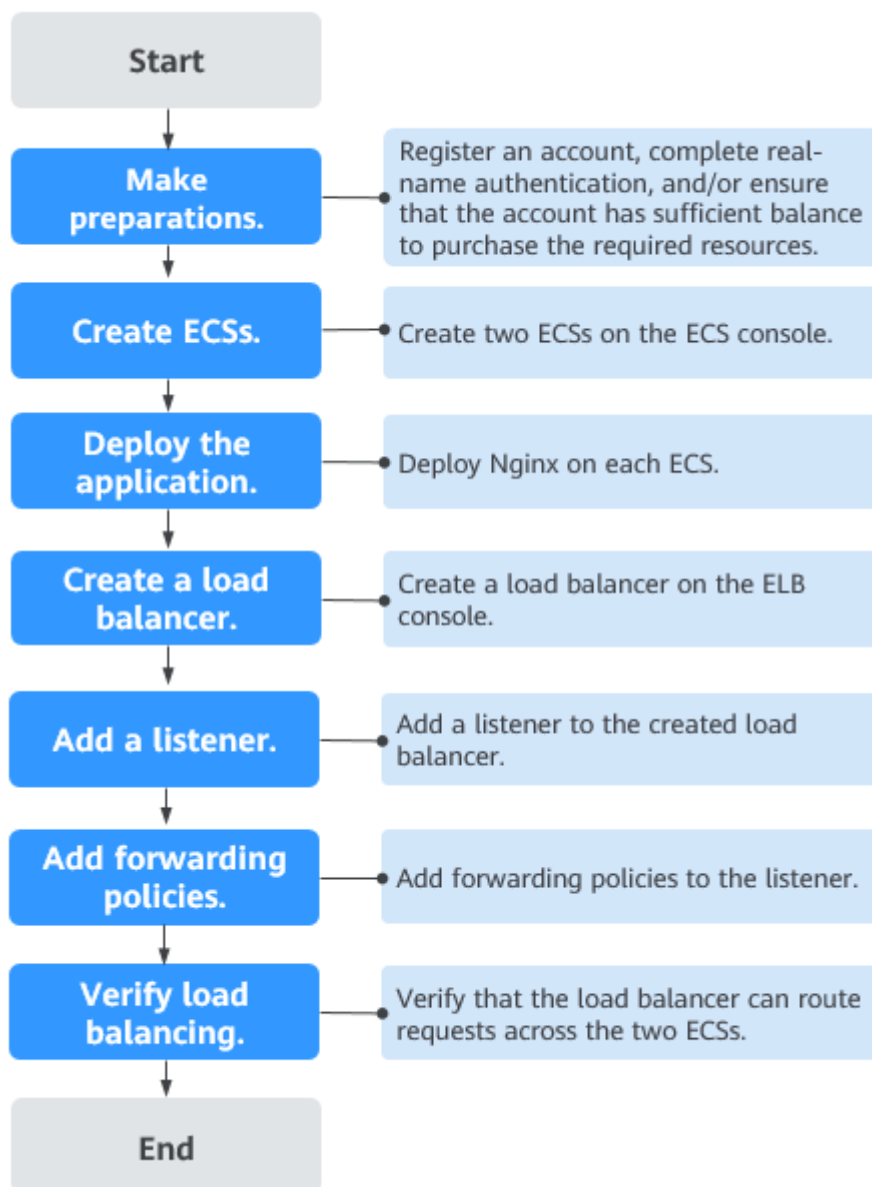
# 2 Process Flowchart

**Figure 1** shows how you can use basic functions of ELB to route requests when you are still new to ELB, and **Figure 2** shows how you can use ELB to route requests based on domain names or URLs more efficiently.

**Figure 2-1** Getting started (entry level)



**Figure 2-2** Getting started (advanced level)





# 3 Preparations for Using ELB

---

Before using ELB to route requests, you need to make some preparations.

- [Signing Up with Huawei Cloud and Completing Real-Name Authentication](#)
- [Topping Up Your Account](#)
- [Creating an IAM User](#)

## Signing Up with Huawei Cloud and Completing Real-Name Authentication

If you already have a Huawei Cloud account and completed real-name authentication, skip this part. If you do not have a Huawei Cloud account, perform the following steps to register one:

1. Visit the [Huawei Cloud official website](#) and click **Sign Up**.
2. Register an account as prompted.  
If the registration is successful, the system automatically redirects you to your personal information page.
3. Complete real-name authentication by following the instructions in [Individual Real-Name Authentication](#).

## Topping Up Your Account

For billing details about shared load balancers, [Billing \(Shared Load Balancers\)](#)

Dedicated load balancers are not free and can be purchased in the pay-per-use billing mode. For details, see [Billing \(Dedicated Load Balancers\)](#).

Ensure that your account has sufficient balance to buy the resources you need.

To top up an account, see [Topping Up an Account](#).

## Creating an IAM User

If you want to allow multiple users to manage your resources without sharing your password or private key, you can create IAM users and grant permissions to the users. These users can use specified links and their own accounts to access the cloud platform and help you manage resources efficiently. You can also configure security policies to ensure the security of these accounts.

If you have registered an account, you can create an IAM user on the IAM console.

For example, to create an ELB administrator, perform the following steps:

1. Enter your account name and password to log in to the management console.
2. In the upper right corner of the page, hover the mouse over the username and select **Identity and Access Management**.
3. In the navigation pane on the left, choose **Users**. In the right pane, click **Create User**.
4. Enter the user information on the **Create User** page.
  - **Username**: Enter **elb\_administrator**.
  - **Email Address**: Email address bound to the IAM user. This parameter is mandatory if the access type is specified as **Set by user**.
  - (Optional) **Mobile Number**: Mobile number of the IAM user.
  - (Optional) **Description**: Enter a description, for example, **ELB administrator**.
5. Select **Management console access** for **Access Type** and **Set now** for **Password**. Enter a password, and click **Next**.

**Figure 3-1** Selecting the access type

★ Access Type You are advised to select only one access type for security purposes.

☐ Programmatic access

Access HUAWEI CLOUD services using development tools (including APIs, CLI, and SDKs) that support key authentication. An access key will be automatically generated for each user. [Learn more](#)

☒ Management console access

Log in to HUAWEI CLOUD management console using the username and a password.

Console Password

☐ Set by user

A one-time login URL will be emailed to the user. The user can then click on the link to set a password.

☐ Automatically generated

A password will be automatically generated by the system. If you choose not to download the password, you can set a new password later.

☒ Set now

Set a password now.

\*\*\*\*\*

☒ Require password reset at first login

Login Protection

☐ Enable

☒ Disable

#### NOTE

An ELB administrator can log in to the management console and manage users. It is good practice to select **Set now** for **Password** when you create an ELB administrator for yourself. If you create an ELB administrator for another user, select **Set by user** for **Password** so that the user can set their own password.

6. (Optional) Add the user to the **admin** user group and click **Create**.

User group **admin** has all the permissions. If you want to grant fine-grained permissions to IAM users, see [Creating a User and Granting Permissions](#).

Check whether the IAM user is displayed in the user list. You can click the IAM user login link above the list and use the created user to log in to the console.

# 4 Using Shared Load Balancers (Entry Level)

---

## Scenarios

You have a web application, which often needs to handle heavy traffic and is deployed on two ECSs for load balancing.

You can create a load balancer to distribute traffic evenly across the two ECSs, which eliminates SPOFs and makes your application more available.

## Prerequisites

To ensure normal communications between the load balancer and backend servers, you need to check the security group rules and network ACL rules configured for the backend servers.

When backend servers receive requests from the load balancer, source IP addresses are translated into those in 100.125.0.0/16.

- Security group rules must allow traffic from the 100.125.0.0/16 to backend servers.
- Network ACL rules are optional for subnets. If network ACL rules are configured for the backend subnet of the load balancer, the network ACL rules must allow traffic from the backend subnet of the load balancer to the backend servers.

### NOTE



If **Transfer Client IP Address** is enabled for the TCP or UDP listeners, network ACL rules and security group rules will not take effect. You can use access control to limit which IP addresses are allowed to access the listener.

## Creating ECSs

ECSs are used as backend servers.

Each ECS needs an EIP for accessing the Internet, and the EIP is used for configuring the application on the ECS. You can determine whether to bind an EIP to the load balancer based on your requirements.

For details, see [Purchasing an ECS](#).

1. Log in to the management console.
2. In the upper left corner of the page, click  and select the desired region and project.
3. Hover on  in the upper left corner to display **Service List** and choose **Compute > Elastic Cloud Server**.
4. Click **Buy ECS**, configure the parameters, and click **Next**.

The following table lists the specifications of the two ECSs.

**Table 4-1** ECS specifications

Item	Example Value
Name	ECS01 and ECS02
OS	CentOS 7.2 64bit
vCPUs	2
Memory	4 GiB
System disk	40 GiB
Data disk	100 GiB
Bandwidth	5 Mbit/s

5. Submit your request.

## Deploying the Application

Deploy Nginx on the two ECSs and edit two HTML pages so that a page with message "Welcome to ELB test page one!" is returned when ECS01 is accessed, and the other page with message "Welcome to ELB test page two!" is returned when ECS02 is accessed.

1. [Log in to the ECSs](#).
2. Install and start Nginx.
  - a. Run the **wget** command to download the Nginx installation package for your operating system in use. CentOS 7.6 is used as an example here.

```
wget http://nginx.org/packages/centos/7/noarch/RPMS/nginx-release-centos-7-0.el7ngx.noarch.rpm
```
  - b. Run the following command to create the Nginx yum repository:

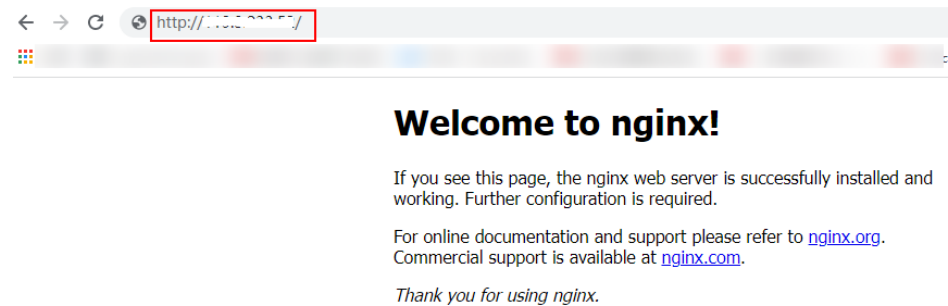
```
rpm -ivh nginx-release-centos-7-0.el7ngx.noarch.rpm
```
  - c. Run the following command to install Nginx:

```
yum -y install nginx
```
  - d. Run the following commands to start Nginx and configure automatic Nginx enabling upon ECS startup:

```
systemctl start nginx  
systemctl enable nginx
```

- e. Enter **http://EIP bound to the ECS** in the address box of your browser. If the following page is displayed, Nginx has been installed.

**Figure 4-1** Nginx installed successfully



3. Modify the HTML page of ECS01.

Modify the **index.html** file in the default root directory of Nginx **/usr/share/nginx/html** to identify access to ECS01.

- a. Open the **index.html** file.  
**vim /usr/share/nginx/html/index.html**
- b. Press **i** to enter editing mode.
- c. Modify the **index.html** file to be as follows:

```
...
<body>
  <h1>Welcome to <strong>ELB</strong> test page one!</h1>

  <div class="content">
    <p>This page is used to test the <strong>ELB</strong>!</p>

    <div class="alert">
      <h2>ELB01</h2>
      <div class="content">
        <p><strong>ELB test (page one)!</strong></p>
        <p><strong>ELB test (page one)!</strong></p>
        <p><strong>ELB test (page one)!</strong></p>
      </div>
    </div>
  </div>
</body>
```

- d. Press **Esc** to exit editing mode. Then, enter **:wq** to save the settings and exit the file.

4. Modify the HTML page of ECS02.

Modify the **index.html** file in the default root directory of Nginx **/usr/share/nginx/html** to identify access to ECS02.

- a. Open the **index.html** file.  
**vim /usr/share/nginx/html/index.html**
- b. Press **i** to enter editing mode.
- c. Modify the **index.html** file to be as follows:

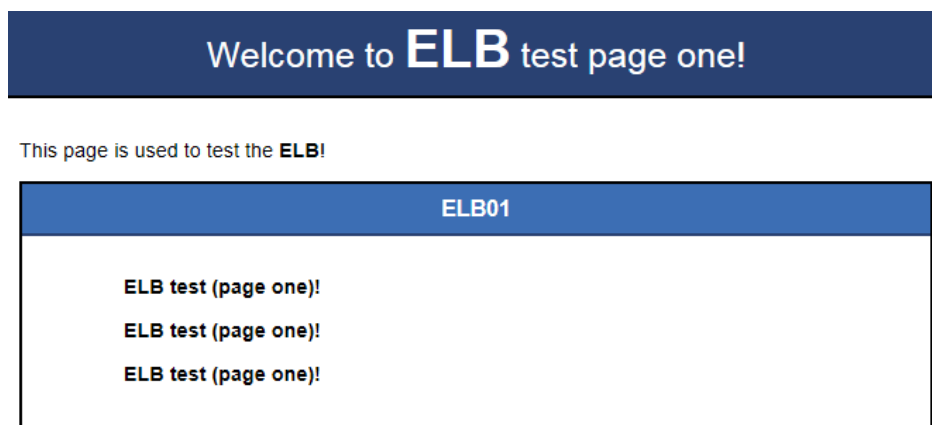
```
...
<body>
  <h1>Welcome to <strong>ELB</strong> test page two!</h1>

  <div class="content">
    <p>This page is used to test the <strong>ELB</strong>!</p>
```

```
<div class="alert">
  <h2>ELB02</h2>
  <div class="content">
    <p><strong>ELB test (page two)!</strong></p>
    <p><strong>ELB test (page two)!</strong></p>
    <p><strong>ELB test (page two)!</strong></p>
  </div>
</div>
</div>
</body>
```


- d. Press **Esc** to exit editing mode. Then, enter **:wq** to save the settings and exit the file.
5. Use your browser to access **http://ECS01 EIP** and **http://ECS02 EIP** to verify that Nginx has been deployed.  
If the modified HTML pages are displayed, Nginx has been deployed.
  - HTML page of ECS01

**Figure 4-2** Nginx successfully deployed on ECS01



- HTML page of ECS02



**Figure 4-3** Nginx successfully deployed on ECS02



## Creating a Load Balancer

The load balancer needs an EIP to access the application deployed on the ECSs over the Internet. You can determine whether to bind an EIP to the load balancer

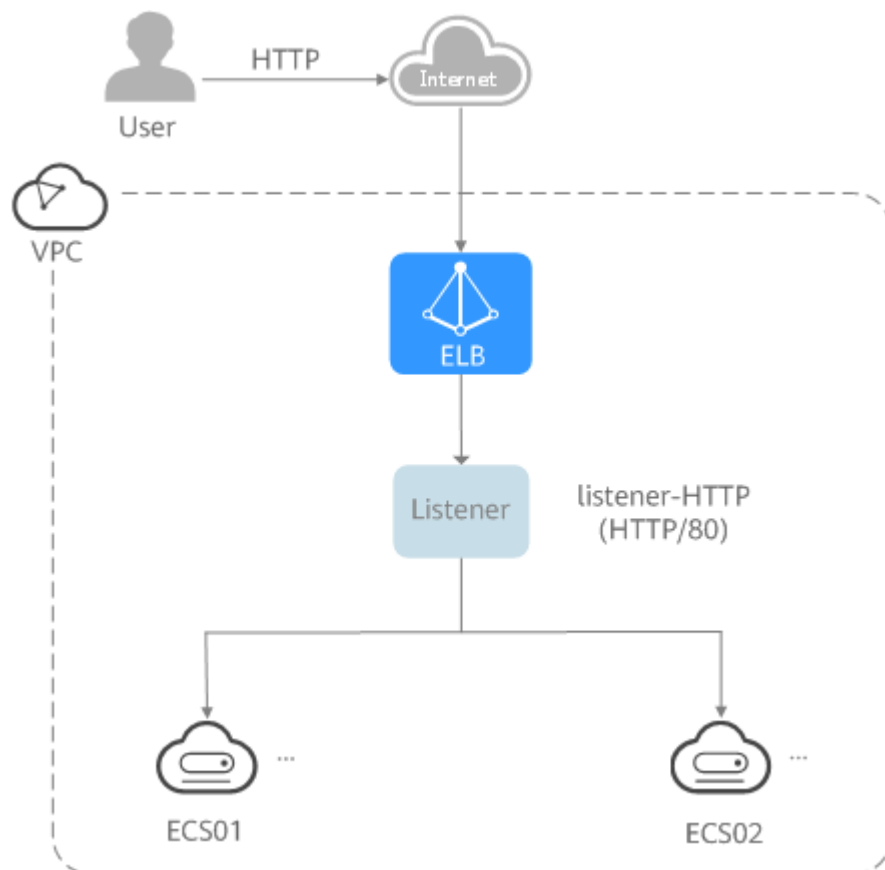
based on your requirements. For details, see [Load Balancing on a Public or Private Network](#).


1. In the upper left corner of the page, click  and select the desired region and project.
2. Hover on  in the upper left corner to display **Service List** and choose **Networking > Elastic Load Balance**.
3. Click **Buy Elastic Load Balancer** and then configure the parameters.
4. Click **Next**.
5. Confirm the configuration and submit your request.
6. View the newly created load balancer in the load balancer list.

## Adding a Listener

Add a listener to the created load balancer. When you add the listener, create a backend server group, configure a health check, and add the two ECSs to the created backend server group.

**Figure 4-4** Traffic forwarding



1. Hover on  in the upper left corner to display **Service List** and choose **Networking > Elastic Load Balance**.



2. Locate the created load balancer (**elb-01**) and click its name.
3. Under **Listeners**, click **Add Listener**.
4. Configure the parameters as follows:
  - **Name**: Enter a name, for example, **listener-HTTP**.
  - **Frontend Protocol**: Select a protocol, for example, **HTTP**.
  - **Frontend Port**: Enter a port, for example, **80**.
5. Click **Next: Configure Request Routing Policy**, select or create a backend group, and select a load balancing algorithm.
  - **Backend Server Group**: Select **Use existing** or **Create new**.  
Here we create a backend server group named **server\_group-ELB**.
  - **Load Balancing Algorithm**: Select an algorithm that the load balancer will use to route requests, for example, **Weighted round robin**.
6. Click **Next: Add Backend Server** and enable the health check.  
Configure the health check as follows:
  - **Protocol**: Select a protocol for the load balancer to perform health checks on backend servers. If the load balancer uses TCP, HTTP, or HTTPS to receive requests, the health check protocol can be TCP or HTTP. Here we use HTTP as an example.
  - **Domain Name**: Enter a domain name that will be used for health checks, for example, **www.example.com**.
  - **Health Check Port**: Enter a port for the load balancer to perform health checks on backend servers, for example, **80**.  
If you do not specify a health check port, the backend port will be used for health checks by default. If you specify a port, it will be used for health check.
7. Click **Next: Confirm**, confirm the configurations, and click **Submit**.
8. On the **Listeners** tab, locate the target listener. In the **Default Backend Server Group** column, click **View/Add Backend Server**.
9. On the **Backend Servers** tab, click the **Backend Servers** tab and click **Add** on the top right.
10. Select the servers you want to add, set the backend port, and click **Finish**.
  - Backend servers: Select **ECS01** and **ECS02**.
  - Backend port: Set it to **80**. Backend servers will use this port to communicate with the load balancer.

## Verifying Load Balancing

After the load balancer is configured, you can access the domain name to check whether the two ECSs are accessible.

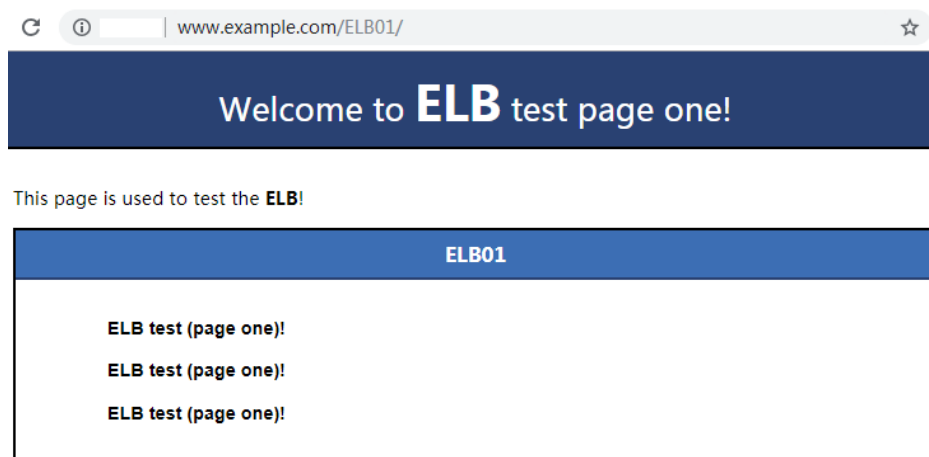
1. Modify the **C:\Windows\System32\drivers\etc\hosts** file on your PC to map the domain name to the load balancer EIP.

View the load balancer EIP on the **Summary** page of the load balancer.

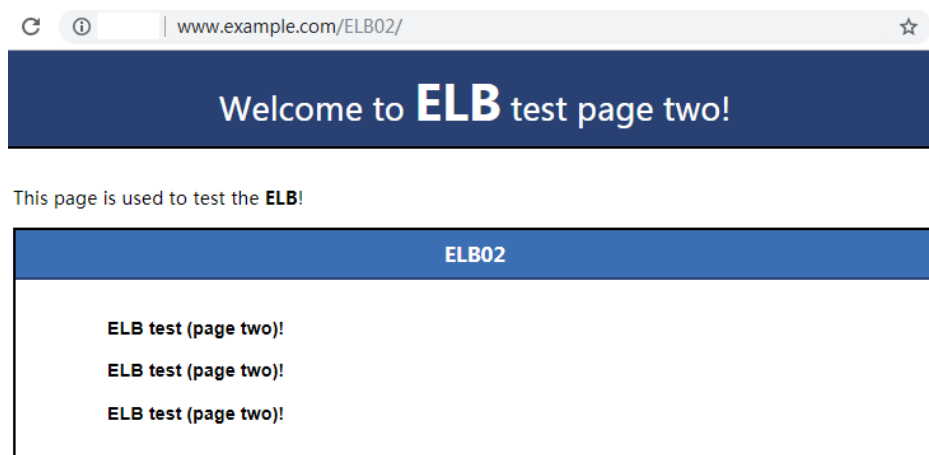
**Figure 4-5** hosts file on your PC

```
# localhost name resolution is handled within DNS itself.  
# 127.0.0.1      localhost  
# ::1           localhost  
  
11.11.11.14 www.example.com
```

2. On the CLI of your PC, run the following command to check whether the domain name is mapped to the load balancer EIP:  
**ping www.example.com**  
If data packets are returned, the domain name has been mapped to the load balancer EIP.
3. Use your browser to access **http://www.example.com**. If the following page is displayed, the load balancer has routed the request to ECS01.

**Figure 4-6** Accessing ECS01

4. Use your browser to access **http://www.example.com**. If the following page is displayed, the load balancer has routed the request to ECS02.

**Figure 4-7** Accessing ECS02

# 5 Using Shared Load Balancers (Advanced Level)

---

## Scenarios

You have two web applications that are deployed on separated ECSs but use the same domain name for access. You can set different URLs to process requests.

To forward requests based on URLs, you need to create a load balancer, add an HTTP or HTTPS listener, and add forwarding policies to specify the URLs.

An HTTP listener is used as an example to describe how to route requests from two URLs (**/ELB01** and **/ELB02**) of the same domain name (www.example.com) to different backend servers.

## Prerequisites

To ensure normal communications between the load balancer and backend servers, you need to check the security group rules and network ACL rules configured for the backend servers.

When backend servers receive requests from the load balancer, source IP addresses are translated into those in 100.125.0.0/16.



- Security group rules must allow traffic from the 100.125.0.0/16 to backend servers.
- Network ACL rules are optional for subnets. If network ACL rules are configured for the backend subnet of the load balancer, the network ACL rules must allow traffic from the backend subnet of the load balancer to the backend servers.

### NOTE

If **Transfer Client IP Address** is enabled for the TCP or UDP listeners, network ACL rules and security group rules will not take effect. You can use access control to limit which IP addresses are allowed to access the listener.

## Creating ECSs

ECSs are used as backend servers to process requests. Each ECS needs an EIP for accessing the Internet and configuring the application on the ECS.

1. Log in to the management console.
2. In the upper left corner of the page, click  and select the desired region and project.
3. Hover on  in the upper left corner to display **Service List** and choose **Compute > Elastic Cloud Server**.
4. Click **Buy ECS**, configure the parameters, and click **Next**.

The following table lists the specifications of the two ECSs.

**Table 5-1** ECS specifications

Item	Example Value
Name	ECS01 and ECS02
OS	CentOS 7.2 64bit
vCPUs	2
Memory	4 GiB
System disk	40 GiB
Data disk	100 GiB
Bandwidth	5 Mbit/s

5. Submit your request.

## Deploying the Application

Deploy Nginx on the two ECSs and edit two HTML pages so that a page with message "Welcome to ELB test page one!" is returned when ECS01 is accessed, and the other page with message "Welcome to ELB test page two!" is returned when ECS02 is accessed.

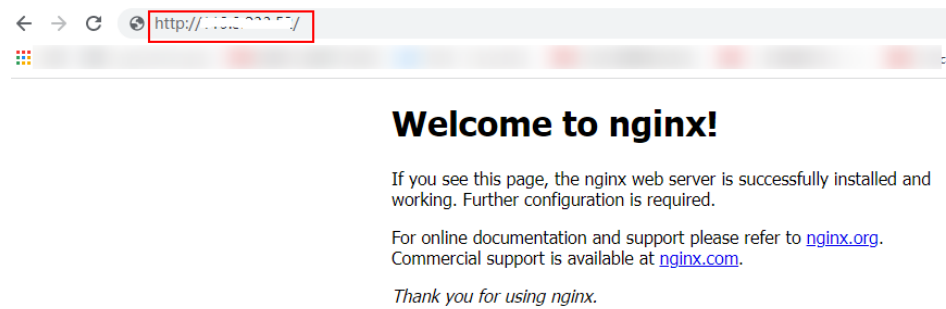
1. **Log in to the ECSs.**
2. Install and start Nginx.
  - a. Run the **wget** command to download the Nginx installation package for your operating system in use. CentOS 7.6 is used as an example here.

```
wget http://nginx.org/packages/centos/7/noarch/RPMS/nginx-release-centos-7-0.el7ngx.noarch.rpm
```
  - b. Run the following command to create the Nginx yum repository:

```
rpm -ivh nginx-release-centos-7-0.el7ngx.noarch.rpm
```
  - c. Run the following command to install Nginx:

```
yum -y install nginx
```
  - d. Run the following commands to start Nginx and configure automatic Nginx enabling upon ECS startup:

```
systemctl start nginx  
systemctl enable nginx
```
  - e. Enter **http://EIP bound to the ECS** in the address box of your browser. If the following page is displayed, Nginx has been installed.

**Figure 5-1** Nginx installed successfully

3. Modify the HTML page of ECS01.

Move the **index.html** file from the default root directory of Nginx **/usr/share/nginx/html** to the **ELB01** directory and modify the file to identify access to ECS01.

a. Create the **ELB01** directory and copy the **index.html** file to this directory:

```
mkdir /usr/share/nginx/html/ELB01
```

```
cp /usr/share/nginx/html/index.html /usr/share/nginx/html/ELB01/
```

b. Open the **index.html** file.

```
vim /usr/share/nginx/html/ELB01/index.html
```

c. Press **i** to enter editing mode.

d. Modify the **index.html** file to be as follows:

```
...
<body>
  <h1>Welcome to <strong>ELB</strong> test page one!</h1>

  <div class="content">
    <p>This page is used to test the <strong>ELB</strong>!</p>

    <div class="alert">
      <h2>ELB01</h2>
      <div class="content">
        <p><strong>ELB test (page one)!</strong></p>
        <p><strong>ELB test (page one)!</strong></p>
        <p><strong>ELB test (page one)!</strong></p>
      </div>
    </div>
  </div>
</body>
```

e. Press **Esc** to exit editing mode. Then, enter **:wq** to save the settings and exit the file.

4. Modify the HTML page of ECS02.

Move the **index.html** file from the default root directory of Nginx **/usr/share/nginx/html** to the **ELB02** directory and modify the file to identify access to ECS02.

a. Create the **ELB02** directory and copy the **index.html** file to this directory:

```
mkdir /usr/share/nginx/html/ELB02
```

```
cp /usr/share/nginx/html/index.html /usr/share/nginx/html/ELB02/
```

b. Open the **index.html** file.

```
vim /usr/share/nginx/html/ELB02/index.html
```

c. Press **i** to enter editing mode.

- d. Modify the **index.html** file to be as follows:


```
...
<body>
  <h1>Welcome to <strong>ELB</strong> test page two!</h1>

  <div class="content">
    <p>This page is used to test the <strong>ELB</strong>!</p>

    <div class="alert">
      <h2>ELB02</h2>
      <div class="content">
        <p><strong>ELB test (page two)!</strong></p>
        <p><strong>ELB test (page two)!</strong></p>
        <p><strong>ELB test (page two)!</strong></p>
      </div>
    </div>
  </div>
</body>
```


- e. Press **Esc** to exit editing mode. Then, enter **:wq** to save the settings and exit the file.
5. Use your browser to access **http://ECS01 EIP/ELB01/** and **http://ECS02 EIP/ELB02/** to verify that Nginx has been deployed.
- If the modified HTML pages are displayed, Nginx has been deployed.
- HTML page of ECS01

**Figure 5-2** Nginx successfully deployed on ECS01



- HTML page of ECS02

**Figure 5-3** Nginx successfully deployed on ECS02



## Creating a Load Balancer

The load balancer needs an EIP to access the application deployed on the ECSs over the Internet. You can determine whether to bind an EIP to the load balancer based on your requirements. For details, see [Load Balancing on a Public or Private Network](#).

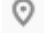

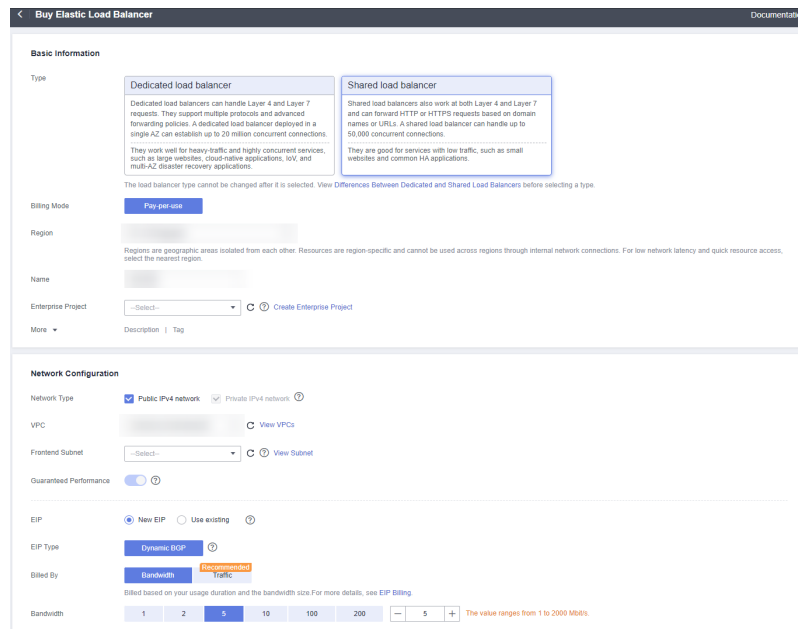
1. In the upper left corner of the page, click  and select the desired region and project.
2. Hover on  in the upper left corner to display **Service List** and choose **Networking > Elastic Load Balance**.
3. Click **Buy Elastic Load Balancer** and then configure the parameters.

Figure 5-4 Buy Elastic Load Balancer



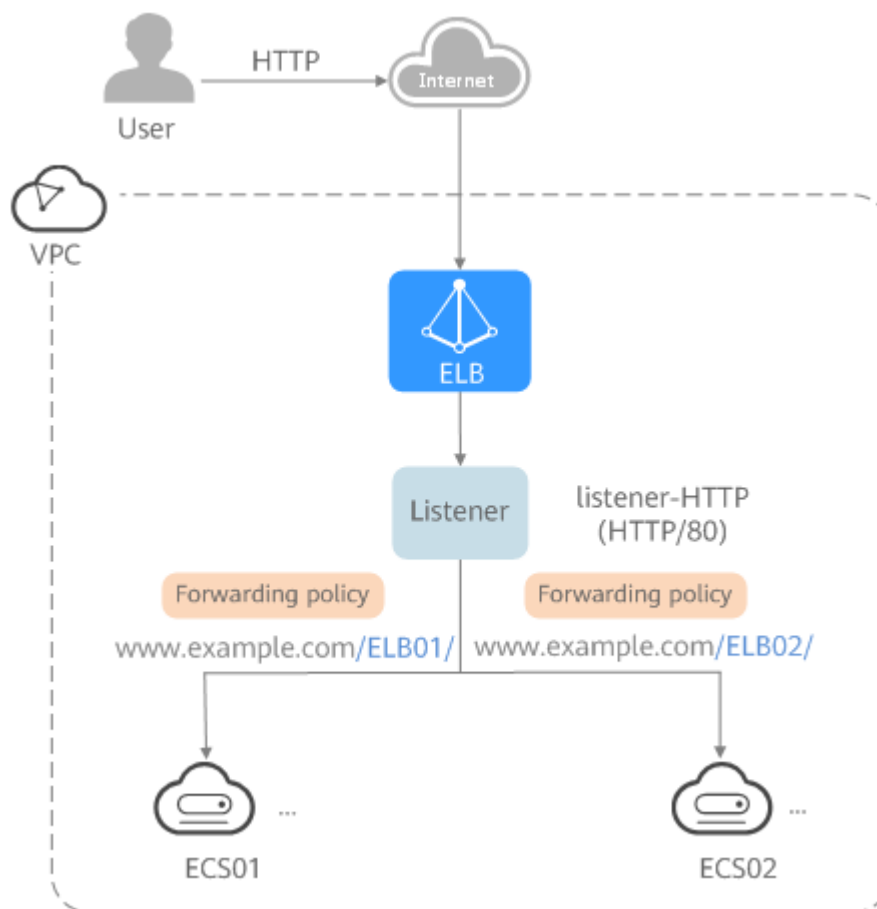
4. Click **Next**.
5. Confirm the configuration and submit your request.
6. View the newly created load balancer in the load balancer list.


## Adding a Listener

Add a listener to the created load balancer. When you add the listener, create a backend server group, configure a health check, and add the two ECSs to the created backend server group.

Configure two forwarding policies to forward HTTP requests to the two ECSs, for example, requests from **www.example.com/ELB01/** to ECS01, and those from **www.example.com/ELB02/** to ECS02.

**Figure 5-5** Traffic forwarding



1. Hover on  in the upper left corner to display **Service List** and choose **Networking > Elastic Load Balance**.
2. Locate the created load balancer and click its name.
3. Under **Listeners**, click **Add Listener**.
4. Configure the parameters as follows:
  - **Name:** Enter a name, for example, **listener-HTTP**.
  - **Frontend Protocol:** Select a protocol, for example, **HTTP**.
  - **Frontend Port:** Enter a port, for example, **80**.
5. Create a backend server group, configure a health check, and click **Finish**.
  - Backend server group
    - **Name:** Enter a name, for example, **server\_group-ELB**.
    - **Load Balancing Algorithm:** Select an algorithm that the load balancer will use to route requests, for example, **Weighted round robin**.
  - Health check
    - **Protocol:** Select a protocol for the load balancer to perform health checks on backend servers. If the load balancer uses TCP, HTTP, or HTTPS to receive requests, the health check protocol can be TCP or



HTTP. Here we use HTTP as an example. Note that the protocol cannot be changed after the listener is added.

- **Domain Name:** Enter a domain name that will be used for health checks, for example, **www.example.com**.
- **Health Check Port:** Enter a port for the load balancer to perform health checks on backend servers, for example, **80**.

## Add a Forwarding Policy

1. Locate the newly added listener and click **Add/Edit Forwarding Policy** in the **Forwarding Policies** column.
2. Click **Add Forwarding Policy** and configure a forwarding policy.
  - **Name:** Enter a forwarding policy name, for example, **forwarding\_policy-ELB01**.
  - **Domain name:** Enter a domain name that will be used to forward the requests, for example, **www.example.com**. The domain name in the request must exactly match that specified in the forwarding policy.
  - **URL:** You can also specify a URL to forward the requests, for example, **/ELB01/**.
  - **URL matching rule:** Select a rule for matching the specified URL string with the URL in the request. Three options are available, **Exact match**, **Prefix match**, and **Regular expression match**. **Exact match** enjoys the highest priority, and **Regular expression match** the lowest priority. Select **Exact match** here.
  - **Action:** Select **Forward to a backend server group**.
  - **Backend Server Group:** Select **Create Backend Server Group**.
3. Create a backend server group and configure a health check.
  - Backend server group
    - **Name:** Enter a name, for example, **server\_group-ELB01**.
    - **Load Balancing Algorithm:** Select an algorithm that the load balancer will use to route requests, for example, **Weighted round robin**.
  - Health check
    - **Protocol:** Select a protocol for the load balancer to perform health checks on backend servers. If the load balancer uses TCP, HTTP, or HTTPS to receive requests, the health check protocol can be TCP or HTTP. Here we use HTTP as an example. Note that the protocol cannot be changed after the listener is added.
    - **Domain Name:** Enter a domain name that will be used for health checks, for example, **www.example.com**.
    - **Health Check Port:** Enter a port for the load balancer to perform health checks on backend servers, for example, **80**.

**Figure 5-6** Configuring a health check

**Configure Health Check**

Health Check ☒

Health check detects the running of backend servers and ensures that requests are routed only to healthy backend servers. [Learn more](#)

The backend server group is associated with a dedicated load balancer. Security group rules need to be configured to allow traffic from the backend subnet where the load balancer is running to backend servers. If security group rules are not configured, backend servers cannot receive requests from the load balancer and will be identified as unhealthy. [Learn how to configure a security group.](#)

★ Health Check Protocol ⓘ HTTP

Domain Name

☐ Private IP address of the backend server

☒ Specified domain name

www.example.com

Enter at least two character strings separated by periods (.). Use only letters, digits, and hyphens (-). Do not start or end strings with a hyphen. Max total: 100 characters. Max string: 63 characters.

Health Check Port

☒ Default backend server port

☐ Specified port

★ Path ⓘ /

Start the path with a slash (/). The path can contain 1 to 80 characters, including letters, digits, and the following characters: / & % ?

OK Cancel

- Click the name of the backend server group configured for the newly added forwarding policy.
- On the **Backend Servers** tab, click the **Backend Servers** tab and click **Add**.
- Select the server you want to add, set the backend port, and click **Finish**.
  - Backend server: ECS01
  - Backend port: Set it to **80**. Backend servers will use this port to communicate with the load balancer.
- Repeat the preceding steps to add another forwarding policy, create a backend server group, and add ECS02 to the backend server group.

## Verifying Load Balancing

After the load balancer is configured, you can access the domain name or the specified URL to check whether the two ECSs are accessible.

- Modify the **C:\Windows\System32\drivers\etc\hosts** file on your PC to map the domain name to the load balancer EIP.

View the load balancer EIP on the **Summary** page of the load balancer.

**Figure 5-7** hosts file on your PC

```
# localhost name resolution is handled within DNS itself.
# 127.0.0.1      localhost
# ::1           localhost

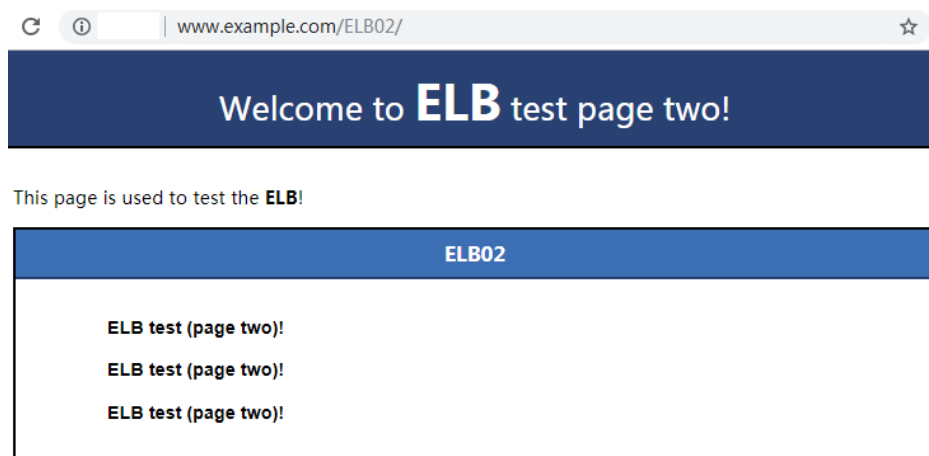
11.11.11.114 www.example.com
```

- On the CLI of your PC, run the following command to check whether the domain name is mapped to the load balancer EIP:  
**ping www.example.com**  
If data packets are returned, the domain name has been mapped to the load balancer EIP.
- Use your browser to access **http://www.example.com/ELB01/**. If the following page is displayed, the load balancer has routed the request to ECS01.

**Figure 5-8** Accessing ECS01**NOTE**

**ELB01/** indicates that the default directory named **ECS01** is accessed, while **ELB01** indicates the file name. Therefore, the slash (/) following **ELB01** must be retained.

- Use your browser to access **http://www.example.com/ELB02/**. If the following page is displayed, the load balancer has routed the request to ECS02.

**Figure 5-9** Accessing ECS02

# 6 Change History

---

Released On	Description
2022-02-11	<p>This issue is the fourth official release, which incorporates the following changes:</p> <p>Added the following sections:</p> <ul style="list-style-type: none"><li>• <a href="#">Process Flowchart</a></li><li>• <a href="#">Preparations for Using ELB</a></li></ul>
2020-05-30	<p>This issue is the third official release, which incorporates the following changes:</p> <p>Changed the name of enhanced load balancers to shared load balancers.</p>
2018-12-30	<p>This issue is the second official release, which incorporates the following changes:</p> <p>Modified some parameters based on the latest console.</p>
2018-11-21	<p>This issue is the first official release.</p>