



CHAPTER 6

Configuring Virtual Servers

Date: 2/17/11

This chapter describes how to configure virtual servers for load balancing on the Cisco Application Control Engine (ACE) using Cisco Application Networking Manager (ANM).



Note

When naming ACE objects (such as a real server, virtual server, parameter map, class map, health probe, and so on), enter an alphanumeric string of 1 to 64 characters, which can include the following special characters: underscore (_), hyphen (-), dot (.), and asterisk (*). Spaces are not allowed.

If you are using ANM with an ACE module or ACE appliance and you configure a named object at the ACE CLI, keep in mind that ANM does not support all of the special characters that the ACE CLI allows you to use when configuring a named object. If you use special characters that ANM does not support, you may not be able to import or manage the ACE using ANM.

This chapter includes the following sections:

- [Information About Load Balancing, page 6-1](#)
- [Configuring Virtual Servers, page 6-2](#)
- [Managing Virtual Servers, page 6-67](#)
- [Deploying Virtual Servers, page 6-76](#)

Information About Load Balancing

Server load balancing (SLB) is the process of deciding to which server a load balancer should send a client request for service. For example, a client request can consist of an HTTP GET for a web page or an FTP GET to download a file. The load balancer selects the server that can successfully fulfill the client request and in the shortest amount of time without overloading either the server or the server farm as a whole.

Depending on the load-balancing algorithm or predictor that you configure, the ACE performs a series of checks and calculations to determine the server that can best service each client request. The ACE bases server selection on several factors, including the server with the fewest connections with respect to load, source or destination address, cookies, URLs, or HTTP headers.

ANM allows you to configure load balancing using:

- Virtual servers—See [Configuring Virtual Servers, page 6-2](#).

- Real servers—See [Configuring Real Servers, page 7-5](#).
- Server farms—See [Configuring Server Farms, page 7-22](#).
- Predictor methods—See [Configuring the Predictor Method for Server Farms, page 7-31](#)
- Health probes—See [Configuring Health Monitoring for Real Servers, page 7-42](#)
- Sticky groups—See [Configuring Sticky Groups, page 8-7](#).
- Parameter maps—See [Configuring Parameter Maps, page 9-1](#).

Configuring Virtual Servers

In a load-balancing environment, a virtual server is a construct that allows multiple physical servers to appear as one for load-balancing purposes. A virtual server is bound to physical services running on real servers in a server farm and uses IP address and port information to distribute incoming client requests to the servers in the server farm according to a specified load-balancing algorithm.

You use class maps to configure a virtual server address and definition. The load-balancing predictor algorithms (for example, round-robin, least connections, and so on) determine the servers to which the ACE sends connection requests.

This section includes the following topics:

- [Virtual Server Configuration and ANM, page 6-2](#)
- [Using ANM to Configure Virtual Servers, page 6-4](#)
- [Virtual Server Usage Guidelines, page 6-5](#)
- [Virtual Server Testing and Troubleshooting, page 6-5](#)
- [Virtual Server Configuration Procedure, page 6-7](#)

Virtual Server Configuration and ANM

This section identifies the constraints and framework used by ANM for virtual server configuration.

In ANM, a virtual server has the following attributes:

- A single Layer 3/Layer 4 match condition
You can specify only a single IP address (or single IP address range if a netmask is used) with only a single port (or port range). A single match condition greatly simplifies and aids virtual server configuration.
- A default Layer 7 action
- A Layer 7 policy map
- A Layer 3/Layer 4 class map
- A single multimatch policy map, a class-map match, and an action

Virtual server attributes also include the following:

- The virtual server multimatch policy map is associated with an interface or is global.
- The name of the virtual server is derived from the name of the Layer 3/Layer 4 class map.

Example 6-1 shows the minimum configuration statements required for a virtual server.

Example 6-1 Minimum Configuration Required for a Virtual Server

```
class-map match-all Example_VIP
  2 match virtual-address 10.10.10.10 tcp eq www
policy-map type loadbalance first-match Example_VIP-l7slb
  class class-default
    forward
policy-map multi-match int10
  class Example_VIP
    loadbalance policy Example_VIP-l7slb

interface vlan 10
  ip address 192.168.65.37 255.255.255.0
  service-policy input int10
  no shutdown
```

Note also the following items regarding the ANM and virtual servers:

- **Additional configuration options**

The Virtual Server configuration window allows you to configure additional items for a functional VIP. These items include server farms, sticky groups, real servers, probes, parameter maps, inspection, class maps, and inline match conditions. Because too many items on a window can be overwhelming, not all configuration options appear on the Virtual Server configuration window, such as sticky statics or backup real servers. These options are available elsewhere in the ANM interface instead of on the Virtual Server configuration window.

- **Configuration options and roles**

To support and maintain the separation of roles, some objects cannot be configured using the Virtual Server configuration window. These objects include SSL certificates, SSL keys, NAT pools, interface IP addresses, and ACLs. Providing these options as separate configuration options in the ANM interface ensures that a user who can view or modify virtual servers or aspects of virtual servers cannot create or delete virtual servers.

- **Changes to virtual servers using the CLI or Expert options can prevent further modifications in the Virtual Server configuration window**

If you create a virtual server using the Virtual Server configuration window, modify it using the CLI or Expert options (Config > Devices > Expert), and then attempt to modify it again using the Virtual Server configuration window, error messages will be displayed and you will not be able to modify the virtual server.

Related Topics

- [Configuring Virtual Servers, page 6-2](#)
- [Using ANM to Configure Virtual Servers, page 6-4](#)
- [Virtual Server Usage Guidelines, page 6-5](#)
- [Virtual Server Testing and Troubleshooting, page 6-5](#)
- [Virtual Server Configuration Procedure, page 6-7](#)

Using ANM to Configure Virtual Servers

Follow these guidelines when using ANM to configure virtual servers:

- **Virtual server configuration windows**

The ANM Virtual Server configuration windows are designed to aid you in configuring virtual servers by presenting configuration options that are relevant to your choices. For example, the protocols that you select in the Properties configuration subset determine the other configuration subsets that appear.

- **Use the virtual server configuration method that suits you**

The ANM Virtual Server configuration windows simplify the process of creating, modifying, and deploying virtual servers by displaying those options that you are most likely to use. In addition, as you specify attributes for a virtual server, such as protocols, the interface refreshes with related configuration options, such as Protocol Inspection or Application Acceleration and Optimization, which speeds virtual server configuration and deployment.

While Virtual Server configuration windows remove some configuration complexities, they have a few constraints that the Expert configuration options do not. If you are comfortable using the CLI, you can use the Expert options (such as Config > Devices > *context* > Expert > Class Maps or Policy or Config > Devices > *context* > Load Balancing > Parameter Maps to configure more complex attributes of virtual servers, traffic policies, and parameter maps.

- **Synchronizing virtual server configurations**

If you configure a virtual server using the CLI and then use the Sync option (Config > Devices > ACE > Sync) to synchronize configurations, the configuration that appears in ANM for the virtual server might not display all configuration options for that virtual server. The configuration that appears in ANM depends on a number of items, such as the protocols configured in class maps or the rules defined for policy maps.

For example, if you configure a virtual server on the CLI that includes a class map that can match any protocol, you will not see the virtual server Application Acceleration and Optimization configuration subset in ANM.

- **Modifying shared objects**

Modifying an object that is used by multiple virtual servers, such as a server farm, real server, or parameter map, could impact the other virtual servers. See the [“Shared Objects and Virtual Servers” section on page 6-9](#) for more information about modifying objects used by multiple virtual servers.

Related Topics

- [Configuring Virtual Servers, page 6-2](#)
- [Virtual Server Configuration and ANM, page 6-2](#)
- [Virtual Server Usage Guidelines, page 6-5](#)
- [Virtual Server Testing and Troubleshooting, page 6-5](#)
- [Virtual Server Configuration Procedure, page 6-7](#)

Virtual Server Usage Guidelines

The Virtual Server configuration window provides you with numerous configuration options. However, instead of setting every option in one pass, configure your virtual server in stages. The first stage should always be to establish basic “pass through” connectivity with simple load balancing and include minimal additional features. This level of setup should verify that ports, VLANs, interfaces, SSL termination (if applicable), and real servers have been set up properly, enabling basic connectivity.

After you establish this level of connectivity, additional virtual server features will be easier to configure and troubleshoot.

Common features to add to a working basic virtual server include:

- Health monitoring probes
- Session persistence (sticky)
- Additional real servers to a server farm
- Application protocol inspection
- Application acceleration and optimization (ACE appliance only)

[Table 6-1](#) identifies and describes virtual server configuration subsets with links to related topics for configuration information.

Related Topics

- [Configuring Virtual Servers, page 6-2](#)
- [Virtual Server Configuration and ANM, page 6-2](#)
- [Virtual Server Testing and Troubleshooting, page 6-5](#)
- [Virtual Server Configuration Procedure, page 6-7](#)

Virtual Server Testing and Troubleshooting

As outlined in the “[Virtual Server Usage Guidelines](#)” section on [page 6-5](#), first set up a basic virtual server that only enables connectivity and simple load balancing, such as round-robin between two real servers. Next, use a client, such as a web browser, to send a request from the client network to the virtual server's VIP address. If the request is successful, you can now make changes or add virtual server features.

If the request is not successful, begin virtual server troubleshooting as outlined in the following sequence:

1. Wait and retry your request after a minute or two, especially if the existing ACE configuration is large. It can take seconds or even minutes for configuration changes to affect how traffic is handled by ACE.
2. Click the **Details** button in the lower right of the Virtual Server page. The Details button displays the output of the **show service-policy** CLI command.
3. Verify that the VIP State in the **show service-policy** CLI command output is **INSERVICE**. If the VIP state is not **INSERVICE**, this may indicate the following:
 - The virtual server has been manually disabled in the configuration.
 - The real servers are all unreachable from ACE or manually disabled. If all of a virtual server's real servers are out of service due to one of those reasons, the virtual server itself will be marked **Out Of Service**.

4. Verify the Hit Count in the **show service-policy** CLI command output. Hit Count shows the number of requests received by ACE. This value should increase for each request attempted by your client. If the hit count does not increase with each request, this indicates that the request is not reaching your virtual server configuration.

This could be a problem with:

- A physical connection.
- VLAN or VLAN interface configuration.
- Missing or incorrect ACL applied to the client interface.
- Incorrect IP address (that is, a VIP that is not valid on the selected VLANs for the virtual server, or a VIP that is not accessible to your client).

If the Hit Count value increases but no response is received (Server Pkt Count does not increase), the problem is more likely to be in the connectivity between the ACE and the backend real servers. This issue is typically caused by one or more of the following problems:

- You are working on a one-armed configuration (that is, do not plan to change routing for your real servers) and have not selected an appropriate NAT pool for your virtual server to use with source NAT.
- A different routing problem (for example, server traffic does not know how to get back to the ACE).
- Addressing problem (for example, you have an incorrect real server address, or the real server is not accessible to ACE due to network topology).



Note Hit count can increase by more than one, even if you make only a single request from your web browser, because retrieving a typical web page makes many requests from the client to the server.

Related Topics

- [Configuring Virtual Servers, page 6-2](#)
- [Virtual Server Configuration and ANM, page 6-2](#)
- [Virtual Server Usage Guidelines, page 6-5](#)
- [Virtual Server Configuration Procedure, page 6-7](#)

Virtual Server Configuration Procedure

You can add virtual servers to the ANM for load-balancing purposes.

Assumptions

This topic assumes the following:

- Depending on the protocol to be used for the virtual server, parameter maps need to be defined.
- For SSL service, SSL certificates, keys, and chain groups, parameter maps must be configured.

Procedure

Step 1 Choose **Config > Devices > context > Load Balancing > Virtual Servers**.

The Virtual Servers table appears.

Step 2 In the Virtual Servers table, click **Poll Now** to instruct ANM to poll the devices and display the current values.

Step 3 Click **OK** when prompted if you want to poll the devices for data now.

Step 4 Click **Add** to add a new virtual server, or choose an existing virtual server and click **Edit** to modify it.

The Virtual Server configuration window appears with a number of configuration subsets. The subsets that you see depend on whether you use the Basic View or the Advanced View and entries that you make in the Properties subset. Change views by using the View object selector at the top of the configuration pane.

[Table 6-1](#) identifies and describes virtual server configuration subsets with links to related topics for configuration information.



Note The protocols that are available depend on the ACE device that you are configuring. For a list of the protocols available for each ACE device type, see [Table 6-2](#).

Table 6-1 Virtual Server Configuration Subsets

Configuration Subset	Description	Related Topics
Properties	Subset that allows you to specify basic virtual server characteristics, such as the virtual server name, IP address, protocol, port, and VLANs.	Configuring Virtual Server Properties, page 6-11
SSL Termination	Subset that appears when TCP is the selected protocol and Other or HTTPS is the application protocol. This subset allows you to configure the virtual server to act as an SSL proxy server and terminate SSL sessions between it and its clients.	Configuring Virtual Server SSL Termination, page 6-17

Table 6-1 Virtual Server Configuration Subsets (continued)

Configuration Subset	Description	Related Topics
Protocol Inspection	<p>Subset that appears in the Advanced View for:</p> <ul style="list-style-type: none"> TCP with FTP, HTTP, HTTPS, Real Time Streaming Protocol (RTSP), or Session Initiated Protocol (SIP) UDP with Domain Name System (DNS) or SIP <p>This subset appears in the Basic view for TCP with FTP.</p> <p>This subset allows you to configure the virtual server so that it can verify protocol behavior and identify unwanted or malicious traffic passing through the ACE on selected application protocols.</p>	Configuring Virtual Server Protocol Inspection, page 6-18
Application Acceleration And Optimization	<p>Subset that appears only for ACE appliances. It appears in the Advanced View when HTTP or HTTPS is the selected application protocol.</p> <p>This subset allows you to configure application acceleration and optimization options for HTTP or HTTPS traffic.</p>	Configuring Application Acceleration and Optimization, page 6-53
L7 Load-Balancing	<p>Subset that appears only in the Advanced View for these protocols:</p> <ul style="list-style-type: none"> TCP with Generic, HTTP, HTTPS, RTSP, or SIP UDP with Generic, RADIUS, or SIP <p>This subset allows you to configure Layer 7 load-balancing options, such as:</p> <ul style="list-style-type: none"> Server farms/real servers Health monitoring probes Stickiness SSL initiation 	Configuring Virtual Server Layer 7 Load Balancing, page 6-30
Default L7 Load-Balancing Action	Subset that allows you to establish the default Layer 7 load-balancing actions for all network traffic that does not meet previously specified match conditions including the SSL initiation configuration.	Configuring Virtual Server Default Layer 7 Load Balancing, page 6-51
NAT	Subset that allows you to set up Name Address Translation (NAT) for the virtual server.	Configuring Virtual Server NAT, page 6-64

Step 5 Do one of the following:

- Click **Deploy Now** to deploy the configuration on the ACE and save your entries to the running-configuration and startup-configuration files.
- Click **Cancel** to exit the procedure without saving your entries and to return to the Virtual Servers table.
- Click **Deploy Later** to save your entries and apply them at a later time.

Step 6 From the Virtual Servers table, to display statistics and status information for an existing virtual server, choose a virtual server, and click **Details**.

The **show service-policy *policy_name* detail** CLI command output appears. See the “[Displaying Virtual Server Statistics and Status Information](#)” section on page 6-66 for details.

Related Topics

- [Configuring Virtual Servers](#), page 6-2
- [Virtual Server Configuration and ANM](#), page 6-2
- [Virtual Server Usage Guidelines](#), page 6-5
- [Using ANM to Configure Virtual Servers](#), page 6-4
- [Shared Objects and Virtual Servers](#), page 6-9
- [Displaying Virtual Servers by Context](#), page 6-66
- [Displaying Virtual Server Statistics and Status Information](#), page 6-66
- [Managing Virtual Servers](#), page 6-67
- [Deploying Virtual Servers](#), page 6-76
- [Understanding Roles](#), page 17-6

Shared Objects and Virtual Servers

A shared object is one that is used by multiple virtual servers.

The following examples are shared objects:

- Action lists
- Class maps
- Parameter maps
- Real servers
- Server farms
- SSL services
- Sticky groups

Because these objects are shared, modifying an object’s configuration in one virtual server can impact other virtual servers that use the same object.

Configuring Shared Objects

ANM offers the following options for shared objects in virtual server configuration windows (Config > Devices > *context* > Load Balancing > Virtual Servers):

- View—Displays the object’s configuration. The window refreshes with read-only fields and the following three buttons.
- Cancel—Closes the read-only view and to return to the previous window.
- Edit—Enables you to modify the selected object’s configuration. The window refreshes with fields that can be modified, except for the Name field which remains read-only.



Note Before changing a shared object's configuration, make sure that you understand the effect of the changes on other virtual servers using the same object. As an alternative, consider using the Duplicate option instead.

- Duplicate—Enables you to create a new object with the same configuration as the selected object. The window refreshes with configurable fields. In the Name field, enter a unique name for the new object, and then modify the configuration as desired. This option allows you to create a new object without impacting other virtual servers using the same object.

Deleting Virtual Servers with Shared Objects

If you create a virtual server and include shared objects in its configuration, deleting the virtual server does not delete the associated shared objects. This action ensures that other virtual servers using the same shared objects are not impacted.

Related Topics

- [Managing Virtual Servers, page 6-67](#)
- [Virtual Server Protocols by Device Type, page 6-10](#)
- [Configuring Virtual Server Properties, page 6-11](#)
- [Configuring Virtual Server SSL Termination, page 6-17](#)
- [Configuring Virtual Server Protocol Inspection, page 6-18](#)
- [Configuring Virtual Server Layer 7 Load Balancing, page 6-30](#)
- [Configuring Virtual Server Default Layer 7 Load Balancing, page 6-51](#)
- [Configuring Application Acceleration and Optimization, page 6-53](#)
- [Configuring Virtual Server NAT, page 6-64](#)

Virtual Server Protocols by Device Type

The protocols that are available for a virtual server depend on the ACE device that you are configuring. [Table 6-2](#) lists the protocols available for each device type.

Table 6-2 *Virtual Server Protocols for ACE Modules and Devices*

Protocol	ACE Modules	ACE Appliance
Any	X	X
TCP		
FTP	X	X
Generic	X	X
HTTP	X	X
HTTPS	X	X
Other	X	X
RTSP	X	X
RDP	X	X

Table 6-2 Virtual Server Protocols for ACE Modules and Devices (continued)

Protocol	ACE Modules	ACE Appliance
SIP	X	X
UDP		
DNS	X	X
Generic	X	X
Other	X	X
RADIUS	X	X
SIP	X	X

Related Topics

- [Configuring Virtual Servers, page 6-2](#)
- [Configuring Virtual Server Properties, page 6-11](#)
- [Managing Virtual Servers, page 6-67](#)

Configuring Virtual Server Properties

You can configure virtual server properties.

Procedure

-
- Step 1** Choose **Config > Devices > context > Load Balancing > Virtual Servers**.
The Virtual Servers table appears.
- Step 2** In the Virtual Servers table, click **Poll Now** to instruct ANM to poll the devices and display the current values, and click **OK** when prompted if you want to poll the devices for data now.
- Step 3** Click **Add** to add a new virtual server, or choose an existing virtual server and click **Edit** to modify it.
The Virtual Server configuration window appears. The Properties configuration subset is open by default.
The fields that you see in the Properties configuration subset depend on whether you are using Advanced View or Basic View:
- To configure Advanced View properties, go to [Step 4](#).
 - To configure Basic View properties, go to [Step 5](#).
- Step 4** In the Advanced View, configure the virtual server properties by entering the information in [Table 6-3](#).

Table 6-3 Virtual Server Properties – Advanced View

Field	Description
Virtual Server Name	Name for the virtual server.
Virtual IP Address	IP address for the virtual server.
Virtual IP Mask	Subnet mask to apply to the virtual server IP address.

Table 6-3 Virtual Server Properties – Advanced View (continued)

Field	Description
Transport Protocol	<p>Protocol that the virtual server supports:</p> <ul style="list-style-type: none"> • Any—The virtual server is to accept connections using any IP protocol. • TCP—The virtual server is to accept connections that use TCP. • UDP—The virtual server is to accept connections that use UDP.
Application Protocol	<p>Field that appears if TCP or UDP is selected. The application protocols that are available depend on the type of ACE being configured.</p> <p>Choose the application protocol to be supported by the virtual server. Table 6-2 identifies the available protocols for each ACE device type.</p> <p>Note This field is read-only if you are editing an existing virtual server. ANM does not allow changes between protocols that require a change to the Layer 7 server load-balancing policy map. You need to delete the virtual server and create a new one with the desired application protocol.</p>
Port	<p>Field that appears for any TCP or UDP protocol.</p> <p>Enter the port to be used for the specified protocol. Valid entries are from 0 to 65535 or a range of integers, such as 10-20. Enter 0 (zero) to indicate all ports.</p> <p>For a complete list of protocols and ports, see the Internet Assigned Numbers Authority available at www.iana.org/numbers/</p>
All VLANs	<p>Check box that enables support of incoming traffic from all VLANs. Uncheck the check box to support incoming traffic from specific VLANs only.</p>
VLAN	<p>Field appears if the All VLANs check box is unchecked.</p> <p>In the Available Items list, choose the VLANs to use for incoming traffic, and click Add. The items appear in the Selected Items list.</p> <p>To remove VLANs, choose them in the Selected Items lists, and click Remove. The items appear in the Available Items list.</p> <p>Note You cannot change the VLAN for a virtual server once it is specified. Instead, delete the virtual server and create a new one with the desired VLAN.</p>
Connection Parameter Maps	<p>Field that appears if TCP is the selected protocol.</p> <p>Choose an existing connection parameter map or click *New* to create a new one as follows:</p> <ul style="list-style-type: none"> • If you chose an existing parameter map, you can view, modify, or duplicate the existing configuration. See the “Shared Objects and Virtual Servers” section on page 6-9 for more information about modifying shared objects. • If you click *New*, the Connection Parameter Maps configuration pane appears. Configure the connection parameter map as described in Table 9-2. <p>Note Click More Settings to access the additional Connection Parameter Maps configuration attributes. By default, ANM hides the default Connection Parameter Maps configuration attributes and the attributes which are not commonly used.</p>

Table 6-3 Virtual Server Properties – Advanced View (continued)

Field	Description
DNS Parameter Maps	<p>Field that appears if DNS is the selected protocol over UDP.</p> <p>Choose an existing DNS parameter map or click *New* to create a new one as follows:</p> <ul style="list-style-type: none"> • If you chose an existing parameter map, you can view, modify, or duplicate the existing configuration. See the “Shared Objects and Virtual Servers” section on page 6-9 for more information about modifying shared objects. • If you click *New*, the DNS Parameter Maps configuration pane appears. Configure the DNS parameter map as described in Table 9-11.
Generic Parameter Maps	<p>Field that appears if Generic is the selected application protocol over TCP or UDP.</p> <p>Choose an existing Generic parameter map or click *New* to create a new one as follows:</p> <ul style="list-style-type: none"> • If you chose an existing parameter map, you can view, modify, or duplicate the existing configuration. See the “Shared Objects and Virtual Servers” section on page 6-9 for more information about modifying shared objects. • If you click *New*, the Generic Parameter Maps configuration pane appears. Configure the Generic parameter map as described in Table 9-4.
HTTP Parameter Maps	<p>Field appears if HTTP or HTTPS is the selected application protocol.</p> <p>Choose an existing HTTP parameter map or click *New* to create a new one as follows:</p> <ul style="list-style-type: none"> • If you chose an existing parameter map, you can view, modify, or duplicate the existing configuration. See the “Shared Objects and Virtual Servers” section on page 6-9 for more information about modifying shared objects. • If you click *New*, the HTTP Parameter Maps configuration pane appears. Configure the HTTP parameter map as described in Table 9-5.
RTSP Parameter Maps	<p>Field that appears if RTSP is the selected application protocol over TCP.</p> <p>Choose an existing RTSP parameter map or click *New* to create a new one as follows:</p> <ul style="list-style-type: none"> • If you chose an existing parameter map, you can view, modify, or duplicate the existing configuration. See the “Shared Objects and Virtual Servers” section on page 6-9 for more information about modifying shared objects. • If you click *New*, the RTSP Parameter Maps configuration pane appears. Configure the RTSP parameter map as described in Table 9-8.

Table 6-3 Virtual Server Properties – Advanced View (continued)

Field	Description
KAL-AP-TAG Name	<p>Feature that is supported only for the ACE module software version A2(2.0), ACE appliance software version A4(1.0), and later versions for both device types. The KAL-AP-TAG feature allows the Cisco Global Site Selector (GSS) proprietary KAL-AP protocol to extract load and availability information from the ACE when a firewall is positioned between the GSS and the ACE. This feature allows you to configure a tag (name) per VIP for a maximum of 4096 tags on an ACE. This feature does not replace the tag per domain feature. For more information about this feature, see the <i>Release Note for the Cisco Application Control Engine Module (Software Version A2(2.0))</i> or the <i>Cisco Application Control Engine Module Server Load-Balancing Configuration Guide (Software Version A2(3.0))</i>, the Configuring Health Monitoring chapter.</p> <p>In the KAL-AP-TAG Name field, enter the name as an unquoted text string with no spaces and a maximum of 76 alphanumeric characters.</p> <p>The following scenarios are not supported and will result in an error:</p> <ul style="list-style-type: none"> • You cannot configure a tag name for a VIP that already has a tag configuration as part of a different policy configuration. • You cannot associate the same tag name with more than one VIP. • You cannot associate the same tag name with a domain and a VIP. • You cannot assign two different tags to two different Layer 3 class maps that have the same VIP, but different port numbers. The KAL-AP protocol considers these class maps to have the same VIP and calculates the load for both Layer 3 rules together when the GSS queries the VIP.
KAL-AP-Primary-Out-Of-Service	<p>Feature that is supported only for ACE module software version A2(3.1), ACE appliance software version A4(1.0), and later versions of either device type. Check the checkbox to enable the ACE to notify a Global Site Selector (GSS) that the primary server farm is down when the backup server farm is in use. Uncheck the checkbox to disable this feature.</p> <p>By default, when you configure a redirect server farm as a backup server farm on the ACE and the primary server farm fails, the backup server farm redirects client requests to another data center; however, the VIP remains in the INSERVICE state.</p> <p>When you configure the ACE to communicate with a GSS, it provides information for server availability. When a backup server is in use after the primary server farm is down, this feature enables the ACE to inform the GSS that the VIP for the primary server farm is out of service by returning a load value of 255. The GSS recognizes that the primary server farm is down and sends future DNS requests with the IP address of the other data center.</p>
ICMP Reply	<p>Virtual server response to ICMP ECHO requests as follows:</p> <ul style="list-style-type: none"> • None—The virtual server is not to send ICMP ECHO-REPLY responses to ICMP requests. • Active—The virtual server is to send ICMP ECHO-REPLY responses only if the configured VIP is active. • Always—The virtual server is always to send ICMP ECHO-REPLY responses to ICMP requests. • Primary Inservice—The virtual server is to reply to an ICMP ping only if the primary server farm state is UP, regardless of the state of the backup server farm. If this option is selected and the primary server farm state is DOWN, the ACE discards the ICMP request and the request times out.

Table 6-3 Virtual Server Properties – Advanced View (continued)

Field	Description
VIP Advertise	<p>Field that appears for ACE modules only.</p> <p>This option allows the ACE to advertise the IP address of the virtual server as the host route.</p> <p>Choose the desired VIP advertise option as follows:</p> <ul style="list-style-type: none"> • None—The ACE does not advertise the IP address of the virtual server as the host route. • Active—The ACE advertises the IP address of the virtual server as the host route only if there is at least one active real server in the server farm. • Always—The ACE always advertises the IP address of the virtual server as the host route. • Active-Metric—The ACE advertises the IP address of the virtual server as the host route if the following occurs: <ul style="list-style-type: none"> • There is at least one active real server in the server farm. • A distance metric is specified for the route in the Distance field. • Always-Metric—The ACE advertises the IP address of the virtual server as the host route, using the distance metric in the Distance field.
Distance	<p>Field that appears for ACE modules only.</p> <p>This field appears if you chose Active-Metric or Always-Metric in the VIP Advertise field.</p> <p>Enter the administrative distance to be included in the routing table. Valid entries are integers from 1 to 254.</p>
Status	<p>Operating state of the virtual server as follows:</p> <ul style="list-style-type: none"> • In Service—Enables the virtual server for load-balancing operations. • Out Of Service—Disables the virtual server for load-balancing operations.

Step 5 In the Basic View, configure virtual server properties by entering the information in [Table 6-4](#).

Table 6-4 Virtual Server Properties – Basic View

Field	Description
Virtual Server Name	Name for the virtual server.
Virtual IP Address	IP address for the virtual server.
Transport Protocol	<p>Protocol that the virtual server supports as follows:</p> <ul style="list-style-type: none"> • Any—The virtual server accepts connections using any IP protocol. • TCP—The virtual server accepts connections that use TCP. • UDP—The virtual server accepts connections that use UDP.

Table 6-4 Virtual Server Properties – Basic View (continued)

Field	Description
Application Protocol	<p>Field that appears if TCP or UDP is selected. The application protocols that are available depend on the type of ACE being configured.</p> <p>Choose the application protocol to be supported by the virtual server. Table 6-2 identifies the available protocols for each ACE device type.</p> <p>Note This field is read-only if you are editing an existing virtual server. ANM does not allow changes between protocols that require a change to the Layer 7 server load-balancing policy map. You need to delete the virtual server and create a new one with the desired application protocol.</p>
Port	<p>Field that appears for any specific TCP or UDP protocol.</p> <p>Enter the port to be used for the specified protocol. Valid entries are from 0 to 65535 or a range of integers, such as 10-20. Enter 0 (zero) to indicate all ports.</p> <p>For a complete list of all protocols and ports, see the Internet Assigned Numbers Authority available at www.iana.org/numbers/</p>
All VLANs	<p>Check box that enables support of incoming traffic from all VLANs. Uncheck the check box to support incoming traffic from specific VLANs only.</p>
VLAN	<p>Field that appears if the All VLANs check box is unchecked.</p> <p>In the Available Items list, choose the VLANs to use for incoming traffic, and click Add. The items appear in the Selected Items list.</p> <p>To remove VLANs, choose them in the Selected Items lists, and click Remove. The items appear in the Available Items list.</p> <p>Note You cannot change the VLAN for a virtual server once it is specified. Instead, delete the virtual server and create a new one with the desired VLAN.</p>

Step 6 Do one of the following:

- Click **Deploy Now** to deploy the configuration on the ACE and save your entries to the running-configuration and startup-configuration files.
- Click **Cancel** to exit the procedure without saving your entries.
- Click **Deploy Later** to save your entries and apply them at a later time.

Related Topics

- [Configuring Virtual Servers, page 6-2](#)
- [Configuring Virtual Server SSL Termination, page 6-17](#)

Configuring Virtual Server SSL Termination

You can configure virtual server SSL termination service, which allows the virtual server to act as an SSL proxy server and terminate SSL sessions between it and its clients.

Assumption

Make sure that a virtual server has been configured for HTTPS over TCP or Other over TCP in the Properties configuration subset. For more information, see the [“Configuring Virtual Server Properties” section on page 6-11](#).

Procedure

-
- Step 1** Choose **Config > Devices > context > Load Balancing > Virtual Servers**.
The Virtual Servers table appears.
- Step 2** In the Virtual Servers table, choose the virtual server that you want to configure for SSL termination, and click **Edit**.
The Virtual Server configuration window appears.
- Step 3** In the Virtual Server configuration window, click **SSL Termination**.
The Proxy Service Name field appears.
- Step 4** In the Proxy Service Name field, choose an existing SSL termination service, or choose ***New*** to create a new SSL proxy service, and do one of the following:
- If you chose an existing SSL service, the window refreshes and allows you to view, modify, or duplicate the existing configuration. See the [“Shared Objects and Virtual Servers” section on page 6-9](#) for more information about modifying shared objects.
 - If you chose ***New***, the Proxy Service configuration subset appears.
- Step 5** Configure the SSL service using the information in [Table 6-5](#).
For more information about SSL, see the [“Configuring SSL” section on page 10-1](#).

Table 6-5 Virtual Server SSL Attributes

Field	Description
Name	Name for this SSL proxy service. Valid entries are alphanumeric strings with a maximum of 26 characters.
Keys	SSL key pair to use during the SSL handshake for data encryption.
Certificates	SSL certificate to use during the SSL handshake.
Chain Groups	Chain group to use during the SSL handshake.
Auth Groups	SSL authentication group to associate with this proxy server service.
CRL Best-Effort	Option that appears if you chose an authentication group in the Auth Groups field. Check the check box to allow the ANM to search client certificates for the service to determine if it contains a CRL in the extension and retrieve the value, if it exists. Uncheck the check box to disable this feature.

Table 6-5 Virtual Server SSL Attributes

Field	Description
CRL Name	Option that appears if the CRL Best-Effort check box is clear. Choose the Certificate Revocation List the ANM is to use for this proxy service.
Parameter Maps	SSL parameter map to associate with this proxy server service.

Step 6 Do one of the following:

- Click **Deploy Now** to deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files.
- Click **Cancel** to exit this procedure without saving your entries.
- Click **Deploy Later** to save your entries and apply them at a later time.

Related Topics

- [Configuring Virtual Servers, page 6-2](#)
- [Configuring Virtual Server Properties, page 6-11](#)

Configuring Virtual Server Protocol Inspection

You can configure protocol inspection on a virtual server, which allows the virtual server to verify protocol behavior and identify unwanted or malicious traffic passing through the ACE.

In the Advanced View, protocol inspection configuration is available for the following virtual server protocol configurations:

- TCP with FTP, HTTP, HTTPS, RTSP, or SIP
- UDP with DNS or SIP

In the Basic View, protocol inspection configuration is available for TCP with FTP.

See [Table 6-2](#) for a list of protocols by ACE device type.

Assumption

Make sure that a virtual server has been configured to use one of the protocols that supports protocol inspection in the Properties configuration subset. See the “[Configuring Virtual Server Properties](#)” section on [page 6-11](#) for information on configuring these protocols.

Procedure

Step 1 Choose the item to configure:

- To configure a virtual server, choose **Config > Devices > context > Load Balancing > Virtual Servers**.
- To configure a configuration building block, choose **Config > Global > All Building Blocks > building_block > Load Balancing > Virtual Servers**.

The Virtual Servers table appears.

- Step 2** In the Virtual Servers table, choose the virtual server that you want to configure for protocol inspection, and click **Edit**.
The Virtual Server configuration window appears.
- Step 3** Click **Protocol Inspection**.
The Enable Inspect check box appears.
- Step 4** Check the **Enable Inspect** check box to enable inspection on the specified traffic or uncheck it to disable inspection on this traffic.
By default, the ACE allows all request methods.
- Step 5** (Optional) If you checked the Enable Inspect check box, configure additional inspection options using the information in [Table 6-6](#).

Table 6-6 Protocol Inspection Configuration Options

Protocol	Action
DNS	In the length field, enter the maximum length of the DNS packet in bytes as defined in the Length field. If you do not enter a value in this field, the DNS packet size is not checked.
FTP	<ol style="list-style-type: none"> a. Check the Use Strict check box to specify that the virtual server is to perform enhanced inspection of FTP traffic and enforce compliance with RFC standards. Uncheck the check box to specify that the virtual server is not to perform enhanced FTP inspection. b. (Optional) If you checked the Use Strict check box, in the Blocked FTP Commands field, identify the commands that are to be denied by the virtual server. See Table 13-8 for more information about the FTP commands. <ul style="list-style-type: none"> • Choose the commands that are to be blocked by the virtual server in the Available Items list, and click Add. The commands appear in the Selected Items list. • To remove commands that you do not want to be blocked, choose them in the Selected Items list, and click Remove. The commands appear in the Available Items list.

Table 6-6 Protocol Inspection Configuration Options (continued)

Protocol	Action
HTTP or HTTPS	<p>a. Check the Logging Enabled check box to enable monitoring of Layer 3 and Layer 4 traffic. When enabled, this feature logs every URL request that is sent in the specified class of traffic, including the source or destination IP address and the URL that is accessed. Uncheck the check box to disable monitoring of Layer 3 and Layer 4 traffic.</p> <p>b. In the Policy subset, click Add to add a new match condition and action, or choose an existing match condition and action and click Edit to modify it. The Policy configuration pane appears.</p> <p>c. In the Matches field, choose an existing class map or *New* or *Inline Match* to configure new match criteria for protocol inspection.</p> <p>If you chose an existing class map, the window refreshes and allows you to view, modify, or duplicate the selected class map. See the “Shared Objects and Virtual Servers” section on page 6-9 for more information about modifying shared objects.</p> <p>d. Configure match criteria and related actions using the information in Table 6-7.</p> <p>e. Do one of the following:</p> <ul style="list-style-type: none"> • Click OK to save your entries. The Conditions table refreshes with the new entry. • Click Cancel to exit the Policy subset without saving your entries. <p>f. In the Default Action field, choose the default action that the virtual server is to take when specified match conditions for protocol inspection are not met:</p> <ul style="list-style-type: none"> • Permit—The specified HTTP traffic is to be received by the virtual server. • Reset—The specified HTTP traffic is to be denied by the virtual server.

Table 6-6 Protocol Inspection Configuration Options (continued)

Protocol	Action
RTSP	There are no protocol-specific inspection options for RTSP.
SIP	<p>a. In the Actions subset, click Add to add a new match condition and action, or choose an existing match condition and action, and click Edit to modify it. The Actions configuration pane appears.</p> <p>b. In the Matches field, choose an existing class map or *New* or *Inline Match* to configure new match criteria for protocol inspection.</p> <p>If you chose an existing class map, the window refreshes and allows you to view, modify, or duplicate the selected class map. See the “Shared Objects and Virtual Servers” section on page 6-9 for more information about modifying shared objects.</p> <p>c. Configure match criteria and related actions using the information in Table 6-9.</p> <p>d. In the Action field, choose the action that the virtual server is to take when the specified match conditions are met:</p> <ul style="list-style-type: none"> – Drop—The specified SIP traffic is discarded by the virtual server. – Permit—The specified SIP traffic is received by the virtual server. – Reset—The specified SIP traffic is denied by the virtual server. <p>e. Do one of the following:</p> <ul style="list-style-type: none"> – Click OK to save your entries. The Conditions table refreshes with the new entry. – Click Cancel to exit the Conditions subset without saving your entries and to return to the Conditions table. <p>f. In the SIP Parameter Map field, choose an existing parameter map or choose *New* to configure a new one.</p> <p>If you chose an existing parameter map, the window refreshes and allows you to view, modify, or delete the selected parameter map. See the “Shared Objects and Virtual Servers” section on page 6-9 for more information about modifying shared objects.</p> <p>g. Configure SIP parameter map options using the information in Table 9-9.</p> <p>h. In the Secondary Connection Parameter Map field, choose an existing parameter map or choose *New* to configure a new one.</p> <p>If you chose an existing parameter map, the window refreshes and allows you to view, modify, or delete the selected parameter map. See the “Shared Objects and Virtual Servers” section on page 6-9 for more information about modifying shared objects.</p> <p>i. Configure secondary connection parameter map options using the information in Table 9-2.</p> <p>j. In the Default Action field, choose the default action that the virtual server is to take when specified match conditions for SIP protocol inspection are not met:</p> <ul style="list-style-type: none"> – Drop—The specified SIP traffic is discarded by the virtual server. – Permit—The specified SIP traffic is received by the virtual server. – Reset—The specified SIP traffic is denied by the virtual server. <p>k. Check the Logging Enabled check box to enable monitoring of Layer 3 and Layer 4 traffic. When enabled, this feature logs every URL request that is sent in the specified class of traffic, including the source or destination IP address and the URL that is accessed. Uncheck the check box to disable monitoring of Layer 3 and Layer 4 traffic.</p>

Table 6-7 HTTP and HTTPS Protocol Inspection Match Criteria Configuration

Selection	Action
Existing class map	<p>a. Click View to review the match condition information for the selected class map.</p> <p>b. Do one of the following:</p> <ul style="list-style-type: none"> – Click Cancel to continue without making changes and to return to the previous window. – Click Edit to modify the existing configuration. – Click Duplicate to create a new class map with the same attributes without affecting other virtual servers using the same class map. <p>See the “Shared Objects and Virtual Servers” section on page 6-9 for information about modifying shared objects.</p> <p>c. In the Action field, choose the action that the virtual server is to perform on the traffic if it matches the specified match criteria:</p> <ul style="list-style-type: none"> – Permit—The specified traffic is received by the virtual server if it meets the specified deep inspection match criteria. – Reset—The specified traffic is denied by the virtual server, which then sends a TCP reset message to the client or server to close the connection.
New	<p>a. In the Name field, specify a unique name for this class map.</p> <p>b. In the Match field, choose the method to be used to evaluate multiple match statements when multiple match conditions exist:</p> <ul style="list-style-type: none"> – Any—A match exists if at least one of the match conditions is satisfied. – All—A match exists only if all match conditions are satisfied. <p>c. In the Conditions table, click Add to add a new set of conditions, or choose an existing entry, and click Edit to modify it. The Type field appears.</p> <p>d. In the Type field, choose the type of condition that is to be met for protocol inspection.</p> <p>e. Provide condition-specific criteria using the information in Table 6-8.</p> <p>f. In the Action field, choose the action that the virtual server is to perform on the traffic if it matches the specified match criteria:</p> <ul style="list-style-type: none"> – Permit—The specified traffic is received by the virtual server if it meets the specified deep inspection match criteria. – Reset—The specified traffic is denied by the virtual server, which then sends a TCP reset message to the client or server to close the connection.
Inline Match	<p>a. In the Conditions Type field, choose the type of inline match condition that is to be met for protocol inspection.</p> <p>b. Provide condition-specific criteria using the information in Table 6-8.</p> <p>c. In the Action field, choose the action that the virtual server is to perform on the traffic if it matches the specified match criteria:</p> <ul style="list-style-type: none"> – Permit—The specified traffic is received by the virtual server if it meets the specified deep inspection match criteria. – Reset—The specified traffic is denied by the virtual server, which then sends a TCP reset message to the client or server to close the connection.

Table 6-8 HTTP and HTTPS Protocol Inspection Conditions and Options

Condition	Description
Content	<p>Specific content contained within the HTTP entity-body to be used for application inspection decisions.</p> <ul style="list-style-type: none"> a. In the Content Expression field, enter the content that is to be matched. Valid entries are alphanumeric strings from 1 to 255 characters. b. In the Content Offset field, enter the number of bytes to be ignored starting with the first byte of the Message body, after the empty line (CR,LF,CR,LF) between the headers and the body of the message. Valid entries are from 1 to 255 bytes.
Content Length	<p>Content parse length is used for application inspection decisions.</p> <ul style="list-style-type: none"> a. In the Content Length Operator field, choose the operand to use to compare content length: <ul style="list-style-type: none"> – Equal To—The content length must equal the number in the Content Length Value field. – Greater Than—The content length must be greater than the number in the Content Length Value field. – Less Than—The content length must be less than the number in the Content Length Value field. – Range—The content length must be within the range specified in the Content Length Lower Value field and the Content Length Higher Value field. b. Enter values to apply for content length comparison: <ul style="list-style-type: none"> – If you chose Equal To, Greater Than, or Less Than in the Content Length Operator field, the Content Length Value field appears. In the Content Length Value field, enter the number of bytes for comparison. Valid entries are from 0 to 4294967295. – If you chose Range in the Content Length Operator field, the Content Length Lower Value and the Content Length Higher Value fields appear: <ol style="list-style-type: none"> 1. In the Content Length Lower Value field, enter the lowest number of bytes to be used for this match condition. Valid entries are from 0 to 4294967295. The number in this field must be less than the number entered in the Content Length Higher Value field. 2. In the Content Length Higher Value field, enter the highest number of bytes to be used for this match condition. Valid entries are from 0 to 4294967295. The number in this field must be greater than the number entered in the Content Length Lower Value field.
Content Type Verification	<p>Verification of MIME-type messages with the header MIME-type is to be used for application inspection decisions. This option verifies that the header MIME-type value is in the internal list of supported MIME-types and that the header MIME-type matches the content in the data or body portion of the message.</p>

Table 6-8 HTTP and HTTPS Protocol Inspection Conditions and Options (continued)

Condition	Description
Header	<p>Name and value in an HTTP header are used for application inspection decisions.</p> <ul style="list-style-type: none"> a. In the Header field, choose one of the predefined HTTP headers to match, or choose HTTP Header to specify a different HTTP header. b. If you chose HTTP Header, in the Header Name field, enter the name of the HTTP header to match. Valid entries are unquoted text strings with no spaces and a maximum of 64 alphanumeric characters. c. In the Header Value field, enter the header-value expression string to compare against the value in the specified field in the HTTP header. Valid entries are text strings with a maximum of 255 alphanumeric characters. The ACE supports regular expressions for matching. Header expressions allow spaces, provided that the spaces are escaped or quoted. All headers in the header map must be matched. See Table 13-34 for a list of the supported characters that you can use in regular expressions.
Header Length	<p>Length of the header in the HTTP message used for application inspection decisions.</p> <ul style="list-style-type: none"> a. In the Header Length Type field, specify whether HTTP header request or response messages are to be used for application inspection decisions: <ul style="list-style-type: none"> – Request—HTTP header request messages are to be checked for header length. – Response—HTTP header response messages are to be checked for header length. b. In the Header Length Operator field, choose the operand to be used to compare header length: <ul style="list-style-type: none"> – Equal To—The header length must equal the number in the Header Length Value field. – Greater Than—The header length must be greater than the number in the Header Length Value field. – Less Than—The header length must be less than the number in the Header Length Value field. – Range—The header length must be within the range specified in the Header Length Lower Value field and the Header Length Higher Value field. c. Enter values to apply for header length comparison: <ul style="list-style-type: none"> – If you chose Equal To, Greater Than, or Less Than in the Header Length Operator field, the Header Length Value field appears. In the Header Length Value field, enter the number of bytes for comparison. Valid entries are from 0 to 255. – If you chose Range in the Header Length Operator field, the Header Length Lower Value and the Header Length Higher Value fields appear: <ol style="list-style-type: none"> 1. In the Header Length Lower Value field, enter the lowest number of bytes to be used for this match condition. Valid entries are from 0 to 255. The number in this field must be less than the number entered in the Header Length Higher Value field. 2. In the Header Length Higher Value field, enter the highest number of bytes to be used for this match condition. Valid entries are from 1 to 255. The number in this field must be greater than the number entered in the Header Length Lower Value field.
Header MIME Type	<p>Multipurpose Internet Mail Extension (MIME) message types are used for application inspection decisions.</p> <p>In the Header MIME Type field, choose the MIME message type to use for this match condition.</p>

Table 6-8 HTTP and HTTPS Protocol Inspection Conditions and Options (continued)

Condition	Description
Port Misuse	<p>Misuse of port 80 (or any other port running HTTP) to be used for application inspection decisions. Choose the application category to use for this match condition as follows:</p> <ul style="list-style-type: none"> • IM—Instant messaging applications are to be checked. • P2P—Peer-to-peer applications are to be checked. • Tunneling—Tunneling applications are to be checked.
Request Method	<p>A request method is to be used for application inspection decisions.</p> <ol style="list-style-type: none"> a. Choose the type of request method to use for this match condition: <ul style="list-style-type: none"> – Ext—An HTTP extension method is to be used. – RFC—The request method defined in RFC 2616 is to be used. b. In the Request Method field, choose the request method that is to be inspected.
Strict HTTP	Compliance with HTTP RFC 2616 to be used for application inspection decisions.
Transfer Encoding	<p>An HTTP transfer-encoding type to be used for application inspection decisions. The transfer-encoding general-header field indicates the type of transformation, if any, that has been applied to the HTTP message body to safely transfer it between the sender and the recipient.</p> <p>In the Transfer Encoding field, choose the type of encoding that is to be checked:</p> <ul style="list-style-type: none"> • Chunked—The message body is transferred as a series of chunks. • Compress—The encoding format that is produced by the UNIX file compression program <i>compress</i>. • Deflate—The .zlib format that is defined in RFC 1950 in combination with the DEFLATE compression mechanism described in RFC 1951. • Gzip—The encoding format that is produced by the file compression program GZIP (GNU zip) as described in RFC 1952. • Identity—The default (identity) encoding which does not require the use of transformation.

Table 6-8 HTTP and HTTPS Protocol Inspection Conditions and Options (continued)

Condition	Description
URL	<p>URL names to be used for application inspection decisions.</p> <p>In the URL field, enter a URL or a portion of a URL to match. Valid entries are URL strings from 1 to 255 alphanumeric characters and include only the portion of the URL following <code>www.hostname.domain</code>. For example, in the URL <code>www.anydomain.com/latest/whatsnew.html</code>, include only <code>/latest/whatsnew.html</code>.</p>
URL Length	<p>URL length to be used for application inspection decisions.</p> <p>a. In the URL Length Operator field, choose the operand to use to compare URL length:</p> <ul style="list-style-type: none"> - Equal To—The URL length must equal the number in the URL Length Value field. - Greater Than—The URL length must be greater than the number in the URL Length Value field. - Less Than—The URL length must be less than the number in the URL Length Value field. - Range—The URL length must be within the range specified in the URL Length Lower Value field and the URL Length Higher Value field. <p>b. Enter values to apply for URL length comparison:</p> <ul style="list-style-type: none"> - If you chose Equal To, Greater Than, or Less Than in the URL Length Operator field, the URL Length Value field appears. In the URL Length Value field, enter the value for comparison. Valid entries are from 1 to 65535 bytes. - If you chose Range in the URL Length Operator field, the URL Length Lower Value and the URL Length Higher Value fields appear: <ol style="list-style-type: none"> 1. In the URL Length Lower Value field, enter the lowest number of bytes to be used for this match condition. Valid entries are from 1 to 65535. The number in this field must be less than the number entered in the URL Length Higher Value field. 2. In the URL Length Higher Value field, enter the highest number of bytes to be used for this match condition. Valid entries are from 1 to 65535. The number in this field must be greater than the number entered in the URL Length Lower Value field.

Table 6-9 SIP Protocol Inspection Match Criteria Configuration

Selection	Action
Existing class map	<p>a. Click View to review the match condition information for the selected class map.</p> <p>b. Do one of the following:</p> <ul style="list-style-type: none"> – Click Cancel to continue without making changes and to return to the previous window. – Click Edit to modify the existing configuration. – Click Duplicate to create a new class map with the same attributes without affecting other virtual servers using the same class map. <p>See the “Shared Objects and Virtual Servers” section on page 6-9 for more information about modifying shared objects.</p> <p>c. In the Action field, choose the action that the virtual server is to perform on the traffic if it matches the specified match criteria:</p> <ul style="list-style-type: none"> – Drop—The specified traffic is to be dropped by the virtual server. – Permit—The specified traffic is to be received by the virtual server. – Reset—The specified traffic is to be denied by the virtual server, which then sends a TCP reset message to the client or server to close the connection.
New	<p>a. In the Name field, specify a unique name for this class map.</p> <p>b. In the Conditions table, click Add to add a new set of conditions, or choose an existing entry, and click Edit to modify it. The Type field appears.</p> <p>c. In the Type field, choose the type of condition that is to be met for protocol inspection.</p> <p>d. Provide condition-specific criteria using the information in Table 6-10.</p> <p>e. In the Action field, choose the action that the virtual server is to perform on the traffic if it matches the specified match criteria:</p> <ul style="list-style-type: none"> – Drop—The specified traffic is to be dropped by the virtual server. – Permit—The specified traffic is to be received by the virtual server. – Reset—The specified traffic is to be denied by the virtual server, which then sends a TCP reset message to the client or server to close the connection.
Inline Match	<p>a. In the Conditions Type field, choose the type of inline match condition that is to be met for protocol inspection.</p> <p>Table 6-10 describes the types of conditions and their related configuration options.</p> <p>b. Provide condition-specific criteria using the information in Table 6-10.</p> <p>c. In the Action field, choose the action that the virtual server is to perform on the traffic if it matches the specified match criteria:</p> <ul style="list-style-type: none"> – Drop—The specified traffic is to be dropped by the virtual server. – Permit—The specified traffic is to be received by the virtual server. – Reset—The specified traffic is to be denied by the virtual server, which then sends a TCP reset message to the client or server to close the connection.

Table 6-10 SIP Protocol Inspection Conditions and Options

Condition	Description
Called Party	<p>Destination or called party specified in the URI of the SIP To header used for SIP protocol inspection decisions.</p> <p>In the Called Party field, enter a regular expression that identifies the called party in the URI of the SIP To header for this match condition. Valid entries are unquoted text strings with no spaces and a maximum of 255 alphanumeric characters. The ACE supports regular expressions for matching string expressions. Table 13-34 lists the supported characters that you can use for matching string expressions.</p>
Calling Party	<p>Source or caller specified in the URI of the SIP From header used for SIP protocol inspection decisions.</p> <p>In the Calling Party field, enter a regular expression that identifies the calling party in the URI of the SIP From header for this match condition. Valid entries are unquoted text strings with no spaces and a maximum of 255 alphanumeric characters. The ACE supports regular expressions for matching string expressions. Table 13-34 lists the supported characters that you can use for matching string expressions.</p>
IM Subscriber	<p>IM (instant messaging) subscriber used for application inspection decisions.</p> <p>In the IP Subscriber field, enter a regular expression that identifies the IM subscriber for this match condition. Valid entries are unquoted text strings with no spaces and a maximum of 255 alphanumeric characters. The ACE supports regular expressions for matching string expressions. Table 13-34 lists the supported characters that you can use for matching string expressions.</p>
Message Path	<p>SIP inspection that allows you to filter messages coming from or transiting through certain SIP proxy servers. The ACE maintains a list of the unauthorized SIP proxy IP addresses or URLs in the form of regular expressions and checks this list against the VIA header field in each SIP packet.</p> <p>In the Message Path field, enter a regular expression that identifies the SIP proxy server for this match condition. Valid entries are unquoted text strings with no spaces and a maximum of 255 alphanumeric characters. The ACE supports regular expressions for matching string expressions. Table 13-34 lists the supported characters that you can use for matching string expressions.</p>
SIP Content Length	<p>SIP message body content length used for SIP protocol inspection decisions.</p> <p>To specify SIP traffic based on SIP message body length:</p> <ol style="list-style-type: none"> <li data-bbox="391 1444 1190 1476">a. In the Content Operator field, confirm that Greater Than is selected. <li data-bbox="391 1486 1414 1612">b. In the Content Length field, enter the maximum size of a SIP message body in bytes that the ACE is to allow without performing SIP protocol inspection. If a SIP message exceeds the specified value, the ACE performs SIP protocol inspection as defined in an associated policy map. Valid entries are from 0 to 65534 bytes.
SIP Content Type	<p>Content type in the SIP message body used for SIP protocol inspection decisions.</p> <p>In the Content Type field, enter a regular expression that identifies the content type in the SIP message body to use for this match condition. Valid entries are unquoted text strings with no spaces and a maximum of 255 alphanumeric characters. The ACE supports regular expressions for matching string expressions. Table 13-34 lists the supported characters that you can use for matching string expressions.</p>

Table 6-10 SIP Protocol Inspection Conditions and Options (continued)

Condition	Description
SIP Request Method	SIP request method used for application inspection decisions. In the Request Method field, choose the request method that is to be inspected.
Third Party	Condition that indicates that the SIP is to allow users to register other users on their behalf by sending REGISTER messages with different values in the From and To header fields. This process can pose a security threat if the REGISTER message is actually a Deregister message. A malicious user could cause a DoS (denial-of-service) attack by deregistering all users on their behalf. To prevent this security threat, you can specify a list of privileged users who can register or unregister someone else on their behalf. The ACE maintains the list as a regex table. If you configure this policy, the ACE drops REGISTER messages with mismatched From and To headers and a From header value that does not match any of the privileged user IDs. In the Third Party Registration Entities field, enter a regular expression that identifies a privileged user who is authorized for third-party registrations. Valid entries are unquoted text strings with no spaces and a maximum of 255 alphanumeric characters. The ACE supports regular expressions for matching string expressions. Table 13-34 lists the supported characters that you can use for matching string expressions.
URI Length	Condition that indicates that the ACE is to validate the length of SIP URIs or Tel URIs. A SIP URI is a user identifier that a calling party (source) uses to contact the called party (destination). A Tel URI is a telephone number that identifies the endpoint of a SIP connection. For more information about SIP URIs and Tel URIs, see RFC 2534 and RFC 3966, respectively. To filter SIP traffic based on URIs, do the following: <ul style="list-style-type: none"> a. In the URI Type field, choose the type of URI to be used: <ul style="list-style-type: none"> – SIP URI—The calling party URI is to be used for this match condition. – Tel URI—A telephone number is to be used for this match condition. b. In the URI Operator field, confirm that Greater Than is selected. c. In the URI Length field, enter the maximum length of the SIP URI or Tel URI in bytes. Valid entries are from 0 to 254 bytes.

Step 6 Do one of the following:

- Click **Deploy Now** to deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files.
- Click **Cancel** to exit this procedure without saving your entries.
- Click **Deploy Later** to save your entries and deploy the configuration at a later time.

Related Topics

- [Configuring Virtual Server Properties, page 6-11](#)
- [Configuring Virtual Server SSL Termination, page 6-17](#)
- [Configuring Virtual Server Layer 7 Load Balancing, page 6-30](#)
- [Managing Virtual Servers, page 6-67](#)

Configuring Virtual Server Layer 7 Load Balancing

You can configure Layer 7 load balancing on a virtual server. In the Advanced View, Layer 7 load balancing is available for virtual servers configured with one of the following protocol combinations:

- TCP with Generic, FTP, HTTP, HTTPS, RDP, RTSP, or SIP
- UDP with Generic, DNS, RADIUS, or SIP

See the “[Configuring Virtual Server Properties](#)” section on page 6-11 for information about configuring these protocols.

[Table 6-2](#) identifies the protocols that are available for each type of ACE device.

Assumption

Make sure that a virtual server has been configured with one of the following protocol combinations:

- TCP with Generic, FTP, HTTP, HTTPS, RDP, RTSP, or SIP
- UDP with Generic, DNS, RADIUS, or SIP

For more information, see the “[Configuring Virtual Server Properties](#)” section on page 6-11.

Procedure

-
- Step 1** Choose **Config > Devices > context > Load Balancing > Virtual Servers**.
The Virtual Servers table appears.
- Step 2** In the Virtual Servers table, choose the virtual server that you want to configure for Layer 7 load balancing, and click **Edit**.
The Virtual Server configuration window appears.
- Step 3** In the Virtual Server configuration window, click **L7 Load-Balancing**.
The Layer 7 Load-Balancing Rule Match table appears.
- Step 4** In the Rule Match table, click **Add** to add a new match condition and action, or choose an existing match condition and action, and click **Edit** to modify it.
The Rule Match configuration pane appears.
- Step 5** In the Rule Match field of the Rule Match configuration pane, choose an existing class map or ***New*** or ***Inline Match*** to configure new match criteria for Layer 7 load balancing, and do one of the following:
- If you chose an existing class map, click **View** to review, modify, or duplicate the existing configuration. See the “[Shared Objects and Virtual Servers](#)” section on page 6-9 for more information about modifying shared objects.
 - If you click ***New*** or ***Inline Match***, the Rule Match configuration pane appears.

Step 6 Configure match criteria using the information in [Table 6-11](#).

Table 6-11 Layer 7 Load-Balancing Match Criteria Configuration

Selection	Action
Existing class map	<p>a. Click View to review the match condition information for the selected class map.</p> <p>b. Do one of the following:</p> <ul style="list-style-type: none"> – Click Cancel to continue without making changes and to return to the previous window. – Click Edit to modify the existing configuration. – Click Duplicate to create a new class map with the same attributes without affecting other virtual servers using the same class map. <p>See the “Shared Objects and Virtual Servers” section on page 6-9 for more information about modifying shared objects.</p>
New	<p>a. In the Name field, enter a unique name for this class map.</p> <p>b. In the Match field, choose the method to be used to evaluate multiple match statements when multiple match conditions exist:</p> <ul style="list-style-type: none"> – match-any—A match exists if at least one of the match conditions is satisfied. – match-all—A match exists only if all match conditions are satisfied. <p>c. In the Conditions table, click Add to add a new set of conditions, or choose an existing entry and click Edit to modify it.</p> <p>d. In the Type field, choose the match condition and configure any of these protocol-specific options:</p> <ul style="list-style-type: none"> – For Generic protocol options, see Table 13-9. – For HTTP and HTTPS protocol options, see Table 6-12. – For RADIUS protocol options, see Table 13-10. – For RTSP protocol options, see Table 13-11. – For SIP protocol options, see Table 13-12. <p>e. Do one of the following:</p> <ul style="list-style-type: none"> – Click OK to accept your entries and to return to the Conditions table. – Click Cancel to exit this procedure without saving your entries and to return to the Conditions table.
Inline Match	<p>In the Conditions Type field, choose the type of inline match condition and configure any protocol-specific options:</p> <ul style="list-style-type: none"> • For Generic protocol options, see Table 13-9. • For HTTP and HTTPS protocol options, see Table 6-12. • For RADIUS protocol options, see Table 13-10. • For RTSP protocol options, see Table 13-11. • For SIP protocol options, see Table 13-12.

Table 6-12 Layer 7 HTTP/HTTPS Load-Balancing Conditions and Options

Match Condition	Action
Class Map	<p>Existing class map used for the match condition.</p> <p>In the Class Map field, choose the class map to be used.</p>
HTTP Content	<p>Specific content contained within the HTTP entity-body used to establish a match condition.</p> <ol style="list-style-type: none"> <li data-bbox="391 468 1464 531">a. In the Content Expression field, enter the content that is to be matched. Valid entries are alphanumeric strings from 1 to 255 characters. <li data-bbox="391 541 1464 636">b. In the Content Offset field, enter the number of bytes to be ignored starting with the first byte of the Message body, after the empty line (CR,LF,CR,LF) between the headers and the body of the message. Valid entries are from 1 to 255.
HTTP Cookie	<p>HTTP cookies used for the match condition.</p> <ol style="list-style-type: none"> <li data-bbox="391 699 1464 762">a. In the Cookie Name field, enter a unique cookie name. Valid entries are unquoted text strings with no spaces and a maximum of 64 alphanumeric characters. <li data-bbox="391 772 1464 898">b. In the Cookie Value field, enter a unique cookie value expression. Valid entries are unquoted text strings with no spaces and a maximum of 255 alphanumeric characters. The ACE supports regular expressions for matching string expressions. Table 13-34 lists the supported characters that you can use for matching string expressions. <li data-bbox="391 909 1464 1035">c. Check the Secondary Cookie Matching check box to indicate that the ACE is to use both the cookie name and the cookie value to satisfy this match condition. Clear this check box to indicate that the ACE is to use either the cookie name or the cookie value to satisfy this match condition.
HTTP Header	<p>HTTP header and corresponding value used to establish match conditions.</p> <ol style="list-style-type: none"> <li data-bbox="391 1098 1464 1308">a. In the Header Name field, specify the header in one of the following ways: <ul style="list-style-type: none"> <li data-bbox="443 1140 1464 1234">– To specify an HTTP header that is not one of the standard HTTP headers, click the first radio button and enter the HTTP header name in the Header Name field. Enter an unquoted text string with no spaces and a maximum of 64 characters. <li data-bbox="443 1245 1464 1308">– To specify one of the standard HTTP headers, click the second radio button and choose the desired HTTP header from the list. <li data-bbox="391 1329 1464 1507">b. In the Header Value field, enter the header-value expression string to compare against the value in the specified field in the HTTP header. Valid entries are text strings with a maximum of 255 alphanumeric characters. The ACE supports regular expressions for matching. Header expressions allow spaces, provided that the spaces are escaped or quoted. All headers in the header map must be matched. Table 13-34 lists the supported characters that you can use in regular expressions.

Table 6-12 Layer 7 HTTP/HTTPS Load-Balancing Conditions and Options (continued)

Match Condition	Action
HTTP URL	<p>Condition that indicates that the ACE is to perform regular expression matching against the received packet data from a particular connection based on the HTTP URL string.</p> <ol style="list-style-type: none"> In the URL Expression field, enter a URL, or portion of a URL, to match. Valid entries are URL strings from 1 to 255 alphanumeric characters. Include only the portion of the URL following <code>www.hostname.domain</code> in the match statement. For example, in the URL <code>www.anydomain.com/latest/whatsnew.html</code>, include only <code>/latest/whatsnew.html</code>. To match the <code>www.anydomain.com</code> portion, the URL string can take the form of a URL regular expression. The ACE supports regular expressions for matching URL strings. Table 13-34 lists the supported characters that you can use in regular expressions. In the Method Expression field, enter the HTTP method to match. Valid entries are unquoted text strings with no spaces and a maximum of 15 alphanumeric characters. The method can either be one of the standard HTTP 1.1 method names (OPTIONS, GET, HEAD, POST, PUT, DELETE, TRACE, or CONNECT) or a text string that must be matched exactly (for example, CORVETTE).
Source Address	<p>Client source IP address used for the match condition.</p> <ol style="list-style-type: none"> In the Source Address field, enter the source IP address of the client. Enter the IP address in dotted-decimal notation (for example, 192.168.11.2). In the Source Netmask field, choose the subnet mask to apply to the source IP address.

- Step 7** In the Primary Action field, choose the action that the virtual server is to perform on the traffic if it matches the specified match criteria:
- **Drop**—Client requests for content are to be discarded when match conditions are met. Continue with [Step 11](#).
 - **Forward**—Client requests for content are to be forwarded without performing load balancing on the requests when match conditions are met. Continue with [Step 11](#).
 - **Load Balance**—Client requests for content are to be directed to a server farm when match conditions are met. Continue with [Step 8](#).
 - **Sticky**—Client requests for content are to be handled by a sticky group when match conditions are met. Continue with [Step 9](#).

Step 8 (Optional) If you chose Load Balance as the primary action, do the following:

- In the Server Farm field, choose the primary server farm to use for load balancing, or choose ***New*** to configure a new server farm (see [Table 6-13](#)).

If you chose an existing object in this field, you can view, modify, or duplicate the selected object's existing configuration. See the “[Shared Objects and Virtual Servers](#)” section on page 6-9 for more information about modifying shared objects in virtual servers.



Note To display statistics and status information for an existing server farm, choose a server farm in the list, and click **Details**. ANM accesses the `show serverfarm name detail` CLI command to display detailed server farm information. See the “[Displaying Server Farm Statistics and Status Information](#)” section on page 7-39.

- b. In the Backup Server Farm field, choose the server farm to act as the backup server farm for load balancing if the primary server farm is unavailable, or choose ***New*** to configure a new backup server farm (see [Table 6-13](#)).

If you chose an existing object in this field, you can view, modify, or duplicate the selected object's existing configuration. See the [“Shared Objects and Virtual Servers” section on page 6-9](#) for more information about modifying shared objects in virtual servers.

Table 6-13 **New Server Farm Attributes**

Field	Description
Name	Unique name for the server farm. Valid entries are unquoted text strings with no spaces and a maximum of 64 characters.
Type	Type of server farm: <ul style="list-style-type: none"> • Host—A typical server farm that consists of real servers that provide content and services to clients. By default, if you configure a backup server farm and all real servers in the primary server farm go down, the primary server farm fails over to the backup server farm. Use the following options to specify thresholds for failover and returning to service. <ol style="list-style-type: none"> 1. In the Partial-Threshold Percentage field, enter the minimum percentage of real servers in the primary server farm that must remain active for the server farm to stay up. If the percentage of active real servers falls below this threshold, the ACE takes the server farm out of service. Valid entries are from 0 to 99. 2. In the Back Inservice field, enter the percentage of real servers in the primary server farm that must be active again for the ACE to place the server farm back into service. Valid entries are from 0 to 99. The value in this field should be larger than the value in the Partial Threshold Percentage field. • Redirect—A server farm that consists only of real servers that redirect client requests to alternate locations specified in the real server configuration.
Fail Action	Action that the ACE takes if any real server in the server farm fails: <ul style="list-style-type: none"> • N/A—Indicates that the ACE is to take no action if any server in the server farm fails. • Purge—Indicates that the ACE is to remove connections to a real server if that real server in the server farm fails. The ACE sends a reset command to both the client and the server that failed. • Reassign—Indicates that the ACE reassign the existing server connections to the backup real server (if configured) if the real server fails after you enter this command. If a backup real server has not been configured for the failing server, this selection leaves the existing connections untouched in the failing real server.

Table 6-13 *New Server Farm Attributes (continued)*

Field	Description
Failaction Reassign Across Vlans	<p>Option that is available only for the ACE module A2(3.0), ACE appliance A4(1.0), and later releases of either device type. This field appears only when the L7 Load-Balancing Action parameters are set as follows: Primary Action: LoadBalance; ServerFarm: New; Fail Action: Reassign.</p> <p>Check the check box to specify that the ACE reassigns the existing server connections to the backup real server on a different VLAN interface (commonly referred to as a bypass VLAN) if the real server fails. If a backup real server has not been configured for the failing server, this option has no effect and leaves the existing connections untouched in the failing real server.</p> <p>Note the following configuration requirements and restrictions when you enable this option:</p> <ul style="list-style-type: none"> • Enable the Transparent option (see the next Field) to instruct the ACE not to use NAT to translate the ACE VIP address to the server IP address. The Failaction Reassign Across Vlans option is intended for use in stateful firewall load balancing (FWLB) on your ACE, where the destination IP address for the connection coming in to the ACE is for the end-point real server, and the ACE reassigns the connection so that it is transmitted through a different next hop. • Enable the MAC Sticky option on all server-side interfaces to ensure that packets that are going to and coming from the same server in a flow will traverse the same firewalls or stateful devices (see the “Configuring VLAN Interfaces” section on page 11-5). • Configure the Predictor Hash Address option. See Table 6-14 for the supported predictor methods and configurable attributes for each predictor method. • You must configure identical policies on the primary interface and the backup-server interface. The backup interface must have the same feature configurations as the primary interface. • If you configure a policy on the backup-server interface that is different from the policies on the primary-server interface, that policy will be effective only for new connections. The reassigned connection will always have only the primary-server interface policies. • Interface-specific features (for example, NAT, application protocol inspection, outbound ACLs, or SYN cookie) are not supported. • You cannot reassign connections to the failed real server after it comes back up. This restriction also applies to same-VLAN backup servers. • Real servers must be directly connected to the ACE. This requirement also applies to same-VLAN backup server. • You must disable sequence number randomization on the firewall (see the “Configuring Connection Parameter Maps” section on page 9-3). • Probe configurations should be similar on both ACEs and the interval values should be low. For example, if you configure a high interval value on ACE-1 and a low interval value on ACE-2, the reassigned connections may become stuck because of the probe configuration mismatch. ACE-2 with the low interval value will detect the primary server failure first and will reassign all its incoming connections to the backup-server interface VLAN. ACE-1 with the high interval value may not detect the failure before the primary server comes back up and will still point to the primary server. <p>To minimize packet loss, we recommend the following probe parameter values on both ACEs: Interval: 2, Faildetect: 2, Passdetect interval: 2, and Passdetect count: 5.</p>

Table 6-13 New Server Farm Attributes (continued)

Field	Description
Transparent	<p>Field that appears only for real servers identified as host servers.</p> <p>Specify whether network address translation from the VIP address to the server IP is to occur. Check the check box to specify that network address translation from the VIP address to the server IP address is to occur. Uncheck the check box to specify that network address translation from the VIP address to the server IP address is not to occur.</p>
Dynamic Workload Scaling	<p>Option that is available only with ACE software version A4(2.0) or later release on either device type (appliance or module). Field that appears only for host server farms.</p> <p>Allows the ACE to burst traffic to remote VMs when the average CPU usage, memory usage, or both of the local VMs has reached it's specified maximum threshold value. The ACE stops bursting traffic to the remote VMs when the average CPU and/or memory usage of the local VMs has dropped to it's specified minimum threshold value. This option requires that you have the ACE configured for Dynamic Workload Scaling using a Nexus 7000, VM Controller, and VM probe (see the “Configuring Dynamic Workload Scaling” section on page 7-18).</p> <p>Click one of the following radio button options:</p> <ul style="list-style-type: none"> • N/A—Not applicable (default). • Local—The ACE can use the VM Controller local VMs only for load balancing (bursting is not allowed). • Burst—Enables the ACE to burst traffic to a remote VMs when needed. <p>When you choose Burst, the VM Probe Name field displays along with a list of available VM probes. Choose an available VM probe or click Add to display the Health Monitoring pop-up window and create a new VM probe or edit an existing one (see the “Configuring Health Monitoring” section on page 7-40).</p>
Fail-On-All	<p>Field that appears for host server farms only.</p> <p>By default, real servers that you configure in a server farm inherit the probes that you configure directly on that server farm. When you configure multiple probes on a server farm, the real servers in the server farm use an OR logic with respect to the probes, which means that if one of the probes configured on the server farm fails, all the real servers in that server farm fail and enter the PROBE-FAILED state.</p> <p>With AND logic, if one server farm probe fails, the real servers in the server farm remain in the OPERATIONAL state. If all the probes associated with the server farm fail, then all the real servers in that server farm fail and enter the PROBE-FAILED state. You can also configure AND logic for probes that you configure directly on real servers in a server farm. For more information, see the command in server farm host real server configuration mode.</p> <p>Check this check box to configure the real servers in a server farm to use AND logic with respect to multiple server farm probes.</p> <p>The Fail On All function is applicable to all probe types.</p>

Table 6-13 *New Server Farm Attributes (continued)*

Field	Description
Inband-Health Check	<p>Option that is available only for the ACE module A4(1.0), ACE appliance A4(1.0), and later releases of either device type. Field that appears only for host server farms.</p> <p>By default, the ACE monitors the health of all real servers in a configuration through the use of ARPs and health probes. However, there is latency period between when the real server goes down and when the ACE becomes aware of the state. The inband health monitoring feature allows the ACE to monitor the health of the real servers in the server farm through the following connection failures:</p> <ul style="list-style-type: none"> • For TCP, resets (RSTs) from the server or SYN timeouts. • For UDP, ICMP Host, Network, Port, Protocol, and Source Route unreachable messages. <p>When you configure the failure-count threshold and the number of these failures exceeds the threshold within the reset-time interval, the ACE immediately marks the server as failed, takes it out of service, and removes it from load balancing. The server is not considered for load balancing until the optional resume-service interval expires.</p> <p>The Inband-Health Check attributes are as follows:</p> <ul style="list-style-type: none"> • Count—Tracks the total number of TCP or UDP failures, and increments the counters. • Log—Logs a syslog error message when the number of events reaches the threshold value that you set for the Connection Failure Threshold Count attribute. • Remove—Logs a syslog error message when the number of events reaches the configured threshold and removes the real server from service.
Connection Failure Threshold Count	<p>This field appears only when the Inband-Health Check is set to Log or Remove.</p> <p>Enter the maximum number of connection failures that a real server can exhibit in the reset-time interval before ACE marks the real server as failed. Valid entries are as follows:</p> <ul style="list-style-type: none"> • ACE appliance—Integers from 1 to 4294967295 • ACE module—Integers from 4 to 4294967295
Reset Timeout (Milliseconds)	<p>This field appears only when the Inband-Health Check is set to Log or Remove.</p> <p>Enter the number of milliseconds for the reset-time interval. Valid entries are integers from 100 to 300000. The default interval is 100.</p> <p>This interval starts when the ACE detects a connection failure. If the connection failure threshold is reached during this interval, the ACE generates a syslog message. If you configure the Remove attribute, the ACE also removes the real server from service.</p> <p>Changing the setting of this option affects the behavior of the real server, as follows:</p> <ul style="list-style-type: none"> • When the real server is in the OPERATIONAL state, even if several connection failures have occurred, the new reset-time interval takes effect the next time that a connection error occurs. • When the real server in the INBAND-HM-FAILED state, the new reset-time interval takes effect the next time that a connection error occurs after the server transitions to the OPERATIONAL state.

Table 6-13 New Server Farm Attributes (continued)

Field	Description
Resume Service (Seconds)	<p>Field that appears only when the Inband-Health Check is set to Remove.</p> <p>Enter the number of seconds after a server has been marked as failed to reconsider it for sending live connections. Valid entries are integers from 30 to 3600. The default setting is 0. The setting of this option affects the behavior of the real server in the inband failed state, as follows:</p> <ul style="list-style-type: none"> • When this field is not configured and has the default setting of 0, the real server remains in the failed state until you manually suspend and then reactivate it. • When this field is not configured and has the default setting of 0 and then you configure this option with an integer between 30 and 3,600, the failed real server immediately transitions to the Operational state. • When you configure this field and then increase the value, the real server remains in the failed state for the duration of the previously-configured value. The new value takes effect the next time the real server transitions to the failed state. • When you configure this field and then decrease the value, the failed real server immediately transitions to the Operational state. • When you configure this field with an integer between 30 and 3,600 and then reset it to the default of 0, the real server remains in the failed state for the duration of the previously-configured value. The default setting takes effect the next time the real server transitions to the failed state. Then the real server remains in the failed state until you manually suspend and then reactivate it. • When you change this field within the reset-time interval the real server in the OPERATIONAL with several connection failures, the new threshold interval takes effect the next time that a connection error occurs, even if it occurs within the current reset-time interval.
Predictor	<p>Method for selecting the next server in the server farm to respond to client requests. Round Robin is the default predictor method for a server farm.</p> <p>See Table 6-14 for the supported predictor methods and configurable attributes for each predictor method.</p>

Table 6-13 New Server Farm Attributes (continued)

Field	Description
Probes	<p>Health monitoring probes to use:</p> <ul style="list-style-type: none"> To include a probe that you want to use for health monitoring, choose it in the Available list, and click Add. The probe appears in the Selected list. <p>The redirect real server probe list contains only configured probes of the type Is Routed, which means that the ACE routes the probe address according to the ACE internal routing table (see the “Configuring Health Monitoring” section on page 7-40</p> <p> Note The list of available probes does not include VM health monitoring probes. To choose a VM probe for monitoring local VM usage, see the Dynamic Workload Scaling field.</p> <ul style="list-style-type: none"> To remove a probe that you do not want to use for health monitoring, choose it in the Selected list, and click Remove. The probe appears in the Available list. To specify a sequence for probe use, choose probes in the Selected list, and click Up or Down until you have the desired sequence. To view the configuration for an existing probe, choose a probe in the list on the right, and click View to review its configuration. To display statistics and status information for an existing probe, choose a probe in the list on the right, and click Details. ANM accesses the show probe name detail CLI command to display detailed probe information. See the “Displaying Health Monitoring Statistics and Status Information” section on page 7-67. <p>To add a new probe, click Create. See the “Configuring Health Monitoring for Real Servers” section on page 7-42 for details on adding a new health monitoring probe and defining attributes for the specific probe type. In addition to the probe attributes that you set as described in the Configuring Health Monitoring for Real Servers section, set the following probe configuration parameters in the Probes section under Server Farm as described below:</p> <ul style="list-style-type: none"> Expect Addresses—To configure expect addresses for a DNS probe, in the Expect Addresses field enter the IP address that the ACE is to expect as a server response to a DNS request. Valid entries are unique IP addresses in dotted-decimal notation, such as 192.168.11.1. Probe Headers—To configure probe headers for either an HTTP or HTTPS probe, in the Probe Headers field enter the name of the HTTP header and the value to be matched using the format <i>header_name=header_value</i> where: <ul style="list-style-type: none"> <i>header_name</i> represents the HTTP header name the probe is to use. Valid entries are unquoted text strings with no spaces and a maximum of 64 alphanumeric characters. You can specify predefined header or any custom header name provided that it does not exceed the maximum length limit. <i>header_value</i> represents the string to assign to the header field. Valid entries are text strings with a maximum of 255 characters. If the string includes spaces, enclose the string with quotes.

Table 6-13 *New Server Farm Attributes (continued)*

Field	Description
Probes (continued)	<ul style="list-style-type: none"> • Probe Expect Status—To configure probe expect status for an FTP, HTTP, HTTPS, RTSP, SIP-TCP, SIP-UDP, or SMTP probe, in the Probe Expect Status field enter the following information: <ul style="list-style-type: none"> • To configure a single expect status code, enter the minimum expect status code for this probe followed by the same expect status code that you entered as the minimum. Valid entries are from 0 to 999. • To configure a range of expect status codes, enter the lower limit of the range of status codes followed by the upper limit of the range of status codes. The maximum expect status code must be greater than or equal to the value specified for the minimum expect status code. Valid entries are from 0 to 999. • SNMP OID Table—To configure the SNMP OID for an SNMP probe, see the “Configuring an OID for SNMP Probes” section on page 7-66. <p>After you add a probe, you can modify the attributes for a health probe from the Health Monitoring table (Config > Virtual Contexts > <i>context</i> > Load Balancing > Health Monitoring) as described in the “Configuring Health Monitoring for Real Servers” section on page 7-42. You can also delete an existing health probe from the Health Monitoring table.</p>

Table 6-13 New Server Farm Attributes (continued)

Field	Description
Real Servers	<p>Table allows you to add, modify, remove, or change the order of real servers.</p> <ol style="list-style-type: none"> a. Choose an existing server, or click Add to add a server to the server farm and do one of the following: <ul style="list-style-type: none"> – If you chose an existing server, you can view, modify, or duplicate the server’s existing configuration. See the “Shared Objects and Virtual Servers” section on page 6-9 for more information about modifying shared objects. – If you click Add, the window refreshes so you can enter server information. b. In the Name field, specify the name of the real server in one of the following ways: <ul style="list-style-type: none"> – To identify a new real server, click the first radio button, and then enter the name of the real server in the adjoining field. – To specify an existing real server, click the second radio button, and then choose one of the real servers listed. c. In the IP Address field, enter the IP address of the real server in dotted-decimal format. d. In the Port field, enter the port number to be used for server port address translation (PAT). Valid entries are from 1 to 65535. e. In the Weight field, enter the weight to assign to this server in the server farm. Valid entries are from 1 to 100, and the default is 8. f. In the Redirection Code field, choose the appropriate redirection code. This field appears only for real servers identified as redirect servers. <ul style="list-style-type: none"> – N/A—Indicates that the webhost redirection code is not defined. – 301—Indicates that the requested resource has been moved permanently. For future references to this resource, the client should use one of the returned URIs. – 302—Indicates that the requested resource has been found, but has been moved temporarily to another location. For future references to this resource, the client should use the request URI because the resource may be moved to other locations from time to time. g. In the Web Host Redirection field, enter the URL string used to redirect requests to another server. This field appears only for real servers identified as redirect servers. Enter the URL and port used to redirect requests to another server. Valid entries are in the form <code>http://host.com:port</code> where host is the name of the server and port is the port to be used. Valid host entries are unquoted text strings with no spaces and a maximum of 255 characters. Valid port numbers are from 1 to 65535. <p>The relocation string supports the following special characters:</p> <ul style="list-style-type: none"> – <code>%h</code>—Inserts the hostname from the request Host header – <code>%p</code>—Inserts the URL path string from the request h. In the Rate Bandwidth field, enter the real server bandwidth limit in bytes per second. Valid entries are from 1 to 300000000 bytes. i. In the Rate Connection field, enter the limit for connections per second (valid entries are from 1 to 350000) and do one of the following: <ul style="list-style-type: none"> – Click OK to accept your entries and add this real server to the server farm. The table refreshes with updated information. – Click Cancel to exit this procedure without saving your entries and to return to the Real Servers table.

Table 6-13 New Server Farm Attributes (continued)

Field	Description
	<p>j. In the State field, choose the administrative state of this server as follows:</p> <ul style="list-style-type: none"> – In Service—The server is to be placed in use as a destination for server load balancing. – In Service Standby—The server is a backup server and remains inactive unless the primary server fails. If the primary server fails, the backup server becomes active and starts accepting connections. – Out Of Service—The server is not to be placed in use by a server load balancer as a destination for client connections. <p>k. In the Fail-On-All field, check this check box to configure a real server to remain in the OPERATIONAL state unless all probes associated with it fail (AND logic). The Fail-On-All function is applicable to all probe types. Fail-On-All is applicable only for host real servers.</p> <p>l. Do one of the following:</p> <ul style="list-style-type: none"> – Click OK to accept your entries and add this real server to the server farm. The table refreshes with updated information. – Click Cancel to exit this procedure without saving your entries and to return to the Real Servers table. <p>To display statistics and status information for an existing real server, choose a real server in the list, and then click Details. ANM accesses the show rserver name detail CLI command to display detailed real server information. See the “Displaying Real Server Statistics and Status Information” section on page 7-8.</p>

Table 6-14 Predictor Methods and Attributes

Predictor Method	Description / Action
Hash Address	<p>Method that indicates that the ACE is to select the server using a hash value based on the source or destination IP address.</p> <p>To configure the hash address predictor method, do the following:</p> <p>a. In the Mask Type field, indicate whether server selection is based on the source IP address or the destination IP address:</p> <ul style="list-style-type: none"> – N/A—Indicates that this option is not defined. – Destination—Indicates that the server is selected based on the destination IP address. – Source—Indicates that the server is selected based on the source IP address. <p>b. In the IP Netmask field, choose the subnet mask to apply to the address. If none is specified, the default is 255.255.255.255.</p>

Table 6-14 Predictor Methods and Attributes (continued)

Hash Content	<p>Method that indicates that the ACE is to select the server by using a hash value based on the specified content string of the HTTP packet body.</p> <p>a. In the Begin Pattern field, enter the beginning pattern of the content string and the pattern string to match before hashing. If you do not specify a beginning pattern, the ACE starts parsing the HTTP body immediate following the offset byte. You cannot configure different beginning and ending patterns for different server farms that are part of the same traffic classification.</p> <p>Valid entries are unquoted text strings with no spaces and a maximum of 255 alphanumeric characters. The ACE supports regular expressions for matching string expressions. Table 13-34 lists the supported characters that you can use for matching string expressions.</p> <p>b. In the End Pattern field, enter the pattern that marks the end of hashing. If you do not specify either a length or an end pattern, the ACE continues to parse the data until it reaches the end of the field or the end of the packet, or until it reaches the maximum body parse length. You cannot configure different beginning and ending patterns for different server farms that are part of the same traffic classification.</p> <p>Valid entries are unquoted text strings with no spaces and a maximum of 255 alphanumeric characters. The ACE supports regular expressions for matching string expressions. Table 13-34 lists the supported characters that you can use for matching string expressions.</p> <p>c. In the Length (Bytes) field, enter the length in bytes of the portion of the content (starting with the byte after the offset value) that the ACE uses for sticking the client to the server. Valid entries are from 1 to 1000 bytes.</p> <p>The offset and length can vary from 0 to 1000 bytes. If the payload is longer than the offset but shorter than the offset plus the length of the payload, the ACE sticks the connection based on that portion of the payload starting with the byte after the offset value and ending with the byte specified by the offset plus the length. The total of the offset and the length cannot exceed 1000.</p> <p> Note You cannot specify both the length and the end-pattern options for a Hash Content predictor.</p> <p>d. In the HTTP Content Offset (Bytes) field, enter the portion of the content that the ACE uses to stick the client on a particular server by indicating the bytes to ignore starting with the first byte of the payload. Valid entries are from 0 to 999 bytes. The default is 0, which indicates that the ACE does not exclude any portion of the content.</p>
Hash Cookie	<p>Method that indicates that the ACE is to select the server by using a hash value based on the cookie name.</p> <p>In the Cookie Name field, enter a cookie name in the form of an unquoted text string with no spaces and a maximum of 64 characters.</p>

Table 6-14 Predictor Methods and Attributes (continued)

Hash Header	<p>Method that indicates that the ACE is to select the server by using a hash value based on the header name.</p> <p>In the Header Name field, choose the HTTP header to be used for server selection as follows:</p> <ul style="list-style-type: none"> To specify an HTTP header that is not one of the standard HTTP headers, click the first radio button and enter the HTTP header name in the Header Name field. Valid entries are unquoted text strings with no spaces and a maximum of 64 characters. To specify one of the standard HTTP headers, click the second radio button, and then choose one of the HTTP headers from the list.
Hash Layer 4	<p>Method that indicates that the ACE is to select the server by using a Layer 4 generic protocol load-balancing method. Use this predictor to load balance packets from protocols that are not explicitly supported by the ACE.</p> <p>a. In the Begin Pattern field, enter the beginning pattern of the Layer 4 payload and the pattern string to match before hashing. If you do not specify a beginning pattern, the ACE starts parsing the HTTP body immediate following the offset byte. You cannot configure different beginning and ending patterns for different server farms that are part of the same traffic classification.</p> <p>Valid entries are unquoted text strings with no spaces and a maximum of 255 alphanumeric characters. The ACE supports regular expressions for matching string expressions. Table 13-34 lists the supported characters that you can use for matching string expressions.</p> <p>b. In the End Pattern field, enter the pattern that marks the end of hashing. If you do not specify either a length or an end pattern, the ACE continues to parse the data until it reaches the end of the field or the end of the packet, or until it reaches the maximum body parse length. You cannot configure different beginning and ending patterns for different server farms that are part of the same traffic classification.</p> <p>Valid entries are unquoted text strings with no spaces and a maximum of 255 alphanumeric characters. The ACE supports regular expressions for matching string expressions. Table 13-34 lists the supported characters that you can use for matching string expressions.</p> <p>c. In the Length (Bytes) field, enter the length in bytes of the portion of the payload (starting with the byte after the offset value) that the ACE uses for sticking the client to the server. Valid entries are from 1 to 1000 bytes.</p> <p>The offset and length can vary from 0 to 1000 bytes. If the payload is longer than the offset but shorter than the offset plus the length of the payload, the ACE sticks the connection based on that portion of the payload starting with the byte after the offset value and ending with the byte specified by the offset plus the length. The total of the offset and the length cannot exceed 1000.</p> <p> Note You cannot specify both the length and end-pattern options for a Hash Layer 4 predictor.</p> <p>d. In the HTTP Content Offset (Bytes) field, enter the portion of the content that the ACE uses to stick the client on a particular server by indicating the bytes to ignore starting with the first byte of the payload. Valid entries are from 0 to 999 bytes. The default is 0, which indicates that the ACE does not exclude any portion of the content.</p>

Table 6-14 *Predictor Methods and Attributes (continued)*

Hash URL	<p>Method that indicates that the ACE is to select the server using a hash value based on the URL. Use this method to load balance firewalls.</p> <p>Enter values in one or both of the pattern fields:</p> <ul style="list-style-type: none"> • In the URL Begin Pattern field, enter the beginning pattern of the URL and the pattern string to parse. • In the URL End Pattern field, enter the ending pattern of the URL and the pattern string to parse. <p>Valid entries for these fields are unquoted text strings with no spaces and a maximum of 255 alphanumeric characters for each pattern you configure.</p>
Least Bandwidth	<p>Method that indicates that the ACE is to select the server with the least amount of network traffic over a specified sampling period.</p> <ol style="list-style-type: none"> a. In the Assess Time field, enter the number of seconds for which the ACE is to collect traffic information. Valid entries are from 1 to 10 seconds. b. In the Least Bandwidth Samples field, enter the number of samples over which you want to weight and average the results of the probe query to calculate the final load value. Valid entries are 1, 2, 4, 8, and 16 (values from 1 to 16 that are also a power of 2).
Least Connections	<p>Method that indicates that the ACE is to select the server with the fewest number of connections.</p> <p>In the Slowstart Duration field, enter the slow-start value to be applied to this predictor method. Valid entries are from 1 to 65535, where 1 is the slowest ramp-up value.</p> <p>The slow-start mechanism is used to avoid sending a high rate of new connections to servers that you have just put into service.</p>

Table 6-14 Predictor Methods and Attributes (continued)

Least Loaded	<p>Method that indicates that the ACE is to select the server with the lowest load based on information from SNMP probes.</p> <ol style="list-style-type: none"> a. In the SNMP Probe Name field, choose the name of the SNMP probe to use. b. In the Auto Adjust field, configure the autoadjust feature to assign a maximum load value of 16000 to that server to prevent it from being flooded with new incoming connections. The ACE periodically adjusts this load value based on feedback from the server's SNMP probe and other configured options. Options include: <ul style="list-style-type: none"> – Average—Instructs the ACE to apply the average load of the server farm to a real server whose load reaches zero. The average load is the running average of the load values across all real servers in the server farm. This is the default setting. – Maxload—Instructs the ACE to apply the maximum load of the server farm to a real server whose load reaches zero. <p>The maxload option requires the following ACE software versions:</p> <ul style="list-style-type: none"> - ACE appliance—A3(2.7) or A4(1.0) or later - ACE module—A2(2.4), A2(3.2), or A4(1.0) or later <p>If you choose the maxload option and the ACE does not support the option, ANM issues a command parse error message.</p> <ul style="list-style-type: none"> – Off—Instructs the ACE to send all new connections to the server that has a load of zero until the next load update arrives from the SNMP probe for this server. There may be times when you want the ACE to send all new connections to a real server whose load is zero. c. In the Weight Connection field, check the check box to instruct the ACE to use the current connection count in the final load calculation for a real server. When you configure this option, the ACE includes the current connection count in the total load calculation for each real server in a server farm. Uncheck the check box to reset the behavior of the ACE to the default of excluding the current connection count from the load calculation.
Response	<p>Method that indicates that the ACE is to select the server with the lowest response time for a requested response-time measurement.</p> <ol style="list-style-type: none"> a. In the Response Type field, choose the type of measurement to use: <ul style="list-style-type: none"> – App-Req-To-Resp—The response time from when the ACE sends an HTTP request to a server to the time that the ACE receives a response from the server for that request. – Syn-To-Close—The response time from when the ACE sends a TCP SYN to a server to the time that the ACE receives a CLOSE from the server. – Syn-To-Synack—The response time from when the ACE sends a TCP SYN to a server to the time that the ACE receives a SYN-ACK from the server. b. In the Response Samples field, enter the number of samples over which you want to average the results of the response-time measurement. Valid entries are 1, 2, 4, 8, and 16 (values from 1 to 16 that are also a power of 2). c. In the Weight Connection field, check the check box to instruct the ACE to use the current connection count in the final load calculation for a real server. When you configure this option, the ACE includes the current connection count in the total load calculation for each real server in a server farm. Uncheck the check box to reset the behavior of the ACE to the default of excluding the current connection count from the load calculation.

Table 6-14 Predictor Methods and Attributes (continued)

Round Robin	Method that indicates that the ACE is to select the next server in the list of servers based on server weight. This is the default predictor method.
-------------	--

- Step 9** (Optional) If you chose Sticky as the primary action, in the Sticky Group field, choose an existing sticky group or click ***New*** to add a new sticky group ([Table 6-15](#)).



Note If you chose an existing sticky group, you can view, modify, or duplicate the selected object's existing configuration. See the [“Shared Objects and Virtual Servers” section on page 6-9](#) for more information about modifying shared objects in virtual servers.

Table 6-15 Sticky Group Attributes

Field	Description
Group Name	Unique identifier for the sticky group. You can either accept the automatically incremented entry that was provided or you can enter your own. Valid entries are unquoted text strings with no spaces and a maximum of 64 alphanumeric characters.
Type	<p>Method to be used when establishing sticky connections and configure any type-specific attributes:</p> <p>Note The available selections listed in the Type drop-down list will vary depending on your selection for Application Protocol in the Properties configuration subset (see Table 6-2). For example, if you chose HTTP or HTTPS as the application protocol, only IP Netmask, HTTP Cookie, HTTP Header, and HTTP Content appear as selections in the Type drop-down.</p> <ul style="list-style-type: none"> • HTTP Content—The virtual server is to stick client connections to the same real server based on a string in the data portion of the HTTP packet. See Table 8-2 for additional configuration options. • HTTP Cookie—The virtual server is either to learn a cookie from the HTTP header of a client request or to insert a cookie in the Set-Cookie header of the response from the server to the client, and then use the learned cookie to provide stickiness between the client and server for the duration of the transaction. See Table 8-3 for additional configuration options. • HTTP Header—The virtual server is to stick client connections to the same real server based on HTTP headers. See Table 8-4 for additional configuration options. • IP Netmask—The virtual server is to stick a client to the same server for multiple subsequent connections as needed to complete a transaction using the client source IP address, the destination IP address, or both. See Table 8-5 for additional configuration options. <p>Note If an organization uses a megaproxy to load balance client requests across multiple proxy servers when a client connects to the Internet, the source IP address is no longer a reliable indicator of the true source of the request. In this situation, you can use cookies or another sticky method to ensure session persistence.</p> <ul style="list-style-type: none"> • Layer 4 Payload—The virtual server is to stick client connections to the same real server based on a string in the payload portion of the Layer 4 protocol packet. See Table 8-6 for additional configuration options. • RADIUS—The virtual server is to stick client connections to the same real server based on a RADIUS attribute. • RTSP Header—The virtual server is to stick client connections to the same real server based on the RTSP Session header field. Table 8-8 for additional configuration options. • SIP Header—The virtual server is to stick client connections to the same real server based on the SIP Call-ID header field.
Sticky Server Farm	Existing server farm that is to act as the primary server farm for this sticky group. You can choose *New* to create a new server farm. If you chose *New* , configure the server farm using the information in Table 6-13 .
Backup Server Farm	Existing server farm that is to act as the backup server farm this sticky group. You can choose *New* to create a new server farm. If you chose *New* , configure the server farm using the information in Table 6-13 .

Table 6-15 *Sticky Group Attributes (continued)*

Field	Description
Aggregate State	Check box to indicate that the state of the primary server farm is to be tied to the state of all real servers in the server farm and in the backup server farm, if configured. The ACE declares the primary server farm down if all real servers in the primary server farm and all real servers in the backup server farm are down. Uncheck the check box if the state of the primary server farm is not to be tied to all real servers in the server farm and in the backup server farm.
Sticky Enabled On Backup Server Farm	Check box to indicate that the backup server farm is sticky. Uncheck the check box if the backup server farm is not sticky.
Replicate On HA Peer	Check box to indicate that the virtual server is to replicate sticky table entries on the backup server farm. If a failover occurs and this option is selected, the new active server farm can maintain the existing sticky connections. Uncheck the check box to indicate that the virtual server is not to replicate sticky table entries on the backup server farm.
Timeout (Minutes)	Number of minutes that the virtual server keeps the sticky information for a client connection in the sticky table after the latest client connection terminates. Valid entries are from 1 to 65535; the default is 1440 minutes (24 hours).
Timeout Active Connections	Check box to specify that the virtual server is to time out sticky table entries even if active connections exist after the sticky timer expires. Uncheck the check box to specify that the virtual is not to time out sticky table entries even if active connections exist after the sticky timer expires. This behavior is the default.

Step 10 (Optional) If you are using the ACE appliance (all versions) or ACE module version A4(1.0) and later, in the Compression Method field, choose the HTTP compression method to indicate how the ACE appliance is to compress packets when a client request indicates that the client browser is capable of packet compression.

By default, HTTP compression is disabled in the ACE. When you configure HTTP compression using the ACE, the appliance compresses data in the HTTP GET responses from the real servers. The ACE does not compress HTTP requests from clients or the HTTP headers in the server responses.



Note By default, the ACE appliance supports HTTP compression at rates of 100 megabits per second (Mbps). Installing an optional HTTP compression license allows you to increase this value to a maximum of 2 Gbps. See the *Cisco 4700 Series Application Control Engine Appliance Administration Guide* for information on ACE licensing options.

Options are as follows:

- **Gzip**—Specifies the gzip compression format as the method to use when the client browser supports both the deflate and gzip compression methods. Gzip is the file format for compression described in RFC1952.
- **Deflate**—Specifies the deflate compression format as the method to use when the client browser supports both the deflate and gzip compression methods. Deflate is the data format for compression described in RFC1951.
- **N/A**—HTTP compression is disabled.

When configuring HTTP compression, we recommend that you exclude the following MIME types from HTTP compression: “.gif”, “.css”, “.js”, “.class”, “.jar”, “.cab”, “.txt”, “.ps”, “.vbs”, “.xsl”, “.xml”, “.pdf”, “.swf”, “.jpg”, “.jpeg”, “.jpe”, or “.png”.

When you enable HTTP compression, the ACE compresses the packets using the following default compression parameter values:

- Mime type—All text formats (text/*).
- Minimum size—512 bytes.
- User agent—None.

Step 11 In the SSL Initiation field, choose an existing service or choose ***New*** to create a new service, and do one of the following:

- If you chose an existing SSL service, you can view, modify, or duplicate the existing configuration. See [“Shared Objects and Virtual Servers” section on page 6-9](#) for more information about modifying shared objects.
- If you chose ***New***, configure the service using the information in [Table 6-5](#). For more information about SSL, see the [“Configuring SSL” section on page 10-1](#).

Step 12 In the Insert HTTP Headers field, enter the name of the HTTP header and the value to be matched using the *header_name=header_value* format where:

- *header_name* represents the name of the HTTP header to insert in the client HTTP request. Valid entries are unquoted text strings with no spaces and a maximum of 64 alphanumeric characters. You can specify predefined header or any custom header name provided that it does not exceed the maximum length limit.
- *header_value* represents the expression string to compare against the value in the specified field in the HTTP header. Valid entries are text strings with a maximum of 255 alphanumeric characters. The ACE supports regular expressions for matching. Header expressions allow spaces, provided that the spaces are escaped or quoted. All headers in the header map must be matched. [Table 13-34](#) lists the supported characters that you can use in regular expressions.

For example, you might enter **Host=www.cisco.com**.

Step 13 Do one of the following:

- Click **OK** to save your entries and to return to the Rule Match table.
- Click **Cancel** to exit this procedure without saving your entries and to return to the Rule Match table.

Step 14 If you are adding Rule Match entries for a new virtual server and you want to modify the sequence of rules in the L7 Load Balancing section of the Virtual Server configuration page, click **Up** or **Down** to change the order of the entries in the Rule Match table.



Note The Up and Down buttons are not available for an existing virtual server, only for a new virtual server. To reorder the entries in the Rule Match table for an existing virtual server, go to Config > Expert > Policy Maps and choose the Layer 7 load balancing policy map, delete the rule entry that you want to reorder, and then add it again by using the Insert Before option to put it in the correct order. See [“Configuring Rules and Actions for Policy Maps” section on page 13-34](#) for details.

Step 15 Do one of the following:

- Click **Deploy Now** to deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files.
- Click **Cancel** to exit this procedure without saving your entries.

- Click **Deploy Later** to save your entries and apply them at a later time.
-

Related Topics

- [Configuring Virtual Servers](#), page 6-2
- [Configuring Virtual Server Properties](#), page 6-11
- [Configuring Virtual Server SSL Termination](#), page 6-17
- [Configuring Virtual Server Protocol Inspection](#), page 6-18

Configuring Virtual Server Default Layer 7 Load Balancing

You can configure default Layer 7 load-balancing actions for all network traffic that does not meet previously specified match conditions.

Assumption

Make sure that a virtual server has been configured in the Properties configuration subset. For more information, see the “[Configuring Virtual Server Properties](#)” section on page 6-11. See the “[Configuring Virtual Servers](#)” section on page 6-2 for information on configuring a virtual server.

Procedure

- Step 1** Choose **Config > Devices > context > Load Balancing > Virtual Servers**.
The Virtual Servers table appears.
- Step 2** In the Virtual Servers table, choose the virtual server that you want to configure for default Layer 7 load balancing, then click **Edit**.
The Virtual Server configuration window appears.
- Step 3** In the Virtual Server configuration window, click **Default L7 Load-Balancing Action**.
The Default L7 Load-Balancing Action configuration pane appears.
- Step 4** In the Primary Action field of the Default L7 Load-Balancing Action configuration pane, choose the default action that the virtual server is to take in response to client requests for content when specified match conditions are not met:
- **Drop**—Client requests that do not meet specified match conditions are to be discarded. Continue with [Step 8](#).
 - **Forward**—Client requests that do not meet specified match conditions are to be forwarded without performing load balancing on the requests. Continue with [Step 8](#).
 - **Load Balance**—Client requests for content are to be directed to a server farm. Continue with [Step 5](#).
 - **Sticky**—Client requests for content are to be handled by a sticky group when match conditions are met. Continue with [Step 6](#).
- Step 5** (Optional) If you chose Load Balance as the primary action, do the following:
- a. In the Server Farm field, choose the primary server farm to use for load balancing, or choose ***New*** to configure a new server farm (see [Table 6-13](#)).



Note To display statistics and status information for an existing server farm, choose a server farm in the list, and then click **Details**. ANM accesses the **show serverfarm name detail** CLI command to display detailed server farm information. See the “[Displaying Server Farm Statistics and Status Information](#)” section on page 7-39.

- b. In the Backup Server Farm field, choose the server farm to act as the backup server farm for load balancing if the primary server farm is unavailable, or choose ***New*** to configure a new backup server farm (see [Table 6-13](#)).



Note If you chose an existing object in either field, you can view, modify, or duplicate the selected object’s existing configuration. See the “[Shared Objects and Virtual Servers](#)” section on page 6-9 for more information about modifying shared objects in virtual servers.

- Step 6** (Optional) If you chose Sticky as the primary action, in the Sticky Group field, choose an existing sticky group or click ***New*** to add a new sticky group (see [Table 6-15](#)).



Note If you chose an existing sticky group, you can view, modify, or duplicate the selected object’s existing configuration. See the “[Shared Objects and Virtual Servers](#)” section on page 6-9 for more information about modifying shared objects in virtual servers.

- Step 7** (Optional) If you are using the ACE appliance (all versions) or ACE module version A4(1.0) and later, in the Compression Method field, choose the HTTP compression method to indicate how the ACE appliance is to compress packets when a client request indicates that the client browser is capable of packet compression.

By default, HTTP compression is disabled in the ACE. When you configure HTTP compression using the ACE, the appliance compresses data in the HTTP GET responses from the real servers. The ACE does not compress HTTP requests from clients or the HTTP headers in the server responses.



Note By default, the ACE appliance supports HTTP compression at rates of 100 megabits per second (Mbps). Installing an optional HTTP compression license allows you to increase this value to a maximum of 2 Gbps. See the *Cisco 4700 Series Application Control Engine Appliance Administration Guide* for information on ACE licensing options.

Options are as follows:

- **Deflate**—Specifies the deflate compression format as the method to use when the client browser supports both the deflate and gzip compression methods. deflate, the data format for compression described in RFC1951.
- **Gzip**—Specifies the gzip compression format as the method to use when the client browser supports both the deflate and gzip compression methods. Gzip is the file format for compression described in RFC1952.
- **N/A**—HTTP compression is disabled.



Note If you enable the Gzip or Deflate compression format, ANM automatically inserts a L7 Load Balance Primary Action to exclude the MIME types listed above. However, if you disable HTTP compression later on, you will need to remove the auto-inserted Load Balance Primary Action.

When you enable HTTP compression, the ACE compresses the packets using the following default compression parameter values:

- Mime type—All text formats (text/*).
- Minimum size—512 bytes.
- User agent—None.

- Step 8** In the SSL Initiation field, choose an existing service or choose ***New*** to create a new service:
- If you chose an existing SSL service, you can view, modify, or duplicate the existing configuration. See the “[Shared Objects and Virtual Servers](#)” section on page 6-9 for more information about modifying shared objects.
 - If you chose ***New***, configure the service using the information in [Table 6-5](#). For more information about SSL, see the “[Configuring SSL](#)” section on page 10-1.
- Step 9** In the Insert HTTP Headers field, enter the name of the HTTP header and the value to be matched using the *header_name=header_value* format where:
- *header_name* represents the name of the HTTP header to insert in the client HTTP request. Valid entries are unquoted text strings with no spaces and a maximum of 64 alphanumeric characters. You can specify predefined header or any custom header name provided that it does not exceed the maximum length limit.
 - *header_value* represents the expression string to compare against the value in the specified field in the HTTP header. Valid entries are text strings with a maximum of 255 alphanumeric characters. The ACE supports regular expressions for matching. Header expressions allow spaces, provided that the spaces are escaped or quoted. All headers in the header map must be matched. [Table 13-34](#) lists the supported characters that you can use in regular expressions.

For example, you might enter **Host=www.cisco.com**.

- Step 10** Do one of the following:
- Click **Deploy Now** to deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files.
 - Click **Cancel** to exit this procedure without saving your entries and to return to the Virtual Servers table.
 - Click **Deploy Later** to save your entries and apply the configuration at a later time.

Related Topics

- [Configuring Virtual Server Properties](#), page 6-11
- [Configuring Virtual Server SSL Termination](#), page 6-17
- [Configuring Virtual Server Protocol Inspection](#), page 6-18
- [Configuring Virtual Server Layer 7 Load Balancing](#), page 6-30

Configuring Application Acceleration and Optimization



Note

This option is available only for ACE appliances and only in the Advanced View.

You can configure acceleration and optimization on virtual servers that are configured on ACE appliances. The ACE appliance includes configuration options that allow you to accelerate enterprise applications, resulting in increased employee productivity, enhanced customer retention, and increased online revenues. The application acceleration functions of the ACE appliance apply several optimization technologies to accelerate Web application performance. This application acceleration functionality enables enterprises to optimize network performance and improve access to critical business information. It also accelerates the performance of Web applications, including customer relationship management (CRM), portals, and online collaboration by up to 10 times.

See the “[Configuring Application Acceleration and Optimization](#)” section on page 14-1 or the *Cisco 4700 Series Application Control Engine Appliance Application Acceleration and Optimization Configuration Guide* for more information about application acceleration and optimization.

Assumption

Make sure that a virtual server has been configured on an ACE appliance with HTTP or HTTPS as the application protocol. See the “[Configuring Virtual Servers](#)” section on page 6-2 for information about configuring a virtual server.

Procedure

-
- Step 1** Choose **Config > Devices > context > Load Balancing > Virtual Servers**.
- The Virtual Servers table appears.
- Step 2** In the Virtual Servers table, choose the virtual server that you want to configure for optimization, and click **Edit**.
- The Virtual Server configuration window appears.
- Step 3** In the Virtual Server configuration window, click **Application Acceleration And Optimization**.
- The Application Acceleration And Optimization configuration pane appears.
- Step 4** In the Configuration field of the Application Acceleration And Optimization configuration pane, choose the method that you want to use to configure application acceleration and optimization:
- **EZ**—Use standard acceleration and optimization options. Continue with [Step 5](#).
 - **Custom**—Associate specific match criteria, actions, and parameter maps for application acceleration and optimization for the virtual server. If you choose this option, continue with [Step 6](#) through [Step 14](#).
- Step 5** (Optional) If you chose EZ, the Latency Optimization (FlashForward) and Bandwidth Optimization (Delta) fields appear.
- Do the following:
- a. Check the **Latency Optimization (FlashForward)** check box to specify that the ACE appliance is to use bandwidth reduction and download acceleration techniques to objects embedded within HTML pages. Uncheck the check box to specify that the ACE appliance is not to employ these techniques to objects embedded within HTML pages. Latency optimization corresponds to FlashForward functionality. For more information about FlashForward functionality, see the “[Optimization Overview](#)” section on page 14-2.
 - b. Check the **Bandwidth Optimization (Delta)** check box to specify that the ACE appliance is to dynamically update client browser caches with content differences, or deltas. Uncheck the check box to specify that the ACE appliance is not to dynamically update client browser caches. Bandwidth optimization corresponds to action list Delta optimization. For more information about configuring Delta optimization, see the “[Optimization Overview](#)” section on page 14-2 and the “[Configuring an HTTP Optimization Action List](#)” section on page 14-3.

- c. Continue with [Step 14](#).
- Step 6** (Optional) If you chose Custom, the Actions configuration pane appears with a table listing match criteria and actions.
- Click **Add** to add an entry to this table or choose an existing entry, and click **Edit** to modify it. The configuration pane refreshes with the available configuration options.
- Step 7** In the Apply Building Block field, choose one of these configuration building blocks for the type of optimization that you want to configure, or leave the field blank to configure optimization without a building block:
- **Bandwidth Optimization**—Maximizes bandwidth for Web-based traffic.
 - **Latency Optimization for Embedded Objects**—Reduces the latency associated with embedded objects in Web-based traffic.
 - **Latency Optimization for Embedded Images**—Reduces the latency associated with embedded images in Web-based traffic.
 - **Latency Optimization for Containers**—Reduces the latency associated with Web containers.
- If you chose one of the building blocks, the Rule Match configuration subset displays the configuration options with selections based on the building block chosen. You can accept the entries as they are or modify them.
- If you do not choose a building block, additional configuration options appear depending on the features you enable.
- Step 8** In the Rule Match field, choose an existing class map or click ***New*** to specify new match criteria, and do one of the following:
- If you chose an existing class map, you can view, modify, or duplicate the existing configuration. See the “[Shared Objects and Virtual Servers](#)” section on page 6-9 for more information about modifying shared objects.
 - If you click ***New***, the window refreshes so that you can enter new match criteria.
- Step 9** Configure match criteria using the information in [Table 6-16](#).

Table 6-16 Optimization Match Criteria Configuration

Field	Description/Action
Name	Unique name for this match criteria rule.
Match	Method to be used to evaluate multiple match statements when multiple match conditions exist: <ul style="list-style-type: none"> • match-any—A match exists if at least one of the match conditions is satisfied. • match-all—A match exists only if all match conditions are satisfied.
Conditions	Field that allows you to add a new set of conditions or choose an existing entry. Click Add to add a new set of conditions, or choose an existing entry and click Edit to modify it: <ol style="list-style-type: none"> In the Type field, choose the match condition to be used, then configure any condition-specific options using the information in Table 6-12. Click OK to save your entries, or Cancel to exit this procedure without saving your entries.

- Step 10** In the Actions field, choose an existing action list to use for optimization or click ***New*** to create a new action list, and do one of the following:
- If you chose an existing action list, you can view, modify, or duplicate the existing configuration. See the “[Shared Objects and Virtual Servers](#)” section on page 6-9 for more information about modifying shared objects.
 - If you click ***New***, the window refreshes so you can configure an action list.
- Step 11** Configure the action list using the information in [Table 6-17](#).

Table 6-17 Optimization Action List Configuration Options

Field	Description
Action List Name	Unique name for the action list. Valid entries are unquoted text strings with a maximum of 64 alphanumeric characters.
Enable Delta	<p>Check box that enables delta optimization for the specified URLs. Delta optimization that dynamically updates client browser caches directly with content differences, or deltas, resulting in faster page downloads.</p> <p>Uncheck the check box to disable this feature.</p> <p>If you are configuring optimization without a building block, additional options appear. Configure these options using the information in Table 6-18.</p>
Enable AppScope	<p>Check box that enables AppScope performance monitoring for use with the ACE appliance. AppScope runs on the Management Console of the optional Cisco AVS 3180A Management Station and measures end-to-end application performance.</p> <p>Uncheck the check box to disable this feature.</p> <p>If you are configuring optimization without a building block, additional options appear. Configure these options using the information in Table 6-18.</p>
Flash Forward	<p>Feature that reduces bandwidth usage and accelerates embedded object downloading by combining local object storage with dynamic renaming of embedded objects, which enforces object freshness within the parent HTML page.</p> <p>Choose how the ACE appliance is to implement FlashForward:</p> <ul style="list-style-type: none"> • N/A—This feature is not enabled. • Flash Forward—FlashForward is to be enabled for the specified URLs and embedded objects are to be transformed. • Flash Forward Object—FlashForward static caching is to be enabled for the objects that the corresponding URLs refer to, such as Cascading Style Sheets (CSS), JPEG, and GIF files. <p>If you are configuring without a building block and chose either FlashForward or FlashForward Object, an addition option appears. Configure this option using the information in Table 6-18.</p>
Cache Dynamic	<p>Check box that enables Adaptive Dynamic Caching for the specified URLs even if the expiration settings in the response indicate that the content is dynamic. The expiration of cache objects is controlled by the cache expiration settings based on time or server load.</p> <p>Uncheck the check box to disable this feature.</p>

Table 6-17 Optimization Action List Configuration Options (continued)

Field	Description
Cache Forward	<p>Field that specifies how the ACE appliance is to implement cache forwarding:</p> <ul style="list-style-type: none"> • N/A—This feature is not enabled. • With Wait—Cache forwarding is enabled with the wait option for the specified URLs. If the object has expired but the maximum cache TTL time period has not yet expired, the ACE appliance sends a request to the origin server for the object. Users requesting this page continue to receive content from the cache during this time but must wait for the object to be updated before their request is satisfied. When the fresh object is returned, it is sent to the requesting user and the cache is updated. • Without Wait—Cache forwarding is enabled without the wait option.
Dynamic Entity Tag	<p>Check box that specifies that the ACE appliance is to implement just-in-time object acceleration for embedded objects not able to be cached. This feature enables the acceleration of embedded objects not able to be cached, which results in improved application response time. When enabled, this feature eliminates the need for users to download objects not able to be cached on each request.</p> <p>Uncheck the check box to disable this feature.</p>

Step 12 (Optional) If you are configuring optimization without a building block, additional options appear when you enable specific features.

Configure the additional options using the information in [Table 6-18](#).

Table 6-18 Application Acceleration and Optimization Additional Configuration Options

Field	Description
Response Codes To Ignore (Comma Separated)	Comma-separated list of HTTP response codes for which the response body must not be read. For example, an entry of 302 indicates that the ACE is to ignore the response body of a 302 (redirect) response from the origin server. Valid entries are unquoted text strings with a maximum of 64 alphanumeric characters.
Set Browse Freshness Period	Method that the ACE is to use to determine the freshness of objects in the client's browser: <ul style="list-style-type: none"> • N/A—This option is not configured. • Disable Browser Object Freshness Control—Browser freshness control is not to be used. • Set Freshness Similar To Flash Forward Objects—The ACE is to set freshness similar to that used for FlashForwarded objects, and to use the values specified in the <i>Maximum Time for Cache Time-To-Live</i> and <i>Minimum Time For Cache Time-To-Live</i> fields.
Duration For Browser Freshness (Seconds)	Field that appears if the Set Browse Freshness Period option is not configured. Enter the number of seconds that objects in the client's browser are considered fresh. Valid entries are 0 to 2147483647 seconds.
Enable Delta Options	
Max. For Post Data To Scan For Logging (kBytes)	Maximum number of kilobytes of POST data the ACE is to scan for parameters for the purpose of logging transaction parameters in the statistics log. Valid entries are 0 to 1000 KB.
Base File Anonymous Level	Feature that enables the ACE to create and deliver condensed base files that contain only information that is common to a large set of users. No information unique to a particular user, or across a very small subset of users, is included in anonymous base files. Information that is common to a large set of users is generally not confidential or user-specific. Conversely, information that is unique to a specific user or a small set of users is generally confidential or user-specific. Enter the value for base file anonymity for the all-user condensation method. Valid entries are from 0 to 50; the default value of 0 disables the base file anonymity feature.
Cache-Key Modifier Expression	Unique identifier that is used to identify a cached object to be served to a client, replacing a trip to the origin server. The cache key modifier feature allows you to modify the canonical form of a URL; that is, the portion before "?" in a URL. For example, the canonical URL of <code>http://www.xyz.com/somepage.asp?action=browse&level=2</code> is <code>http://www.xyz.com/somepage.asp</code> . Enter a regular expression containing embedded variables as described in Table 6-19 . The ACE transforms URLs specified in class maps for this virtual server with the expression and variable entered here. Valid entries are unquoted text strings with no spaces and a maximum of 255 alphanumeric characters. If the string includes spaces, enclose the string with quotation marks ("").

Table 6-18 Application Acceleration and Optimization Additional Configuration Options (continued)

Field	Description
Min. Time For Cache Time-To-Live (Seconds)	Minimum number of seconds that an object without an explicit expiration time should be considered fresh in the ACE cache. This value specifies the minimum time that content can be cached. If the ACE is configured for FlashForward optimization, this value should normally be 0. If the ACE is configured for dynamic caching, this value should indicate how long the ACE should cache the page. (See Table 6-17 for information about these configuration options.) Valid entries are 0 to 2147483647 seconds.
Max. Time For Cache Time-To-Live (Seconds)	Maximum number of seconds that an object without an explicit expiration time should be considered fresh in the ACE cache. Valid entries are 0 to 2147483647 seconds.
Cache Time-To-Live Duration (%)	Percent of an object's age at which an embedded object without an explicit expiration time is considered fresh. Valid entries are 0 to 100 percent.
Expression To Modify Cache Key Query Parameter	Feature that allows you to modify the query parameter of a URL; that is, the portion after “?” in a URL. For example, the query parameter portion of <code>http://www.xyz.com/somepage.asp?action=browse&level=2</code> is <code>action=browse&level=2</code> . Enter a regular expression containing embedded variables as described in Table 6-19 . The ACE transforms URLs specified in class maps for this virtual server with the expression and variable entered here. If no string is specified, the query parameter portion of the URL is used as the default value for this portion of the cache key. Valid entries are unquoted text strings with no spaces and a maximum of 255 alphanumeric characters.
Canonical URL Expressions	Canonical URL feature to eliminate the “?” and any characters that follow to identify the general part of the URL. This general URL is then used to create the base file. In this way, the ACE maps multiple URLs to a single canonical URL. Enter a comma-separated list of parameter expander functions as defined in Table 6-19 to identify the URLs to associate with this parameter map. Valid entries are unquoted text strings with a maximum of 255 alphanumeric characters.
Enable Cacheable Content Optimization	Check box that enables delta optimization of content that can be cached. This feature allows the ACE to detect content that can be cached and perform delta optimization on it. Uncheck the check box to disable this feature.
Enable Delta Optimization On First Visit To Web Page	Check box that enables condensation on the first visit to a Web page. Uncheck the check box to disable this feature.
Min. Page Size For Delta Optimization (Bytes)	Minimum page size, in bytes, that can be condensed. Valid entries are from 1 to 250000 bytes.
Max. Page Size For Delta Optimization (Bytes)	Maximum page size, in bytes, that can be condensed. Valid entries are from 1 to 250000 bytes.

Table 6-18 Application Acceleration and Optimization Additional Configuration Options (continued)

Field	Description
Set Default Client Script	Scripting language that the ACE is to recognize on condensed content pages: <ul style="list-style-type: none"> • N/A—Indicates that this option is not configured. • Javascript—Indicates that the default scripting language is JavaScript. • Visual Basic Script—Indicates that the default scripting language is Visual Basic.
Exclude Iframes From Delta Optimization	Check box to specify that delta optimization is not to be applied to IFrames (inline frames). Uncheck the check box to indicate that delta optimization is to be applied to IFrames.
Exclude Non-ASCII Data From Delta Optimization	Check box to specify that delta optimization is not to be applied to non-ASCII data. Uncheck the check box to indicate that delta optimization is to be applied to non-ASCII data.
Exclude JavaScripts From Delta Optimization	Check box to specify that delta optimization is not to be applied to JavaScript. Uncheck the check box to indicate that delta optimization is to be applied to JavaScript.
MIME Types To Exclude From Delta Optimization	<p>a. In the first field, enter a comma-separated list of the MIME (Multipurpose Internet Mail Extension) type messages that are not to have delta optimization applied, such as image/Jpeg, text/html, application/msword, or audio/mpeg. See the “Supported MIME Types” section on page 9-26 for a list of supported MIME types.</p> <p>b. Click Add to add the entry to the list box on the right. You can position the entries in the list box by using the Up and Down buttons.</p>
Remove HTML META Elements From Documents	Check box to specify that HTML META elements are to be removed from documents to prevent them from being condensed. Uncheck the check box to indicate that HTML META elements are not to be removed from documents.
Rebase Delta Optimization Threshold (%)	Delta threshold, expressed as a percent, when rebasing is to be triggered. This entry represents the size of a page delta relative to total page size, expressed as a percent. This entry triggers rebasing when the delta response size exceeds the threshold as a percentage of base file size. Valid entries are 0 to 10000 percent.
Rebase Flash Forward Threshold (%)	Threshold, expressed as a percent, when rebasing is to be triggered based on the percent of FlashForwarded URLs in the response. This entry triggers rebasing when the difference between the percentages of FlashForwarded URLs in the delta response and the base file exceeds the threshold. Valid entries are 0 to 10000 percent.
Rebase History Size (Pages)	Number of pages to be stored before the ACE resets all rebase control parameters to zero and starts over. This option prevents the base file from becoming too rigid. Valid entries are 10 to 2147483647.
Rebase Modify Cool-Off Period (Seconds)	Number of seconds after the last modification before performing a rebase. Valid entries are 1 to 14400 seconds (4 hours).
Rebase Reset Period (Seconds)	Period of time, in seconds, for performing a meta data refresh. Valid entries are 1 to 900 seconds (15 minutes).

Table 6-18 Application Acceleration and Optimization Additional Configuration Options (continued)

Field	Description
UTF-8 Character Set Threshold	<p>Number of 8-bit Unicode Transformation Format (UTF-8) characters that need to appear on a page to create a UTF-8 character set page. The UTF-8 character set is an international standard that allows Web pages to display non-ASCII or non-English multibyte characters. It can represent any universal character in the Unicode standard and is backwards compatible with ASCII.</p> <p>Valid entries are from 1 to 1,000,000.</p>
Server Load Threshold Trigger (%)	<p>Threshold, expressed as a percent, at which the TTL for cached objects is to be changed. The server load threshold trigger indicates that the time-to-live (TTL) period for cached objects is to be based dynamically on server load. With this method, TTL periods increase if the current response time from the origin sever is greater than the average response time and decrease if the current response time from the origin server is less than the average response time when the difference in response times exceeds a specified threshold amount.</p> <p>Valid entries are from 0 to 100 percent.</p>
Server Load Time-To-Live Change (%)	<p>Percentage by which the cache TTL is to be increased or decreased when the server load threshold trigger is met. This option specifies the percentage by which the cache TTL is increased or decreased in response to a change in server load. For example, if this value is set to 20 and the current TTL for a response is 300 seconds, and if the current server response times exceeds the trigger threshold, the cache TTL for the response is raised to 360 seconds.</p> <p>Valid entries are from 0 to 100 percent.</p>
Delta Optimization Mode	<p>Method by which delta optimization is to be implemented:</p> <ul style="list-style-type: none"> • N/A—Indicates that a delta optimization mode is not configured. • Enable The All-User Mode For Delta Optimization—Indicates that the ACE is to generate the delta against a single base file that is shared by all users of the URL. This option is usable in most cases if the structure of a page is common across all users, and the disk space overhead is minimal. • Enable The Per-User Mode For Delta Optimization—Indicates that the ACE is to generate the delta against a base file that is created specifically for that user. This option is useful when page contents, including layout elements, are different for each user, and delivers the highest level of condensation. However, this increases disk space requirements because a copy of the base page that is delivered to each user is cached. This option is useful when privacy is required because base pages are not shared among users.
Enable Appscope Options	
Appscope Optimize Rate (%)	<p>Percentage of all requests or sessions to be sampled for performance with acceleration (or optimization) applied. All applicable optimizations for the class will be performed. Valid entries are from 0 to 100 percent, with a default of 10 percent. The sum of this value and the value entered in the Passthrough Rate Percent field must not exceed 100.</p>
Appscope Passthrough Rate (%)	<p>Percentage of all requests or sessions to be sampled for performance without optimization. No optimizations for the class will be performed. Valid entries are from 0 to 100, with a default of 10 percent. The sum of this value and the value entered in the Optimize Rate Percent field must not exceed 100.</p>

Table 6-18 Application Acceleration and Optimization Additional Configuration Options (continued)

Field	Description
Max Number For Parameter Summary Log (Bytes)	Maximum number of bytes that are to be logged for each parameter value in the parameter summary of a transaction log entry in the statistics log. If a parameter value exceeds this limit, it is truncated at the specified limit. Valid entries are 0 to 10,000 bytes.
Specify String For Grouping Requests	<p>String that the ACE is to use to sort requests for AppScope reporting. The string can contain a URL regular expression that defines a set of URLs in which URLs that differ only by their query parameters are to be treated as separate URLs in AppScope reports.</p> <p>For example, to define a string that is used to identify the URLs <code>http://server/catalog.asp?region=asia</code> and <code>http://server/catalog.asp?region=america</code> as two separate reporting categories, you would enter <code>http_query_param(region)</code>.</p> <p>Valid entries contain 1 to 255 characters and can contain the parameter expander functions listed in Table 6-19.</p>

[Table 6-19](#) lists the parameter expander functions that you can use.

Table 6-19 Parameter Expander Functions

Variable	Description
<code>\$(number)</code>	<p>Expands to the corresponding matching subexpression (by <i>number</i>) in the URL pattern. Subexpressions are marked in a URL pattern using parentheses (). The numbering of the subexpressions begins with 1 and is the number of the left-parenthesis “(“ counting from the left. You can specify any positive integer for the number. <code>\$(0)</code> matches the entire URL. For example, if the URL pattern is <code>((http://server/.*)(.*)/a.jsp)</code>, and the URL that matches it is <code>http://server/main/sub/a.jsp?category=shoes&session=99999</code>, then the following are correct:</p> <p><code>\$(0)</code> = <code>http://server/main/sub/a.jsp</code> <code>\$(1)</code> = <code>http://server/main/sub/</code> <code>\$(2)</code> = <code>http://server/main</code> <code>\$(3)</code> = <code>sub</code></p> <p>If the specified subexpression does not exist in the URL pattern, then the variable expands to the empty string.</p>
<code>\$http_query_string()</code>	<p>Expands to the value of the whole query string in the URL. For example, if the URL is <code>http://myhost/dohis?param1=value1&param2=value2</code>, then the following is correct:</p> <p><code>\$http_query_string()</code> = <code>param1=value1&param2=value2</code></p> <p>This function applies to both GET and POST requests.</p>

Table 6-19 Parameter Expander Functions (continued)

Variable	Description
<p>\$http_query_param(<i>query-param-name</i>)</p> <p>The obsolete syntax is also supported: \$param(<i>query-param-name</i>)</p>	<p>Expands to the value of the named query parameter (case-sensitive). For example, if the URL is <code>http://server/main/sub/a.jsp?category=shoes&session=99999</code>, then the following are correct:</p> <pre>\$http_query_param(category) = shoes \$http_query_param(session) = 99999</pre> <p>If the specified parameter does not exist in the query, then the variable expands to the empty string. This function applies to both GET and POST requests.</p>
\$http_cookie(<i>cookie-name</i>)	Evaluates to the value of the named cookie. For example, \$http_cookie(cookiexyz). The cookie name is case-sensitive.
\$http_header(<i>request-header-name</i>)	Evaluates to the value of the specified HTTP request header. In the case of multivalued headers, it is the single representation as specified in the HTTP specification. For example, \$http_header(user-agent). The HTTP header name is not case-sensitive.
\$http_method()	Evaluates to the HTTP method used for the request, such as GET or POST.
<p>Boolean Functions:</p> <p>\$http_query_param_present(<i>query-param-name</i>)</p> <p>\$http_query_param_notpresent(<i>query-param-name</i>)</p> <p>\$http_cookie_present(<i>cookie-name</i>)</p> <p>\$http_cookie_notpresent(<i>cookie-name</i>)</p> <p>\$http_header_present(<i>request-header-name</i>)</p> <p>\$http_header_notpresent(<i>request-header-name</i>)</p> <p>\$http_method_present(<i>method-name</i>)</p> <p>\$http_method_notpresent(<i>method-name</i>)</p>	<p>Evaluates to a Boolean value: True or False, depending on the presence or absence of the element in the request. The elements are a specific query parameter (<i>query-param-name</i>), a specific cookie (<i>cookie-name</i>), a specific request header (<i>request-header-name</i>), or a specific HTTP method (<i>method-name</i>). All identifiers are case-sensitive except for the HTTP request header name.</p>
\$regex_match(<i>param1</i> , <i>param2</i>)	<p>Evaluates to a Boolean value: True if the two parameters match and False if they do not match. The two parameters can be any two expressions, including regular expressions, that evaluate to two strings. For example, this function:</p> <pre>\$regex_match(\$http_query_param(URL), .*Store\.asp.*)</pre> <p>compares the query URL with the regular expression string <code>.*Store\.asp.*</code>.</p> <p>If the URL matches this regular expression, this function evaluates to True.</p>

- Step 13** When you finish configuring match criteria and actions, do one of the following:
- Click **OK** to save your entries and to return to the Rule Match and Actions table.
 - Click **Cancel** to exit this procedure without saving your entries and to return to the Rule Match and Actions table.
- Step 14** When you finish configuring virtual server properties, do one of the following:
- Click **Deploy Now** to deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files. The ACE appliance validates the action list configuration and deploys it.
 - Click **Cancel** to exit this procedure without saving your entries and to return to the Virtual Servers table.
 - Click **Deploy Later** to save your entries and apply the configuration at a later time.

Related Topics

- [Optimization Traffic Policies and Typical Configuration Flow, page 14-2](#)
- [Configuring Traffic Policies for HTTP Optimization, page 14-7](#)
- [Configuring Virtual Server Protocol Inspection, page 6-18](#)
- [Configuring Virtual Server Layer 7 Load Balancing, page 6-30](#)
- [Configuring Virtual Server Default Layer 7 Load Balancing, page 6-51](#)

Configuring Virtual Server NAT

You can configure Name Address Translation (NAT) for virtual servers.

Assumptions

This topic assumes the following:

- Make sure that a virtual server has been configured in the Properties configuration subset. For more information, see the [“Configuring Virtual Server Properties”](#) section on page 6-11
- Make sure that a VLAN has been configured. See the [“Configuring VLAN Interfaces”](#) section on page 11-5 for information on configuring a VLAN interface.
- Make sure that at least one NAT pool has been configured on a VLAN interface. See the [“Configuring VLAN Interface NAT Pools”](#) section on page 11-16 for information on configuring a NAT pool.

Procedure

- Step 1** Choose **Config > Devices > context > Load Balancing > Virtual Servers**.
The Virtual Servers table appears.
- Step 2** In the Virtual Servers table, choose the virtual server you want to configure for NAT, and click **Edit**.
The Virtual Server configuration window appears.
- Step 3** In the Virtual Server configuration window, click **NAT**.
The NAT table appears.

- Step 4** In the NAT table, click **Add** to add an entry, or choose an existing entry and click **Edit** to modify it.
- Step 5** In the VLAN drop-down list, choose the VLAN that you want to use for NAT.
VLANs that have previously been defined for NAT do not appear in this list. VLAN numbers provide an indication of available NAT pools.
- Step 6** In the NAT Pool ID drop-down list, choose the NAT pool that you want to associate with the selected VLAN.
Note the following about the NAT pool ID selections:
NAT Pool IDs (Begin IP - End IP: Netmask: PAT) appear in a format that provides the details of the beginning and ending IP address range, netmask, and the PAT enabled or disabled setting. For example:
2 (10.77.241.2 - 10.77.241.15: 255.255.255.192: PAT Enabled).
If the NAT pool had previously been associated but is no longer defined, then it appears as “<NAT_POOL_ID> (Warning: Undefined NAT Pool)”. For example:
2 (Warning: Undefined NAT Pool)
For more information about NAT pools, see the [“Configuring VLAN Interface NAT Pools”](#) section on page 11-16.
- Step 7** Do one of the following:
- Click **OK** to save your entries and to return to the NAT table. The NAT table refreshes with the new entry.
 - Click **Cancel** to exit the procedure without saving your entries and to return to the NAT table.
- Step 8** When you finish configuring virtual server properties, do one of the following:
- Click **Deploy Now** to deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files.
 - Click **Cancel** to exit this procedure without saving your entries and to return to the Virtual Servers table.
 - Click **Deploy Later** to save your entries and apply the configuration at a later time.
-

Related Topics

- [Configuring Virtual Servers, page 6-2](#)
- [Configuring Virtual Server Properties, page 6-11](#)
- [Configuring Virtual Server SSL Termination, page 6-17](#)
- [Configuring Virtual Server Protocol Inspection, page 6-18](#)
- [Configuring Virtual Server Layer 7 Load Balancing, page 6-30](#)
- [Configuring Virtual Server Default Layer 7 Load Balancing, page 6-51](#)

Displaying Virtual Servers by Context

You can display all virtual servers associated with a virtual context.

Procedure

Step 1 Choose **Config > Devices**.

The device tree appears.

Step 2 In the device tree, choose the context associated with the virtual servers that you want to display, and choose **Load Balancing > Virtual Servers**.

The Virtual Servers table appears with the following information:

- Virtual server name
 - Configured state, such as Inservice or Out of service
 - Virtual IP address
 - Port
 - Associated VLANs
 - Associated server farms
 - The owner, and context in which the virtual server was created
-

Related Topics

- [Configuring Virtual Servers, page 6-2](#)
- [Managing Virtual Servers, page 6-67](#)
- [Displaying Detailed Virtual Server Information, page 6-71](#)
- [Displaying Virtual Servers, page 6-72](#)

Displaying Virtual Server Statistics and Status Information

You can display virtual server statistics and status information for a particular virtual server by using the **Details** button. ANM accesses the **show service-policy *policy_name* detail** CLI command to display detailed virtual server information.

Procedure

Step 1 Choose **Config > Devices > *context* > Load Balancing > Virtual Servers**.

The Virtual Servers table appears.

Step 2 In the Virtual Servers table, choose a virtual server from the Virtual Servers table, and click **Details**.

The **show service-policy *policy_name* detail** CLI command output appears. For details on the displayed output fields, see either the *Cisco ACE Module Server Load-Balancing Configuration Guide* or the *Cisco ACE 4700 Series Appliance Server Load-Balancing Configuration Guide*.

Step 3 Click **Update Details** to refresh the output for the **show service-policy *policy_name* detail** CLI command.

Step 4 Click **Close** to return to the Virtual Servers table.

Related Topics

- [Configuring Virtual Servers, page 6-2](#)
- [Managing Virtual Servers, page 6-67](#)
- [Displaying Detailed Virtual Server Information, page 6-71](#)
- [Displaying Virtual Servers, page 6-72](#)

Managing Virtual Servers

This section shows how to display and manage the virtual servers from the Virtual Servers window (Config > Operations > Virtual Servers). This window provides you with information about each virtual server configured on ANM (see the “[Displaying Virtual Servers](#)” section on [page 6-72](#)) and provides access to function buttons that allow you to perform tasks such as activate or suspend a virtual server, display a virtual server topology map, or display connection statistics graphs.

This section includes the following topics:

- [Activating Virtual Servers, page 6-67](#)
- [Suspending Virtual Servers, page 6-68](#)
- [Managing GSS VIP Answers, page 6-69](#)
- [Activating and Suspending DNS Rules Governing GSS Load Balancing, page 6-70](#)
- [Displaying Detailed Virtual Server Information, page 6-71](#)
- [Displaying Virtual Servers, page 6-72](#)
- [Using the Virtual Server Connection Statistics Graph, page 6-74](#)
- [Using the Virtual Server Topology Map, page 6-74](#)
- [Understanding CLI Commands Sent from Virtual Server Table, page 6-75](#)

Activating Virtual Servers

You can activate a virtual server.



Note

A missing operation or Admin state on a CSM or CSS device most likely means that the community string was not enabled on those devices. If the community string is not enabled on a CSM or CSS device, and any kind of operation is performed on those devices, it will not succeed, and ANM will not provide any kind of indication.

- For CSM devices, you must enable the community string of the Catalyst 6500 series chassis.
 - For CSS devices, you must enable the community string of the CSS device itself.
-

Procedure

-
- Step 1** Choose **Config > Operations > Virtual Servers**.
The Virtual Servers table appears.
- Step 2** In the Virtual Servers table, choose the virtual server that you want to activate, and click **Activate**.
The server is activated and the window refreshes with updated information in the Configured State column.
-

Related Topics

- [Managing Virtual Servers, page 6-67](#)
- [Displaying Virtual Servers, page 6-72](#)
- [Suspending Virtual Servers, page 6-68](#)

Suspending Virtual Servers

You can suspend a virtual server.

**Note**

A missing operation or Admin state on a CSM or CSS device most likely means that the community string was not enabled on those devices. If the community string is not enabled on a CSM or CSS device, and any kind of operation is performed on those devices, it will not succeed, and ANM will not provide any kind of indication.

- For CSM devices, you must enable the community string of the Catalyst 6500 series chassis.
 - For CSS devices, you must enable the community string of the CSS device itself.
-

Procedure

-
- Step 1** Choose **Config > Operations > Virtual Servers**.
The Virtual Servers table appears.
- Step 2** In the Virtual Servers table, choose the virtual server that you want to suspend, and click **Suspend**.
The Suspend Virtual Server window appears.
- Step 3** In the Reason field of the Suspend Virtual Server window, enter the reason for this action.
You might enter a trouble ticket, an order ticket, or a user message.



Note Do not enter a password in this field.

Related Topics

- [Managing Virtual Servers, page 6-67](#)

- [Displaying Virtual Servers, page 6-72](#)
- [Activating Virtual Servers, page 6-67](#)

Managing GSS VIP Answers

This section describes how to manage GSS VIP answers. In a GSS network, the term *answers* refers to resources that respond to content queries. When you create an answer using the primary Global Site Selector Manager (PGSSM), you are simply identifying a resource on your GSS network to which queries can be directed and that can provide your user's D-proxy with the address of a valid host to serve their request.

Virtual IP (VIP) addresses associated with an SLB such as the Cisco CSS, Cisco CSM, Cisco IOS-compliant SLB, LocalDirector, or a Web server are types of answers that are specified in the ANM user interface in the GSS VIP Answers table found in ANM under Configuration > Operations. Use this procedure to poll, activate, or suspend GSS VIP answers.

Assumption

Make sure that you have established GSS VIP answers using the PGSSM.

Procedure

Step 1 Choose **Config > Operations > GSS VIP Answers**.

The GSS Answers table appears. For a list of fields available, see [Table 6-20](#).

Table 6-20 GSS Answer Table

Field	Description
Multiple Row Selection Checkbox	Check box that selects all entries at the same time, or you can check line items individually.
IP Address	VIP answer IP address.
Name	VIP answer name.
Config State	VIP answer configured status.
PGSSM Oper State	Operational status as shown on the primary GSS manager (PGSSM).
Answer Group	Answer group names to which the VIP answer belong.
Location	Logical groupings for GSS resources that correspond to geographical entities such as a city, data center, or content site.
Device	Primary GSS device name on ANM.
PGSSM Time	Last operational status update time on the primary GSS.

Step 2 In the GSS Answers table, check the check boxes to the left of the servers that you want to poll, activate, or suspend.

Step 3 Do one of the following:

- Click **Active/Suspended** hyperlink to view the VIP answer details across the GSS node(s). A popup window appears listing all nodes associated with the VIP, operational state, hit count, and timestamp for each node.
- Click **Poll Now** to query the chosen resource to verify it is still active.

**Note**

If you click **Poll Now** immediately after you click **Activate** or **Suspend**, you might not get the VIP answer operational status on the PGSSM that reflects your most recent configuration. It might be necessary to click Poll Now two or three times in succession to get an accurate result.

The ability of Cisco License Manager to update the VIP answer operational status and statistics accurately in detailed GSS statistics window might depend on the polling interval that has been configured on the GSS. The polling interval can be configured directly on the GSS device. (The default is 5 minutes.) Depending on the interval, it can take 5 minutes or more for the ANM server to show an accurate result.

- Click **Activate** to reactivate a GSS answer.
- Click **Suspend** to temporarily stop the GSS from using an associated answer.

If you clicked **Activate** or **Suspend**, a dialog box prompts for a Reason. Acceptable text consists of any characters or nothing at all.

Step 4 Do one of the following:

- Click **Deploy Now** to complete Activation or Suspension.
- Click **Cancel** to cancel the Activation or Suspension operation.

Related Topics

- [Information About Load Balancing, page 6-1](#)
- [Activating and Suspending DNS Rules Governing GSS Load Balancing, page 6-70](#)

Activating and Suspending DNS Rules Governing GSS Load Balancing

You can activate or suspend DNS rules associated with your GSS VIP answers table. The DNS rules table in Configuration > Operations navigation tree specifies actions for the GSS to take when it receives a request from a known source (a member of a source address list) for a known hosted domain (a member of a domain list).

The DNS rule specifies which response (answer) is given to the requesting user's local DNS host (D-proxy) and how that answer is chosen. One of a variety of balance methods is used to determine the best response to the request, based on the status and load of the GSS host devices.

Assumption

Make sure that you have established GSS VIP answers and DNS rules using the PGSSM.

Procedure

Step 1 Choose **Config > Operations > DNS Rules**.

The DNS Rules table appears. For a list of fields available, see [Table 6-21](#).

Table 6-21 **DNS Rules Table**

Field	Description
Multiple Row Selection Checkbox	Check box that selects all entries at the same time, or you can check line items individually.
Name	Name of the DNS rule.
Source Address	Collection of IP addresses or address blocks for known client DNS proxies (or D-proxies).
Domains	Domain list name containing one or more domain names that point to content for which the GSS is acting as the authoritative DNS server and for which you wish to use the GSS technology to balance traffic and user requests.
Config State	DNS rules configured status, either Active or Suspended.
Answer Group	Lists of GSS resources that are candidates to respond to DNS queries received from a user for a hosted domain.
Owner	Owner names, providing a simple way to organize and identify groups of related GSS resources.
Device	Primary GSS device name on ANM.
PGSSM Time	Last operational status update time on the GSS.

Step 2 In the DNS Rules table, check the checkbox to the left of the servers that you want to activate or suspend.

Step 3 Click the **Activate** or **Suspend** button.

A dialog box prompts for a Reason. Acceptable text consists of any characters or none at all.

Step 4 Do one of the following:

- Click **Deploy Now** to complete Activation or Suspension.
- Click **Cancel** to cancel the Activation or Suspension operation.

Related Topics

- [Information About Load Balancing, page 6-1](#)
- [Managing GSS VIP Answers, page 6-69](#)

Displaying Detailed Virtual Server Information

You can display detailed information about the state of a virtual server.

Procedure

Step 1 Choose **Config > Operations > Virtual Servers**.

The Virtual Servers table appears.

Step 2 In the Virtual Servers table, choose the virtual server whose configuration details that you want to display.

Click the hyperlinked entry for that virtual server that appears in the Operational State column.

The Details window appears with the following information:

- Current operational status
- Description, if one was entered
- Configured interfaces, such as VLANs
- Configured service policies including:
 - Configured class maps, detailed by type (such as load balancing or inspection)
 - States of configured options, indicated by word (ACTIVE, DISABLED, OUTFSERVICE) and color (green, orange/yellow, and red)
 - Associated policy maps with details on their type and action (L7 loadbalance, serverfarm)
 - Statistics regarding connections and counts

Related Topics

- [Configuring Virtual Servers, page 6-2](#)
- [Displaying Virtual Servers by Context, page 6-66](#)
- [Displaying Virtual Server Statistics and Status Information, page 6-66](#)
- [Managing Virtual Servers, page 6-67](#)

Displaying Virtual Servers

You can display all virtual servers by choosing **Config > Operations > Virtual Servers**. The Virtual Servers table appears. [Table 6-22](#) describes the Virtual Servers table information.

Table 6-22 Virtual Server Table Fields

Item	Description
Name	Server farm name sorted by virtual context.
Policy Map	Associated policy map.
IP Address:Protocol:Port	Server farm IP address, protocol, and port used for communications.
HA	Indicator that specifies that the virtual server is part of a high availability pair. The table displays HA pair virtual servers together in the same row so that they remain together no matter how you sort the information.
SLB Device	Associated ACE IP address and context.
Admin	Administrative state of the virtual server: Up or Down.
	 <p>Note For a CSM device, the virtual server Admin State is derived from the Operational State. In this case, the Operational State may display an Out of Service condition when the virtual server is configured to be Inservice (if all of the real servers are out of service).</p>

Table 6-22 Virtual Server Table Fields (continued)

Item	Description
Oper	Operational state of the virtual server: Up or Down.  Note For an ACE appliance, this column is populated for that device. When you click on the value in this column (irrespective of ACE version), detailed information about the virtual server displays in a pop-up window.
DWS	Operating state of Dynamic Workload Scaling for the virtual server, which can be: <ul style="list-style-type: none"> • N/A—Not applicable; the server farms associated with the virtual server are not configured to use Dynamic Workload Scaling. • Local—At least one server farm associated the virtual server is configured to use Dynamic Workload Scaling, but the ACE is sending traffic to the VM Controller’s local VMs only. • Expanded—At least one server farm associated the virtual server is configured to use Dynamic Workload Scaling and the ACE is sending traffic to the VM Controller’s local and remote VMs.
Conn	Number of active connections.  Note This column is populated for ACE appliances. For ACE devices, the Active Connections column displays N/A for older versions of the ACE appliance and module.
Stat Age	Age of the statistical information.
Serverfarms	Associated server farms.  Note If you have the Details pop-up window feature enabled, click the value in this column to open the Details pop-up window and display detailed information about the server farm. By default, this feature is disabled. For information about enabling or disabling this feature, see the “Enabling the ACE Server Farm Details Pop-up Window Option for Virtual Servers” section on page 17-89.
VLANs	Associated VLANs.

Use the display toggle button () located above the table to control which virtual servers ANM displays as follows:

- Show ANM recognized Virtual Servers—Displays only virtual servers that match ANM’s virtual server definition.
- Show all Virtual Servers—Displays virtual servers that match ANM’s virtual server definition and those that do not match ANM’s virtual server definition but that ANM can recognize as virtual servers using SNMP polling.

You can activate or suspend virtual servers from this table and obtain additional information about the state of the virtual server.

Related Topics

- [Activating Virtual Servers, page 6-67](#)
- [Suspending Virtual Servers, page 6-68](#)
- [Displaying Detailed Virtual Server Information, page 6-71](#)
- [Displaying Virtual Server Statistics and Status Information, page 6-66](#)
- [Displaying Virtual Servers by Context, page 6-66](#)

Using the Virtual Server Connection Statistics Graph

You can display real time and historical statistical information about the connections of a virtual server. ANM displays the information in graph or chart form. This feature also allows you to compare similar connection information across multiple virtual servers.

Procedure

-
- Step 1** Choose **Config > Operations > Virtual Servers**.
- The Virtual Servers table appears.
- Step 2** In the Virtual Servers table, check the check box next to server whose connection information you want to display, and click **Graph**.
- You can choose up to four virtual servers if you want to compare statistical data.
- The Virtual Server Graph window appears, displaying the default graph for each selected virtual server. For details about using the graph feature, see “[Configuring Historical Trend and Real Time Graphs for Devices](#)” section on page 16-46.
- Step 3** Click **Exit** to return to the Virtual Server widow.
-

Related Topics

- [Configuring Historical Trend and Real Time Graphs for Devices, page 16-46](#)
- [Activating Virtual Servers, page 6-67](#)
- [Suspending Virtual Servers, page 6-68](#)
- [Displaying Detailed Virtual Server Information, page 6-71](#)
- [Displaying Virtual Servers, page 6-72](#)
- [Using the Virtual Server Topology Map, page 6-74](#)
- [Displaying Virtual Server Statistics and Status Information, page 6-66](#)
- [Displaying Virtual Servers by Context, page 6-66](#)

Using the Virtual Server Topology Map

You can display the nodes on your network based on the virtual server that you select.

Procedure

Step 1 Choose **Config > Operations > Virtual Servers**.

The Virtual Servers table appears.

Step 2 Use the display toggle button () to ensure that the Virtual Servers table is set to Show ANM Recognized Virtual Servers.



Note The topology map feature is not available when the Virtual Server table is set to Show All Virtual Servers (for more information, see [“Displaying Virtual Servers” section on page 6-72](#)).

Step 3 In the Virtual Servers table, choose the server whose topology map you want to display, and click **Topology**.

The ANM Topology map appears. The map includes several tools for navigating the network map and zooming in and out. For details about using the map tools, see the [“Displaying Network Topology Maps” section on page 16-63](#).

Step 4 Click **Exit** to return to the Virtual Server widow.

Related Topics

- [Suspending Virtual Servers, page 6-68](#)
- [Displaying Detailed Virtual Server Information, page 6-71](#)
- [Displaying Virtual Servers, page 6-72](#)
- [Using the Virtual Server Connection Statistics Graph, page 6-74](#)
- [Displaying Virtual Server Statistics and Status Information, page 6-66](#)
- [Displaying Virtual Servers by Context, page 6-66](#)

Understanding CLI Commands Sent from Virtual Server Table

[Table 6-23](#) displays the CLI commands dispatched to the device for a given Virtual Servers table option, and is sorted by device.

Table 6-23 *CLI Commands Deployed from Virtual Servers Table*

Command	Sample CLI Sent
ACE Modules and Appliances	
Virtual Server Activate	<pre>policy-map multi-match int25 class VIP3 loadbalance vip inservice</pre>
Virtual Server Suspend	<pre>policy-map multi-match int25 class VIP3 no loadbalance vip inservice</pre>
CSMs	
Virtual Server Activate	<pre>vserver APP1 inservice</pre>

Table 6-23 CLI Commands Deployed from Virtual Servers Table (continued)

Command	Sample CLI Sent
Virtual Server Suspend	<pre>vserver APPl no inservice</pre>
CSS Devices	
Virtual Server Activate	<pre>owner hm content LB active</pre>
Virtual Server Suspend	<pre>owner hm content LB suspend</pre>

Deploying Virtual Servers

You can deploy virtual servers on your network at times that are convenient and appropriate for your environment. For example, if your site prefers to make changes to the network during a specific time each night, you can modify and save virtual server configurations during the day and then deploy them when appropriate.

This section includes the following topics:

- [Deploying a Virtual Server, page 6-76](#)
- [Displaying All Staged Virtual Servers, page 6-77](#)
- [Modifying Deployed Virtual Servers, page 6-77](#)
- [Modifying Staged Virtual Servers, page 6-78](#)

Deploying a Virtual Server

You can deploy virtual servers on your network at times that are convenient and appropriate for your environment. For example, if your site prefers to make changes to the network during a specific time each night, you can modify and save virtual server configurations during the day and then deploy them when appropriate.

Procedure

-
- Step 1** Choose **Config > Deploy**.
The Staged Objects table appears.
- Step 2** From the Staged Objects table, choose the virtual server that you want to deploy on your network, and click **Deploy**.
The virtual server is deployed and the table refreshes with updated information.
-

Related Topics

- [Configuring Virtual Servers, page 6-2](#)

- [Displaying All Staged Virtual Servers, page 6-77](#)
- [Modifying Staged Virtual Servers, page 6-78](#)

Displaying All Staged Virtual Servers

You can display all objects that have been configured but have not yet been deployed on your network.

Procedure

- Step 1** Do one of the following:
- Choose **Config > Deploy**.
The Staged Objects table appears listing the following:
 - Virtual server name
 - Device ID and virtual context
 - Time the virtual server was created
 - User who last modified the object
 - Time the object was last updated
 - Choose **Config > Devices > context > Load Balancing > Virtual Servers**.
The Virtual Servers table appears. Virtual servers with configurations that have not been deployed appear with the status Not Deployed in the Configured State column.
-

Related Topics

- [Configuring Virtual Servers, page 6-2](#)
- [Deploying a Virtual Server, page 6-76](#)
- [Modifying Staged Virtual Servers, page 6-78](#)
- [Modifying Deployed Virtual Servers, page 6-77](#)

Modifying Deployed Virtual Servers

You can modify the configuration of a deployed virtual server.

Procedure

- Step 1** Choose **Config > Devices > context > Load Balancing > Virtual Servers**.
The Virtual Servers table appears.
- Step 2** In the Virtual Servers table, choose the virtual server you want to modify, then click **Edit**.
The Virtual Server configuration window appears.
- Step 3** In the Virtual Server configuration window, modify the virtual server's configuration as desired.
See [Table 6-1](#) for virtual server configuration options.

- Step 4** When you are done modifying the configuration, do one of the following:
- Click **Deploy Now** to immediately deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files.
 - Click **Cancel** to exit this procedure without saving your entries and to return to the Virtual Servers table.
-

Related Topics

- [Managing Virtual Servers, page 6-67](#)
- [Displaying All Staged Virtual Servers, page 6-77](#)
- [Activating Virtual Servers, page 6-67](#)
- [Suspending Virtual Servers, page 6-68](#)

Modifying Staged Virtual Servers

You can modify the configuration of a staged virtual server.

Procedure

-
- Step 1** Choose **Config > Deploy**.
- The Staged Objects table appears, listing those virtual servers that have not yet been deployed in the network.
- Step 2** From the Staged Objects table, choose the virtual server you want to modify, and click **Edit**.
- The Virtual server configuration window appears.
- Step 3** In the Virtual server configuration window, modify the virtual server configuration as desired.
- See [Table 6-1](#) for virtual server configuration options.
- Step 4** When you are done modifying the configuration, do one of the following:
- Click **Deploy Now** to immediately deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files.
 - Click **Cancel** to exit this procedure without saving your entries and to return to the Virtual Servers table.
 - Click **Deploy Later** to save your entries and apply this configuration at a later time.
-

Related Topics

- [Deploying a Virtual Server, page 6-76](#)
- [Displaying All Staged Virtual Servers, page 6-77](#)
- [Activating Virtual Servers, page 6-67](#)