

Server Load Balancer

Best Practices

Best Practices

Troubleshoot Health Check Exceptions

You can [Click Here](#) to learn the implementation mechanism of health check in Server Load Balancer.

For a Layer-7 (HTTP protocol) service, when a listener has “exception” in health check, go through the following steps for troubleshooting:

- First, make sure that you can access your application service through ECS directly.
- As communication between Layer-7 (HTTP protocol) service and backend ECS uses intranet, you need to log in to the server to ensure the application server port listens to the intranet address.
- Make sure that the backend server has enabled corresponding port and that port must be consistent with the backend port for your Server Load Balancer listening.
- Check whether the backend ECS has a firewall or other security protection software. This type of software may easily block local IP address of the Server Load Balancer system, and thus disable communication between the Server Load Balancer system and the backend server.
- Check whether the Server Load Balancer health check parameters are correctly set. It is recommended to keep the default values.
- It is recommended to use static pages to conduct health check. If the page for health check is not the default page on the backend ECS application server, you need to specify the URI of the page in health check configuration.
- Check whether there is a high load on the backend ECS resources that causes ECS to respond slowly in providing service.

In addition, you can also use the following methods for troubleshooting.

Because the Server Load Balancer and backend ECS communicate with each other over the intranet, it is necessary for ECS to monitor intranet ports or all ports.

1.If the IDs of the Server Load Balancer frontend port and the ECS backend port are both 80 and the ECS intranet IP address is 10.11.192.1

For Windows ECS, use the following command:

```
netstat -ano | findstr :80
```

For Linux ECS, use the following command:

```
netstat -anp | grep :80
```

If you can see the listening of 10.11.192.1:80 or 0.0.0.0:80, it means that this part is normal.

2. Check whether the server intranet firewall allows port 80. The firewall can be temporarily closed for testing.

For Windows system, the command "firewall.cpl" can be entered to close the firewall.

For Linux system, the command "/etc/init.d/iptables stop" can be entered to close the firewall.

3. For Layer-4 Server Load Balancer, it is okay if the telnet command of the backend port responds. Telnet 10.11.192.1 80 for testing. For Layer-7 Server Load Balancer, the HTTP status code must be 200 or other codes that represent normal. Testing methods are as follows:

For Windows system, the intranet IP address can be entered in ECS browser directly to see whether it is normal, for example, <http://http://10.11.192.1>

For Linux system, the "curl -I" command can be used to check whether the status is HTTP/1.1 200 OK, for example, `curl -I 10.11.192.1`

4. It is recommended that simple pages in HTML format are specified for health check. It's only used to check returned results. Dynamic scripting languages like PHP are not recommended.

If the problem persists, contact Alibaba Cloud After-Sales Technical Support.

Obtain a real visitor IP

Note: The following is for Server Load Balancer services under the Layer-7 protocol (HTTP).

General methods of an application server to obtain a real visitor IP:

Because service under the Layer-4 protocol (TCP) can obtain a real visitor IP address at the backend ECS directly, the extra configuration is not needed.

To obtain a real visitor IP address, the function "Obtain a real IP" needs to be enabled in the Server Load Balancer service listener configuration. The Server Load Balancer system provides the X-Forwarded-For to obtain the real visitor IP.

General configuration schemes of application server are shown below:

IIS 6 configuration scheme

To obtain the real visitor IP addresses through the IIS 6 log, the plug-in F5XForwardedFor.dll needs to be installed first. Click to download the F5XForwardedFor.dll installation package and install it as

follows.

In accordance with the operating system version of your server, copy the F5XForwardedFor.dll from directory x86\Release or x64\Release into a directory such as C:\ISAPIFilters and the IIS process should have the right to read the directory.

Open the IIS manager, find out the current website and right click "Attribute" on the website to open an attribute page.

Switch into "ISAPI filter" on the attribute page and click "Add" for addition window.

In the addition window: input "F5XForwardedFor" in the "Filter name", the full path of F5XForwardedFor.dll in the "Executable file" and then click OK.

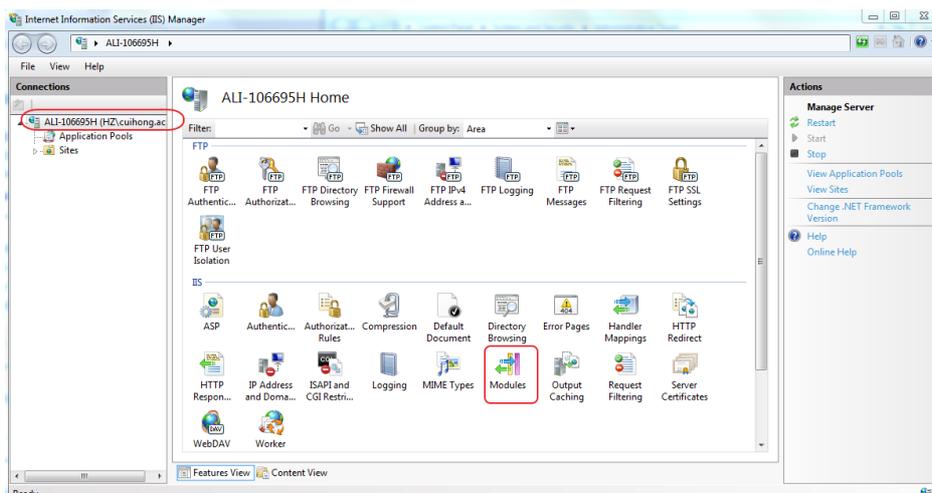
Reboot the IIS server and wait for the configuration to take effect.

IIS 7 configuration scheme

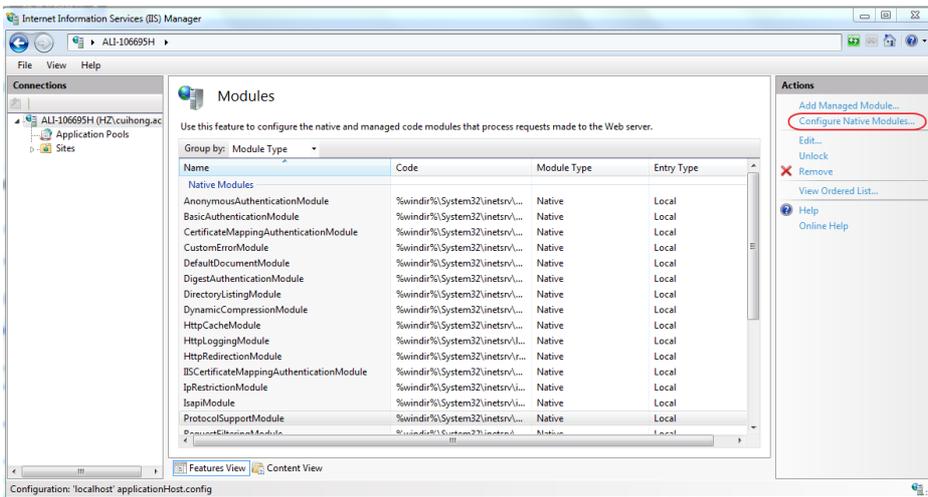
To obtain the real visitor IP addresses through the F5XForwardedFor module, the plug-in F5XForwardedFor module needs to be downloaded and installed first. Click to download the F5XForwardedFor installation package and install it as follows.

1. In accordance with the operating system version of your server, copy the F5XFFHttpModule.dll and F5XFFHttpModule.ini under directory x86\Release or x64\Release into a directory such as C:\F5XForwardedFor\ and the IIS process should have the right to read the directory.

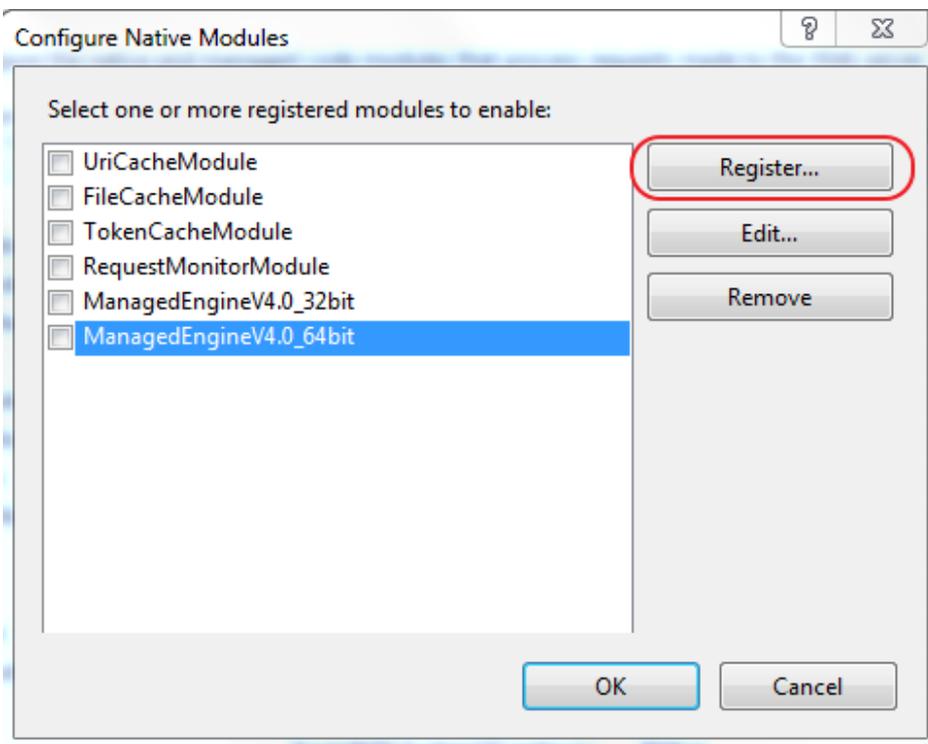
2. Select "IIS server" to choose "Module" as shown in the figure below:



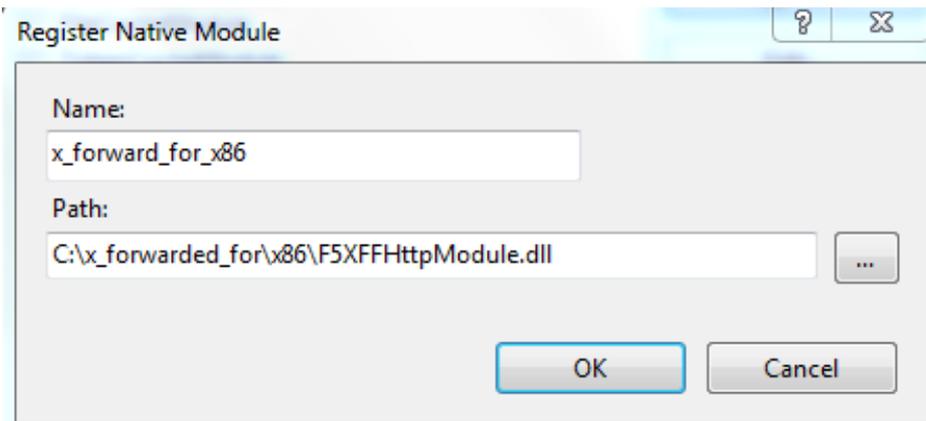
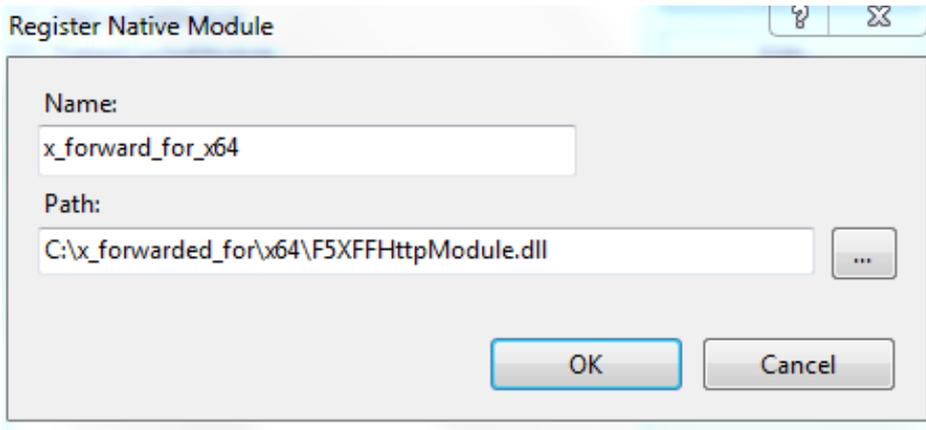
3. Double click "Module" and click "Configure the module".



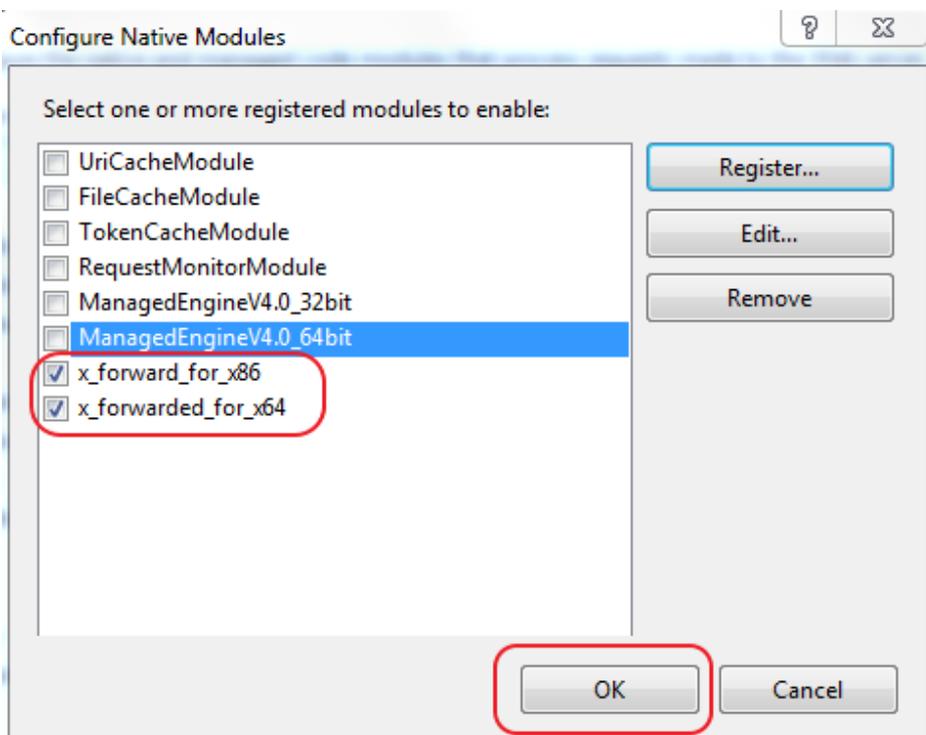
4. Click the "Register" button in the pop-up box:



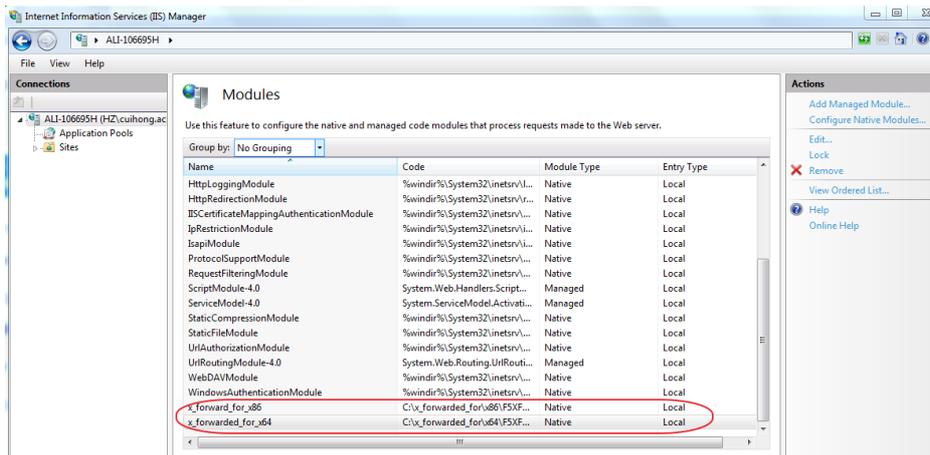
5. Add a DLL file downloaded as shown below:



6. Then check and click "OK" :



7. Add two DLL files on the "API and CGI restriction" and change into Allow:



8.Reboot the IIS server and wait for the configuration to take effect.

Apache configuration scheme

Install a third-party module "mod_rpaf" of apache from the following website.

```
wget http://stderr.net/apache/rpaf/download/mod_rpaf-0.6.tar.gz
tar zxvf mod_rpaf-0.6.tar.gz
cd mod_rpaf-0.6
/alidata/server/httpd/bin/apxs -i -c -n mod_rpaf-2.0.so mod_rpaf-2.0.c
```

Modify the configuration /alidata/server/httpd/conf/httpd.conf of apache and add at the end:

```
LoadModule rpaf_module modules/mod_rpaf-2.0.so
RPAFenable On
RPAFsethostname On
RPAFproxy_ips ip address
RPAFheader X-Forwarded-For
```

//RPAFproxy_ips IP address (This IP address is not a public network IP provided by the Server Load Balancer and the specific IP is shown in the apache log. Generally, there are two IP addresses and both shall be written)

Reboot the apache after adding.

```
/alidata/server/httpd/bin/apachectl restart
```

Specific case

```
LoadModule rpfaf_module modules/mod_rpfaf-2.0.so
RPAFenable On
RPAFsethostname On
RPAFproxy_ips 10.242.230.65 10.242.230.131
RPAFheader X-Forwarded-For
```

Nginx configuration scheme

Nginx is used by the Server Load Balancer to obtain a real IP via `http_realip_module`. There is no `http_realip_module` in Nginx installed by the defaulted “one-click” installation package and you need to recompile the Nginx and add `--with-http_realip_module`.

```
wget http://soft.phpwind.me/top/nginx-1.0.12.tar.gz
tar zxvf nginx-1.0.12.tar.gz
cd nginx-1.0.12
./configure --user=www --group=www --prefix=/alidata/server/nginx --with-http_stub_status_module --
without-http-cache --with-http_ssl_module --with-http_realip_module
make
make install
kill -USR2 `cat /alidata/server/nginx/logs/nginx.pid`
kill -QUIT `cat /alidata/server/nginx/logs/nginx.pid.oldbin`
```

After that, modify the `nginx.conf`.

```
vi /alidata/server/nginx/conf/nginx.conf
```

under

```
fastcgi connect_timeout 300;
fastcgi send_timeout 300;
fastcgi read_timeout 300;
fastcgi buffer_size 64k;
fastcgi buffers 4 64k;
fastcgi busy_buffers_size 128k;
fastcgi temp_file_write_size 128k;
```

add

```
set_real_ip_from IP address (This IP address is not a public network IP provided by the Server Load Balancer and the specific IP is shown in the previous nginx log. If there are multiple IP addresses, all of them shall be written))
```

```
set_real_ip_from IP address (This IP address is not a public network IP provided by the Server Load Balancer and the specific IP is shown in the previous nginx log. If there are multiple IP addresses, all of them shall be written)) real_ip_header X-Forwarded-For;
```

```
real_ip_header X-Forwarded-For;
```

Reboot the Nginx

```
/alidata/server/nginx/sbin/nginx -s reload
```

Configure different session persistence rules

This part introduces the implementation method to configure different session persistence rules for multiple domain names.

For example, your VIP has two domain names , vip.a.com and img.a.com. You hope that vip.a.com has session persistence and img.a.com does not have session persistence. Open the session persistence configuration, select "Override a service cookie" such as inputting the name cookie field and set key as the cookie of name for the domain name vip.a.com on your own application server without any settings for the domain name img.a.com to meet your needs.

Here are three mainstream methods for the application server to configure the cookie for domain name, for your reference.

Apache

Open the httpd.conf to ensure this line is not commented.

```
LoadModule usertrack_module modules/mod_usertrack.so
```

Configure in the virtual host.

```
CookieName name  
CookieExpires "1 days"  
CookieStyle Cookie  
CookieTracking on
```

Nginx

```
server {  
listen 8080;
```

```
server_name wqwq.example.com;
location / {
    add_header Set-Cookie name=xxxx;
    root html;
    index index.html index.htm;
}
}
```

Lighttpd

```
server.modules = ( "mod_setenv" )
$http["host"] == "test.example.com" {
    server.document-root = "/var/www/html/"
    setenv.add-response-header = ( "Set-Cookie" => "name=XXXXXX" )
}
```

But the cookie of a main domain name may be inherited by its subdomain name. If you hope that the session persistence is provided for the primary domain names and not for the secondary domain names, you can apply for more VIP later to solve this problem.

Configure Server Load Balancer through Open API

How to configure Server Load Balancer service through the Server Load Balancer Open API?

URL is used as a request parameter in the following example and no any common request parameter is included. Refer to the [Alibaba Cloud ECS API Reference](#) to form a complete request URL. To facilitate reading, the parameter values of request URL in the example are not encoded in format of URL.

Preconditions : There is one or multiple ECS instances. ECS instance is created and it gives access right to SSH and Web ports. For more details, refer to ECS introduction. It is assumed that an ECS instance has been created and configured with Nginx. To check the Server Load Balancer, change the Body characters in the `/usr/share/nginx/www/index.html` into: "Welcome to nginx on mytesthost1!"

1. "Copy" an Instance

"Copy" the second instance from the first instance.

It is assumed that the first instance is configured and can be used to create a custom image.

1、 Create a snapshot for the system disk

1) Query the system disk ID of the instance:

Request:

```
https://ecs.aliyuncs.com/?Action=DescribeInstanceDisks&InstanceId=id5ab1760-3498-4d95-9687-a91545ef90b3
```

Return:

```
{
  "RequestId" : "9F2188AC-AFAC-4F43-B452-C88463B9F069",
  "Disks" : {
    "Disk" : [
      {
        "DiskId" : "1008-27930",
        "Size" : 20,
        "Type" : "system"}]
    }
  }
  System Disk ID is: "1008-27930"
```

2) Create a snapshot for the system disk:

Request:

```
https://ecs.aliyuncs.com/?Action=CreateSnapshot&InstanceId=id5ab1760-3498-4d95-9687-a91545ef90b3&DiskId=1008-27930&SnapshotName=mytesthost1-init
```

Return:

```
{
  "RequestId" : "5CA4F9E6-81D2-42E1-A317-4C25284C6939",
  "SnapshotId" : "1008-27930-1097358"
}
```

Snapshot ID created is "1008-27930-1097358" .

3) Query the snapshot process and when it is 100, the snapshot has been created.

Request:

```
https://ecs.aliyuncs.com/?Action=DescribeSnapshotAttribute&RegionId=cn-hangzhou-dg-a01&SnapshotId=1008-27930-1097358
```

Return:

```
{
  "RequestId" : "8307863A-1415-40EF-9520-8974871E651C",
```

```
"SnapshotId" : "1008-27930-1097358",
"SnapshotName" : "mytesthost1-snp-init",
"Progress" : "100",
"CreationTime" : "2013-05-19T03:19Z"
}
```

2、 Make a custom image with the snapshot created.

Request:

```
https://ecs.aliyuncs.com/?Action=CreateImage&RegionId=cn-hangzhou-dg-a01&SnapshotId=1008-27930-1097358&Description=for creating test instances
```

Return:

```
{
"RequestId" : "38C930E9-5CE9-4E24-A392-8538FC20D503",
"ImageId" : "m8a1f80fe-ed9d-4156-a7a8-432f66305c36"
}
```

3、 “Copy” the machine

With the custom image, a machine with the same configuration can be “copied” and the second ECS instance will be created by ImageId=m8a1f80fe-ed9d-4156-a7a8-432f66305c36.

It is assumed that the following is attributes of the ECS instance created:

```
{
"RequestId" : "850ED7ED-A4D5-40A1-A7EF-C33B74B1296B",
"InstanceId" : "i6b47cd72-843f-4558-b911-2776acae06fb",
"ImageId" : "m8a1f80fe-ed9d-4156-a7a8-432f66305c36",
"RegionId" : "cn-hangzhou-dg-a01",
"ZoneId" : "cn-hangzhou-gy002-a",
"InstanceType" : "ecs.t1.small",
"HostName" : "mytesthost2",
"Status" : "Stopped",
"SecurityGroupIds" : {
"SecurityGroupId" : [
"g1f91e6e8-3c4b-4923-98dd-78aacbd09d17"
]
},
"PublicIpAddress" : {
"IpAddress" : [
"10.10.10.173"
]
},
"InnerIpAddress" : {
"IpAddress" : [
"10.32.148.152"
]
},
"InternetMaxBandwidthIn" : 2,
```

```
"InternetMaxBandwidthOut" : 2,  
"SerialNumber" : "1fec6c01-7186-2c3e-fa10-a672b8c300ec"  
}
```

To facilitate demonstration, distinguish from the first ECS instance during access to indicate the Server Load Balancer configuration succeeds and change the Body characters in the /usr/share/nginx/www/default/index.html into: "Welcome to nginx on mytesthost2! "

2.Create a LoadBalancer of Server Load Balancer

Request:

```
https://slb.aliyuncs.com/?Action=CreateLoadBalancer&RegionId=cn-hangzhou-dg-a01
```

Return: { "RequestId" : "3DE96B24-E2AB-4DFA-9910-1AADD60E13A5" , "LoadBalancerId" : "13ebb82ceaa-cn-hangzhou-dg-a01" , "Address" : "10.10.10.77" }

LoadBalancer ID created: "13ebb82ceaa-cn-hangzhou-dg-a01" ; add the HTTP Listener to listen to the HTTP 80 port for the LoadBalancer. When the Server Load Balancer under the Layer-4 protocol (TCP) is used, the similar method can be adopted.

Request:

```
https://slb.aliyuncs.com/?Action=CreateLoadBalancerHttpListener&LoadBalancerId=13ebb82ceaa-cn-hangzhou-dg-a01&ListenerPort=80&BackendServerPort=80&ListenerStatus=active
```

Activate the LoadBalancer into the operating condition.

Request:

```
https://slb.aliyuncs.com/?Action=SetLoadBalancerStatus&LoadBalancerId=13ebb82ceaa-cn-hangzhou-dg-a01&LoadBalancerStatus=active
```

3.Add a Backend Server

Add the backend server of Server Load Balancer through the AddBackendServers interface. After adding the backend server, the Server Load Balancer forwards requests into the backend server according to the configuration.

Request:

```
https://slb.aliyuncs.com/?Action=AddBackendServers&LoadBalancerId=13ebb82ceaa-cn-hangzhou-dg-a01&BackendServers=[{"ServerId":"id5ab1760-3498-4d95-9687-a91545ef90b3"}]
```

Return:

```
{
  "RequestId" : "FA2F2172-63F2-409D-927C-86BD1D536F13",
  "LoadBalancerId" : "13ebb82ceaa-cn-hangzhou-dg-a01",
  "BackendServers" : {
    "BackendServer" : [
      {
        "ServerId" : "id5ab1760-3498-4d95-9687-a91545ef90b3",
        "Weight" : 100
      }
    ]
  }
}
```

Add the second instance (or add multiple instances for one time):

Request:

```
https://slb.aliyuncs.com/?Action=AddBackendServers&LoadBalancerId=13ebb82ceaa-cn-hangzhou-dg-a01&BackendServers=[{"ServerId":"i6b47cd72-843f-4558-b911-2776acae06fb"}]
```

Return:

```
{
  "RequestId" : "C61FAD0A-2E87-4D0C-80B0-95AB758FCA70",
  "LoadBalancerId" : "13ebb82ceaa-cn-hangzhou-dg-a01",
  "BackendServers" : {
    "BackendServer" : [
      {
        "ServerId" : "id5ab1760-3498-4d95-9687-a91545ef90b3",
        "Weight" : 100
      },
      {
        "ServerId" : "i6b47cd72-843f-4558-b911-2776acae06fb",
        "Weight" : 100
      }
    ]
  }
}
```

Query the configuration information of LoadBalancer.

Request:

```
https://slb.aliyuncs.com/?Action=DescribeLoadBalancerAttribute&LoadBalancerId=13ebb82ceaa-cn-hangzhou-dg-a01
```

Return:

```
{
  "RequestId" : "4747E9AE-ADFD-412D-B523-C1CBD45A2154",
```

```
"LoadBalancerId" : "13ebb82ceaa-cn-hangzhou-dg-a01",
"Address" : "10.10.10.77",
"IsPublicAddress" : "true",
"ListenerPorts" : {
"ListenerPort" : [
80
]
},
"BackendServers" : {
"BackendServer" : [
{
"ServerId" : "id5ab1760-3498-4d95-9687-a91545ef90b3",
"Weight" : 100
},
{
"ServerId" : "i6b47cd72-843f-4558-b911-2776acae06fb",
"Weight" : 100
}
]
}
}
```

Now, two backend servers have been added into the LoadBalancer of Server Load Balancer. To test you can access <http://10.10.10.77/> and continuously refresh it so that the page shows alternately: "Welcome to nginx on mytesthost1!" and "Welcome to nginx on mytesthost2!". If the two messages are displayed accordingly, requests are forwarded correctly to two servers configured.