



CHAPTER 6

Configuring Server Load Balancing

This chapter describes how to configure server load balancing (SLB) on the Cisco Application Control Engine (ACE) module.

This chapter contains the following sections:

- [Information About Server Load Balancing](#)
- [Configuring Server Load Balancing](#)
- [Configuration Example for Configuring Server Load Balancing](#)
- [Where to Go Next](#)

Information About Server Load Balancing

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

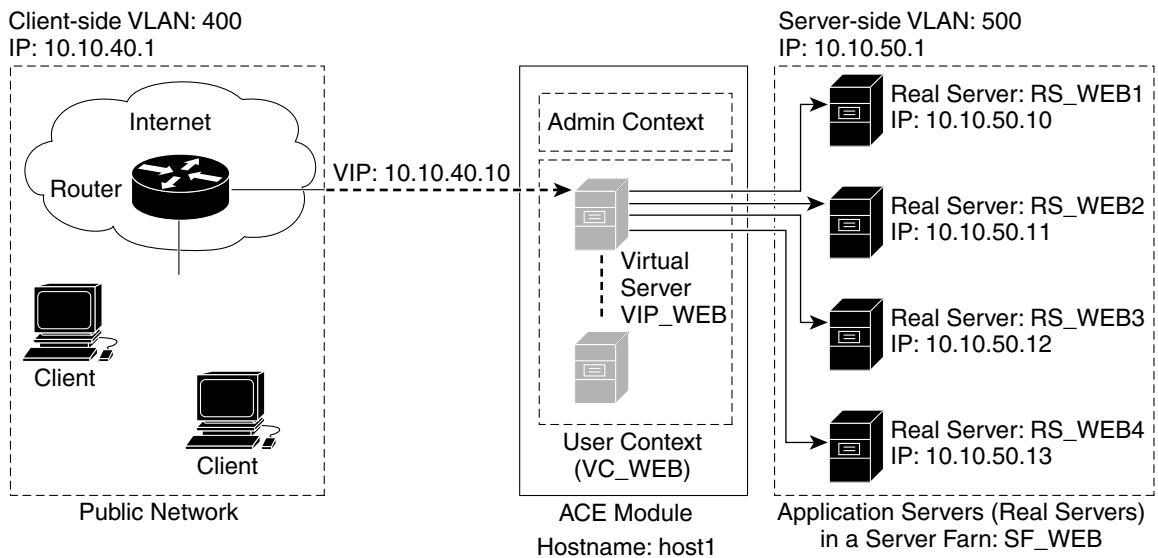
When there is a client request for web services, a load-balancing device decides to which server it should send the request. For example, a client request may consist of an HTTP GET for a web page or an FTP GET to download a file. The ACE, as a server load balancer, selects a server that can successfully 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 TCP), and the port address.

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

This chapter describes how to configure a virtual server using the network example in [Figure 6-1](#).

Figure 6-1 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 Server Load Balancing

This section describes how to configure server load balancing. It contains the following topics:

- [Task Flow for Configuring Server Load Balancing](#)
- [Configuring Real Servers](#)
- [Creating a Server Farm](#)
- [Creating a Virtual Server Traffic Policy](#)

Task Flow for Configuring Server Load Balancing

Follow these steps to configure server load balancing:

-
- Step 1** Configure real servers.
- Step 2** Create a server farm.
- Step 3** Create a virtual server traffic policy.
-

Configuring Real Servers

Procedure

	Command	Purpose
Step 1	changeto <i>context</i> Example: host1/Admin# changeto VC_WEB host1/VC_WEB#	Changes to the correct context if necessary. Check the CLI prompt to verify that you are operating in the desired context.
Step 2	config Example: host1/VC_WEB# config host1/VC_WEB(config)#	Enters configuration mode.
Step 3	rserver <i>name</i> Example: host1/VC_WEB(config)# rserver RS_WEB1 host1/VC_WEB(config-rserver-host)#	Creates a real server named RS_WEB1 as type host (the default).
Step 4	description <i>string</i> Example: host1/VC_WEB(config-rserver-host)# description content server web-one	Enters a description of the real server.
Step 5	ip address <i>address</i> Example: host1/VC_WEB(config-rserver-host)# ip address 10.10.50.10	Assigns the real server with an IP address of 10.10.50.10.
Step 6	inservice Example: host1/VC_WEB(config-rserver-host)# inservice	Places the real server in service.
Step 7	exit Example: host1/VC_WEB(config-rserver-host)# exit host1/VC_WEB(config)#	Exits real server host configuration mode.

Command	Purpose
<p>Step 8 <code>rserver name</code></p> <p>Example: <pre>host1/VC_WEB(config)# rserver RS_WEB2 host1/VC_WEB(config-rserver-host)#</pre></p> <p>description string</p> <p>Example: <pre>host1/VC_WEB(config-rserver-host)# description content server web-two</pre></p> <p>ip address address</p> <p>Example: <pre>host1/VC_WEB(config-rserver-host)# ip address 10.10.50.11</pre></p> <p>inservice</p> <p>Example: <pre>host1/VC_WEB(config-rserver-host)# inservice</pre></p> <p>exit</p> <p>Example: <pre>host1/VC_WEB(config-rserver-host)# exit host1/VC_WEB(config)#</pre></p>	<p>Add three more real servers by repeating Steps 3 through 7, using the following real server names, descriptions, and IP addresses.</p> <p>For RS_WEB2, enter:</p> <ul style="list-style-type: none"> • Name: RS_WEB2 • Description: content server web-two • IP Address: 10.10.50.11 <p>For RS_WEB3, enter:</p> <ul style="list-style-type: none"> • Name: RS_WEB3 • Description: content server web-three • IP Address: 10.10.50.12 <p>For RS_WEB4, enter:</p> <ul style="list-style-type: none"> • Name: RS_WEB4 • Description: content server web-four <p>IP Address: 10.10.50.13</p>
<p>Step 9 <code>do show running-config rserver</code></p> <p>Example: <pre>host1/VC_WEB(config)# do show running-config rserver</pre></p>	<p>Displays the configuration of the real servers.</p>
<p>Step 10 <code>do copy running-config startup-config</code></p> <p>Example: <pre>host1/Admin(config)# do copy running-config startup-config</pre></p>	<p>(Optional) Copies the running configuration to the startup configuration.</p>

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.

Procedure

	Command	Purpose
Step 1	serverfarm <i>name</i> Example: host1/VC_WEB(config)# serverfarm SF_WEB host1/VC_WEB(config-sfarm-host)#	Creates a server farm of type host (the default) named SF_WEB.
Step 2	rserver <i>name</i> [<i>port</i>] Example: host1/VC_WEB(config-sfarm-host)# rserver RS_WEB1 80 host1/VC_WEB(config-sfarm-host-rs)#	Associates real server RS_WEB1 with the server farm through port 80. Specifying a port number is optional. If you do not specify a port number, the ACE does not perform PAT and the destination port that was used from the client to the VIP will also be used from the VIP to the real server. When you specify a port number, it is the only destination port that the ACE uses from the VIP to the real server.
Step 3	inservice Example: host1/VC_WEB(config-sfarm-host-rs)# inservice	Places the real server in service within the server farm. 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.
Step 4	exit Example: host1/VC_WEB(config-sfarm-host-rs)# exit host1/VC_WEB(config-sfarm-host)#	Exits server farm host real server configuration mode.

	Command	Purpose
<p>Step 5</p> <pre>rserver name [port] inservice</pre> <p>Example:</p> <pre>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</pre>	<p>Similarly, associates the RS_WEB2, RS_WEB3, and RS_WEB4 real servers with the SF_WEB server farm and places the real server in service. Repeat steps 2 through 4 to configure the remaining real servers in the server farm.</p>	
<p>Step 6</p> <pre>exit</pre> <p>Example:</p> <pre>host1/VC_WEB(config-sfarm-host)# exit host1/VC_WEB(config)#</pre>	<p>Exits server farm host configuration mode.</p>	
<p>Step 7</p> <pre>do show rserver name</pre> <p>Example:</p> <pre>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</pre>	<p>Displays the information for the real servers.</p> <p>Note The real server status is shown as ARP_FAILED because network connectivity has not been established yet.</p>	
<p>Step 8</p> <pre>do copy running-config startup-config</pre> <p>Example:</p> <pre>host1/VC_WEB(config)# do copy running-config startup-config</pre>	<p>(Optional) Copies the running configuration to the startup configuration.</p>	

Creating a Virtual Server Traffic Policy

Procedure

	Command	Purpose
Step 1	<pre>policy-map type loadbalance first-match name Example: host1/VC_WEB(config)# policy-map type loadbalance first-match PM_LB host1/VC_WEB(config-pmap-lb)#</pre>	<p>Creates 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.</p> <p>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.</p>
Step 2	<pre>class class-default Example: host1/VC_WEB(config-pmap-lb)# class class-default host1/VC_WEB(config-pmap-lb-c)#</pre>	<p>For a simple load-balancing policy, assigns the ACE default class map which contains an implicit match any statement to match any traffic classification.</p>
Step 3	<pre>serverfarm name Example: host1/VC_WEB(config-pmap-lb-c)# serverfarm SF_WEB</pre>	<p>Adds the server farm SF_WEB to the Layer 7 server load-balancing policy map and exits configuration mode.</p>
Step 4	<pre>exit Example: host1/VC_WEB(config-pmap-c)# exit host1/VC_WEB(config-pmap)# exit host1/VC_WEB(config)#</pre>	<p>Exits policy map class configuration mode. Exits policy map configuration mode.</p>
Step 5	<pre>class-map {match-all match-any type} name Example: host1/VC_WEB(config)# class-map VS_WEB host1/VC_WEB(config-cmap)#</pre>	<p>Creates a Layer 3 and Layer 4 load-balancing class map VS_WEB. The default is match-all.</p>
Step 6	<pre>match virtual-address address netmask tcp eq port Example: host1/VC_WEB(config-cmap)# match virtual-address 10.10.40.10 255.255.255.255 tcp eq 80</pre>	<p>Defines a match statement for the IP address 10.10.40.10, the TCP IP protocol, and port 80.</p>
Step 7	<pre>exit Example: host1/VC_WEB(config-cmap)# exit host1/VC_WEB(config)#</pre>	<p>Exits class map configuration mode.</p>

	Command	Purpose
Step 8	<p>policy-map multi-match <i>name</i></p> <p>Example: host1/VC_WEB(config)# policy-map multi-match PM_MULTI_MATCH host1/VC_WEB(config-pmap)#</p>	Creates a Layer 3 and Layer 4 multi-match policy map to direct classified incoming requests to the load-balancing policy map.
Step 9	<p>class <i>name</i></p> <p>Example: host1/VC_WEB(config-pmap)# class VS_WEB host1/VC_WEB(config-pmap-c)#</p>	Associates the VS_WEB Layer 3 and Layer 4 class map that you created in Step 4 with the PM_MULTI_MATCH policy map.
Step 10	<p>loadbalance policy <i>name</i></p> <p>Example: host1/VC_WEB(config-pmap-c)# loadbalance policy PM_LB host1/VC_WEB(config-pmap-lb-c)#</p>	Associates the PM_LB Layer 7 load-balancing policy map with the PM_MULTI_MATCH Layer 3 and Layer 4 policy map.
Step 11	<p>loadbalance vip inservice</p> <p>Example: host1/VC_WEB(config-pmap-lb-c)# loadbalance vip inservice</p>	Enables a VIP for load-balancing operations.
Step 12	<p>exit</p> <p>Example: host1/VC_WEB(config-pmap-c)# exit host1/VC_WEB(config-pmap)# exit host1/VC_WEB(config)#</p>	Exits policy map class configuration mode. Exits policy map configuration mode.
Step 13	<p>interface vlan <i>vlan_id</i></p> <p>Example: host1/VC_WEB(config)# interface vlan 400 host1/VC_WEB(config-if)#</p>	Accesses the interface to which you want to apply the multi-match policy map.
Step 14	<p>service-policy input <i>policy_name</i></p> <p>Example: host1/VC_WEB(config-if)# service-policy input PM_MULTI_MATCH</p>	Applies the PM_MULTI_MATCH Layer 3 and Layer 4 policy map.
Step 15	<p>exit</p> <p>Example: host1/VC_WEB(config-if)# exit host1/VC_WEB(config)# exit host1/VC_WEB#</p>	Exits interface configuration mode. Exits configuration mode.
Step 16	<p>show service-policy <i>policy_name</i></p> <p>Example: host1/VC_WEB# show service-policy PM_MULTI_MATCH</p>	Displays the service policy state for the PM_MULTI_MATCH policy map.
Step 17	<p>copy running-config startup-config</p> <p>Example: host1/VC_WEB# copy running-config startup-config</p>	(Optional) Copies the running configuration to the startup configuration.

Configuration Example for Configuring Server Load Balancing

The following example shows how to configure server load balancing. The commands that you have configured in this chapter appear in bold text.

```
switch/VC_WEB(config)# do show running config
Generating configuration...

access-list INBOUND line 8 extended permit ip any any

rserver host RS_WEB1
  description content server web-one
  ip address 10.10.50.10
  inservice
rserver host RS_WEB2
  description content server web-two
  ip address 10.10.50.11
  inservice
rserver host RS_WEB3
  description content server web-three
  ip address 10.10.50.12
  inservice
rserver host RS_WEB4
  description content server web-four
  ip address 10.10.50.13
  inservice

serverfarm host SF_WEB
  rserver RS_WEB1 80
  inservice
  rserver RS_WEB2 80
  inservice
  rserver RS_WEB3 80
  inservice
  rserver RS_WEB4 80
  inservice

class-map type management match-any REMOTE_ACCESS
  description Remote access traffic match
  2 match protocol ssh any
  3 match protocol telnet any
  4 match protocol icmp any
class-map match-all VS_WEB
  2 match virtual-address 10.10.40.10 tcp eq www

policy-map type management first-match REMOTE_MGMT_ALLOW_POLICY
  class REMOTE_ACCESS
  permit
policy-map type loadbalance first-match PM_LB
  class class-default
  serverfarm SF_WEB
policy-map multi-match PM_MULTI_MATCH
  class VS_WEB
  loadbalance vip inservice
  loadbalance policy PM_LB

service-policy input REMOTE_MGMT_ALLOW_POLICY

interface vlan 400
  description Client connectivity on VLAN 400
  ip address 10.10.40.1 255.255.255.0
  access-group input INBOUND
  service-policy input PM_MULTI_MATCH
  no shutdown
```

```
interface vlan 500
  description Server connectivity on VLAN 500
  ip address 10.10.50.1 255.255.255.0
  no shutdown

ip route 0.0.0.0 0.0.0.0 172.25.91.1

domain DOMAIN1
add-object all

username USER1 password 5 $1$vAN9gQDI$MmbmjQgJPj45lxbtzXPpB1 role SLB-Admin domain
DOMAIN1
```

Where to Go Next

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