

This section describes how to configure server load balancing on the Cisco 4700 Series Application Control Engine (ACE) appliance.

Guide Contents

Overview
Setting Up an ACE Appliance
Creating a Virtual Context
Configuring Access Control Lists
Configuring Role-Based Access Control
Configuring Server Load Balancing (this section)
Configuring a Load-Balancing Predictor
Configuring Server Persistence Using Stickiness
Configuring SSL Security
Configuring Health Monitoring Using Health Probes

Contents

- [1 Overview](#)
- [2 Configuring Layer 7 Server Load Balancing Using the Device Manager GUI](#)
- [3 Configuring Layer 7 Server Load Balancing Using the CLI](#)
 - ◆ [3.1 Configuring Real Servers](#)
 - ◆ [3.2 Creating a Server Farm](#)
 - ◆ [3.3 Creating a Virtual Server Traffic Policy](#)

Overview

After reading this section, you should have an understanding of the basic server load-balancing capabilities provided by the ACE appliance. You should also be able to configure a virtual server for Layer 7 load-balancing purposes.

When there is a client request for web services, a load-balancing device decides to which server it should send the request. For example, a client request may consist of an HTTP GET for a web page or an FTP GET to download a file. The ACE, as a server load balancer, selects a server that can successfully fulfill the client request in the shortest amount of time without overloading either the server or the server farm as a whole.

The ACE uses a virtual server to intercept web traffic to a website. A virtual server allows multiple real

servers to appear as one for load-balancing purposes. A virtual server, also called a Virtual IP (VIP), is defined by its IP address, the protocol used (for example, UDP or TEC), and the port address.

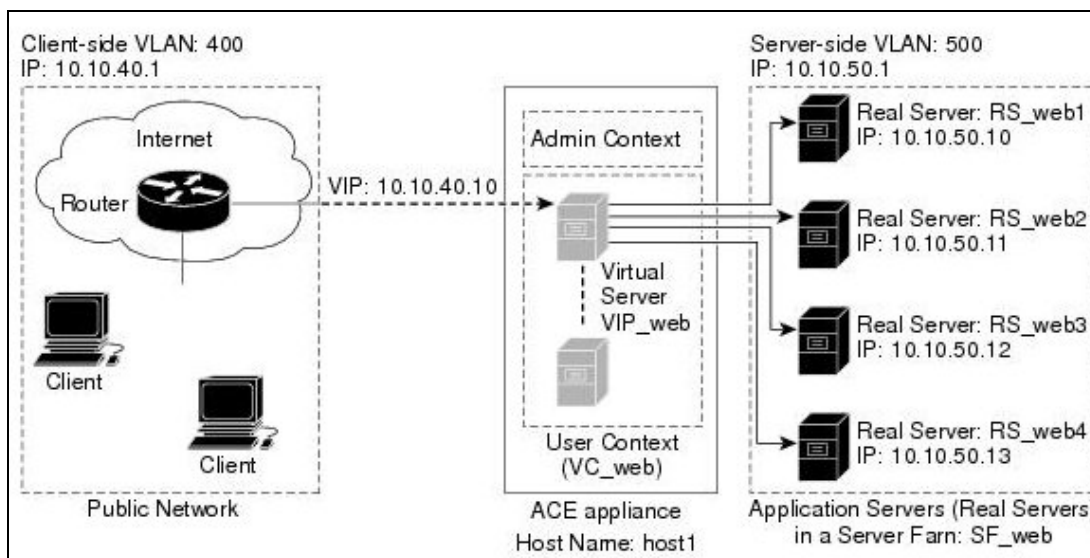
Multiple servers grouped together in server farms are assigned to each virtual server and the ACE appliance carries out load balancing across them. Real servers are dedicated servers that provide services to clients?for example, delivery of HTTP or XML content. Server farms contain the same content and typically reside in the same physical location in a data center.

You can configure the ACE for server load balancing by following these steps:

1. Create a virtual server.
2. Configure the real servers and associate them with a server farm.
3. Assign the server farm to the virtual server.
4. Deploy the configuration.

This section describes how to configure a virtual server using either the Device Manager GUI or the CLI, using the network setup example illustrated in Figure 1.

Figure 1 Example Server Load-Balancing Setup



The configuration of the example setup is as follows:


- A virtual server VS_web is created with a virtual IP address 10.10.40.10 to forward the client traffic from VLAN 400 to the application servers in VLAN 500.
- There are four real servers grouped into the server farm SF_web.
- The virtual server uses a round-robin predictor to forward the client requests to one of the real servers in the server farm.

Configuring Layer 7 Server Load Balancing Using the Device Manager GUI

You can configure Layer 7 server load balancing using the Device Manager GUI by following these steps:


1. Choose **Load Balancing > Virtual Servers**. The Virtual Servers pane appears (Figure 2). Choose the user context **VC_web**.

Figure 2 Virtual Servers Pane



2. Click **Add** to add a new virtual server. The Virtual Server configuration window appears (Figure 3).

Figure 3 Properties in the Virtual Server Configuration Window



By default, the Basic View configuration option is selected and the Properties section is open.

3. In Properties, enter the following virtual server attributes. Leave the remaining attributes blank or with their default values.

◇ VIP Name: VS_web

◇ VIP IP: 10.10.40.10

Note A client request targeted at a website (a URL) is translated to an IP address according to the Domain Name System (DNS). A virtual IP address assigned to a virtual server is the IP address that corresponds to the URL of the website from which the client requests services.

◇ Protocol: TCP

◇ Application Protocol: HTTP

◇ Port: 80

◇ VLAN: 400

4. In the Default L7 Load-Balancing Action section, choose **loadbalance** from the Primary Action drop-down list.

5. Choose ***New*** from the Server Farm drop-down list to configure a new server farm.

6. Enter the following server farm attributes. Leave the remaining attributes blank or with their default values.

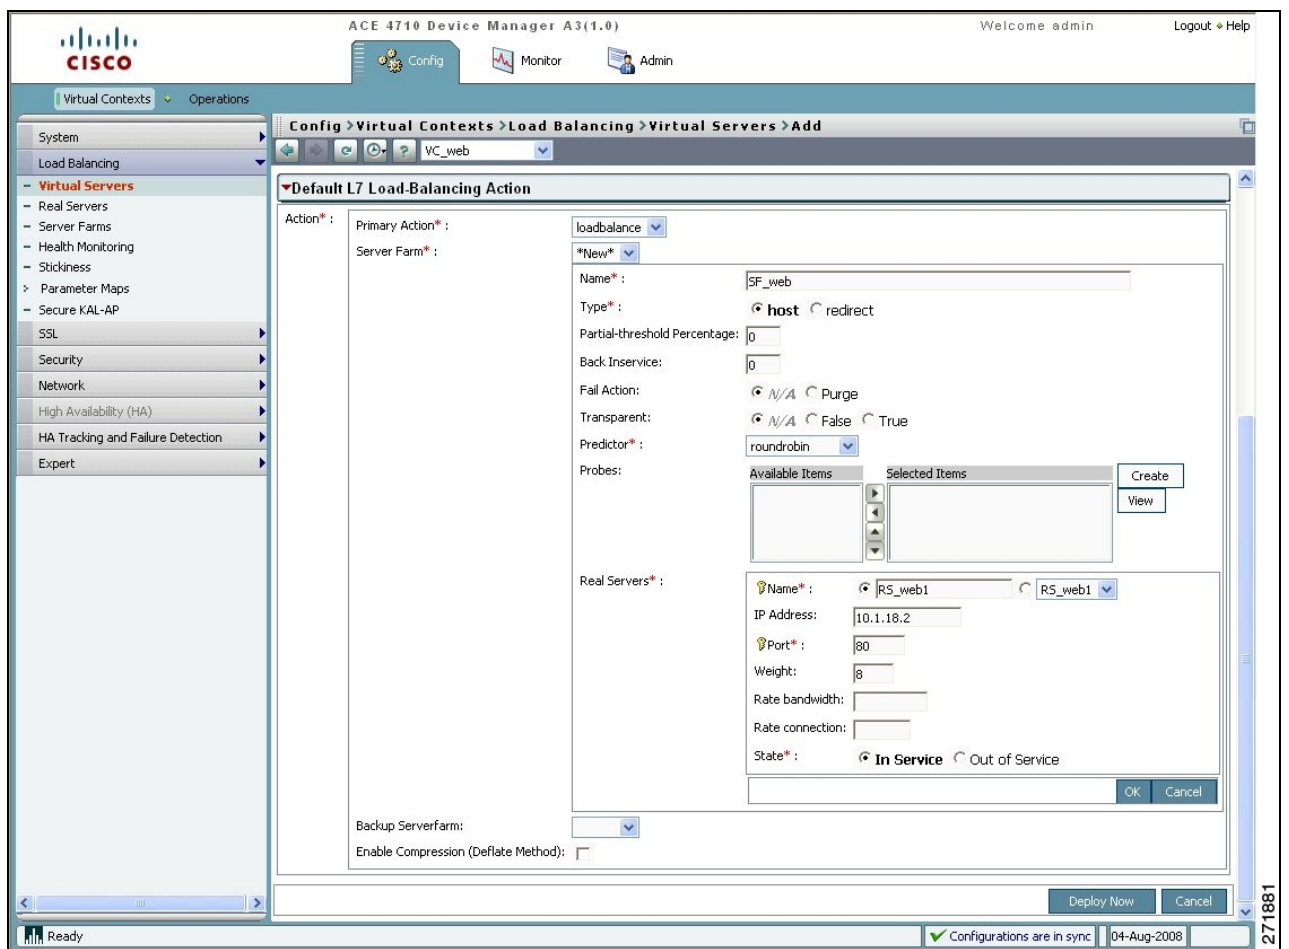
◇ Name: SF_web

◇ Type: host

◇ Predictor: roundrobin

7. Click **Add** to add a new entry to the Real Servers pane. A new entry appears in the Real Servers pane (Figure 4).

Figure 4 Real Servers Pane in the Virtual Server Configuration Window



8. Enter the following attributes for the first real server to be configured. Leave the remaining attributes blank or with their default values.

◇ Name: RS_web1

◇ IP Address: 10.10.50.10

◇ Port: 80

◇ Weight: 8

◇ State: In Service

Click **OK** to save the attributes of the first real server.

Note For information on how to configure a health probe, see the [Configuring Health Monitoring Using Health Probes](#) section.

9. Add three more entries to the Real Servers pane by repeating Steps 7 and 8 with the following real server names and corresponding IP addresses. Leave the remaining attributes with their default values.

For RS_web2, enter:

◇ Name: RS_web2

◇ IP Address: 10.10.50.11

◇ Port: 80

For RS_web3, enter:

◇ Name: RS_web3

◇ IP Address: 10.10.50.12

◇ Port: 80

For RS_web4, enter:


◇ Name: RS_web4

◇ IP Address: 10.10.50.13

◇ Port: 80

10. Click **Deploy Now** at the bottom of the window to save your settings for the virtual server. The Virtual Servers pane reappears (Figure 5). The newly configured virtual server appears in the pane and is in the Inservice state, which means that the virtual server is in use as a destination for server load balancing.

Figure 5 Virtual Servers Pane with a Virtual Server Created



Configuring Layer 7 Server Load Balancing Using the CLI

You can configure Layer 7 server load balancing using the command-line interface (CLI). This section contains the following subsections:

- [Configuring Real Servers](#)
- [Creating a Server Farm](#)
- [Creating a Virtual Server Traffic Policy](#)

Configuring Real Servers

Configure real servers on the ACE using the CLI by following these steps:

1. Verify that you are operating in the desired context by checking the CLI prompt. If necessary, change to the correct context.

```
host1/Admin# changeto VC_web
```

```
host1/VC_web#
```

2. Enter configuration mode.

```
host1/VC_web# config
```

3. Create a real server named RS_web1 as type host (the default).

```
host1/VC_web(config)# rserver RS_web1
```

```
host1/VC_web(config-rserver-host)#
```

4. Enter a description of the real server.

```
host1/VC_web(config-rserver-host)# description content server web-one
```

5. Assign the real server with an IP address of 10.10.50.10.

```
host1/VC_web(config-rserver-host)# ip address 10.10.50.10
```

6. Place the real server in service and exit configuration mode.

```
host1/VC_web(config-rserver-host)# inservice
```

```
host1/VC_web(config-rserver-host)# exit
```

```
host1/VC_web(config)#
```

7. Add three more real servers by repeating Steps 3 through 6, using the following real server names, descriptions, and IP addresses.

For RS_web2, enter:

◇ Name: RS_web2

◇ Description: content server web-two

◇ IP Address: 10.10.50.11

For RS_web3, enter:

◇ Name: RS_web3

◇ Description: content server web-three

◇ IP Address: 10.10.50.12

For RS_web4, enter:

◇ Name: RS_web4

◇ Description: content server web-four

◇ IP Address: 10.10.50.13

8. Display the configuration of the real servers.

```
host1/VC_web(config)# do show running-config rserver
```


Creating a Server Farm

After you create and configure the real servers, you can create a server farm and associate the real servers with it. Create a server farm by following these steps:

1. Create a server farm of type host (the default) named SF_web.

```
host1/VC_web(config)# serverfarm SF_web
```

```
host1/VC_web(config-sfarm-host)#
```

2. Associate real server RS_web1 to the server farm through port 80.

```
host1/VC_web(config-sfarm-host)# rserver RS_web1 80
```

```
host1/VC_web(config-sfarm-host-rs)#
```

3. Place the real server in service within the server farm and exit configuration mode.

```
host1/VC_web(config-sfarm-host-rs)# inservice
```

```
host1/VC_web(config-sfarm-host-rs)# exit
```

```
host1/VC_web(config-sfarm-host)#
```

Note Before you can start sending connections to a real server in a server farm, you must place it in service. Otherwise, the ACE considers it out of service and the server farm cannot receive or respond to client requests.

4. Similarly, associate the RS_web2, RS_web3, and RS_web4 real servers with the SF_web server farm.

```
host1/VC_web(config-sfarm-host)# rserver RS_web2 80
```

```
host1/VC_web(config-sfarm-host-rs)# inservice
```

```
host1/VC_web(config-sfarm-host-rs)# exit
```

```
host1/VC_web(config-sfarm-host)# rserver RS_web3 80
```

```
host1/VC_web(config-sfarm-host-rs)# inservice
```

```
host1/VC_web(config-sfarm-host-rs)# exit
```

```
host1/VC_web(config-sfarm-host)# rserver RS_web4 80
```

```
host1/VC_web(config-sfarm-host-rs)# inservice
```

```
host1/VC_web(config-sfarm-host-rs)# exit
```

5. Exit server farm configuration mode.

```
host1/VC_web(config-sfarm-host)# exit
```

```
host1/VC_web(config)#
```

6. Display the information for the real servers and verify that the real servers appear as operational (even though network connectivity has not been established).

```
host1/VC_web(config)# do show rserver RS_web1
```

```
host1/VC_web(config)# do show rserver RS_web2
```

```
host1/VC_web(config)# do show rserver RS_web3
```

```
host1/VC_web(config)# do show rserver RS_web4
```

7. Display how the ACE populates the ARP table with the real servers.

```
host1/VC_web(config)# do show arp
```

Creating a Virtual Server Traffic Policy

You can create a virtual server traffic policy on the ACE by following these steps:

1. Create a Layer 7 server load-balancing policy map named PM_LB to match the class maps in the order in which they occur for load balancing.

```
host1/VC_web(config)# policy-map type loadbalance first-match PM_LB
```

```
host1/VC_web(config-pmap-lb)#
```

Note The ACE uses a class map to specify a series of flow match criteria (traffic classifications). The ACE uses a policy map to define a series of actions (functions) that you want applied to a set of classified inbound traffic.

2. For a simple load-balancing policy, assign the ACE default class map which contains an implicit match any statement to match any traffic classification.

```
host1/VC_web(config-pmap-lb)# class class-default
```

```
host1/VC_web(config-pmap-lb-c)#
```

3. Add the server farm SF_web to the Layer 7 server load-balancing policy map and exit configuration mode.

```
host1/VC_web(config-pmap-lb-c)# serverfarm SF_web
```

```
host1/VC_web(config-pmap-c)# exit
```

```
host1/VC_web(config-pmap)# exit
```

```
host1/VC_web(config)#
```

4. Create a Layer 3 and Layer 4 load-balancing class map VS_web.

```
host1/VC_web(config)# class-map VS_web
```

```
host1/VC_web(config-cmap)#
```

5. Define a match statement for the IP address 10.10.40.10 for any IP protocol and exit configuration mode.

```
host1/VC_web(config-cmap)# match virtual-address 10.10.40.10 255.255.255.0 tcp eq 80
```

```
host1/VC_web(config-cmap)# exit
```

```
host1/VC_web(config)#
```

6. Create a Layer 3 and Layer 4 multi-match policy map to direct classified incoming requests to the load-balancing policy map.

```
host1/VC_web(config)# policy-map multi-match PM_multi_match  
  
host1/VC_web(config-pmap)#
```

7. Associate the Layer 3 and Layer 4 class map VS_web with the policy map.

```
host1/VC_web(config-pmap)# class VS_web  
  
host1/VC_web(config-pmap-c)#
```

8. Associate the Layer 7 load-balancing policy map PM_LB with the Layer 3 and Layer 4 policy map.

```
host1/VC_web(config-pmap-c)# loadbalance policy PM_LB  
  
host1/VC_web(config-pmap-lb-c)#
```

9. Enable a VIP for load-balancing operations and exit configuration mode.

```
host1/VC_web(config-pmap-lb-c)# loadbalance vip inservice  
  
host1/VC_web(config-pmap-c)# exit  
  
host1/VC_web(config-pmap)# exit  
  
host1/VC_web(config)#
```

10. Access the interface to which you want to apply the multi-match policy map.

```
host1/VC_web(config)# interface vlan 400  
  
host1/VC_web(config-if)#
```

11. Apply the multi-match policy map PM_multi_match.

```
host1/VC_web(config-if)# service-policy input PM_multi_match  
  
host1/VC_web(config-if)# exit  
  
host1/VC_web(config)#
```

12. Save the running configuration to the startup configuration.

```
host1/VC_web(config)# do copy running-config startup-config
```

13. Display the service policy state for the PM_multi_match policy map.

```
host1/VC_web(config)# do show service-policy PM_multi_match
```

In this section, you have configured a virtual server for load-balancing HTTP traffic. In the next section, you will configure a load-balancing predictor to forward client requests to the appropriate real servers.