

Server Farm Host Configuration Mode Commands

Serverfarm host configuration mode commands allow you to create and configure host server farms and associate host real servers with the server farm. Host server farms are clusters of real servers that provide web content or services in a data center. You must configure a real server using the **(config) rserver** command in configuration mode before you can associate it with a server farm.

To create a host server farm and access serverfarm host configuration mode, use the **serverfarm** command. Note that host is the default server-farm type, so you do not need to enter the **host** option. The CLI prompt changes to (config-sfarm-host). For information about the commands in this mode, see the following commands.

Use the **no** form of this command to remove a server farm from the configuration.

serverfarm [**host**] *name*

no serverfarm *name*

Syntax Description

host	(Optional) Specifies a server farm of mirrored real servers that provide web content or services.
<i>name</i>	Unique identifier of the server farm. Enter an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.

Command Modes

Configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.

ACE Appliance Release	Modification
A1(7)	This command was introduced.

Usage Guidelines

The commands in this mode require the server-farm feature in your user role. For details about role-based access control (RBAC) and user roles, see the *Virtualization Guide, Cisco ACE Application Control Engine*.

Examples

To create a host server farm named SFARM1, enter:

```
host1/Admin(config) # serverfarm SFARM1
host1/Admin(config-sfarm-host) #
```

To delete the server farm named SFARM1, enter:

```
host1/Admin(config) # no serverfarm SFARM1
```

Related Commands

[show serverfarm](#)
[show running-config](#)
[\(config\) rserver](#)

(config-sfarm-host) description

To configure the description of a server farm, use the **description** command. Use the **no** form of this command to delete the description of a server farm.

description *text*

no description

Syntax Description

<i>text</i>	Text description of a server farm. Enter an unquoted text string with a maximum of 240 alphanumeric characters.
-------------	---

Command Modes

Serverfarm host configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.
ACE Appliance Release	Modification
A1(7)	This command was introduced.

Usage Guidelines

This command has no usage guidelines.

Examples

To configure a description of a server farm, enter:

```
host1/Admin(config-sfarm-host)# description CURRENT EVENTS ARCHIVE
```

To delete the description of a server farm, enter:

```
host1/Admin(config-sfarm-host)# no description
```

Related Commands

This command has no related commands.

(config-sfarm-host) dws

To enable a server farm for the dynamic workload scaling (DWS) feature, use the **dws** command. Use the **no** form of this command to disable the DWS feature on a server farm.

```
dws {local | burst probe name}
```

```
no dws {local | burst probe name}
```

Syntax Description	local	Specifies that only the local pool of VMs are taken into account for load balancing decisions
	burst	Specifies that remote VMs are taken into account for load balancing decisions when the configured threshold for the load of the local pool of VMs is reached or exceeded.
	probe <i>name</i>	Existing VM probe associated with this server farm. For details about configuring a VM probe, see the (config-sfarm-host) probe command.

Command Modes
Serverfarm host configuration mode
Admin and user contexts

Command History	ACE Module/Appliance Release	Modification
	A4(2.0)	This command was introduced.

Usage Guidelines
This command has no usage guidelines.

Examples
To enable DWS on a server farm, enter the following command:

```
host1/Admin(config-sfarm-host)# dws burst probe VM_PROBE
```

 To disable DWS on a server farm, enter the following command:

```
host1/Admin(config-sfarm-host)# no dws burst probe VM_PROBE
```

Related Commands
This command has no related commands.

(config-sfarm-host) failaction

To configure the action that the ACE takes if a real server in a server farm goes down, use the **failaction** command. Use the **no** form of this command to reset the ACE to its default of taking no action when a server fails.

failaction { **purge** | **reassign** [**across-interface**] }

no failaction

When the **failaction purge** command is present in a server farm, and if the probe fails or the real server is taken out of service with the "**no inservice**" configuration command, the ACE immediately purges all connections to this real server by sending a reset (RST) both to the client and server.

Syntax Description

purge	Specifies that the ACE remove the connections to a real server if that real server in the server farm fails after you configure this command. The ACE sends a reset (RST) both to the client and to the server that failed.
reassign	Specifies that the ACE reassigns existing server connections to the backup real server, if a backup real server is configured. If no backup real server is configured, this keyword has no effect.
across-interface	(Optional) Instructs the ACE to reassign all connections from the failed real server to a backup real server on a different VLAN that is commonly referred to as a bypass VLAN. By default, this feature is disabled.

Command Modes

Serverfarm host configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.
A2(1.0)	This command was revised (reassign keyword added).
A2(3.0)	The across-interface option was added.

ACE Appliance Release	Modification
A1(7)	This command was introduced.
A3(1.0)	This command was revised (reassign keyword added).
A4(1.0)	The across-interface option was added.

Usage Guidelines

If you do not configure this command, the ACE takes the real server out of rotation for new connections and allows existing connections to complete. The ACE does not send the connections to a backup server in the server farm or to a backup server farm if all servers in the primary server farm fail. To clear connections to servers that have failed prior to entering the **failaction** command, use the **clear conn** command.

This feature is required for stateful firewall load balancing (FWLB). For details about FWLB, see the *Server Load-Balancing Guide, Cisco ACE Application Control Engine*.

The use of the **failaction reassign** command requires that you enable the **transparent** command (see **(config-sfarm-host) transparent**) to instruct the ACE not to use NAT to translate the ACE VIP address to the server IP address. The **failaction reassign** command is intended for use in FWLB 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.

Follow these configuration requirements and restrictions when you use the **across-interface** option:

- 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.
- You must connect real servers directly to the ACE. This requirement also applies to same-VLAN backup servers.
- You must disable sequence number randomization on the firewall.
- 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 ACE1 and a low interval value on ACE2, the reassigned connections may become stuck because of the probe configuration mismatch. ACE2 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. ACE1 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.

To minimize packet loss, we recommend the following probe parameter values on both ACEs:

- Interval: 2
- Faildetect: 2
- Passdetect interval: 2
- Passdetect count: 5

Examples

To instruct the ACE to remove connections from a failed server in the server farm, enter:

```
host1/Admin(config-sfarm-host) # failaction purge
```

To specify that the ACE reassign the existing server connections to a backup real server in a different VLAN, enter:

```
host1/Admin(config-sfarm-host) # failaction reassign across-interface
host1/Admin(config-sfarm-host) # transparent
```

To specify that the ACE reassign the existing server connections to the backup real server, enter:

```
host1/Admin(config-sfarm-host)# failaction reassign  
host1/Admin(config-sfarm-host)# transparent
```

To reset the ACE to its default of taking no action if a real server fails, enter:

```
host1/Admin(config-sfarm-host)# no failaction
```

Related Commands [\(config-sfarm-host\) transparent](#)

(config-sfarm-host) fail-on-all

To configure the real servers in a server farm to use AND logic with respect to multiple server farm probes, use the **fail-on-all** command in server farm host configuration mode. This command is applicable to all probe types. The syntax of this command is:

```
fail-on-all
```

```
no fail-on-all
```

Syntax Description This command has no keywords or arguments.

Command Modes Server farm host configuration mode
Admin and user contexts

Command History	ACE Module Release	Modification
	A2(1.0)	This command was introduced.
	ACE Appliance Release	Modification
	A3(1.0)	This command was introduced.

Usage Guidelines 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. This 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.

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.

Examples To configure the SERVER1 real server to remain in the OPERATIONAL state unless all associated probes fail, enter the following commands:

```
host1/Admin(config)# rserver SERVER1
host1/Admin(config-rserver-host)# ip address 192.168.12.15
host1/Admin(config-rserver-host)# probe HTTP_PROBE
host1/Admin(config-rserver-host)# probe ICMP_PROBE
host1/Admin(config-rserver-host)# fail-on-all
```

To remove the AND probe logic from the real server and return the behavior to the default of OR logic, enter the following command:

```
host1/Admin(config-rserver-host)# no fail-on-all
```


Related Commands This command has no related commands.

(config-sfarm-host) inband-health check

To enable inband health monitoring for each real server in a server farm, use the **inband-health check** command. Use the **no** form of this command to disable inband health monitoring.

```
inband-health check {count | {log fail_threshold_count [reset milliseconds]} | {remove
fail_threshold_count [reset milliseconds] [resume-service seconds]}}
```

```
no inband-health
```

Syntax Description		
count		Tracks the total number of TCP or UDP failures, and increments the counters as displayed by the show serverfarm name inband command.
log		Logs a syslog error message when the number of events reaches the configured connection failure threshold.
remove		Specifies a syslog error message when the number of events reaches the threshold specified by the <i>threshold_number</i> argument and the ACE removes the server from service.
<i>fail-threshold</i>		The maximum number of connection failures that a real server can have during the configurable reset-time interval before the ACE marks the real server as failed. <ul style="list-style-type: none"> For the ACE module, enter an integer from 4 to 4294967295. For the ACE appliance, enter an integer from 1 to 4294967295.
reset <i>milliseconds</i>		Specifies the reset-time interval in milliseconds. For the <i>milliseconds</i> argument, enter an integer from 100 to 300000. The default interval is 100. 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 keyword, the ACE also removes the real server from service.
remove		Logs a syslog error message when the number of events reaches the configured threshold and removes the real server from service.
resume-service <i>seconds</i>		(Optional) Specifies the number of seconds after a server has been marked as failed for the ACE to reconsider sending live connections. For the <i>seconds</i> argument, enter an integer from 30 to 3600. The default setting is 0.

Command Modes	
	Server-farm host configuration mode Admin and user contexts

Command History	ACE Module/Appliance Release	Modification
	A4(1.0)	This command was introduced.

Usage Guidelines

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.

When you configure the inband health monitoring feature, it informs the ACE load balancer of connection failures on the real servers in a server farm. These connection failures are as follows:

- For TCP, resets (RSTs) from the server or SYN timeouts
- For UDP, ICMP Host, Network, Port, Protocol, and Source Route unreachable messages

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.

Inband health monitoring has the following considerations and restrictions:

- When you configure inband health monitoring, the setting of the resume-service option for the inband-health check command affects the behavior of the real server in the INBAND-HM-FAILED state.
- Inband health monitoring works with connection reuse only when the ACE to server connection is torn down, not for every request that is sent out on the reused connection.
- The state of the real server is not synchronized to the standby ACE when the state of the real server changes due to inband health monitoring.
- If you configure a different port for probes than what is used for traffic forwarding (for example, when you configure port inheritance or specify the port under the probe configuration), out-of-band and inband health monitoring monitor different ports.
- If a server farm is attached to two different VIPs, one servicing TCP and the other servicing UDP requests, and both TCP and UDP inband health monitoring are enabled on that server farm, the inband probe that goes down first takes the real server down. We recommend that you configure two different server farms, and enable both with inband health monitoring.
- When you configure inband health monitoring with a Layer 7 configuration containing a Layer 4 or Layer 7 class map, you must configure the inactivity timeout using the **set timeout inactivity** command to a time greater than the time to teardown the connection. The teardown time is based on the number of SYN retries configured by the **set tcp syn-retry** command. Otherwise, inband health monitoring does not track the syn-timeout failures. For example, if you configure the **set tcp syn-retry** command to 4, the connection teardown takes 45 seconds. You must configure the **set timeout inactivity** command to greater than 45 seconds.
- You can configure inband health monitoring to work with health probes to monitor a server. If you do, both sets of health checks are required to keep a real server in service within a server farm. If either detects a server is out of service, the ACE does not select the server for load balancing.
- You can configure inband health monitoring with HTTP return codes under the same server farm.

The **reset** 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** keyword, the ACE also removes the real server from service.

Changing the setting of the **reset** option affects the behavior of the real server, as follows:

- 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.

- The default setting is 0. The setting of this option affects the behavior of the real server in the INBAND-HM-FAILED state, as follows:
 - When the **resume-service** option is not configured and has the default setting of 0, the real server remains in the failed state until you manually enter the **no inservice** command followed by the **inservice** command.
 - When this option 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 option 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 the **resume-service** option and then decrease the value, the failed real server immediately transitions to the Operational state.

- When you configure this option 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 enter the **no inservice** command followed by the **inservice** command.
- When you change this option within the reset-time interval and the real server is in the OPERATIONAL state 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.

Examples

To track the total number of TCP or UDP failures for the real servers on a server farm and increment the **show serverfarm name inband** command counters, enter:

```
host1/Admin(config)# serverfarm host SF1
host1/Admin(config-sfarm-host)# inband-health check count
```

To configure the ACE to remove a real server at a failure threshold of 400, and resume service to it after 300 seconds, enter:

```
host1/Admin(config-sfarm-host)# inband-health check remove 400 resume-service 300
```

To disable inband health monitoring, enter:

```
host1/Admin(config-sfarm-host)# no inband-health
```

Related Commands

[show serverfarm](#)

(config-sfarm-host) partial-threshold

By default, if you configured 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. Partial server farm failover allows you to specify a failover threshold. If the percentage of active real servers in a server farm falls below the specified threshold, the primary server farm fails over to the backup server farm (if configured).

To enable partial server farm failover, use the **partial-threshold** command in server farm host configuration mode. Use the **no** form of this command to disable partial server farm failover.

partial-threshold *percentage1* **back-in-service** *percentage2*

no partial-threshold

Syntax Description		
	<i>percentage1</i>	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. Enter an integer from 0 to 99.
	back-in-service <i>percentage2</i>	Specifies 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. Enter an integer from 0 to 99. The percentage configured with the back-in-service keyword must be greater than or equal to the <i>percentage1</i> value.

Command Modes	
	Server-farm host configuration mode Admin and user contexts

Command History	ACE Module Release	Modification
	A2(1.0)	This command was introduced.
	ACE Appliance Release	Modification
	A3(1.0)	This command was introduced.

Usage Guidelines	
	Each time that a server is taken out of service (for example, by an administrator using the CLI, because of a probe failure, or because the retcode threshold is exceeded), the ACE is updated. If the percentage of active real servers in a server farm falls below the specified threshold, the primary server farm fails over to the backup server farm (if a backup server farm is configured).
	With partial server farm failover configured, the ACE allows current connections on the remaining active servers in the failed primary server farm to complete. The ACE redirects any new connection requests to the backup server farm.

Examples	
	To configure partial server farm failover, enter: <pre>host1/Admin(config-sfarm-host)# partial-threshold 40 back-in-service 60</pre>

To disable partial server farm failover, enter:

```
host1/Admin(config-sfarm-host) # no partial-threshold
```

Related Commands [show serverfarm](#)

(config-sfarm-host) predictor

To configure the load-balancing algorithm for the server farm, use the **predictor** command. Use the **no** form of this command to revert to the default load-balancing algorithm (the round-robin algorithm).

```
predictor {hash {address [destination | source] [v6-prefix prefix-length | netmask]} | {content
  [offset number1] [length number2] [begin-pattern expression1] [end-pattern expression2]} |
  {cookie [secondary] name1} | {header name2} | {layer4-payload [offset number3] [length
  number4] [begin-pattern expression3] [end-pattern expression4]} | {url [begin-pattern
  expression5] [end-pattern expression6]} | {least-bandwidth [samples number5]
  [assess-time seconds]} | {least-loaded probe name3 [samples number6]} | {leastconns
  [slowstart seconds]} | {response {app-req-to-req | syn-to-close | syn-to-synack} [samples
  number7]} | {roundrobin}
```

no predictor

Syntax Description

hash address	Selects the server using a hash value based on the source and destination IP addresses. Use the hash address source and hash address destination methods for firewall load balancing (FWLB).
destination	(Optional) Selects the server using a hash value based on the destination IP address.
source	(Optional) Selects the server using a hash value based on the source IP address.
v6-prefix <i>prefix-length</i>	(Optional) Specifies how many of the most significant bits (MSBs) of the IPv6 address are used for the network identifier. Enter an integer from 1 to 128.
<i>netmask</i>	(Optional) Bits in the IP address to use for the hash. If not specified, the default is 255.255.255.255.
hash content	Selects the server using a hash value based on the specified content string of the HTTP packet body.
offset <i>number1</i>	(Optional) Specifies 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. Enter an integer from 0 to 999. The default is 0, which indicates that the ACE does not exclude any portion of the content.

length <i>number2</i>	<p>(Optional) Specifies the length 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. Enter an integer from 1 to 1000. The default is the entire payload.</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 in the same hash content command.</p>
begin-pattern <i>expression1</i>	<p>(Optional) Specifies 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 immediately 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>Enter an unquoted text string with no spaces and a maximum of 255 alphanumeric characters for each pattern that you configure. You can enter a text string with spaces if you enclose the entire string in quotation marks (“”). The ACE supports the use of regular expressions for matching string expressions.</p> <p>When matching data strings, note that the period (.) and question mark (?) characters do not have a literal meaning in regular expressions. Use brackets ([]) to match these symbols (for example, enter <code>www[.]xyz[.]com</code> instead of <code>www.xyz.com</code>). You can also use a backslash (\) to escape a dot (.) or a question mark (?).</p>
end-pattern <i>expression2</i>	<p>(Optional) Specifies 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>Enter an unquoted text string with no spaces and a maximum of 255 alphanumeric characters for each pattern that you configure. You can enter a text string with spaces if you enclose the entire string in quotation marks (“”). The ACE supports the use of regular expressions for matching string expressions.</p> <p>Note: You cannot specify both the length and the end-pattern options in the same hash content command.</p>
hash cookie	<p>Selects the server using a hash value based on the cookie name or based on the name in the cookie name of the URL query string.</p>
secondary	<p>(Optional) Selects the server by using the hash value based on the specified name in the cookie name in the URL query string, not the cookie header. If you do not include this option, the ACE selects a real server using the hash value of the cookie name.</p>
<i>name1</i>	<p>Cookie name. Enter a cookie name as an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.</p>

hash header <i>name2</i>	<p>Selects the server using a hash value based on the header name. Enter a header name as an unquoted text string with no spaces and a maximum of 64 alphanumeric characters, or enter one of the following standard headers:</p> <ul style="list-style-type: none"> • Accept • Accept-Charset • Accept-Encoding • Accept-Language • Authorization • Cache-Control • Connection • Content-MD5 • Expect • From • Host • If-Match • Pragma • Referrer • Transfer-Encoding • User-Agent • Via
hash layer4-payload	<p>Specifies 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>
offset <i>number3</i>	<p>(Optional) Specifies the portion of the payload 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. Enter an integer from 0 to 999. The default is 0, which indicates that the ACE does not exclude any portion of the payload.</p>
length <i>number4</i>	<p>(Optional) Specifies the length 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. Enter an integer from 1 to 1000. The default is the entire payload.</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 in the same hash layer4-payload command.</p>

begin-pattern <i>expression3</i>	<p>(Optional) Specifies 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 immediately 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>Enter an unquoted text string with no spaces and a maximum of 255 alphanumeric characters for each pattern that you configure. You can enter a text string with spaces if you enclose the entire string in quotation marks (“”). The ACE supports the use of regular expressions for matching string expressions.</p> <p>When matching data strings, note that the period (.) and question mark (?) characters do not have a literal meaning in regular expressions. Use brackets ([]) to match these symbols (for example, enter www[.]xyz[.]com instead of www.xyz.com). You can also use a backslash (\) to escape a dot (.) or a question mark (?).</p>
end-pattern <i>expression4</i>	<p>(Optional) Specifies 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>Enter an unquoted text string with no spaces and a maximum of 255 alphanumeric characters for each pattern that you configure. You can enter a text string with spaces if you enclose the entire string in quotation marks (“”). The ACE supports the use of regular expressions for matching string expressions.</p> <p>Note: You cannot specify both the length and the end-pattern options in the same hash layer4-payload command.</p>
hash url	<p>Selects the server using a hash value based on the requested URL. Use this predictor method to load balance cache servers. Cache servers perform better with the URL hash method because you can divide the contents of the caches evenly if the traffic is random enough. In a redundant configuration, the cache servers continue to work even if the active ACE switches over to the standby ACE. For information about configuring redundancy, see the <i>Administration Guide, Cisco ACE Application Control Engine</i>.</p>
begin-pattern <i>expression5</i>	<p>(Optional) Specifies the beginning pattern of the URL and the pattern string to match before hashing. You cannot configure different beginning and ending patterns for different server farms that are part of the same traffic classification. Enter an unquoted text string with no spaces and a maximum of 255 alphanumeric characters for each pattern that you configure. If you want to match a URL that contains spaces, you must use \x20 for each space character.</p>

end-pattern <i>expression6</i>	(Optional) Specifies the pattern that marks the end of hashing. You cannot configure different beginning and ending patterns for different server farms that are part of the same traffic classification. Enter an unquoted text string with no spaces and a maximum of 255 alphanumeric characters for each pattern that you configure. If you want to match a URL that contains spaces, you must use <code>\x20</code> for each space character.
least-bandwidth	Selects the server that processed the least amount of network traffic over a specified sampling period. Use this predictor for heavy traffic use, such as downloading a video clip. The ACE measures traffic statistics between itself and the real servers in the server farm in both directions and calculates the bandwidth over the sampling period. Then, it creates an ordered list of real servers based on the sampling results and selects the server that used the least amount of bandwidth during the sampling period.
samples <i>number5</i>	(Optional) Specifies the number of samples over which you want to weight and average the results of the probe query to calculate the final load value. Enter an integer from 1 to 16. Each value must be a power of 2, so the valid values are as follows: 1, 2, 4, 8, and 16. The default is 8.
assess-time <i>seconds</i>	(Optional) Specifies the sampling period over which the ACE measures traffic for all the servers in the server farm. Enter an integer from 1 to 10. The default is 2 seconds.
least-loaded	Selects the server with the lowest load based on information obtained from SNMP probes. To use this predictor, you must associate an SNMP probe with the server farm. The ACE queries one user-specified OID (for example, CPU utilization or memory utilization). The ACE uses the retrieved value directly to determine the server with the lowest load.
probe <i>name3</i>	Specifies the name of the SNMP probe that you want to query. Enter an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.
samples <i>number6</i>	(Optional) Specifies the number of samples over which you want to weight and average the results of the probe query to calculate the final load value. Enter an integer from 1 to 16. Each value must be a power of 2, so the valid values are as follows: 1, 2, 4, 8, and 16. The default is 8.
leastconns	Selects the real server with the fewest number of active connections based on the server weight. Use this predictor for processing light user requests (for example, browsing simple static web pages). For information about setting real server weight, see the (config-sfarm-host-rs) weight section.
slowstart <i>seconds</i>	(Optional) Specifies that the connections to the real server be in a slow-start mode for the duration indicated by the <i>seconds</i> value. Use the slow-start mechanism to avoid sending a high rate of new connections to servers that you have recently put into service. Enter an integer from 1 to 65535, where 1 is the slowest ramp-up value. By default, slowstart is disabled.

response	Selects the server with the lowest response time for the requested response-time measurement. If you do not specify a response-time measurement method, the ACE uses the HTTP app-req-to-response method.
app-req-to-resp	(Default) Measures 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. The ACE does not allow you to configure this predictor response in a generic load-balancing policy map.
syn-to-close	Measures 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	Measures the response time from when the ACE sends a TCP SYN to a server to the time that the ACE receives the SYN-ACK from the server.
samples <i>number</i>	(Optional) Number of samples over which you want to average the results of the response time measurement. Enter an integer from 1 to 16 in powers of 2. Valid values are: 1, 2, 4, 8, and 16. The default is 8.
roundrobin	(Default) Selects the next server in the list of real servers based on server weight (weighted round-robin). For information about setting real server weight, see the (config-sfarm-host-rs) weight section.

Command Modes

Server-farm host configuration mode

Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.
A2(1.0)	This command was revised.
A2(1.4) and A2(2.1)	This secondary option for the hash cookie keywords was added.

ACE Appliance Release	Modification
A1(7)	This command was introduced.
A3(1.0)	This command was revised.
A3(2.2)	This secondary option for the hash cookie keywords was added.

Usage Guidelines

Use this command to specify the load-balancing algorithm that the ACE uses in choosing a real server in the server farm. If you do not specify the **predictor** command, the default algorithm is **roundrobin**. Using the **no** form of this command changes the configured predictor algorithm to the default algorithm.

The weight assigned to the real servers is used only in the **roundrobin** and **leastconns** predictor methods. The **hash** and the **response** predictor methods do not recognize the weight for the real servers. For information about setting real server weight, see the [\(config-sfarm-host-rs\) weight](#) section.

If you configure the **leastconns** predictor, you can use a **slowstart** mechanism (ramp-up) to avoid sending a high rate of new connections to the servers that have just been put in service. The real server with the fewest number of active connections will get the next connection request for the server farm with the **leastconns** predictor. The ramp-up stops when the duration timer that you specify expires.

The only time that the sequence of servers starts over at the beginning (with the first server) is when there is a configuration or server state change (for example, a probe failure).

Server weights take effect only when there are open connections to the servers. When there are no sustained connections to any of the servers, the leastconns predictor method behaves like the roundrobin method.

The **secondary** option allows the ACE to correctly load balance in cases when the query string identifies the actual resource, instead of the URL.

Examples

To configure the ACE to select the real server with the lowest number of connections in the server farm, enter:

```
host1/Admin(config-sfarm-host) # predictor leastconns slowstart 300
```

To reset the load-balancing algorithm to the default of roundrobin, enter:

```
host1/Admin(config-sfarm-host) # no predictor
```

Related Commands

[\(config-sfarm-host-rs\) weight](#)

(config-sfarm-host) probe

Use probes to monitor the health of real servers in a server farm. To associate a probe with a server farm, use the **probe** command. Use the **no** form of this command to dissociate a probe from a server farm.

probe *probe-name*

no probe *probe-name*

Syntax Description

<i>probe-name</i>	Identifier of an existing probe that you want to associate with a server farm. Enter an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.
-------------------	--

Command Modes

Serverfarm host configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.

ACE Appliance Release	Modification
A1(7)	This command was introduced.

Usage Guidelines

The probe must already exist. (To create a probe, see the [\(config\) probe](#) command.) You can associate multiple probes of the same or different protocols with each server farm.

Examples

To associate a probe with a server farm, enter:

```
host1/Admin(config-sfarm-host)# probe TCP1
```

To dissociate a probe from a server farm, enter:

```
host1/Admin(config-sfarm-host)# no probe TCP1
```

Related Commands

[\(config\) probe](#)

(config-sfarm-host) retcode

To configure HTTP return-code checking (retcode map) for a server farm, use the **retcode** command. Use the **no** form of this command to dissociate a return code map. You can specify a single return code number or a range of return code numbers. For example, you can instruct the ACE to check for and count the number of occurrences of such return codes as HTTP/1.1 200 OK, HTTP/1.1 100 Continue, or HTTP/1.1 404 Not Found.

```
retcode number1 number2 check {count | {log threshold_number reset seconds1
| {remove threshold_number reset seconds1 [resume-service seconds2]}}
```

```
no retcode number1 number2
```

Syntax Description

<i>number1</i>	Minimum value for an HTTP return code. Enter an integer from 100 to 599. The minimum value must be less than or equal to the maximum value.
<i>number2</i>	Maximum value for an HTTP return code. Enter an integer from 100 to 599. The maximum value must be greater than or equal to the minimum value.
check	Checks for HTTP return codes associated with the server farm.
count	Tracks the total number of return codes received for each return code number that you specify.
log	Specifies a syslog error message when the number of events reaches the threshold specified by the <i>threshold_number</i> argument.
remove	Specifies a syslog error message when the number of events reaches the threshold specified by the <i>threshold_number</i> argument and the ACE removes the server from service.
<i>threshold_number</i>	Threshold for the number of events that the ACE receives before it performs the log or remove action. <ul style="list-style-type: none"> For the ACE module, enter an integer from 4 to 4294967295. For the ACE appliance, enter an integer from 1 to 4294967295.
reset <i>seconds1</i>	Specifies the time interval in seconds over which the ACE checks for the return code for the log or remove action. <ul style="list-style-type: none"> For the ACE module, enter an integer from 1 to 4294967295. For the ACE appliance, enter an integer from 1 to 2147483647.
resume-service <i>seconds2</i>	(Optional) Specifies the number of seconds that the ACE waits before it resumes service for the real server automatically after taking the real server out of service because the remove option is configured. Enter an integer from 30 to 3600. The default setting is 0.

Command Modes

Server-farm host configuration mode

Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.
A2(1.0)	This command was revised.
A4(1.0)	The lowest integer for the <code>threshold_number</code> argument was changed from 2 to 4.

ACE Appliance Release	Modification
A1(7)	This command was introduced.
A3(1.0)	This command was revised.

Usage Guidelines

You can configure multiple return code maps on each server farm. You can view hitcounts for return code checking by using the **show serverfarm** command.

The setting of the **remove** option affects the behavior of the real server in the failed state, as follows:

- When the **resume-service** option is not configured and has the default setting of 0, the real server remains in the failed state until you manually enter the **no inservice** command followed by the **inservice** command.
- When this option is not configured and has the default setting of 0 and then you configure this option with an integer between 30 and 3600, the failed real server transitions to the Operational state.
- When you configure this option 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 option and then decrease the value, the failed real server transitions to the Operational state.
- When you configure this option with an integer between 30 and 3600 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 enter the **no inservice** command followed by the **inservice** command.

The ACE performs the log or remove actions only if the *threshold_number* value for a particular retcode is reached within a specified period of time. The time period is defined from the receipt of a retcode until the next reset time.

Examples

To check for and count the number of return code hits for all return codes from 200 to 500 inclusive, enter:

```
host1/Admin(config-sfarm-host)# retcode 200 500 check count
```

To remove the HTTP return-code map from the configuration, enter:

```
host1/Admin(config-sfarm-host)# no retcode 200 500
```

Related Commands

[show serverfarm](#)

(config-sfarm-host) rserver

To associate one or more existing host real servers with a server farm and access serverfarm host real server configuration mode, use the **rserver** command. The CLI prompt changes to (config-sfarm-host-rs). For information on commands in serverfarm host real server configuration mode, see the “[Server Farm Host Real Server Configuration Mode Commands](#)” section. Use the **no** form of this command to dissociate the real server from the server farm.

```
rserver name [port]
```

```
no rserver name [port]
```

Syntax Description

<i>name</i>	Unique identifier of the real server. Enter an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.
<i>port</i>	(Optional) Port number used for the real server Port Address Translation (PAT). Enter an integer from 1 to 65535.

Command Modes

Serverfarm host configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.

ACE Appliance Release	Modification
A1(7)	This command was introduced.

Usage Guidelines

The real server must already exist. To create a real server, see the [\(config\) rserver](#) command. You can associate a maximum of 16,384 real servers with a server farm.

You can configure a combination of IPv6 and IPv4 servers in a server farm (mixed mode), but a mixed mode server farm will have limited feature support.



Caution

Do not configure under a real server a peer IPv6 address that is calculated from EUI64. In a redundant configuration, if you configure a peer IPv6 address as EUI64 on an interface, the address will not be learned by the active member of an FT group because the address is calculated only on the peer. If you then configure the same calculated IPv6 address on the active under a real server, the CLI accepts it because it does not calculate it. This IPv6 address is not synced to the standby because it conflicts with the interface address. If you subsequently apply a probe to the real server, the state of the real server is PROBE-FAILED on the active and OUTFSERVICE on the the standby. This same check applies to VIPs, routes, interfaces, and probes.

If you choose not to assign a port number for the real server association with the server farm, the default behavior by the ACE is to automatically assign the same destination port that was used by the inbound connection to the outbound server connection. For example, if the incoming connection to the ACE is a

secure client HTTPS connection, the connection is typically made on port 443. If you do not assign a port number to the real server, the ACE will automatically use port 443 to connect to the server, which results in the ACE making a clear-text HTTP connection over port 443. In this case, you would typically define an outbound destination port of 80, 81, or 8080 for the backend server connection.

Examples

To associate a real server with a server farm, enter:

```
host1/Admin(config-sfarm-host)# rserver server1 80
```

To dissociate a real server from a server farm, enter:

```
host1/Admin(config-sfarm-host)# no rserver server1 80
```

Related Commands [\(config\) rserver](#)

(config-sfarm-host) transparent

To configure the ACE not to use Network Address Translation (NAT) to translate the ACE VIP address to the server IP address, use the **transparent** command. Use the **no** form of this command to reset the ACE to its default of using NAT to translate the VIP address to the server IP address.

transparent

no transparent

Syntax Description

This command has no keywords or arguments.

Command Modes

Serverfarm host configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.

ACE Appliance Release	Modification
A1(7)	This command was introduced.

Usage Guidelines

Use this command in firewall load balancing (FWLB) when you configure the insecure and secure sides of the firewall as a server farm. For details about FWLB, see the *Server Load-Balancing Guide, Cisco ACE Application Control Engine*.

Examples

To prevent the ACE from using NAT to translate the ACE VIP address to the server IP address, enter:

```
host1/Admin(config-sfarm-host) # transparent
host1/Admin(config-sfarm-host) #
```

To reset the ACE to its default of using NAT to translate the VIP address to the server IP address, enter:

```
host1/Admin(config-sfarm-host) # no transparent
host1/Admin(config-sfarm-host) #
```

Related Commands

This command has no related commands.

Serverfarm Host Predictor Configuration Mode Commands

Serverfarm host predictor configuration mode commands allow you to configure additional parameters for some of the server farm predictor methods.

To configure these additional predictor parameters, use the **predictor least-loaded** or the **predictor response** command in serverfarm host configuration mode. The CLI prompt changes to (config-sfarm-host-predictor). For information about the commands in this mode, see the following commands. Use the **no** form of this command to remove the predictor from the server farm.

```
predictor {least-loaded probe name} | {response {app-req-to-resp | syn-to-close |
syn-to-synack}[samples number]}}
```

```
no predictor
```

Syntax Description

least-loaded	Selects the server with the lowest load based on information obtained from SNMP probes. To use this predictor, you must associate an SNMP probe with the server farm. The ACE queries one user-specified OID (for example, CPU utilization or memory utilization). The ACE uses the retrieved value directly to determine the server with the lowest load.
probe name	Specifies the name of the SNMP probe that you want to query. Enter an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.
response	Selects the server with the lowest response time for the requested response-time measurement. If you do not specify a response-time measurement method, the ACE uses the HTTP app-req-to-response method.
app-req-to-resp	(Default) Measures 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. The ACE does not allow you to configure this predictor response in a generic load-balancing policy map.
syn-to-close	Measures 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	Measures the response time from when the ACE sends a TCP SYN to a server to the time that the ACE receives the SYN-ACK from the server.
samples <i>number</i>	(Optional) Number of samples over which you want to average the results of the response time measurement. Enter an integer from 1 to 16 in powers of 2. Valid values are: 1, 2, 4, 8, and 16. The default is 8.

Command Modes

Serverfarm host configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
A2(1.0)	This command was introduced.

ACE Appliance Release	Modification
A3(1.0)	This command was introduced.

Usage Guidelines

The commands in this mode require the server-farm feature in your user role. For details about role-based access control (RBAC) and user roles, see the *Virtualization Guide, Cisco ACE Application Control Engine*.

Examples

To specify the least-loaded predictor method with a probe called SNMP_PROBE for the server farm, enter:

```
host1/Admin(config-sfarm-host)# predictor least-loaded probe SNMP_PROBE
host1/Admin(config-sfarm-host-predictor)#
```

To remove the least-loaded predictor from the server farm, enter:

```
host1/Admin(config-sfarm-host)# no predictor
```

To specify the response predictor method that measures 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, enter:

```
host1/Admin(config-sfarm-host)# predictor response app-req-to-req
host1/Admin(config-sfarm-host-predictor)#
```

To remove the response predictor from the server farm, enter:

```
host1/Admin(config-sfarm-host)# no predictor
```

Related Commands

[show serverfarm detail](#)
[\(config-sfarm-host\) predictor](#)
[\(config-sfarm-host-predictor\) autoadjust](#)
[\(config-sfarm-host-predictor\) weight connection](#)

(config-sfarm-host-predictor) autoadjust

After you specify the **predictor least-loaded** command, use the **autoadjust** command to instruct the ACE to apply the maximum load of 16000 to a real server whose load reaches zero or override the default behavior. Use the **no** form of this command to return the ACE behavior to the reset the behavior of the ACE to the default of average load of the server farm to a real server whose load is zero.

autoadjust { **average** | **maxload** | **off** }

no autoadjust

Syntax Description

average	Applies the average load of the server farm to a real server whose load is zero. This setting allows the server to participate in load balancing, while preventing it from being flooded by new connections. This is the default setting.
maxload	Instructs the ACE to apply the maximum load of 16000 to a real server whose load reaches zero.
off	Overrides the default behavior of the ACE of applying the average load of the server farm to a real server whose load is zero. When you configure this command, the ACE sends 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. If two servers have the same lowest load (either zero or nonzero), the ACE load balances the connections between the two servers in a round-robin manner.

Command Modes

Serverfarm host predictor configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
A2(1.0)	This command was introduced.
A2(2.4)	The average load became the default autoadjust setting for the least-loaded predictor. Previously, the default setting was maximum load. The maxload keyword was added to set the least-loaded predictor to maximum load.

ACE Appliance Release	Modification
A3(1.0)	This command was introduced.
A4(1.0)	The average load became the default autoadjust setting for the least-loaded predictor. Previously, the default setting was maximum load. The maxload keyword was added to set the least-loaded predictor to maximum load.

Usage Guidelines

Whenever a server's load reaches zero, by default, the ACE uses the autoadjust feature to assign an average load value 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.

Using the least-loaded predictor with the configured server weight and the current connection count option enabled, the ACE calculates the final load of a real server as follows:

$$\text{final load} = \text{weighted load} \times \text{static weight} \times \text{current connection count}$$

where:

- *weighted load* is the load reported by the SNMP probe
- *static weight* is the configured weight of the real server
- *current connection count* is the total number of active connections to the real server

The ACE recalculates the final load whenever the connection count changes, provided that the **weight connection** command is configured. If the **weight connection** command is not configured, the ACE updates the final load when the next load update arrives from the SNMP probe.

Examples

To instruct the ACE to apply the maximum load of 16000 to a real server whose load reaches zero, enter:

```
host1/Admin(config-sfarm-host-predictor) # autoadjust maxload
```

To turn off the autoadjust feature for all servers in a server farm so that servers with a load of zero receive all new connections, enter:

```
host1/Admin(config-sfarm-host-predictor) # autoadjust off
```

To reset the behavior of the ACE to the default of applying the average load of the server farm to a real server whose load is zero, enter:

```
host1/Admin(config-sfarm-host-predictor) # no autoadjust
```

You can also reset the behavior of the ACE to the default by entering the following:

```
host1/Admin(config-sfarm-host-predictor) # autoadjust average
```

Related Commands

[show serverfarm detail](#)
[\(config-sfarm-host\) predictor](#)
[\(config-sfarm-host-predictor\) weight connection](#)

(config-sfarm-host-predictor) weight connection

After you specify the **predictor least-loaded** or the **predictor response** command, use the **weight connection** command to instruct the ACE to use the current connection count in the final load calculation for each real server in the server farm. Use the **no** form of this command to reset the behavior of the ACE to the default of excluding the current connection count from the load calculation.

weight connection

no weight connection

Syntax Description This command has no keywords or arguments.

Command Modes Serverfarm host predictor configuration mode
Admin and user contexts

Command History	ACE Module Release	Modification
	A2(1.0)	This command was introduced.

Command History	ACE Appliance Release	Modification
	A3(1.0)	This command was introduced.

Usage Guidelines To see how the **weight connection** command affects the [\(config-sfarm-host-predictor\) autoadjust](#) command for the least-loaded predictor, see the Usage Guidelines section of the [\(config-sfarm-host-predictor\) autoadjust](#) command.

Examples To instruct the ACE to use the current connection count in the final load calculation for each real server in the server farm, enter:

```
host1/Admin(config-sfarm-host-predictor)# weight connection
```

To reset the behavior of the ACE to the default of excluding the current connection count from the load calculation, enter:

```
host1/Admin(config-sfarm-host-predictor)# no weight connection
```

Related Commands [show serverfarm detail](#)
[\(config-sfarm-host\) predictor](#)
[\(config-sfarm-host-predictor\) autoadjust](#)

Server Farm Host Real Server Configuration Mode Commands

Serverfarm host real server configuration mode commands allow you to associate a host real server with a host server farm and configure the real server attributes.

To associate one or more existing host real servers with a host server farm and access serverfarm host real server configuration mode, use the **rserver** command in serverfarm host configuration mode. The CLI prompt changes to (config-sfarm-host-rs). For information about the commands in this mode, see the following commands. Use the **no** form of this command to remove the real server from the server farm.

```
rserver name [port]
```

```
no rserver name
```

Syntax Description

<i>name</i>	Unique identifier of the real server. Enter an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.
<i>port</i>	(Optional) Port number used for the real server Port Address Translation (PAT). Enter an integer from 1 to 65535.

Command Modes

Serverfarm host configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.
ACE Appliance Release	Modification
A1(7)	This command was introduced.

Usage Guidelines

The commands in this mode require the server-farm feature in your user role unless otherwise specified. For details about role-based access control (RBAC) and user roles, see the *Virtualization Guide, Cisco ACE Application Control Engine*.

The real server must already exist. To create a real server, see the **(config) rserver** command. You can associate a maximum of 16,384 real servers with a server farm.

Examples

To associate a real server with a server farm, enter:

```
host1/Admin(config-sfarm-host)# rserver SERVER1
```

To dissociate a real server from a server farm, enter:

```
host1/Admin(config-sfarm-host)# no rserver SERVER1
```

Related Commands This command has no related commands.

(config-sfarm-host-rs) backup-rserver

To configure a backup real server for a real server in a server farm, use the **backup-rserver** command. If a real server associated with a server farm becomes unavailable, the ACE directs flows to the configured backup real server. Use the **no** form of this command to remove a backup real server from the configuration.

backup-rserver *name* [*port*]

no backup-rserver

Syntax Description

<i>name</i>	Unique identifier of an existing real server that you want to configure as a backup server in a server farm. Enter an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.
<i>port</i>	(Optional) Port number used for the backup real server Port Address Translation (PAT). Enter an integer from 0 to 65535.

Command Modes

Serverfarm host real server configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.

ACE Appliance Release	Modification
A1(7)	This command was introduced.
A3(2.2)	This command now supports cyclic backup of real servers in a server farm.

Usage Guidelines

The real server used as a backup server must already exist. To create a real server, see the [\(config\) rserver](#) command.

Examples

To associate a backup real server with a server farm, enter:

```
host1/Admin(config-sfarm-host-rs) # backup-rserver BACKUP_SERVER1 3500
```

To dissociate a backup real server from a server farm, enter:

```
host1/Admin(config-sfarm-host-rs) # no backup-rserver
```

Related Commands

[\(config\) rserver](#)
[\(config-sfarm-host-rs\) inservice](#)

(config-sfarm-host-rs) conn-limit

To configure the maximum and minimum number of connections that you want to allow for a host real server in a server farm, use the **conn-limit** command. Use the **no** form of this command to reset the limits for the real server maximum connections and minimum connections to the default of 4000000.

conn-limit max *maxconns* **min** *minconns*

no conn-limit

Syntax Description

max <i>maxconns</i>	Specifies the maximum number of connections allowed for this real server. Enter an integer from 2 to 4000000. The default is 4000000.
min <i>minconns</i>	Specifies the connection threshold below which the real server will start accepting connections again after the number of connections exceeds the configured maximum number of connections. Enter an integer from 2 to 4000000. The default is <i>minconns</i> equal to <i>maxconns</i> .

Command Modes

Serverfarm host real server configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.

ACE Appliance Release	Modification
A1(7)	This command was introduced.

Usage Guidelines

Use this command to specify the maximum number of connections and the minimum connection threshold for a host real server in a server farm. The *minconns* value must be less than or equal to the *maxconns* value. The ACE uses the *minconns* value as a threshold to start accepting connections again after the *maxconns* limit is exceeded.

Examples

To configure the maximum number of connections and the minimum connection threshold for a host real server, enter:

```
host1/Admin(config-sfarm-host-rs)# conn-limit max 65535 min 40000
```

To reset the maximum number of connections and the minimum connection threshold for a host real server to the default of 4000000, enter:

```
host1/Admin(config-sfarm-host-rs)# no conn-limit
```

Related Commands

[\(config-sfarm-host-rs\) rate-limit](#)

(config-sfarm-host-rs) cookie-string

To configure a cookie string value of the real server for HTTP cookie insertion when establishing a sticky connection, use the **cookie-string** command. Use the **no** form of this command to remove the user-defined cookie string value of the real server for cookie insertion and have the ACE generate the cookie string for the associated real server.

cookie-string *text_string*

no cookie-string

Syntax Description

text_string

Cookie string value for the real server. Enter a text string with a maximum of 32 alphanumeric characters. When you include spaces and special characters in a cookie string value, enter a quoted text string (for example, "test cookie string"). The quotes appear in the running-configuration file.

Command Modes

Serverfarm host real server configuration mode
Admin and user contexts

Command History

ACE Module Release

Modification

A4(1.0)

This command was introduced.

ACE Appliance Release

Modification

A3(2.2)

This command was introduced.

Usage Guidelines

Use cookie insertion when you want to use a session cookie for persistence if the server is not currently setting the appropriate cookie. When you configure a cookie string value, the ACE inserts the cookie in the Set-Cookie header of the response from the server to the client.

If you do not configure a cookie string value, when you enable cookie insertion for a sticky group, the ACE generates the cookie string for each real server after sending a connection to it. The ACE-generated cookie string appears as "Rxxxxxxx" (for example, R2148819051).

When configuring a cookie string value, consider the following:

- You can configure one cookie string for each real server.
- The ACE automatically uses the user-defined cookie string for cookie insertion for a sticky group instead of the ACE-generated cookie string.
- Ensure that there are no duplicate strings configured for real servers. If there are duplicate cookie strings, the old entry will be removed and sticky database will use the latest configured cookie string for the real server.

If you remove the user-defined cookie string from a real server, the ACE generates the cookie string for the associated real server after sending a connection.

Examples

To configure a cookie string value of the real server for HTTP cookie insertion, enter:

```
host1/Admin(config-sfarm-host-rs)# cookie-string ABC123
```

To remove the configured cookie string value, enter:

```
host1/Admin(config-sfarm-host-rs)# no cookie-string
```

Related Commands

[show sticky database](#)

(config-sfarm-host-rs) description

To configure the description of a real server in a server farm, use the **description** command. Use the **no** form of this command to delete the description of a real server.

description *text*

no description

Syntax Description

<i>text</i>	Text description of a server farm. Enter an unquoted text string with a maximum of 240 alphanumeric characters. If the text string includes spaces, enclose the string in quotes/
-------------	---

Command Modes

Serverfarm host real server configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
A2(2.1)	This command was introduced.
ACE Appliance Release	Modification
A3(2.4)	This command was introduced.

Usage Guidelines

This command has no usage guidelines.

Examples

To configure a description of a real server in a server farm, enter:

```
host1/Admin(config-sfarm-host-rs)# description "CURRENT EVENTS: EDITION 1"
```

To delete the description, enter:

```
host1/Admin(config-sfarm-host-rs)# no description
```

Related Commands

This command has no related commands.

(config-sfarm-host-rs) fail-on-all

To configure a real server in a server farm to remain in the OPERATIONAL state unless all probes associated with it fail (AND logic), use the **fail-on-all** command in server farm host real server configuration mode. This command is applicable to all probe types. The syntax of this command is:

fail-on-all

no fail-on-all

Syntax Description This command has no keywords or arguments.

Command Modes Server farm host real server configuration mode
Admin and user contexts

Command History	ACE Module Release	Modification
	A2(1.0)	This command was introduced.

Command History	ACE Appliance Release	Modification
	A3(1.0)	This command was introduced.

Usage Guidelines By default, multiple probes that you configure directly on a real server in a server farm have an OR logic associated with them. This means that, if one of the real server probes fails, then the real server fails and enters the PROBE-FAILED state.

You can selectively configure this command on only certain real servers in the server farm to give those server AND logic. Any real server that you do not configure with the **fail-on-all** command, maintains its default OR logic with respect to probes.

Examples For example, to configure the SERVER1 real server in SFARM1 to remain in the OPERATIONAL state unless all associated probes fail, enter the following commands:

```
host1/Admin(config)# serverfarm SFARM1
host1/Admin(config-sfarm-host)# rserver SERVER1
host1/Admin(config-sfarm-host-rs)# inservice
host1/Admin(config-sfarm-host-rs)# probe HTTP_PROBE
host1/Admin(config-sfarm-host-rs)# probe ICMP_PROBE
host1/Admin(config-sfarm-host-rs)# fail-on-all
```

If either HTTP_PROBE or ICMP_PROBE fails, the SERVER1 real server remains in the OPERATIONAL state. If both probes fail, the real server fails and enters the PROBE-FAILED state.

To remove the AND probe logic from the real server in a server farm and return the behavior to the default of OR logic, enter the following command:

```
host1/Admin(config-rserver-host)# no fail-on-all
```

Related Commands This command has no related commands.

(config-sfarm-host-rs) inservice

To place a real server associated with a server farm in service, use the **inservice** command. Use the **no** form of this command to take a real server out of service.

inservice [**standby**]

no inservice

Syntax Description	standby	(Optional) Used with backup real servers, specifies that a backup real server remain inactive unless the primary real server fails. If the primary fails, the backup server becomes active and starts accepting connections.
--------------------	---------	--

Command Modes	Serverfarm host real server configuration mode Admin and user contexts
---------------	---

Command History	ACE Module Release	Modification
	3.0(0)A1(2)	This command was introduced.

ACE Appliance Release	Modification
A1(7)	This command was introduced.

Usage Guidelines This command requires the real-inservice feature in your user role. For details about role-based access control (RBAC) and user roles, see the *Virtualization Guide, Cisco ACE Application Control Engine*.

To start load balancing connections to a real server in a server farm, you must place the real server in service by using the **inservice** command.

You can modify the attributes of a real server in a server farm without taking the server out of service.

In addition to putting a backup real server in service standby, another use of the **inservice standby** command is to provide the graceful shutdown of primary real servers. Use this command to gracefully shut down servers with sticky connections. When you enter this command for a primary real server, the ACE does the following:

- Tears down existing non-TCP connections to the server
- Allows current TCP connections to complete
- Allows new sticky connections for existing server connections that match entries in the sticky database
- Load balances all new connections (other than the matching sticky connections mentioned above) to the other servers in the server farm
- Eventually takes the server out of service

Examples

To place a real server in service, enter:

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

To take a real server out of service, enter:

```
host1/Admin(config-sfarm-host-rs) # no inservice
```

To perform a graceful shutdown on a primary real server with sticky connections in a server farm, enter:

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

Related Commands

This command has no related commands.

(config-sfarm-host-rs) probe

To configure a probe to monitor the health of a host real server in a host server farm, use the **probe** command. Use the **no** form of this command to remove the probe from the real server.

probe *probe-name*

no probe *probe-name*

Syntax Description

<i>probe-name</i>	Identifier of an existing probe that you want to assign to a real server to monitor its health. Enter an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.
-------------------	---

Command Modes

Serverfarm host real server configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.

ACE Appliance Release	Modification
A1(7)	This command was introduced.

Usage Guidelines

You can associate multiple probes with each real server.

The ACE periodically sends the probes to the real servers. If the ACE receives the appropriate responses from the servers, the ACE includes the servers in load-balancing decisions. If not, the ACE marks the servers as out of service, depending on the configured number of retries.

Examples

To configure a probe for a host real server, enter:

```
host1/Admin(config-sfarm-host-rs)# probe SERVER1_PROBE
```

To remove a probe from a host real server, enter:

```
host1/Admin(config-sfarm-host-rs)# no probe SERVER1_PROBE
```

Related Commands

This command has no related commands.

(config-sfarm-host-rs) rate-limit

To configure a limit for the connection rate and the bandwidth rate of a real server in a host server farm, use the **rate-limit** command. The connection rate is the number of connections per second received by the ACE and destined to a particular real server. The bandwidth rate is the number of bytes per second received by the ACE and destined for a particular real server. Use the **no** form of this command to revert to the ACE default of not limiting the connection rate or bandwidth rate of real servers in a server farm.

```
rate-limit {connection number1 | bandwidth number2}
```

```
no rate-limit {connection | bandwidth}
```

Syntax Description

connection <i>number1</i>	Specifies the real server connection-rate limit in connections per second. <ul style="list-style-type: none"> For the ACE module, enter an integer from 2 to 350000. For the ACE appliance, enter an integer from 1 to 350000. There is no default value.
bandwidth <i>number2</i>	Specifies the real server bandwidth-rate limit in bytes per second. <ul style="list-style-type: none"> For the ACE module, enter an integer from 4 to 300000000. For the ACE appliance, enter an integer from 1 to 300000000. There is no default value.

Command Modes

Serverfarm host real server configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.
A4(1.0)	The lowest real server bandwidth-rate limit was changed from 2 to 4.

ACE Appliance Release	Modification
A3(1.0)	This command was introduced.

Usage Guidelines

If the connection rate or the bandwidth rate of incoming traffic destined for a particular real server exceeds the configured rate for the server, the ACE blocks any further traffic destined to that real server until the connection rate or bandwidth rate drops below the configured limit. Also, the ACE removes the blocked real server from future load-balancing decisions. By default, the ACE does not limit the connection rate or the bandwidth rate of real servers in a server farm.

Examples

To limit the connection rate of a real server to 100,000 connections per second, enter:

```
host1/Admin(config-sfarm-host-rs)# rate-limit connection 100000
```

To revert to the ACE default of not limiting the real-server connection rate, enter:

```
host1/Admin(config-sfarm-host-rs)# no rate-limit connection
```

To limit the real-server bandwidth rate to 5,000,000 bytes per second, enter:

```
host1/Admin(config-sfarm-host-rs)# rate-limit bandwidth 5000000
```

To revert to the ACE default of not limiting real-server bandwidth, enter:

```
host1/Admin(config-sfarm-host-rs)# no rate-limit bandwidth
```

Related Commands [\(config-sfarm-host-rs\) conn-limit](#)

(config-sfarm-host-rs) weight

To configure the capacity of a real server in relation to other servers in a server farm, use the **weight** command. The weight value that you specify for a server is used in the weighted round-robin and least-connections predictor load-balancing methods. Use the **no** form of this command to reset the real server weight to the default.

weight *number*

no weight

Syntax Description	<i>number</i>	Weight value assigned to a real server in a server farm. This value is used in the weighted round-robin and least-connections predictor load-balancing algorithms. Enter an integer from 1 to 100. The default is 8.
--------------------	---------------	--

Command Modes	Serverfarm host real server configuration mode Admin and user contexts
---------------	---

Command History	ACE Module Release	Modification
	3.0(0)A1(2)	This command was introduced.
	ACE Appliance Release	Modification
	A1(7)	This command was introduced.

Usage Guidelines	Servers with higher weight values receive a proportionally higher number of connections than servers with lower weight values. If you do not specify a weight in serverfarm host real server configuration mode, the ACE uses the weight that you configured for the global real server in real server host configuration mode.
------------------	---

To specify different weight values for a host real server in a server farm, you can assign multiple IP addresses to the server. You can also use the same IP address of a real server with different port numbers.

Server weights take effect only when there are open connections to the servers. When there are no sustained connections to any of the servers, the leastconns predictor method behaves like the roundrobin method.

Examples

To configure a weight value for a real server, enter:

```
host1/Admin(config-sfarm-host-rs) # weight 50
```

To reset the weight of a real server to the default of 8, enter:

```
host1/Admin(config-sfarm-host-rs) # no weight
```

Related Commands

([config-rserver-host](#)) **weight**

([config-sfarm-host](#)) **predictor**

Server Farm Redirect Configuration Mode Commands

Serverfarm redirect configuration mode commands allow you to create and configure redirect server farms and associate redirect real servers with the server farm. Redirect server farms are clusters of real servers that redirect users to alternative URLs where content has been moved, either temporarily or permanently. The server farm consists only of real servers that redirect client requests to alternative locations specified by the relocation string or port number in the real server configuration. You must configure a redirect real server using the ([config rserver redirect](#)) command in configuration mode before you can associate it with a server farm.

To create a redirect server farm and access serverfarm redirect configuration mode, use the **serverfarm redirect** command. The CLI prompt changes to (config-sfarm-redirect). For information about the commands in this mode, see the following commands.

Use the **no** form of this command to remove a server farm from the configuration.

```
serverfarm redirect name
```

```
no serverfarm redirect name
```

Syntax Description

<i>name</i>	Unique identifier of the server farm. Enter an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.
-------------	---

Command Modes

Configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.

ACE Appliance Release	Modification
A1(7)	This command was introduced.

Usage Guidelines

The commands in this mode require the server-farm feature in your user role. For details about role-based access control (RBAC) and user roles, see the *Virtualization Guide, Cisco ACE Application Control Engine*.

Examples

To create a redirect server farm named SFARM2, enter:

```
host1/Admin(config)# serverfarm redirect SFARM2
host1/Admin(config-sfarm-redirect)#
```

To delete the redirect server farm named SFARM2, enter:

```
host1/Admin(config)# no serverfarm redirect SFARM2
```

Related Commands

[show serverfarm](#)
[show running-config](#)
[\(config\) rserver](#)

(config-sfarm-redirect) description

To configure the text description of a server farm, use the **description** command. Use the **no** form of this command to delete the description of a server farm.

description *text*

no description

Syntax Description	<i>text</i>	Text description of a server farm. Enter an unquoted text string with a maximum of 240 alphanumeric characters.
--------------------	-------------	---

Command Modes	Serverfarm redirect configuration mode Admin and user contexts
---------------	---

Command History	ACE Module Release	Modification
	3.0(0)A1(2)	This command was introduced.

Command History	ACE Appliance Release	Modification
	A1(7)	This command was introduced.

Usage Guidelines	This command has no usage guidelines.
------------------	---------------------------------------

Examples	<p>To configure a description of a server farm, enter:</p> <pre>host1/Admin(config-sfarm-redirect)# description REDIRECT_NEW_SITE</pre> <p>To delete the description of a server farm, enter:</p> <pre>host1/Admin(config-sfarm-redirect)# no description</pre>
----------	---

Related Commands	This command has no related commands.
------------------	---------------------------------------

(config-sfarm-redirect) failaction

To configure the action that the ACE takes if a real server in a server farm goes down, use the **failaction** command. Use the **no** form of this command to reset the ACE to its default of taking no action when a server fails.

```
failaction {purge | reassign [across-interface]}
```

```
no failaction
```

Syntax Description

purge	Specifies that the ACE removes the connections to a real server in the server farm if that real server fails. The ACE sends a reset (RST) both to the client and to the server that failed.
reassign	Specifies that the ACE reassigns existing server connections to the backup real server if a backup real server is configured. If no backup real server is configured, this keyword has no effect.
across-interface	(Optional) Instructs the ACE to reassign all connections from the failed real server to a backup real server on a different VLAN that is commonly referred to as a bypass VLAN. By default, this feature is disabled.

Command Modes

Serverfarm redirect configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.
A2(1.0)	This command was revised (reassign keyword added).
A2(3.0)	The across-interface option was added.
A5(1.0)	Added IPv6 support.

ACE Appliance Release	Modification
A1(7)	This command was introduced.
A3(1.0)	This command was revised (reassign keyword added).
A4(1.0)	The across-interface option was added.
A5(1.0)	Added IPv6 support.

Usage Guidelines

If you do not configure this command, the ACE takes the real server out of rotation for new connections and allows existing connections to complete. The ACE does not send the connections to a backup server in the server farm or to a backup server farm if all servers in the primary server farm fail.

This feature is required for stateful firewall load balancing (FWLB). For details about FWLB, see the *Server Load-Balancing Guide, Cisco ACE Application Control Engine*.

Examples

To instruct the ACE to remove connections from a failed server in the server farm, enter:

```
host1/Admin(config-sfarm-redirect)# failaction purge
```

To reset the ACE to its default of taking no action if a real server fails, enter:

```
host1/Admin(config-sfarm-redirect)# no failaction
```

Related Commands

This command has no related commands.

(config-sfarm-redirect) predictor

To configure the load-balancing algorithm for the server farm, use the **predictor** command. Use the **no** form of this command to revert to the default load-balancing algorithm (the round-robin algorithm).

```
predictor {hash {address [destination | source] [v6-prefix prefix-length | netmask]} | {content
  [offset number1] [length number2] [begin-pattern expression1] [end-pattern expression2]} |
  {cookie [secondary] name1} | {header name2} | {layer4-payload [offset number3] [length
  number4] [begin-pattern expression3] [end-pattern expression4]} | {url [begin-pattern
  expression5] [end-pattern expression6]} } | {least-bandwidth [samples number5]
  [assess-time seconds]} | {least-loaded probe name3 [samples number6]} | {leastconns
  [slowstart seconds]} | {response {app-req-to-resp | syn-to-close | syn-to-synack} [samples
  number7] [threshold milliseconds] [resume-timer seconds]} } | {roundrobin}
```

no predictor

Syntax Description		
hash address		Selects the server using a hash value based on the source and destination IP addresses. Use the hash address source and hash address destination methods for firewall load balancing (FWLB).
destination		(Optional) Selects the server using a hash value based on the destination IP address.
source		(Optional) Selects the server using a hash value based on the source IP address.
v6-prefix <i>prefix-length</i>		(Optional) Specifies how many of the most significant bits (MSBs) of the IPv6 address are used for the network identifier. Enter an integer from 1 to 128.
<i>netmask</i>		(Optional) Bits in the IP address to use for the hash. If not specified, the default is 255.255.255.255.
hash content		Selects the server using a hash value based on the specified content string of the HTTP packet body.
offset <i>number1</i>		(Optional) Specifies 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. Enter an integer from 0 to 999. The default is 0, which indicates that the ACE does not exclude any portion of the content.
length <i>number2</i>		(Optional) Specifies the length 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. Enter an integer from 1 to 1000. The default is the entire payload. 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. You cannot specify both the length and the end-pattern options in the same hash content command.

begin-pattern <i>expression1</i>	<p>(Optional) Specifies 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 immediately 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>Enter an unquoted text string with no spaces and a maximum of 255 alphanumeric characters for each pattern that you configure. You can enter a text string with spaces if you enclose the entire string in quotation marks (“”). The ACE supports the use of regular expressions for matching string expressions.</p> <p>When matching data strings, note that the period (.) and question mark (?) characters do not have a literal meaning in regular expressions. Use brackets ([]) to match these symbols (for example, enter www[.]xyz[.]com instead of www.xyz.com). You can also use a backslash (\) to escape a dot (.) or a question mark (?).</p>
end-pattern <i>expression2</i>	<p>(Optional) Specifies 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>Enter an unquoted text string with no spaces and a maximum of 255 alphanumeric characters for each pattern that you configure. You can enter a text string with spaces if you enclose the entire string in quotation marks (“”). The ACE supports the use of regular expressions for matching string expressions.</p> <p>You cannot specify both the length and the end-pattern options in the same hash content command.</p>
hash cookie	<p>Selects the server using a hash value based on the cookie name or based on the name in the cookie name of the URL query string.</p>
secondary	<p>(Optional) Selects the server by using the hash value based on the specified name in the cookie name in the URL query string, not the cookie header. If you do not include this option, the ACE selects a real server using the hash value of the cookie name.</p>
<i>name1</i>	<p>Cookie name. Enter a cookie name as an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.</p>

hash header <i>name2</i>	<p>Selects the server using a hash value based on the header name. Enter a header name as an unquoted text string with no spaces and a maximum of 64 alphanumeric characters, or enter one of the following standard headers:</p> <ul style="list-style-type: none"> • Accept • Accept-Charset • Accept-Encoding • Accept-Language • Authorization • Cache-Control • Connection • Content-MD5 • Expect • From • Host • If-Match • Pragma • Referrer • Transfer-Encoding • User-Agent • Via
hash layer4-payload	<p>Specifies 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>
offset <i>number3</i>	<p>(Optional) Specifies the portion of the payload 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. Enter an integer from 0 to 999. The default is 0, which indicates that the ACE does not exclude any portion of the payload.</p>
length <i>number4</i>	<p>(Optional) Specifies the length 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. Enter an integer from 1 to 1000. The default is the entire payload.</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>You cannot specify both the length and the end-pattern options in the same hash layer4-payload command.</p>

begin-pattern <i>expression3</i>	<p>(Optional) Specifies 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 immediately 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>Enter an unquoted text string with no spaces and a maximum of 255 alphanumeric characters for each pattern that you configure. You can enter a text string with spaces if you enclose the entire string in quotation marks (“”). The ACE supports the use of regular expressions for matching string expressions.</p> <p>When matching data strings, note that the period (.) and question mark (?) characters do not have a literal meaning in regular expressions. Use brackets ([]) to match these symbols (for example, enter www[.]xyz[.]com instead of www.xyz.com). You can also use a backslash (\) to escape a dot (.) or a question mark (?).</p>
end-pattern <i>expression4</i>	<p>(Optional) Specifies 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>Enter an unquoted text string with no spaces and a maximum of 255 alphanumeric characters for each pattern that you configure. You can enter a text string with spaces if you enclose the entire string in quotation marks (“”). The ACE supports the use of regular expressions for matching string expressions.</p> <p>Note: You cannot specify both the length and the end-pattern options in the same hash layer4-payload command.</p>
hash url	<p>Selects the server using a hash value based on the requested URL. Use this predictor method to load balance cache servers. Cache servers perform better with the URL hash method because you can divide the contents of the caches evenly if the traffic is random enough. In a redundant configuration, the cache servers continue to work even if the active ACE switches over to the standby ACE. For information about configuring redundancy, see the <i>Administration Guide, Cisco ACE Application Control Engine</i>.</p>
begin-pattern <i>expression5</i>	<p>(Optional) Specifies the beginning pattern of the URL and the pattern string to match before hashing. You cannot configure different beginning and ending patterns for different server farms that are part of the same traffic classification. Enter an unquoted text string with no spaces and a maximum of 255 alphanumeric characters for each pattern that you configure. If you want to match a URL that contains spaces, you must use \x20 for each space character.</p>

end-pattern <i>expression6</i>	(Optional) Specifies the pattern that marks the end of hashing. You cannot configure different beginning and ending patterns for different server farms that are part of the same traffic classification. Enter an unquoted text string with no spaces and a maximum of 255 alphanumeric characters for each pattern that you configure. If you want to match a URL that contains spaces, you must use <code>\x20</code> for each space character.
least-bandwidth	Selects the server that processed the least amount of network traffic over a specified sampling period. Use this predictor for heavy traffic use, such as downloading a video clip. The ACE measures traffic statistics between itself and the real servers in the server farm in both directions and calculates the bandwidth over the sampling period. Then, it creates an ordered list of real servers based on the sampling results and selects the server that used the least amount of bandwidth during the sampling period.
samples <i>number5</i>	(Optional) Specifies the number of samples over which you want to weight and average the results of the probe query to calculate the final load value. Enter an integer from 1 to 16. Each value must be a power of 2, so the valid values are as follows: 1, 2, 4, 8, and 16. The default is 8.
assess-time <i>seconds</i>	(Optional) Specifies the sampling period over which the ACE measures traffic for all the servers in the server farm. Enter an integer from 1 to 10. The default is 4 seconds.
least-loaded	Selects the server with the lowest load based on information obtained from SNMP probes. To use this predictor, you must associate an SNMP probe with the server farm. The ACE queries one user-specified OID (for example, CPU utilization or memory utilization). The ACE uses the retrieved value directly to determine the server with the lowest load. This predictor is not supported with IPv6.
probe <i>name3</i>	Specifies the name of the SNMP probe that you want to query. Enter an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.
samples <i>number6</i>	(Optional) Specifies the number of samples over which you want to weight and average the results of the probe query to calculate the final load value. Enter an integer from 1 to 16. Each value must be a power of 2, so the valid values are as follows: 1, 2, 4, 8, and 16. The default is 8.
leastconns	Selects the real server with the fewest number of active connections based on the server weight. Use this predictor for processing light user requests (for example, browsing simple static web pages). For information about setting real server weight, see the (config-sfarm-redirect-rs) weight section.
slowstart <i>seconds</i>	(Optional) Specifies that the connections to the real server be in a slow-start mode for the duration indicated by the <i>seconds</i> value. Use the slow-start mechanism to avoid sending a high rate of new connections to servers that you have recently put into service. Enter an integer from 1 to 65535, where 1 is the slowest ramp-up value. By default, slowstart is disabled.

response	Selects the server with the lowest response time for the requested response-time measurement. If you do not specify a response-time measurement method, the ACE uses the HTTP app-req-to-response method.
app-req-to-resp	(Default) Measures 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. The ACE does not allow you to configure this predictor response in a generic load-balancing policy map.
syn-to-close	Measures 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	Measures the response time from when the ACE sends a TCP SYN to a server to the time that the ACE receives the SYN-ACK from the server.
samples <i>number</i>	(Optional) Specifies the number of samples that you want to average from the results of the response time measurement. Enter an integer from 1 to 16 in powers of 2. Valid values are 1, 2, 4, 8, and 16. The default is 8.
threshold <i>milliseconds</i>	(Optional) Specifies the required minimum average response time for a server. If the server response time is greater than the specified threshold value, the ACE removes the server from the load-balancing decision process (takes the server out of service). Enter an integer from 1 to 300000 milliseconds (5 minutes). The default is no threshold (servers are not taken out of service).
resume-timer <i>seconds</i>	(Optional) Specifies the number of seconds after which the ACE sends traffic again to a server that was taken out of the load-balancing decision process. The ACE monitors the server's response time. If that response time is less than or equal to the value set with the threshold keyword, the ACE places the server back in service. Enter an integer from 30 to 3600 seconds (1 hour). The default value is 300 seconds (5 minutes) if you configure a threshold without specifying the resume timer.
roundrobin	(Default) Selects the next server in the list of real servers based on server weight (weighted round-robin). For information about setting the real server weight, see the (config-sfarm-redirect-rs) weight section.

Command Modes

Server-farm redirect configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.
A2(1.0)	This command was revised.
A2(1.4) and A2(2.1)	This secondary option for the hash cookie keywords was added.

ACE Appliance Release	Modification
A1(7)	This command was introduced.
A3(1.0)	This command was revised (reassign keyword added).
A3(2.2)	This secondary option for the hash cookie keywords was added.

Usage Guidelines

Use this command to specify the load-balancing algorithm that the ACE uses in choosing a real server in the server farm. If you do not specify the **predictor** command, the default algorithm is **roundrobin**. Using the **no** form of this command changes the configured predictor algorithm to the default algorithm.

The weight assigned to the real servers is used only in the **roundrobin** and **leastconns** predictor methods. The **hash** and the **response** predictor methods do not recognize the weight for the real servers. For information about setting the real server weight, see the [\(config-sfarm-redirect-rs\) weight](#) section.

If you configure the **leastconns** predictor, you can use a **slowstart** mechanism (ramp-up) to avoid sending a high rate of new connections to the servers that have just been put in service. The real server with the fewest number of active connections will get the next connection request for the server farm with the **leastconns** predictor. The ramp-up stops when the duration timer that you specify expires.

The only time that the sequence of servers starts over at the beginning (with the first server) is when there is a configuration or server state change (for example, a probe failure).

The **secondary** option allows the ACE to correctly load balance in cases when the query string identifies the actual resource, instead of the URL.

Examples

To configure the ACE to select the real server with the lowest number of connections in the server farm, enter:

```
host1/Admin(config-sfarm-redirect)# predictor leastconns slowstart 300
```

To reset the load-balancing algorithm to the default round-robin algorithm, enter:

```
host1/Admin(config-sfarm-redirect)# no predictor
```

Related Commands

[\(config-sfarm-redirect-rs\) weight](#)

(config-sfarm-redirect) probe

Use probes to monitor the health of real servers in a server farm. To associate a probe with a server farm, use the **probe** command. Use the **no** form of this command to dissociate a probe from a server farm.

probe *probe-name*

no probe *probe-name*

Syntax Description

<i>probe-name</i>	Identifier of an existing probe that you want to associate with a server farm. Enter an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.
-------------------	--

Command Modes Serverfarm redirect configuration mode
Admin and user contexts

Command History	ACE Module Release	Modification
	A2(3.2). Not applicable for A4(1.0).	This command was introduced.
Command History	ACE Appliance Release	Modification
	A3(2.7). Not applicable for A4(1.0).	This command was introduced.

Usage Guidelines The probe must already exist. (To create a probe, see the [\(config\) probe](#) command.) You can associate multiple probes of the same or different protocols with each server farm.
You can only configure probes with an IP address in routed mode under a redirect server.
You cannot associate a scripted probe with a redirect server.

Examples To associate a probe with a server farm, enter:

```
host1/Admin(config-sfarm-redirect)# probe TCP1
```

 To dissociate a probe from a server farm, enter:

```
host1/Admin(config-sfarm-redirect)# no probe TCP1
```

Related Commands [\(config\) probe](#)
[\(config-probe-probe_type\) ip address](#)

(config-sfarm-redirect) rserver

To associate one or more existing redirect real servers with a server farm and access serverfarm redirect real server configuration mode, use the **rserver** command. The CLI prompt changes to (config-sfarm-redirect-rs). For information on commands in serverfarm redirect real server configuration mode, see the [“Server Farm Redirect Real Server Configuration Mode Commands”](#) section. Use the **no** form of this command to dissociate the real server from the server farm.

rserver *name* [*port*]

no rserver *name* [*port*]

Syntax Description	<i>name</i>	Unique identifier of the real server. Enter an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.
	<i>port</i>	(Optional) Port number used for the real server Port Address Translation (PAT). Enter an integer from 1 to 65535.

Command Modes Serverfarm redirect configuration mode
Admin and user contexts

Command History	ACE Module Release	Modification
		3.0(0)A1(2)

Command History	ACE Appliance Release	Modification
		A1(7)

Usage Guidelines The real server must already exist. To create a real server, see the [\(config\) rserver](#) command. You can associate a maximum of 16,384 real servers with a server farm.

Examples

To associate a real server with a server farm, enter:

```
host1/Admin(config-sfarm-redirect)# rserver server1 4000
host1/Admin(config-sfarm-redirect-rs)#
```

To dissociate a real server from a server farm, enter:

```
host1/Admin(config-sfarm-redirect)# no rserver server1
host1/Admin(config-sfarm-redirect)#
```

Related Commands [\(config\) rserver](#)

Serverfarm Redirect Predictor Configuration Mode Commands

Serverfarm redirect predictor configuration mode commands allow you to configure additional parameters for some of the server farm predictor methods.

To configure these additional predictor parameters, use the **predictor least-loaded** or the **predictor response** command in serverfarm host configuration mode. The CLI prompt changes to (config-sfarm-host-predictor). For information about the commands in this mode, see the following commands. Use the **no** form of this command to remove the predictor from the server farm.

```
predictor {least-loaded probe name} | {response {app-req-to-resp | syn-to-close |
syn-to-synack} [samples number]}}
```

```
no predictor
```

Syntax Description

least-loaded	Selects the server with the lowest load based on information obtained from SNMP probes. To use this predictor, you must associate an SNMP probe with the server farm. The ACE queries one user-specified OID (for example, CPU utilization or memory utilization). The ACE uses the retrieved value directly to determine the server with the lowest load.
probe <i>name</i>	Specifies the name of the SNMP probe that you want to query. Enter an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.
response	Selects the server with the lowest response time for the requested response-time measurement. If you do not specify a response-time measurement method, the ACE uses the HTTP app-req-to-response method.
app-req-to-resp	(Default) Measures 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. The ACE does not allow you to configure this predictor response in a generic load-balancing policy map.
syn-to-close	Measures 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	Measures the response time from when the ACE sends a TCP SYN to a server to the time that the ACE receives the SYN-ACK from the server.
samples <i>number</i>	(Optional) Number of samples over which you want to average the results of the response time measurement. Enter an integer from 1 to 16 in powers of 2. Valid values are: 1, 2, 4, 8, and 16. The default is 8.

Command Modes

Serverfarm redirect configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
A2(1.0)	This command was introduced.

ACE Appliance Release	Modification
A3(1.0)	This command was introduced.

Usage Guidelines

The commands in this mode require the server-farm feature in your user role. For details about role-based access control (RBAC) and user roles, see the *Virtualization Guide, Cisco ACE Application Control Engine*.

Examples

To specify the least-loaded predictor method with a probe called SNMP_PROBE for the server farm, enter:

```
host1/Admin(config-sfarm-redirect)# predictor least-loaded probe SNMP_PROBE
host1/Admin(config-sfarm-redirect-predictor)#
```

To remove the least-loaded predictor from the server farm, enter:

```
host1/Admin(config-sfarm-redirect)# no predictor
```

To specify the response predictor method that measures 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, enter:

```
host1/Admin(config-sfarm-redirect)# predictor response app-req-to-resp
host1/Admin(config-sfarm-redirect-predictor)#
```

To remove the response predictor from the server farm, enter:

```
host1/Admin(config-sfarm-redirect)# no predictor
```

Related Commands

[show serverfarm detail](#)
[\(config-sfarm-redirect\) predictor](#)
[\(config-sfarm-redirect-predictor\) autoadjust](#)
[\(config-sfarm-redirect-predictor\) weight connection](#)

(config-sfarm-redirect-predictor) autoadjust

After you specify the **predictor least-loaded** command, use the **autoadjust** command to instruct the ACE to apply the maximum load of 16000 to a real server whose load reaches zero or override the default behavior. Use the **no** form of this command to return the ACE behavior to the reset the behavior of the ACE to the default of average load of the server farm to a real server whose load is zero.

```
autoadjust { average | maxload | off }
```

```
no autoadjust
```

Syntax Description

average	Applies the average load of the server farm to a real server whose load is zero. This setting allows the server to participate in load balancing, while preventing it from being flooded by new connections. This is the default setting.
maxload	Instructs the ACE to apply the maximum load of 16000 to a real server whose load reaches zero.
off	Overrides the default behavior of the ACE of applying the average load of the server farm to a real server whose load is zero. When you configure this command, the ACE sends 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. If two servers have the same lowest load (either zero or nonzero), the ACE load balances the connections between the two servers in a round-robin manner.

Serverfarm redirect predictor configuration mode

Admin and user contexts

Command History

ACE Module Release	Modification
A2(1.0)	This command was introduced.
A2(2.4)	The average keyword became the default autoadjust setting for the least-loaded predictor. Previously, the default setting was maximum load. The maxload keyword was added to set the least-loaded predictor to maximum load.
ACE Appliance Release	Modification
A3(1.0)	This command was introduced.
A4(1.0)	The average keyword became the default autoadjust setting for the least-loaded predictor. Previously, the default setting was maximum load. The maxload keyword was added to set the least-loaded predictor to maximum load.

Usage Guidelines

Whenever a server's load reaches zero, by default, the ACE uses the `autoadjust` feature to assign an average load value 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.

Using the least-loaded predictor with the configured server weight and the current connection count option enabled, the ACE calculates the final load of a real server as follows:

$$\text{final load} = \text{weighted load} \times \text{static weight} \times \text{current connection count}$$

where:

- *weighted load* is the load reported by the SNMP probe
- *static weight* is the configured weight of the real server
- *current connection count* is the total number of active connections to the real server

The ACE recalculates the final load whenever the connection count changes, provided that the **weight connection** command is configured. If the **weight connection** command is not configured, the ACE updates the final load when the next load update arrives from the SNMP probe.

Examples

To instruct the ACE to apply the maximum load of the server farm to a real server whose load value reaches zero, enter:

```
host1/Admin(config-sfarm-redirect-predictor)# autoadjust maxload
```

To turn off the `autoadjust` feature for all servers in a server farm so that servers with a load of zero receive all new connections, enter:

```
host1/Admin(config-sfarm-redirect-predictor)# autoadjust off
```

To reset the behavior of the ACE to the default of applying the average load value of 16000 to a real server whose load is zero, enter:

```
host1/Admin(config-sfarm-redirect-predictor)# no autoadjust
```

Related Commands

[show serverfarm detail](#)
[\(config-sfarm-redirect\) predictor](#)
[\(config-sfarm-redirect-predictor\) weight connection](#)

(config-sfarm-redirect-predictor) weight connection

After you specify the **predictor least-loaded** or the **predictor response** command, use the **weight connection** command to instruct the ACE to use the current connection count in the final load calculation for each real server in the server farm. Use the **no** form of this command to reset the behavior of the ACE to the default of excluding the current connection count from the load calculation.

weight connection

no weight connection

Syntax Description This command has no keywords or arguments.

Command Modes Serverfarm redirect predictor configuration mode
Admin and user contexts

Command History	ACE Module Release	Modification
	A2(1.0)	This command was introduced.

Command History	ACE Appliance Release	Modification
	A3(1.0)	This command was introduced.

Usage Guidelines To see how the **weight connection** command affects the **(config-sfarm-redirect-predictor) autoadjust** command for the least-loaded predictor, see the Usage Guidelines section of the **(config-sfarm-redirect-predictor) autoadjust** command.

Examples To instruct the ACE to use the current connection count in the final load calculation for each real server in the server farm, enter:

```
host1/Admin(config-sfarm-redirect-predictor)# weight connection
```

To reset the behavior of the ACE to the default of excluding the current connection count from the load calculation, enter:

```
host1/Admin(config-sfarm-redirect-predictor)# no weight connection
```

Related Commands [show serverfarm detail](#)
[\(config-sfarm-redirect\) predictor](#)
[\(config-sfarm-redirect-predictor\) autoadjust](#)

Server Farm Redirect Real Server Configuration Mode Commands

Serverfarm redirect real server configuration mode commands allow you to associate a redirect real server with a redirect server farm and configure the real server attributes.

To associate one or more existing redirect real servers with a redirect server farm and access serverfarm redirect real server configuration mode, use the **rserver** command in serverfarm redirect configuration mode. The CLI prompt changes to (config-sfarm-redirect-rs). For information about the commands in this mode, see the following commands. Use the **no** form of this command to remove the real server from the server farm.

rserver *name*

no rserver *name*

Syntax Description

<i>name</i>	Unique identifier of the real server. Enter an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.
-------------	---

Command Modes

Serverfarm redirect configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.

ACE Appliance Release	Modification
A1(7)	This command was introduced.

Usage Guidelines

The commands in this mode require the server-farm feature in your user role unless otherwise specified. For details about role-based access control (RBAC) and user roles, see the *Virtualization Guide, Cisco ACE Application Control Engine*.

The redirect real server must already exist. To create a real server, see the **(config) rserver redirect** command. You can associate a maximum of 16,384 real servers with a server farm.

Examples

To associate a real server with a server farm, enter:

```
host1/Admin(config-sfarm-redirect)# rserver server1
```

To dissociate a real server from a server farm, enter:

```
host1/Admin(config-sfarm-redirect)# no rserver server1
```


Related Commands This command has no related commands.

(config-sfarm-redirect-rs) backup-rserver

To configure a backup real server for a real server in a server farm, use the **backup-rserver** command. If a real server associated with a server farm becomes unavailable, the ACE directs flows to the configured backup real server. Use the **no** form of this command to remove a backup real server from the configuration.

backup-rserver *name*

no backup-rserver

Syntax Description

<i>name</i>	Unique identifier of an existing real server that you want to configure as a backup server in a server farm. Enter an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.
-------------	--

Command Modes

Serverfarm redirect real server configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.
A5(1.0)	Added support for IPv6.

ACE Appliance Release	Modification
A1(7)	This command was introduced.
A3(2.2)	This command supports cyclic backup of real servers in a server farm.
A5(1.0)	Added support for IPv6.

Usage Guidelines

The real server used as a backup server must already exist. To create a redirect real server, see the [\(config\) rserver redirect](#) command.

IPv6 servers can back up IPv4 servers, but only for the HTTP and the HTTPS protocols.

Examples

To associate a backup real server with a server farm, enter:

```
host1/Admin(config-sfarm-redirect-rs)# backup-rserver BACKUP_SERVER1
```

To dissociate a backup real server from a server farm, enter:

```
host1/Admin(config-sfarm-redirect-rs)# no backup-rserver
```

Related Commands

[\(config\) rserver](#)

(config-sfarm-redirect-rs) conn-limit

To configure the maximum and minimum number of connections that you want to allow for a redirect real server in a server farm, use the **conn-limit** command. Use the **no** form of this command to reset the real server maximum connections and minimum connections threshold to the default of 4000000.

conn-limit **max** *maxconns* **min** *minconns*

no conn-limit

Syntax Description	max <i>maxconns</i>	min <i>minconns</i>
	Specifies the maximum number of connections allowed for this real server. Enter an integer from 2 to 4000000. The default is 4000000.	Specifies the connection threshold below which the real server will start accepting connections again after the number of connections exceeds the configured maximum number of connections. Enter an integer from 2 to 4000000. The default is <i>minconns</i> equal to <i>maxconns</i> .

Command Modes
Serverfarm redirect real server configuration mode
Admin and user contexts

Command History	ACE Module Release	Modification
	3.0(0)A1(2)	This command was introduced.
	ACE Appliance Release	Modification
	A1(7)	This command was introduced.

Usage Guidelines
Use this command to specify the maximum number of connections and the minimum connection threshold for a redirect real server in a server farm. The *minconns* value must be less than or equal to the *maxconns* value. The ACE uses the *minconns* value as a threshold to start accepting connections again after the *maxconns* limit is exceeded.

Examples
To configure the maximum number of connections and the minimum connection threshold for a redirect real server, enter:

```
host1/Admin(config-sfarm-redirect-rs)# conn-limit max 65535 min 40000
```

To reset the maximum number of connections and the minimum connection threshold for a redirect real server to the default of 4000000, enter:

```
host1/Admin(config-sfarm-redirect-rs)# no conn-limit
```

Related Commands [\(config-sfarm-redirect-rs\) rate-limit](#)

(config-sfarm-redirect-rs) inservice

To place a real server associated with a server farm in service, use the **inservice** command. Use the **no** form of this command to take a real server out of service.

inservice [standby]

no inservice

Syntax Description

standby	(Optional) Used with backup real servers, specifies that a backup real server remain inactive unless the primary real server fails. If the primary fails, the backup server becomes active and starts accepting connections.
----------------	--

Command Modes

Serverfarm redirect real server configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
3.0(0)A1(2)	This command was introduced.

ACE Appliance Release	Modification
A1(7)	This command was introduced.

Usage Guidelines

This command requires the real-inservice feature in your user role. For details about role-based access control (RBAC) and user roles, see the *Virtualization Guide, Cisco ACE Application Control Engine*.

To start load-balancing connections to a real server in a server farm, you must place the real server in service by using the **inservice** command.

You can modify the attributes of a real server in a server farm without taking the server out of service.

In addition to putting a backup real server in service standby, another use of the **inservice standby** command is to provide the graceful shutdown of primary real servers. Use this command to gracefully shut down servers with sticky connections. When you enter this command for a primary real server, the ACE does the following:

- Tears down existing non-TCP connections to the server
- Allows current TCP connections to complete
- Allows new sticky connections for existing server connections that match entries in the sticky database
- Load balances all new connections (other than the matching sticky connections mentioned above) to the other servers in the server farm
- Eventually takes the server out of service

Examples

To place a real server in service, enter:

```
host1/Admin(config-sfarm-redirect-rs)# inservice
```

To perform a graceful shutdown on a primary real server with sticky connections in a server farm, enter:

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

To take a real server out of service, enter:

```
host1/Admin(config-sfarm-redirect-rs)# no inservice
```

Related Commands

This command has no related commands.

(config-sfarm-host-rs) probe

To configure a probe to monitor the health of a redirect real server in a redirect server farm, use the **probe** command. Use the **no** form of this command to remove the probe from the real server.

probe *probe-name*

no probe *probe-name*

Syntax Description

<i>probe-name</i>	Identifier of an existing probe that you want to assign to a real server to monitor its health. Enter an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.
-------------------	---

Command Modes

Serverfarm redirect real server configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
A2(3.2). Not applicable for A4(1.0).	This command was introduced.

ACE Appliance Release	Modification
A3(2.7). Not applicable for A4(1.0).	This command was introduced.

Usage Guidelines

You can associate multiple probes with each real server.

The ACE periodically sends the probes to the real servers. If the ACE receives the appropriate responses from the servers, the ACE includes the servers in load-balancing decisions. If not, the ACE marks the servers as out of service, depending on the configured number of retries.

You can only configure probes with an IP address in routed mode under a redirect server.

You cannot associate a scripted probe with a redirect server.

Examples

To configure a probe for a redirect real server, enter:

```
host1/Admin(config-sfarm-host-rs)# probe SERVER1_PROBE
```

To remove a probe from a redirect real server, enter:

```
host1/Admin(config-sfarm-host-rs)# no probe SERVER1_PROBE
```

Related Commands

[\(config-probe-probe_type\) ip address](#)

(config-sfarm-redirect-rs) rate-limit

To configure a limit for the connection rate and the bandwidth rate of a real server in a redirect server farm, use the **rate-limit** command. The connection rate is the number of connections per second received by the ACE and destined to a particular redirect real server. The bandwidth rate is the number of bytes per second received by the ACE and destined for a particular redirect real server. Use the **no** form of this command to revert to the ACE default of not limiting the connection rate or bandwidth rate of real servers in a server farm.

```
rate-limit { connection number1 | bandwidth number2 }
```

```
no rate-limit { connection | bandwidth }
```

Syntax Description

connection <i>number1</i>	Specifies the real server connection-rate limit in connections per second. Enter an integer from 2 to 350000. There is no default value.
bandwidth <i>number2</i>	Specifies the real server bandwidth-rate limit in bytes per second. Enter an integer from 2 to 300000000. There is no default value.

Command Modes

Serverfarm redirect real server configuration mode
Admin and user contexts

Command History

ACE Module Release	Modification
A2(1.0)	This command was introduced.

ACE Appliance Release	Modification
A3(1.0)	This command was introduced.

Usage Guidelines

If the connection rate or the bandwidth rate of incoming traffic destined for a particular real server exceeds the configured rate for the server, the ACE blocks any further traffic destined to that real server until the connection rate or bandwidth rate drops below the configured limit. Also, the ACE removes the blocked real server from future load-balancing decisions. By default, the ACE does not limit the connection rate or the bandwidth rate of real servers in a server farm.

Examples

To limit the connection rate of a real server to 100,000 connections per second, enter:

```
host1/Admin(config-sfarm-redir-rs)# rate-limit connection 100000
```

To revert to the ACE default of not limiting the real-server connection rate, enter:

```
host1/Admin(config-sfarm-redir-rs)# no rate-limit connection
```

To limit the real-server bandwidth rate to 5,000,000 bytes per second, enter:

```
host1/Admin(config-sfarm-redir-rs)# rate-limit bandwidth 5000000
```

To revert to the ACE default of not limiting real-server bandwidth, enter:

```
host1/Admin(config-sfarm-redirect-rs)# no rate-limit bandwidth
```

Related Commands [\(config-sfarm-redirect-rs\) conn-limit](#)

(config-sfarm-redirect-rs) weight

To configure the capacity of a real server in relation to other servers in a server farm, use the **weight** command. The weight value that you specify for a server is used in the weighted round-robin and least-connections predictor load-balancing methods. Use the **no** form of this command to reset the real server weight to the default.

weight *number*

no weight

Syntax Description	<i>number</i>	Weight value assigned to a real server in a server farm. This value is used in the weighted round-robin and least-connections predictor load-balancing algorithms. Enter an integer from 1 to 100. The default is 8.

Command Modes	Serverfarm redirect real server configuration mode Admin and user contexts

Command History	ACE Module Release	Modification
	3.0(0)A1(2)	This command was introduced.
	ACE Appliance Release	Modification
	A1(7)	This command was introduced.

Usage Guidelines	Servers with higher weight values receive a proportionally higher number of connections than servers with lower weight values. To specify different weight values for a redirect real server in a server farm, you can assign multiple IP addresses to the server.

Examples	To configure a weight value for a real server, enter: host1/Admin(config-sfarm-redirect-rs)# weight 50 To reset the weight of a real server to the default of 8, enter: host1/Admin(config-sfarm-redirect-rs)# no weight

Related Commands [\(config-sfarm-redirect\) predictor](#)