Global Configuration Mode Commands

Global configuration mode allows a SuperUser to:

- Configure global CSS parameters.
- Initially access subordinate configuration modes on the CSS. These modes allow you to configure ACLs, boot, circuits and their IP interface addresses, EQLs, physical interfaces, global keepalives, source groups, owners and their content rules, RMON alarm, events and history, and services.

To access global configuration mode, use the **configure** command in SuperUser mode.

This section describes the commands in global configuration mode. For more information on commands for the subordinate configuration modes available on the CSS, see their sections later in this chapter.

For a list of general commands you can use in global configuration mode, see the "General Commands" section.

(config) acl

To access ACL configuration mode, configure an access control list (ACL) on the CSS, and enable or disable all ACLs on the CSS, use the **acl** command. Use the **no** form of this command to delete an ACL.

acl [*index*|**enable**|**disable**]

no acl *index*

Syntax Description

index	Number you want to use to create a new ACL or the number for an existing ACL to access ACL mode. Enter a number from 1 to 99. When you access this mode, the prompt changes to (config-acl [*index*]). For information about commands available in this mode, see the "ACL Configuration Mode Commands" section.

disable	Disables all ACLs on the CSS.
enable	Enables all ACLs on the CSS.

Usage Guidelines

To enable global logging for ACLs, you must enter the **(config) logging subsystem acl level debug-7** command.

Caution When you enable ACL mode, all traffic not configured in an ACL permit clause *will be denied*. ACLs function as a firewall security feature. You must first configure an ACL to permit traffic *before you enable ACL mode*. If you do not permit any traffic, you will lose network connectivity. Note that the console port is not affected.

If you do not configure ACLs on the CSS, all packets passing through the CSS could be allowed onto the entire network. For example, you may want to permit all e-mail traffic, but block Telnet traffic. You can also use ACLs to allow one client to access a part of the network and prevent another client from accessing the same area.

Related Commands

show acl
(config-acl) apply
(config-acl) clause
(config-acl) remove

(config) app

To enable all Application Peering Protocol (APP) sessions, use the **app** command. An APP session is the exchange of content information between a group of configured CSSs. APP provides a guaranteed and private communications channel for this exchange. Use the **no** form of this command to disable all APP sessions.

app

no app

Related Commands

(config) dns-server
(config-owner) dns
(config-owner-content) add dns

(config) app framesz

To set the maximum frame size allowed on an APP channel between CSSs, use the **app framesz** command. Use the **no** form of this command to restore the default frame size to 10240.

app framesz *size*

no app framesz

Syntax Description

size	Maximum frame size. Enter a number from 10240 to 65535. The default is 10240.

(config) app port

To set the TCP port number, use the **app port** command. This port listens for APP connections. Use the **no** form of this command to restore the default port number to 5001.

app port *port_number*

no app port

Syntax Description

port_number	Port number. Enter a number from 1025 to 65535. The default is 5001.

(config) app session

To create an APP session between the CSS and its peer CSS, use the **app session** command. These CSSs are a content domain that share the same content rules, load, and DNS information with each other. Use the **no** form of this command to terminate an APP session.

app session *ip_address* {*ka_freq* {[**authChallenge**|**authNone**] *secret* {[**encryptMd5hash**|**encryptNone**] {[**rcmdEnable**|**rcmdDisable**]}}}}

no app session *ip_address*

Syntax Description

ip_address	IP address for the peer CSS. Enter the address in dotted-decimal notation (for example, 192.168.11.1).
ka_freq	(Optional) Time in seconds between sending keepalive messages to the peer CSS. Enter an integer from 14 to 255. The default is 14.
authChallenge | **authNone**	(Optional) Authentication method for the session. Enter either **authChallenge** for Challenge Handshake Authentication Protocol (CHAP) method or **authNone** for no authentication method. The default is no authentication.

secret	Secret sent with each packet identifier. Enter an unquoted text string with a maximum of 32 characters. If you entered **authNone** for the authentication method, enter any character as the secret.
encryptMd5hash | **encryptNone**	(Optional) Encryption method for the packets. Enter either **encryptMd5hash** for the MD5 base hashing method or **encryptNone** for the no encryption method. The default is no encryption.
rcmdEnable | **rcmdDisable**	(Optional) Setting for sending remote CLI commands to the peer through the **rcmd** command. Enter either **rcmdEnable** to send CLI commands or **rcmdDisable** to not send CLI commands. The default setting is enabled.

Related Commands

show app
show dns-peer
show dns-server

(config) app-udp

To enable Application Peering Protocol-User Datagram Protocol (APP-UDP) datagram messaging, use the **app-udp** command. Messaging is enabled by default. An APP datagram allows an exchange of information between applications resident on the CSS. Use the **no** form of this command to disable APP-UDP messaging.

app-udp

no app-udp

Usage Guidelines

The **app-udp** command is available on a Proximity Database and a DNS CSS.

Related Commands

show app-udp

(config) app-udp options

To configure encryption with an IP address, use the **app-udp options** command. Use the **no** form of this command to delete the options from an IP address.

app-udp options *ip_address* **encrypt-md5hash** *secret*

no app-udp options *ip_address*

Syntax Description

ip_address	IP address that you want to associate with this group of options. Enter the address in dotted-decimal notation (for example, 192.168.11.1).
secret	String used in encryption and decryption of the MD5 hashing method. Enter an unquoted text string with a maximum of 31 characters. There is no default.

Usage Guidelines

The CSS applies encryption to packets sent to this destination address or when the CSS receives datagrams with a matching source IP address. You can set the IP address to 0.0.0.0 to apply encryption to all incoming and outbound datagrams that are not more specifically configured. Use of the 0.0.0.0 IP address allows you to set a global security configuration that may be applied to an arbitrary number of peers.

Examples

The following example shows the application of a specific option set to 10.6.3.21 and a global option set to all other IP addresses. The CSS encrypts datagrams received from 10.6.3.21 and transmitted to 10.6.3.21 with secret *mySecret*. The CSS subjects all other datagrams, received or transmitted, to the default encryption secret *anotherSecret*.

```
(config) # app-udp options 10.6.3.21 encrypt-md5hash mySecret
(config) # app-udp options 0.0.0.0 encrypt-md5hash anotherSecret
```

Related Commands

(config) app-udp secure

(config) app-udp port

To set the UDP port number, use the **app-udp port** command. This port listens for APP datagrams. Use the **no** form of this command to restore the UDP port number to its default value of 5002.

app-udp port *port_number*

no app-udp port

Syntax Description

port_number	UDP port number. Enter a value from 1025 to 65535. The default is 5002.

(config) app-udp secure

To require the encryption of all inbound APP datagrams, use the **app-udp secure** command. This prevents unauthorized messages from entering the CSS. Use the **no** form of this command to restore the default behavior of allowing the CSS to accept all APP datagrams.

app-udp secure

no app-udp secure

Usage Guidelines

Use the **app-udp secure** command with the **(config) app-udp options** command to specify the secure messages that are accepted. If you use this command without the **(config) app-udp options** command, the CSS drops all incoming data.

Examples

The following commands allow only incoming traffic from 10.6.3.21 encrypted with the secret "mySecret."

```
(config) # app-udp secure
(config) # app-udp options 10.6.3.21 encrypt-md5hash mySecret
```

Related Commands

(config) app-udp options

(config) arp

To define a static ARP mapping IP address to Media Access Control (MAC) address translations necessary for the CSS to send data to network nodes, use the **arp** command. Use the **no** form of this command to delete a static mapping address.

arp *ip_or_host mac_address interface* {*vlan*}

no arp *ip_or_host*

Syntax Description

ip_or_host	IP address of the system for static mapping. Enter an IP address in dotted-decimal notation (for example, 192.168.11.1) or in mnemonic host-name format (for example, myhost.mydomain.com).
mac_address	MAC address of the system mapped to the IP address. Enter the MAC address in hyphenated-hexadecimal notation (for example, 00-60-97-d5-26-ab).
interface	CSS interface that you want to configure as the egress logical port. For a CSS 11501, enter the interface name in *interface-port* format (for example, e2). For a CSS 11503 or 11506, the interface format is *slot*/*port* (for example, 3/1). To see a list of interfaces, enter: **arp** *ip_or_host mac_address* **?**
vlan	(Optional) VLAN number configured in a trunked interface on which the ARP address is configured. Enter an integer from 1 to 4094 as the VLAN number.

Usage Guidelines

To show static ARP mapping when you use the **show arp** command, the IP route must exist in the routing table. To view all static ARP entries, use the **show running-config** command.

The CSS discards ARP requests from hosts that are not on the same network as the CSS circuit IP address. Thus, if a CSS and a host are within the same VLAN but configured for different IP networks, the CSS does not respond to ARP requests from the host.

Related Commands

clear
show arp
show running-config
update arp

(config) arp timeout

To set the time in seconds to hold an ARP resolution result in the ARP cache, use the **arp timeout** command. Use the **no** form of this command to restore the default timeout value of 14400 seconds.

arp timeout *timeout_time*

no arp timeout

Syntax Description

timeout_time	Number of seconds to hold an ARP resolution result. To set a timeout period, enter an integer from 60 to 86400 (24 hours). The default is 14400 (4 hours). If you do not want the ARP entries to timeout, enter **none** or 86401.

Usage Guidelines

When you change the timeout value, it only affects new ARP entries. All previous ARP entries retain the old timeout value. To remove all entries with the old timeout value, enter the **clear arp cache** command.

Related Commands

clear arp cache
show arp config

(config) arp wait

To set the time in seconds to wait for an ARP resolution before discarding the packet waiting to be forwarded to the address, use the **arp wait** command. Use the **no** form of this command to restore the default wait time of 5 seconds.

arp wait *wait_time*

no arp wait

Syntax Description

wait_time	Number of seconds to wait for an ARP resolution. Enter an integer from 5 to 30. The default is 5 seconds.

Related Commands **show arp config**

(config) boot

To access boot configuration mode, use the **boot** command. Boot configuration mode contains all commands necessary to manage booting the CSS and to maintain the software revision.

boot

Usage Guidelines

When you use the **boot** command to access boot mode, the prompt changes to (config-boot). For information about commands available in this mode, see the "Boot Configuration Mode Commands" section.

(config) bridge

To configure the spanning-tree bridge parameters that apply to the CSS, use the **bridge** command. The options for this global configuration mode command are:

- **bridge aging-time** - Sets the bridge filtering database aging time
- **bridge bpdu-guard** - Enables or disables the Bridge Protocol Data Unit (BPDU) guard feature on the CSS
- **bridge forward-time** - Sets the bridge forward delay time
- **bridge hello-time** - Sets the bridge hello time interval
- **bridge max-age** - Sets the bridge spanning-tree maximum age
- **bridge priority** - Sets the spanning-tree priority for the root bridge on the network
- **bridge spanning-tree** - Enables or disables the bridge spanning tree

For more information on these options and associated variables, see the following commands.

Note For information on bridge commands you can use in interface mode, see the **(config-if) bridge** command.

Related Commands

show bridge
(config) interface
(config-if) bridge

bridge aging-time

To set the spanning-tree bridge filtering database aging time for the CSS, use the **bridge aging-time** command. Use the **no** form of this command to restore the default aging time of 300.

bridge aging-time *timeout*

no bridge aging-time

Syntax Description

timeout	Timeout period in seconds for aging out dynamically learned forwarding information. Enter an integer from 10 to 1000000. The default is 300.

Command Modes

Global configuration mode

Related Commands

show bridge status

bridge bpdu-guard

To globally enable or disable the Bridge Protocol Data Unit (BPDU) guard feature on the CSS, use the **bridge bpdu-guard** command. The command shuts down PortFast-configured interfaces that receive BPDUs rather than putting the interfaces into the spanning-tree blocking state. By default, the BPDU guard feature is disabled.

bridge bpdu-guard [**enabled**|**disabled**]

Syntax Description

enabled	Enables the BPDU guard feature
disable	Disables the BPDU guard feature (default)

Command Modes

Global configuration

Usage Guidelines

The BPDU guard feature affects interfaces that have the PortFast feature enabled on them. PortFast should only be configured on interfaces that connect to end stations; otherwise, an accidental topology loop could cause a data packet loop and disrupt CSS and network operation. An interface with PortFast mode enabled is moved directly to the spanning-tree forwarding state when linkup occurs, without waiting for the standard forward-time delay.

When properly connected to other devices, PortFast-configured interfaces do not receive BPDUs. If a BPDU is received on a PortFast-configured interface, the interface is connected to an invalid device, such as a switch or router, and the BPDU guard feature disables the interface. The BPDU guard feature provides a secure response to invalid connections because you must manually put the interface back in service.

Related Commands

show bridge
(config-if) bridge port-fast

bridge forward-time

To set the spanning-tree bridge forward delay time, use the **bridge forward-time** command. Use the **no** form of this command to restore the default delay time of 4.

bridge forward-time *delay*

no bridge forward-time

Syntax Description

delay	Delay time in seconds that all bridges use for forward delay when this bridge is acting as the root. Enter an integer from 4 to 30. The default is 4.

Command Modes

Global configuration mode

Usage Guidelines

Make sure that the bridge maximum age is less than or equal to 2 x (bridge forward-time – 1 second) and greater than or equal to 2 x (bridge hello-time + 1 second).

Related Commands

show bridge status
(config) bridge max-age

bridge hello-time

To set the bridge hello time interval, use the **bridge hello-time** command. Use the **no** form of this command to restore the default hello time interval of 1.

bridge hello-time *hello*

no bridge hello-time

Syntax Description

hello	Hello time in seconds that all bridges use when this bridge is acting as the root. Enter an integer from 1 to 10. The default is 1.

Command Modes

Global configuration mode

Usage Guidelines

Make sure that the bridge maximum age is greater than or equal to 2 x (bridge hello-time + 1 second) and less than or equal to 2 x (bridge forward-time – 1 second).

Related Commands

show bridge status
(config) bridge max-age

bridge max-age

To set the bridge spanning-tree maximum age, use the **bridge max-age** command. Use the **no** form of this command to restore the default maximum age of 6.

bridge max-age *age*

no bridge max-age

Syntax Description

age	Maximum age in seconds that all bridges use when this bridge is acting as the root. Enter an integer from 6 to 40. The default is 6.

Command Modes

Global configuration mode

Usage Guidelines

Make sure that the bridge maximum age is greater than or equal to 2 x (bridge hello-time + 1 second) and less than or equal to 2 x (bridge forward-time – 1 second).

Related Commands

show bridge status
(config) bridge forward-time
(config) bridge hello-time

bridge priority

To set the priority used by the spanning-tree protocol to choose the root bridge on the network, use the **bridge priority** command. This command can override the root bridge selection in your network. Use the **no** form of this command to restore the default priority of 32768.

bridge priority *priority*

no bridge priority

Syntax Description

priority	Decimal value for the write portion of the bridge ID; the first two octets of the 8-octet bridge ID. The last 6 octets of the bridge ID come from the base bridge address. Enter an integer from 0 to 65535 (0 to ffff, hexadecimal). The default is 32768 (0x8000, hexadecimal).

Command Modes

Global configuration mode

Related Commands

show bridge status

bridge spanning-tree

To enable or disable the spanning tree, use the **bridge spanning-tree** command.

bridge spanning-tree [**disable**|**enable**]

Syntax Description

disable	Disables the spanning tree.
enabled	Enables the spanning tree. This is the default state.

Command Modes

Global configuration mode

Usage Guidelines

Disabling spanning-tree bridging may make your network susceptible to packet storms. When you disable spanning-tree bridging, the CSS drops Bridge Protocol Data Units (BPDUs), but forwards the Cisco Systems 802.1Q BPDUs (tagged with the proprietary 01-00-0c-cc-cc-cc-cd destination MAC address) on an 802.1Q VLAN trunk. The CSS can still operate in an 802.1Q spanning-tree environment as long as you do not require that the CSS put any of its ports into a blocking state.

Related Commands

show bridge status

(config) bypass persistence

To determine if the CSS performs either a service remapping or HTTP redirection operation to reset a bypassed service when a content request matches on a content rule, but a previous request caused the bypass, use the **bypass persistence** command. By default, **bypass persistence** is enabled.

bypass persistence [**disable**|**enable**]

Syntax Description

disable	Performs remapping or redirection to reset the connection according to the setting of the persistence reset method
enable	Does not perform remapping or redirection to reset the connection, and continues to bypass a service

Usage Guidelines

The **bypass persistence** command affects all flows.

Related Commands

show remap
(config) persistence reset
(config-owner-content) persistent

(config) cdp

To configure the global Cisco Discovery Protocol (CDP) parameters on the CSS, use the **cdp** command. The options for this global configuration mode command are:

- **cdp holdTime** - Defines the period of time to hold the CSS CDP information before discarding it
- **cdp run** - Enables CDP on the CSS and the broadcasting of CDPv1 advertisements by the CSS
- **cdp timer** - Specifies how often the CSS sends CDP advertisements to Cisco CDP-compatible devices

For more information on these options and associated variables, see the following commands.

Usage Guidelines

The Cisco Discovery Protocol (CDP) is a media-independent protocol that runs over Layer 2 (the data link layer) on the CSS and other Cisco-manufactured equipment, such as routers, switches, bridges, and access servers. CDP allows the CSS to advertise itself to all other neighboring Cisco CDP-compatible devices on a network.

Note The CSS only transmits CDP advertisements to other CDP-compatible devices on the network; it does not listen for CDP messages from other CDP-compatible devices.

Any Cisco device with CDP support can learn about the CSS by listening to the periodic advertisements transmitted by the CSS and determine when the CSS is active. Network operators and analysts can use this information for configuration monitoring, topology discovery, and fault diagnosis.

CDP advertisements include the following information about the CSS:

- Device ID (CSS base MAC address)
- IP address (CSS management port IP address)
- Ethernet port ID name
- CSS functional capability flag (router, transparent bridge, or switch)

- CSS software version
- CSS platform

CDP advertisements also include time-to-live, or hold-time information, which defines the length of time the receiving device is to hold CDP information before discarding it.

Related Commands **show cdp**

cdp holdTime

To define the hold time in the CSS CDP advertisement to receiving devices, use the **cdp holdTime** command. The hold time defines how long the CSS wants the device to hold the CSS CDP information before discarding it. If a device does not receive a CSS CDP advertisement before the hold time expires, it drops the CSS as a neighbor. Use the **no** form of this command to reset the hold time to its default of 180 seconds.

cdp holdTime *seconds*

no cdp holdTime

Syntax Description

seconds	Number of seconds for holding the CSS CDP information. The range is from 10 to 255. The default is 180.

Command Modes Global configuration mode

cdp run

To enable CDP transmissions to advertise the CSS in the form of CDPv1 packet broadcasts to neighboring Cisco CDP-compatible devices on the network, use the **cdp run** command. By default, CDP advertisement is disabled for the CSS. Use the **no** version of this command to disable the CSS CDP transmissions.

cdp run

no cdp run

Command Modes Global configuration mode

cdp timer

To specify the interval at which the CSS advertises CDP packets to all receiving CDP-compatible devices, use the **cdp timer** command. Use the **no** form of the command to reset the interval to its default of 60 seconds.

cdp timer *interval*

no cdp timer

Syntax Description

interval	Number of seconds that the CSS advertises CDP packets. The range is from 5 to 254. The default is 60.

Command Modes Global configuration mode

(config) circuit

To access circuit configuration mode and configure a circuit on the CSS, use the **circuit** command. A circuit on the CSS is a logical entity that maps IP interfaces to a logical port or group of logical ports.

circuit *circuit_name*

Syntax Description

circuit_name	Name of the circuit you want to configure. To see a list of available circuits, enter: `circuit ?`

Usage Guidelines

When you use the **circuit** command to access circuit mode, the prompt changes to (config-circuit [*circuit_name*]). For information about commands available in this mode, see the "Circuit Configuration Mode Commands" section.

Related Commands

show circuits

(config) cmd-sched

To enable command scheduling, use the **cmd-sched** command. Use the **no** form of this command to disable command scheduling.

cmd-sched
no cmd-sched

(config) cmd-sched record

To create a configuration record for the scheduled execution of any CLI commands, including the playing of scripts, use the **cmd-sched record** command. Use the **no** form of this command to delete a configuration record.

cmd-sched record *name minute hour day month weekday* **"***command...***"** {*logfile_name*}

no cmd-sched record

Syntax Description

name	Name of the configuration record. Enter an unquoted text string with a maximum of 16 characters. Any of the following time variables can contain one or some combination of the following values: • A single number to define a single or exact value for the specified time variable • A **"*"** wildcard character matching any valid number for the specified time variable • A list of numbers separated by commas, with a maximum of 40 characters, to define multiple values for a time variable • Two numbers separated by a dash (-) character indicating a range of values for a time variable
minute	Minute of the hour to execute the command. Valid numbers are from 0 to 59.
hour	Hour of the day. Valid numbers are from 0 to 23.
day	Day of the month. Valid numbers are from 0 to 31.
month	Month of the year. Valid numbers are from 1 to 12.
weekday	Day of the week. Valid numbers are from 1 to 7. Sunday is 1.

"*command...***"**	The commands you want to execute. Enter a quoted text string with a maximum of 255 characters. Separate multiple commands with a semicolon (:) character. If the command string includes quoted characters, use a single quote character; any single quoted characters not preceded by a **"\"** character is converted to double quotes when the commands string is executed.
logfile_name	(Optional) Defines the name of the log file. Enter a text string with a maximum of 32 characters.

Usage Guidelines

The commands that the **cmd-sched record** command executes are referred to as the command string. To schedule commands, you must create a configuration record including when to execute the commands and the command string.

For example, you can use this command to schedule periodic content replication and configuration changes and gather statistics. At the specified time, the command scheduler executes a command string by creating a pseudo login shell where each string is executed. A cmd-sched record is only scheduled for execution upon completion of its shell. Use the **show lines** command to display information about active pseudo shells.

Note To terminate the execution of a command string, you can use the **disconnect** command.

Related Commands

disconnect
show cmd-sched
show lines

(config) console authentication

To configure the primary, secondary, or tertiary console port authentication of locally-defined usernames and passwords logging into the CSS, use the **console authentication** command. Use the **no** form of this command to disable authentication on the console port allowing users to access the CSS without a username and password.

console authentication [**primary** [**local**|**radius**|**tacacs**|**none**]
|**secondary**|**tertiary** [**local**|**radius**|**tacacs**|**none**|**disallowed**]]

no console authentication

Syntax Description

primary	Defines the first authentication method that the CSS uses. The default primary console authentication method is the local user database.
secondary	Defines the second authentication method that the CSS uses if the first method fails. The default secondary console authentication method is to disallow all user access. If you are configuring a TACACS+ server as the primary authentication method, define a secondary authentication method, such as **local**. If you do not configure a secondary method and use the default of **disallowed**, you have the possibility of being locked out of the CSS.
tertiary	Defines the third authentication method that the CSS uses if the second method fails. The default tertiary console authentication method is to disallow all user access.
local	The CSS uses the local user database for authentication.
radius	The CSS uses the configured RADIUS server for authentication.
tacacs	The CSS uses the configured TACACS+ server for authentication.

none	The CSS uses no authentication method. All users can access the CSS.
disallowed	The CSS does not allows access by all users (secondary or tertiary authentication method only). Entering this keyword does not terminate existing connections. To remove users currently logged into the CSS, use the **disconnect** command.

Usage Guidelines

To control access to the CSS, you can configure the CSS to authenticate console users. The CSS can authenticate users by using the local user database, RADIUS server, or TACACS+ server. You can also allow user access without authenticating or disallowing all remote user access to the CSS.

You can set a maximum of three authentication methods: a primary, secondary, or tertiary authentication method. The primary method is the first authentication method that the CSS tries. If the primary authentication method fails, the CSS tries the secondary method. If the secondary method fails, the CSS tries the tertiary method. In the event that the tertiary method also fails, the CSS displays a message that authentication has failed.

Before you can use RADIUS or TACACS+ as the console authentication method, you must enable communication with the RADIUS or TACACS+ security server. Use either the **(config) radius-server** command or the **(config) tacacs-server** command.

Related Commands

show user-database
(config) restrict console
(config) radius-server
(config) tacacs-server
(config) virtual authentication

(config) date european-date

To change the behavior of the **clock date date** command to accept date input in the format of day, month, and year, use the **date european-date** command. Use the **no** form of this command to reset the format for the **clock date** command to its default format of month, day, and year.

date european-date

no date european-date

Related Commands

clock date
show clock

(config) dfp

To configure a DFP agent listening for DFP connections on an IP address and TCP port combination on a server, and to enable the DFP manager on the CSS, use the **dfp** command. You can configure a maximum of 127 DFP agents for the DFP manager in the CSS. Use the **no** form of this command to disable the DFP agent connection to an IP address.

dfp *ip_or_host* {*port* {**key "***secret***"**|[**des-encrypted** *encrypted_key* |**"***encrypt_key***"**]} {**timeout** *seconds*} {**retry** *count*} {**delay** *time*} {**max-agent-wt** *weight*}

no dfp *ip_or_host* {*port*}

Syntax Description

ip_or_host	IP address or host name of the configured DFP agent. Enter an IP address in dotted-decimal notation (for example, 192.168.11.1) or a mnemonic host name (for example, myhost.mydomain.com).
port	(Optional) Server TCP port that the configured DFP agent uses to listen for connections from the CSS DFP manager. Valid entries are 0 to 65535. The default is 14001.

key "*md5secret*"	(Optional) MD5 (Message Digest Algorithm Version 5) security key used for encryption to provide a secure data exchange between the CSS DFP load-balancing manager and the DFP agents. MD5 encryption is a one-way hash function that provides strong encryption protection. Enter the secret as a case-sensitive quoted text string (maximum of 64 characters). It can include any printable ASCII character except tabs. Ensure that you configure the same key on each DFP agent for MD5 encryption to function properly.
des-encrypted	(Optional) Defines a Data Encryption Standard (DES) encryption key.
encrypted_key	DES encryption key that the CSS had previously encrypted. The CSS does not reencrypt this key and saves it in the running-config as you entered it. Enter an unquoted case-sensitive text string with no spaces and a maximum of 128 characters.
"*encrypt_key*"	DES encryption key that you want the CSS to encrypt. The CSS saves the encrypted key in the running-config as you entered it. Enter a quoted case-sensitive text string with no spaces and a maximum of 64 characters.
timeout *seconds*	(Optional) Maximum inactivity time period (the keepalive time) for the connection between the CSS DFP manager and the server DFP agent. If the inactivity time period exceeds the timeout value, the DFP manager closes the connection. The DFP manager attempts to reopen the connection as often as specified by the value of the **retry** option. The range is from 1 to 10000 seconds. The default is 3600 seconds (1 hour).
retry *count*	(Optional) Number of times the CSS DFP manager tries to reopen a connection with the server DFP agent. The range is 0 (for continuous retries) to 65535. The default is three retry attempts.

delay *time*	Optional. The delay time, in seconds, between each connection reestablishment attempt. Valid entries are 1 (immediately) to 65535 seconds (18 hours). The default value is 5 seconds.
max-agent-wt *weight*	Optional. Maximum value of the weight reported by a DFP agent. A CSS uses this option to scale the reported weight when the weight range of a DFP agent does not match the weight range of the DFP manager. For example, the DFP manager weight range is 0 to 255. If a DFP agent reports weight in the range 0 to 16, the CSS scales up the agent-reported weight to match the weight range of the DFP manager. If an agent reports weight in the range 0 to 65535, the CSS scales down the agent-reported weight to match the weight range of the DFP manager. If a DFP agent reports a weight greater than the maximum configured weight, then the CSS rejects the weight report and does not use the weight in load-balancing decisions. In this case, the CSS also logs an error in SYSLOG. Enter an integer from 1 to 65535. The default is 255.

Related Commands

show dfp
show dfp-reports

(config) dhcp-agent max-hops

To set the maximum allowable number in the hops field of the BOOTP header, use the **dhcp-agent max-hops** command. The CSS does not forward packets with headers that have a larger number. Use the **no** form of this command to reset the maximum allowable number in the hops field to its default of 4.

dhcp-agent max-hops *number*

no dhcp-agent max-hops

Syntax Description

number	Maximum allowable number in the hops field of the BOOTP header. The range is 1 to 15. The default is 4.

Related Commands **show dhcp-relay-agent global**

(config) dns

To enter commands that control the Domain Name System (DNS) client, the facility that translates host names such as myhost.mydomain.com to IP (Internet Protocol) addresses such as 192.168.11.1, use the **dns** command. The options for this global configuration mode command are:

- **dns primary** - Specifies the primary DNS server to use for DNS name resolution
- **dns secondary** - Specifies the secondary DNS server to use for DNS name resolution
- **dns suffix** - Specifies the default suffix to use during a DNS query

For information on these options and associated variables, see the following commands.

Related Commands **show running-config global**
(config) dns-server

dns primary

To specify the primary DNS server to use for DNS queries and resolution, use the **dns primary** command. Use the **no** form of this command to remove the primary DNS server.

dns primary *ip_or_host*

no dns primary

Syntax Description

ip_or_host	Default DNS address to use for DNS queries. Enter the IP address in dotted-decimal notation (for example, 192.168.11.1) or the mnemonic host name (for example, myhost.mydomain.com).

Command Modes Global configuration mode

dns secondary

To specify the secondary DNS server, use the **dns secondary** command. When the primary server fails, the CSS uses the secondary server for DNS name resolution. Use the **no** form of this command to remove a secondary DNS server on a client.

dns secondary *ip_or_host*

no dns secondary *ip_or_host*

Syntax Description

ip_or_host	IP address for the secondary DNS server. Enter the IP address in dotted-decimal notation (for example, 192.168.11.1) or the mnemonic host name (for example, myhost.mydomain.com).

Command Modes Global configuration mode

Usage Guidelines You can specify a maximum of two secondary servers. To specify each additional server, repeat the **dns secondary** command. The order in which you enter them is the order in which they are used if the primary DNS server fails.

dns suffix

To specify the default suffix to use when querying the DNS server to resolve a DNS name, use the **dns suffix** command. Use the **no** form of this command to remove the default suffix.

dns suffix *suffix*

no dns suffix

Syntax Description

suffix	Default suffix. Enter an unquoted text string with no spaces and a maximum length of 64 characters (for example, webhoster.com).

Command Modes Global configuration mode

(config) dns-boomerang client

To configure and enable the Content Routing Agent (CRA) functionality on the CSS, use the **dns-boomerang client** command. The CSS functioning as a CRA improves HTTP response time for a client request. A Cisco Content Router 4430B configured as a Content Routing server redirects a client to the closest (best) replicated-content site represented by a CRA, based on network delay.

The options for this global configuration mode command are:

- **dns-boomerang client cpu-threshold** - Specifies the CPU load threshold for a CSS CRA
- **dns-boomerang client domain** - Creates a client domain record in the CSS CRA domain name server or creates a client alias record
- **dns-boomerang client enable** - Enables the CRA functionality on the CSS

For information on these options and associated variables, see the following commands.

Related Commands **show dns-boomerang client**

dns-boomerang client cpu-threshold

To set the CPU load threshold for a CSS CRA, use the **dns-boomerang client cpu-threshold** command. If the CSS CPU load exceeds the configured threshold value, then the CSS drops incoming DNS requests from the Content Router. Use the **no** form of this command to reset the CSS CPU threshold to the default value of 99.

dns-boomerang client cpu-threshold *number*

no dns-boomerang client cpu-threshold

Syntax Description

number	The load threshold value. Enter a number from 1 to 99. The default value is 99.

Command Modes Global configuration mode

Usage Guidelines The load threshold value is the percentage of CPU utilization shown in the **show system-resources** command.

Related Commands **show system-resources**
(config) dns-boomerang client domain

dns-boomerang client domain

To create a client domain record in the CSS CRA or an alias for the record, use the **dns-boomerang client domain** command. The record maps to each of the domains you associated with the agent when you configured domains on the Content Router. Use the **no** form of this command to remove a client domain or the alias for the domain.

dns-boomerang client domain *dns_name* [**alias** *alias_name*|*ip_or_host* {**"***uri***"**} {**key** [**"***secret***"**|**des-encrypted** *encrypted_key*|**"***encrypt_key***"**]} {**dns-ttl** *number1*} {**ip-ttl** *number2*} {**threshold** *number3*}]

no dns-boomerang client domain *dns_name* {**alias** *alias_name*}

Syntax Description

dns_name	Domain name mapped to the client record. Enter the name as a case-sensitive, unquoted text string with no spaces and a maximum length of 72 characters. For example, www.sample.com.
alias	Creates an alias for an existing client domain. The alias behaves exactly the same as the configured domain.
alias_name	Alias name for the associated DNS name. Enter the name as a case-sensitive unquoted text string with no spaces and a maximum length of 72 characters.

ip_or_host	IP address or host name of the content server or web cache bound to the domain name on the CSS. This address can be a local VIP. Enter the address in dotted-decimal notation (for example, 192.168.11.1) or a mnemonic host name (for example, myhost.mydomain.com).
"*uri***"**	(Optional) Defines the URI that the CSS uses for the keepalive probe to the Content Router for a domain. Enter a quoted text string with a maximum of 255 characters. If you do not prepend the URI with a slash (/) character, the CSS prepends it.
key	(Optional) Defines the clear-text secret or DES encryption key on the Content Router.
"*secret***"**	Clear-text secret for encrypting packets sent between a Content Router and the CSS client. The secret is the same as the secret on the CR. Enter the secret as a case-sensitive quoted text string with a maximum of 64 characters.
des-encrypted	(Optional) Defines a Data Encryption Standard (DES) encryption key.
encrypted_key	DES encryption key that the CSS had previously encrypted. The CSS does not reencrypt this key and saves it in the running-config as you entered it. Enter an unquoted case-sensitive text string with no spaces and a maximum of 64 characters.
"*encrypt_key***"**	DES encryption key that you want the CSS to encrypt. The CSS saves the encrypted key in the running-config as you entered it. Enter a quoted case-sensitive text string with no spaces and a maximum of 16 characters.
dns-ttl *number1*	(Optional) Defines the DNS time-to-live value returned with the DNS responses of the CSS client. This keyword determines the length of time that a domain name server caches the returned information for reuse. Enter an integer from 10 to 2147483647 seconds. The default value is from the Content Router.

ip-ttl *number2*	(Optional) Defines the IP routing time-to-live value in hops that is set in the IP packets for returned CSS client DNS responses. This keyword determines how many router hops a response packet traverses en route to the client's local name server, D-Proxy, before it is discarded. This helps to eliminate the CSS client from longer races. Enter an integer from 1 to 255. The default value is from the Content Router.
threshold *number3*	(Optional) Defines the load threshold for testing the keepalive state of a local VIP. If the load on the associated rule is greater than the threshold, then the CSS drops Content Router requests until the load goes below the threshold. Enter an integer from 2 to 254. The default value is 254.

Command Modes Global configuration mode

Usage Guidelines If the matching domain record keepalive messaging succeeds, the CSS uses this record for DNS resolutions and will respond to the D-Proxy on behalf of the Content Router.

dns-boomerang client enable

To enable the Content Routing Agent (CRA) functionality on a CSS, use the **dns-boomerang client enable** command. Use the **no** form of this command to disable the CRA functionality.

dns-boomerang client enable

no dns-boomerang client enable

Command Modes Global configuration mode

Usage Guidelines

Before you enable the CRA functionality on a CSS, configure a Cisco Content Router 4430B as a Content Routing server and CRAs on the server. For information on configuring the server, refer to the *Cisco Content Router 4430B User Guide*.

(config) dns-peer

To control the DNS peer functionality on the CSS, use the **dns-peer** command. You can configure the CSS as a DNS peer to exchange DNS information over an APP connection to other CSSs. The options for this global configuration mode command are:

- **dns-peer interval** - Sets the time between sending load reports to each CSS DNS peer
- **dns-peer load-variance** - Sets the range of load numbers between peers that a CSS considers to be similar for the least-loaded algorithm in a DNS load-balancing decision
- **dns-peer receive-slots** - Sets the maximum number of DNS names that the CSS can receive from each CSS DNS peer
- **dns-peer send-slots** - Sets the maximum number of DNS names that the CSS can send to each CSS DNS peer

For information on these options and associated variables, see the following commands.

Related Commands

show dns-peer
(config) app
(config) dns
(config-owner) dns
(config-owner-content) add dns

dns-peer interval

To set the time between sending load reports to CSS DNS peers over an APP connection, use the **dns-peer interval** command. Use the **no** form of this command to reset the interval to its default value of 5.

dns-peer interval *number*

no dns-peer interval

Syntax Description

number	Time in seconds between generating load reports. Enter an integer from 5 to 120. The default is 5.

Command Modes Global configuration mode

dns-peer load-variance

To set the range of load numbers between peers that a CSS considers to be similar for the least-loaded algorithm in a DNS load-balancing decision, use the **dns-peer load-variance** command. If the load numbers of all peers are within the specified range, the CSS calculates the minimum response time of each site, then selects the site with the fastest response time. Use the **no** form of this command to reset the **load-variance** to its default value of 50.

dns-peer load-variance *number*

no dns-peer variance

Syntax Description

number	Upper limit of the range of load numbers considered similar. Enter an integer from 0 to 254. The default is 50.

Command Modes Global configuration mode

Usage Guidelines If you configure the absolute load calculation method for GSLB, we recommend that you configure a load variance of 0, regardless of whether you are using zone-based or rule-based DNS load balancing. For information on absolute load calculation, refer to the *Cisco Content Services Switch Content Load-Balancing Configuration Guide*.

dns-peer receive-slots

To set the maximum number of DNS names that the CSS can receive from each CSS DNS peer over an APP connection, use the **dns-peer receive-slots** command. Use the **no** form of this command to reset the maximum number of DNS names received from a peer to its default value of 128.

dns-peer receive-slots *number*

no dns-peer receive-slots

Syntax Description

number	Maximum number of DNS names that can be received from a peer. Enter an integer from 128 to 1024. The default is 128.

Command Modes Global configuration mode

dns-peer send-slots

To set the maximum DNS names that the CSS can send to each CSS DNS peer, use the **dns-peer send-slots** command. Use the **no** form of this command to reset the maximum number of DNS names sent to a peer to its default value of 128.

dns-peer send-slots *number*

no dns-peer send-slots

Syntax Description

number	Maximum number of DNS names sent to a peer. Enter an integer from 128 to 1024. The default is 128.

Command Modes Global configuration mode

(config) dns-record

To create a domain record, use the **dns-record** command and its options. This command is not available on a Proximity Database CSS. The command options are:

- **dns-record a** - Creates a domain record on the CSS Zone Domain Name Server mapped directly to an IP address
- **dns-record accel** - Creates a domain acceleration record on the CSS mapped to a content rule through an IP address
- **dns-record ns** - Creates a domain record on the CSS Zone Domain Name Server mapped to a name server IP address
- **dns-record zero** - Resets the DNS record statistics to zero

For information on these options and associated variables, see the following commands.

Related Commands **show dns-record**
(config) dns-server {**zone**}

dns-record a

To create a domain record on the CSS Zone Domain Name Server that maps the DNS name to an IP address, use the **dns-record a** command. If a domain *can* be directly translated to an IP address, configure it as an A-record. Use the **no** form of this command to delete a domain address record.

dns-record a *dns_name ip_address* {*ttl_value* {**single**|**multiple** {**kal-ap-vip**|**kal-ap**|**kal-icmp**|**kal-none** {*ip_address2* {*threshold* {**sticky-disabled**|**sticky-enabled** {**usedefault**|**weightedrr**|**srcip** |**leastloaded**|**preferlocal**|**roundrobin**|**proximity** {*weight*}}}}}}}}

no dns-record a *dns_name*

Syntax Description

dns_name	DNS name mapped to the address record. Enter the name as a case-sensitive, unquoted text string with no spaces and a maximum of 63 characters.
ip_address	IP address bound to the *dns_name* within the CSS zone. Enter the address in dotted-decimal notation (for example, 192.168.11.1). This is the VIP for which a CSS client sends a **kal-ap-vip** request to itself or another CSS agent for load information.
ttl_value	(Optional) Time to Live (TTL) value in seconds. This value determines how long the DNS client remembers the IP address response to the query. Enter a value from 0 to 65535. The default is 0.
single|**multiple**	(Optional) Number of records to return on a DNS response message. Enter either **single** or **multiple**. By default, the DNS server returns a single a-record. Setting this parameter to **single** ensures that only one a-record is returned.

kal-ap|kal-icmp| kal-none	(Optional) Keepalive message type for this record. The types are: • **kal-ap** - The keepalive message type keyword that specifies the CSS keepalive message. This is the recommended keepalive message type to obtain load information from remote as well as local services based on domains configured on a single content rule. **Note** To use **kal-ap** proximity keepalive messages, lower-level CSSs acting as either data centers or DNS servers must be running the Enhanced feature set. When the Proximity Domain Name Server (PDNS) is directly attached to a server farm, an internal keepalive is used. • **kal-icmp** - The keepalive message type keyword that specifies ICMP echo (ping). To obtain load information from local services only, use the **add dns** *record_name* command in the associated content rule. This is the default setting. • **kal-none** - For no keepalive messaging.
ip_address2	(Optional) IP address of the local interface receiving CSS keepalive messages.
threshold	(Optional) Load threshold used with the CSS proximity keepalive. The CSS considers that this record is in the Down state when the load number is greater than this value. Enter a value from 2 to 254. The default is 254.
sticky-disabled |sticky-enabled	(Optional) Disables or enables DNS sticky for the domain. The **sticky-disabled** option disables DNS sticky for the specified domain. This is the default setting.
usedefault	(Optional) Returns domain records using the default DNS load-balancing method configured for the zone.
weightedrr	(Optional) Returns domain records based on the weighted roundrobin load-balancing method. This method uses the *weight* value to determine the zone from which the record should be requested.

srcip	(Optional) Returns domain records using a source IP address hash. For sticky-enabled domains without a GSDB, the CSS uses the srcip method regardless of the configured balance method. For sticky-enabled domains with a GSDB, a CSS uses the configured balance method when the GSDB does not contain an entry for the requested domain.
leastloaded	(Optional) Returns domain records from the zone with the smallest load.
preferlocal	(Optional) Returns local domain records whenever possible. If no local record exists, the CSS uses the balance method configured for the zone with the lowest zone index.
roundrobin	(Optional) Returns domain records by cycling among records available at the different zones to evenly distribute the load.
proximity	(Optional) Returns domain records based on proximity information. If a PDB is not configured or is unavailable in a zone, the CSS applies the default balance method for the selected zone for DNS resolution. This is the default method.
weight	(Optional) Value assigned to a domain in the local zone to determine how many requests the local zone receives for the specified domain compared with other zones in a peer mesh. A domain with a weight of 10 in the local zone will receive twice as many requests as the same domain in another zone with a weight of 5. Use this parameter with the **weighted** roundrobin DNS load-balancing method. CSSs configured as authoritative DNS servers in a peer mesh share domain weights with each other. Enter an integer from 0 to 10 The default is 1. For details on configuring a DNS record with a *weight* of 0, refer to the *Cisco Content Services Switch Global Server Load-Balancing Configuration Guide*.

Command Modes

Global configuration mode

Usage Guidelines

This command is available on a CSS PDNS.

If you need to modify an existing A-record configuration, you must first remove the record using the **no dns-record a** *domain_name* command. Then, recreate the A-record with the change using the **dns-record a** command.

When you enable DNS Sticky through the **sticky-enabled** option, the CSS makes a decision based on one of the following three scenarios:

- In a global server load-balancing (GSLB) environment without a global sticky database (GSDB), the CSS selects a server based on the srcip hash (regardless of the default load-balancing method) and the availability of the domain in the zone mesh. The use of the srcip hash ensures that the CSS selects a consistent zone for a given source IP address.

  In a GSLB environment with a GSDB, the CSS sends a lookup request to the Global Sticky Database for the requesting client's local DNS server. If the GSDB has an entry in its sticky database for the client's local DNS server IP address, it returns the appropriate zone index to the CSS. The CSS then returns the associated IP address to the client. Otherwise, the CSS selects a zone based on the default load-balancing method and informs the GSDB about the selected zone.

- In a Network Proximity environment, the CSS configured as a Proximity Domain Name Server (PDNS) first consults the GSDB. If a sticky database entry exists for the client's local DNS server IP address, the PDNS sends the appropriate IP address to the client based on the zone index returned by the GSDB. If the GSDB does not contain an entry for the client's local DNS server IP address, the PDNS consults the Proximity Database (PDB).

- If the PDB contains an entry for the client's local DNS server IP address, the PDNS formulates a response to the client based on the ordered zone index returned by the PDB and keepalive information. The PDNS informs the GSDB about the selected zone (performs a "set" function). If the PDB does not have an entry for the client's local DNS server IP address or the sticky zone is unavailable, the CSS selects a zone based on its default load-balancing method and informs the GSDB about the selected zone.

Note If you configure any sticky domains in a particular zone, you must configure all sticky domains participating in the peer mesh in that same zone. Otherwise, the thrashing of the sticky zone index will cause DNS Sticky to fail.

For details on configuring DNS Sticky, refer to the *Cisco Content Services Switch Global Server Load-Balancing Configuration Guide*.

The CSS uses the following guidelines when selecting a DNS load-balancing method on a domain basis:

- If a local record exists, the CSS uses the configured domain balance method to determine local DNS resolutions. This applies regardless of the local record's keepalive state.
- If no local record exists, the CSS uses the balance method configured for the zone with the lowest zone index.

To provide backup sites in a DNS weighted roundrobin configuration when all domain records with weights from 1 to 10 are unavailable, configure domain records with a weight of zero. When a DNS record has a weight of zero, a CSS does not consider that record for selection when using the weighted roundrobin algorithm unless all of the other records, with weights from 1 to 10, are unavailable. This feature is intended especially for use in disaster recovery sites. For details, refer to the *Cisco Content Service Switch Global Server Load-Balancing Configuration Guide*.

Related Commands **show dns-record**

dns-record accel

To create a DNS acceleration record for the domains you want to accelerate on the CSS, use the **dns-record accel** command. Use the **no** form of this command to delete a DNS acceleration record.

dns-record accel *dns_name ip_address* {*ageout*}

no dns-record accel *dns_name*

Syntax Description

dns_name	DNS name you want to map to the acceleration record. Enter a case-sensitive unquoted text string with no spaces and a maximum of 63 characters.
ip_address	IP address of the local content rule that will handle content request for the DNS name during content acceleration.
ageout	(Optional) Number of minutes that the domain remains accelerated. Enter a number from 0 to 525600. The default is 180 minutes. If you enter 0, the accelerated domain record does not age out.

Usage Guidelines

The DNS acceleration record indicates a DNS name that is eligible for content acceleration. The record maps the name to a content rule through an IP address. To enable the acceleration of domains, use the **(config) dns-server accelerate domains** command. The **dns-record accel** command is *not* available on a Proximity Database CSS.

Configure nonaccelerated domains as either A-records or NS-records.

Note If the content rule associated with the acceleration candidate domain is suspended or cannot provide service for content requests, the CSA does not accelerate the domain.

Related Commands

show dns-record accel
(config) dns-server accelerate domains

dns-record ns

To create a domain record on the CSS Zone Domain Name Server that maps the DNS name to a Name Server IP address, use the **dns-record ns** command. If a domain *cannot* be directly translated to an IP address, configure it as an NS-record. Use the **no** form of this command to delete a DNS record.

dns-record ns *dns_name ip_address* {*ttl_value* {**single|multiple** {**kal-ap-vip|kal-ap|kal-icmp|kal-none** {*ip_address2* {*threshold* {**default|forwarder** {**sticky-disabled|sticky-enabled** {*weight* {**usedefault|weightedrr|srcip|leastloaded|preferlocal |roundrobin|proximity**}}}}}}}}}

no dns-record ns *dns_name*

Syntax Description

dns_name	DNS name mapped to the name server record. Enter the name as a case-sensitive, unquoted text string with no spaces and a maximum of 63 characters.
ip_address	IP address of the DNS server bound to the *dns_name* within the CSS zone. Enter the address in dotted-decimal notation (for example, 192.168.11.1).
ttl_value	(Optional) Time to Live (TTL) value in seconds. This value determines how long the DNS client remembers the IP address response to the query. Enter a value from 0 to 65535. The default is 0.
single|multiple	(Optional) Number of records to return on a DNS response message. Enter either **single** or **multiple**. By default, the DNS server returns a single ns-record. Setting this parameter to **single** ensures that only one ns-record is returned.

kal-ap| **kal-icmp**| **kal-none**	(Optional) Keepalive message type for this record. The types are: • **kal-ap** - For the CSS keepalive message. • **kal-icmp** - For an ICMP echo message (ping). This is the default setting. • **kal-none** - For no keepalive messaging.
ip_address2	(Optional) IP address of the local interface receiving CSS keepalive messages.
threshold	(Optional) Load threshold for the record. The CSS considers that the record is in the Down state when the load number is greater than this value. Enter a value from 2 to 254. The default is 254.
default	(Optional) Uses PDB information to return the next most proximate location. When a PDB is not available or configured, the roundrobin method is used.
forwarder	(Optional) Eliminates a potential single point of failure by providing a maximum of two alternative DNS servers called forwarders. A forwarder can be a CSS configured as a DNS server or a fully-functional BIND DNS server. If an optimal miss occurs (the lower-level DNS server indicated in the NS-record is Down), the PDNS sends the DNS request to the primary or secondary forwarder, depending on forwarder health and configuration. An optimal miss occurs when the PDNS cannot return the NS-record for the zone that the PDB indicated was most proximate. For this failover to occur, the local NS-record must be in the Down state, and the PDB has indicated the local zone to be the zone most proximate to the client.
sticky-disabled |**sticky-enable**	(Optional) Disables or enables DNS sticky for the domain. The **sticky-disabled** option disables DNS sticky for the specified domain. This is the default setting.

weight	(Optional) Value assigned to a domain in the local zone to determine how many requests the local zone receives for the specified domain compared with other zones in a peer mesh. A domain with a weight of 10 in the local zone will receive twice as many requests as the same domain in another zone with a weight of 5. Use this parameter with the **weighted** roundrobin DNS load-balancing method. CSSs configured as authoritative DNS servers in a peer mesh share domain weights with each other. Enter an integer from 0 to 10 The default is 1. For details on configuring a DNS record with a *weight* of 0, refer to the *Cisco Content Services Switch Global Server Load-Balancing Configuration Guide*.
usedefault	(Optional) Returns domain records using the default DNS load-balancing method configured for the zone.
weightedrr	(Optional) Returns domain records based on the weighted roundrobin load-balancing method. This method uses the *weight* value to determine the zone from which the record should be requested.
srcip	(Optional) Returns domain records using a source IP address hash. For sticky-enabled domains without a GSDB, the CSS uses the srcip method regardless of the configured balance method. For sticky-enabled domains with a GSDB, a CSS uses the configured balance method when the GSDB does not contain an entry for the requested domain.
leastloaded	(Optional) Returns domain records from the zone with the smallest load.
preferlocal	(Optional) Returns local domain records whenever possible. If no local record exists, the CSS uses the balance method configured for the zone with the lowest zone index.

roundrobin	(Optional) Returns domain records by cycling among records available at the different zones to evenly distribute the load.
proximity	(Optional) Returns domain records based on proximity information. If a PDB is not configured or is unavailable in a zone, the CSS applies the default balance method for the selected zone for DNS resolution. This is the default method.

Command Modes

Global configuration mode

Usage Guidelines

This command is available on a CSS PDNS.

If you need to modify an existing NS-record configuration, you must first remove the record using the **no dns-record ns** *domain_name* command. Recreate the NS-record with the change using the **dns-record ns** command.

When you enable DNS Sticky through the **sticky-enabled** keyword, The CSS makes a decision based on one of the following three scenarios:

- In a global server load-balancing (GSLB) environment without a global sticky database (GSDB), the CSS selects a server based on the srcip hash (regardless of the default load-balancing method) and the availability of the domain in the zone mesh. The use of the srcip hash ensures that the CSS selects a consistent zone for a given source IP address.

  In a GSLB environment with a GSDB, the CSS sends a lookup request to the Global Sticky Database for the requesting client's local DNS server. If the GSDB has an entry in its sticky database for the client's local DNS server IP address, it returns the appropriate zone index to the CSS. The CSS then returns the associated IP address to the client. Otherwise, the CSS selects a zone based on the default load-balancing method and informs the GSDB about the selected zone.

- In a Network Proximity environment, the CSS configured as a Proximity Domain Name Server (PDNS) first consults the GSDB. If a sticky database entry exists for the client's local DNS server IP address, the PDNS sends the appropriate IP address to the client based on the zone index returned by the GSDB. If the GSDB does not contain an entry for the client's local DNS server IP address, the PDNS consults the Proximity Database (PDB).
- If the PDB contains an entry for the client's local DNS server IP address, the PDNS formulates a response to the client based on the ordered zone index returned by the PDB and keepalive information. The PDNS informs the GSDB about the selected zone (performs a "set" function). If the PDB does not have an entry for the client's local DNS server IP address or the sticky zone is unavailable, the CSS selects a zone based on its default load-balancing method and informs the GSDB about the selected zone.

Note If you configure any sticky domains in a particular zone, you must configure all sticky domains participating in the peer mesh in that same zone. Otherwise, the thrashing of the sticky zone index will cause DNS Sticky to fail.

For details on configuring DNS Sticky, refer to the *Cisco Content Services Switch Global Server Load-Balancing Configuration Guide*.

The CSS uses the following guidelines when selecting a DNS load-balancing method on a domain basis:

- If a local record exists, the CSS uses the configured domain balance method to determine local DNS resolutions. This applies regardless of the local record's keepalive state.
- If no local record exists, the CSS uses the balance method configured for the zone with the lowest zone index.

To provide backup sites in a DNS weighted roundrobin configuration when all domain records with weights from 1 to 10 are unavailable, configure domain records with a weight of zero. When a DNS record has a weight of zero, a CSS does not consider that record for selection when using the weighted roundrobin algorithm unless all of the other records, with weights from 1 to 10, are unavailable. This feature is intended especially for use in disaster recovery sites. For details, refer to the *Cisco Content Service Switch Global Server Load-Balancing Configuration Guide*.

Related Commands **show dns-record**
(config) dns-server forwarder

dns-record zero

To reset the statistics or counters displayed by the **show dns-record** command to zero for all domain records or a specific domain name, use the **dns-record zero** command.

dns-record zero [**a/ns** {*domain_name*}|**accel** {*domain_name*}]

Syntax Description

a/ns	Resets the statistics for the domain records displayed by the **show dns-record statistics** command and the **show dns-record proximity** command.
domain_name	(Optional) Specified domain name mapped to the DNS record. To view a list of domain names, enter: `dns-record zero [a/ns|accel] ?`
accel	(Optional) Resets the counters for the acceleration records displayed by the **show dns-record accel** command.

Usage Guidelines The **dns-record zero** command is *not* available on a Proximity Database CSS.

Related Commands **show dns-record**
(config) dns-record

(config) dns-server

To enable the DNS server function on the CSS, use the **dns-server** command. The CSS acts as the authoritative name server for the content domain. Use the **no** form of this command to disable DNS server functionality on the CSS.

dns-server

no dns-server

Related Commands

show dns-server
show zone
(config) app
(config) dns
(config-owner) dns
(config-owner-content) add dns

(config) dns-server accelerate domains

To enable the domain acceleration and configure the Client Side Accelerator (CSA) on the CSS, use the **dns-server accelerate domains** command. Use the **no** form of this command to disable domain acceleration.

dns-server accelerate domains {*threshold interval max_number* [**single-location**|**multi-location**]}

no dns-server accelerate domains

Syntax Description

threshold	(Optional) Hits threshold used to determine whether a domain is accelerated. When the hits on the domain are greater than or equal to the threshold, the CSA accelerates the domain. Enter a number from 0 to 65535. The default is 0, indicating that the CSA always accelerates the candidate domains.
interval	Interval in minutes over which the CSA samples the hits on the domain and compares the hits with the threshold. Enter a number from 1 to 3600. The default is 5.
max_number	Maximum number of domains that the CSA can accelerate. Enter a number from 0 to 4096. The default is 1024.
single-location	Allows CSA peers to share content by maintaining the content on the cache farm of a single CSA.
multi-location	Allows multiple CSAs to accelerate the same domain resulting in multiple cache farms maintaining the same content. This can occur when two or more CSAs (located in different POPs) are configured for multi-location and accelerate the same domain. Each cache farm maintains the same content after: • The CSAs accelerate the same domain. • A cache in each POP retrieves the same content from the origin server.

Usage Guidelines

Use the **dns-server accelerate** command to enable the acceleration of domains configured through the **dns-record accel** command.

Related Commands

show dns-server accelerate domains
(config) dns-record accel

(config) dns-server bufferCount

To change the DNS response buffer count on the CSS, use the **dns-server bufferCount** command. Use the **no** form of this command to set the DNS response buffer count to its default value of 50.

dns-server bufferCount *number*

no dns-server bufferCount

Syntax Description

number	Number of buffers allocated for query responses. Enter an integer from 2 to 1000. The default is 50.

Usage Guidelines

Only use the **dns-server bufferCount** command to tune the CSS if the CSS experiences buffer depletion during normal use. If the name server buffers (NS Buffers) drop below two, increase the buffer count and the responder task with the **(config) dns-server respTasks** command. To view the buffers, use the **show dns-server** command.

Related Commands

show dns-server

(config) dns-server domain-cache

To enable domain caching to track DNS request counts and configure the parameters for the domain cache on the CSA, use the **dns-server domain-cache** command. Use the **no** form of this command to disable domain caching.

dns-server domain-cache {*cache_size ageout*|**purge** {*dns_name*} |**zero** {*dns_name*}}

no dns-server domain-cache

Syntax Description

cache_size	(Optional) Number of domains that the CSA can cache. Enter a number from 1 to 4096. The default is 1024.
ageout	(Optional) Maximum number of seconds that the domain entry remains in cache. Enter a number from 0 to 60. The default is 10 seconds. If you enter 0, the domain entries remain in cache unless they are removed with the **purge** option.
purge	(Optional) Removes all entries or the specified entries in the domain cache.
dns_name	(Optional) DNS entry in the domain cache. To see a list of entries, enter: `dns-server domain-cache [purge|zero] ?`
zero	(Optional) Resets all counters for all entries or the specified entry in the domain cache displayed through the **show dns-server domain-cache** command.

Usage Guidelines

Use the **dns-server domain-cache** command to create the domain cache and enable it. The domain cache records all domains including accelerated domains.

Enabling or disabling the domain cache does not affect domain acceleration. The operation of the domain cache can impact the DNS request/response rate performance. Use the domain cache only when you need to identify potential acceleration candidates.

Related Commands

show dns-server domain-cache

(config) dns-server forwarder

To configure a DNS server forwarder on a CSS, use the **dns-server forwarder** command. The forwarder is an alternative server for resolving DNS requests. In the case of proximity, the forwarder is a CSS in the same zone as the PDB. When the CSS is acting as a CSA, the forwarder is a fully-functional Berkeley Internet Name Domain (BIND) DNS server, not a CSS. Use the **no** form of this command to delete the DNS forwarder.

dns-server forwarder [**primary** *ip_address*|**secondary** *ip_address*|**zero**]

no dns-server forwarder primary|secondary

Syntax Description

primary	Specifies the first choice forwarder. The CSS sends unresolvable requests to the primary forwarder unless it is unavailable, in which case, it uses the secondary forwarder. When the primary forwarder is available again, the CSS resumes sending requests to the primary forwarder.
secondary	Specifies the second choice as the forwarder.
ip_address	IP address for the DNS forwarder. Enter the address in dotted-decimal notation (for example, 192.168.11.1).
zero	Resets the statistics of both forwarders on the CSS. The statistics are displayed through the **show dns-server forwarder** command.

Usage Guidelines

The CSS uses the primary forwarder first. If it is unavailable, the CSS uses the secondary forwarder.

The forwarder receives DNS requests that the CSS cannot resolve, or that contain an unsupported request or record type. The forwarder sends DNS responses to the client transparently through the CSS. To monitor forwarder health, an internal keepalive mechanism sends queries periodically to validate the state of the forwarder.

Related Commands

show dns-server forwarder
(config) dns-record ns

(config) dns-server respTasks

To change the DNS server responder task count, use the **dns-server respTasks** command. These tasks handle responses to incoming DNS query requests. Use the **no** form of this command to set the DNS responder task count to its default value of 2.

dns-server respTasks *number*

no dns-server respTasks

Syntax Description

number	Number of tasks. Enter an integer from 1 to 250. The default is 2.

Usage Guidelines

If you increase the responder task count, also increase the buffer count with the **(config) dns-server bufferCount** command.

(config) dns-server zero

To set the DNS server request and response statistics displayed by the **show dns-server** command to zero, use the **dns-server zero** command.

dns-server zero

Usage Guidelines

The **dns-server zero** command is *not* available on a Proximity Database CSS.

Related Commands

show dns-server
(config) dns-server

(config) dns-server zone

To enable the CSS Zone Domain Name Server (DNS) on a CSS or configure how the CSS handles the least-loaded balance method, use the **dns-server zone** command. This service allows the CSS to respond to DNS requests based upon proximity and shared zone domain availability. Use the **no** form of this command to disable the CSS Proximity Domain Name Server or disable DNS server zone load reporting.

dns-server zone *zoneIndex* {**tier1**|**tier2** {**"***description***"** {**weightedrr**|**srcip**|**leastloaded**|**preferlocal**|**roundrobin**|*ip_address* {**weightedrr**|**srcip**|**leastloaded**|**preferlocal**|**roundrobin**} {*weight*}}}} |**load** [**reporting**|**frequency** *seconds*|**variance** *number*]

no dns-server zone|**load** [**reporting**|**frequency**|**variance**]

Syntax Description

zoneIndex	Numerical identifier of the Proximity Zone of the CSS. This number should match the *zoneIndex* configured on the Proximity Database in a dedicated CSS 11150. Enter an integer from 0 to 15.
tier1|**tier2**	(Optional) Maximum number of zones the CSS expects to participate in its proximity zone mesh. Enter **tier1** for a maximum of 6 zones, numbered 0 to 5. Enter **tier2** for a maximum of 16 zones, numbered 0 to 15. Tier1 is the default. For CSA applications, the tier you select must be the same as the tier for the other CSAs participating in the mesh.
"*description***"**	(Optional) Text description of the CSS zone. Enter a quoted string with a maximum of 20 characters.
ip_address	(Optional) IP address of the PDB. Enter the address in dotted-decimal notation (for example, 192.168.11.1). This enables the DNS server to respond to DNS requests based on proximity. For CSA applications, do not enter an IP address.

weightedrr| roundrobin |srcip |leastloaded |preferlocal	(Optional) Balance method to determine the algorithm that the DNS server uses to choose returned records when a PDB is unavailable or not configured. • **weightedrr** - The CSS gives a zone priority over other zones in a peer mesh according to the assigned domain weights. Each CSS in the mesh maintains an internal list of zones ordered from highest to lowest according to weight. The heaviest zone (the zone with the highest weight number) receives DNS requests until it reaches its maximum number of requests, then the next heaviest zone receives DNS requests until it reaches its maximum, and so on. When all the zones have reached their maximum number of requests, the CSS resets the counters and the cycle starts over again. When you add a new DNS zone, each CSS adds the new zones to its list by weight. In this case, the CSSs do not reset their hit counters. This process prevents flooding of the heaviest zone every time you add or remove a zone. For example, a domain with a weight of 10 in the local zone will receive twice as many hits as the same domain with a weight of 5 in another zone. You assign domain weights using the **dns-record** command. • **roundrobin** - The CSS cycles between records available from different zones. This is the default method. • **srcip** - The CSS uses a source IP address hash to select the zone index to return to the client. • **leastloaded** - The CSS reports loads and selects a record from the zone that has the least traffic. • **preferlocal** - The CSS returns a record from the local zone whenever possible. Otherwise, the server uses the roundrobin method.

weight	(Optional) Default weight applied to all DNS records in the zone if you do not configure a weight for individual records using the **dns-record** command. Enter an integer from 0 to 10. The default is 1. To display the weight that you configured on a record using either the **dns-server zone** command or the **dns-record** command, enter the **show dns-record weight** command.
reporting	Enables the processing of local DNS server zone load information and sharing it with peers. The default is enabled.
frequency *seconds*	Specifies the period of time in seconds between processing local DNS server load information and the subsequent delivery of load information to peers. Enter an integer from 5 and 300 seconds (5 minutes). The default is 30 seconds.
variance *number*	Specifies the range of load numbers between zones that will be considered similar. If the load numbers of all zones are within the specified range, the CSS uses response times to identify the least-loaded site. Enter an integer from 0 to 255. The default is 255.

Usage Guidelines

The **dns-server zone** command is available in the CSS Enhanced feature set.

If you need to modify a **dns-server zone** value, you must first disable the DNS server using the **no dns-server** command and then remove the zone using the **no dns-server zone** command. Restore the DNS server zone with the value change, and then reenable the DNS server. To enable or disable the **dns-server zone load reporting** command, you must first disable the DNS server using the **no dns-server** command, and then enter the **dns-server zone load reporting** or the **no dns-server zone load reporting** command.

Note If you configure the absolute load calculation method for GSLB, we recommend that you configure a load variance of 0, regardless of whether you are using zone-based or rule-based DNS load balancing. For information on absolute load calculation, refer to the *Cisco Content Services Switch Content Load-Balancing Configuration Guide*.

(config) dnsflow

To set up either TCP or UDP traffic to DNS server port 53 as a CSS flow or to forward the traffic, use the **flow-state** commands. The **flow-state** commands replace the following **dnsflow** command and options:

Syntax Description

disable	This command option has been deprecated (obsoleted). If you enter **dnsflow disable** at the CLI or if it already exists in your running-config, the CSS automatically converts it to the following **flow-state** commands: • (config)# **flow-state 53 udp flow-disable nat-enable** • (config)# **flow-state 53 tcp flow-disable**
enable	This **dnsflow** command option has been removed from the CLI. Use the **flow-state** commands instead (see Note below).

Note For details about the **flow-state** commands, see the (config) flow-state port_number command.

Command Modes

Global configuration mode

Related Commands

(config) flow-state port_number

(config) domain hotlist

To enable the domain hot list, use the **domain hotlist** command. The domain hot list is disabled by default. A domain hot list lists the most accessed domains on the CSS during a user-defined period of time. Use the **no** form of this command to disable the domain hot list.

domain hotlist

no domain hotlist

Related Commands **show domain hotlist**

(config) domain hotlist interval

To configure the interval, in minutes, to refresh the domain hot list and start a new list, use the **domain hotlist interval** command. Use the **no** form of this command to reset the interval to its default setting of 1 minute.

domain hotlist interval *minutes*

no domain hotlist interval

Syntax Description

minutes	Interval in minutes. Enter an integer from 1 to 60. The default is 1.

Related Commands **show domain hotlist**

(config) domain hotlist size

To configure the maximum number of domain entries contained in the hot list, use the **domain hotlist size** command. Use the **no** form of this command to reset the maximum size to its default setting of 10 entries.

domain hotlist size *max_entries*

no domain hotlist size

Syntax Description

max_entries	Maximum number of domain hot-list entries. Enter an integer from 1 to 100. The default is 10.

Related Commands **show domain hotlist**

(config) domain hotlist threshold

To configure the threshold (the number of domain hits per interval) that must be exceeded for a domain to be considered hot and added to the list, use the **domain hotlist threshold** command. Use the **no** form of this command to reset the threshold to its default setting of 0.

domain hotlist threshold *number*

no domain hotlist threshold

Syntax Description

number	Threshold number. Enter a number from 0 to 65535. The default is 0 which indicates that the threshold is disabled.

Related Commands **show domain hotlist**

(config) dql

To access and configure a domain qualifier list (DQL), use the **dql** command. A DQL is a collection of domain names that you can assign to a content rule, instead of creating a rule for each domain.

Use the **no** form of this command to remove an existing DQL.

dql *dql_name*

no dql *existing_dql_name*

Syntax Description

dql_name	Name of a new DQL you want to create or of an existing list. Enter an unquoted text string with no spaces and a maximum of 31 characters. To see a list of existing DQL names, enter: `dql ?`

Usage Guidelines

When you use the **dql** command to access DQL mode, the prompt changes to (config-dql [*name*]). You can also use this command from DQL mode to access another DQL. For information about commands available in this mode, see the "DQL Configuration Mode Commands" section.

Related Commands

show dql
(config-owner-content) url

(config) dump

To enable or disable core dumps when the CSS experiences a fatal error, use the **dump** command. Core dumps are enabled by default.

Note Core dump information is for customer support use only.

dump [**disable**|**enable**]

Syntax Description

disable	Disables core dumps. When the CSS experiences a fatal error and core dumps are disabled, the CSS reboots automatically. The CSS does not write information to the hard disk or flash disk.
enable	Enables core dumps. This is the default setting. When the CSS experiences a fatal error and core dumps are enabled, the CSS: • Writes information about the fatal error to the Core directory of the volume root (for example, c:\core) on either the hard or flash disk. On the hard or flash disk stores one dump file per slot per card type until the disk is full. Files can be 10 to 20 MB in size. • Reboots automatically

Usage Guidelines

For a flash disk-based system, if the core dump file is older than 15 minutes, it may be overwritten. If you want to save the core dump file for later examination, archive it to another directory or disk before it is overwritten. To archive a log file, see the **archive log** command.

Related Commands

show core

(config) eql

To access EQL configuration mode and configure an extension qualifier list (EQL), use the **eql** command. This list is a collection of file extensions for content requests joined together through content rules. The CSS uses this list to identify which requests to send to a service.

Use the **no** form of this command to delete an existing extension list.

eql *eql_name*

no eql *existing_eql_name*

Syntax Description

eql_name	Name of a new extension list you want to create or of an existing list. Enter an unquoted text string with no spaces and a maximum length of 31 characters. To see a list of existing EQL names, enter: `eql ?`

Usage Guidelines

When you use the **eql** command to access eql mode, the prompt changes to (config-eql [*name*]). For information about commands available in this mode, see the "EQL Configuration Mode Commands" section.

Related Commands

show eql
(config-owner-content) url

(config) flow-state *port_number*

To set the flow states of TCP and UDP ports in the CSS flow-state table, use the **flow-state** command. Use the **no** form of the command to disable the flow state.

flow-state *port_number* **tcp** [**flow-enable**|**flow-disable**]

flow-state *port_number* **udp** [**flow-enable**|**flow-disable** {**nat-enable**|**nat-disable**}]

no flow-state *number* **tcp**|**udp**

Syntax Description

number	TCP or UDP port number on which you want to configure the flow state.
tcp	Specifies a TCP port.
udp	Specifies a UDP port.
flow-enable	Enables flows on the specified TCP or UDP port. With this option, the CSS performs full content-rule and source-group matching, including Layer 5 (IP address, destination port, and URL) content-based load balancing and sticky.
flow-disable	Disables flows on the specified TCP or UDP port. When you disable flows on a port, the CSS does not perform content rule and source group matching. The benefit is no flow setup overhead.
nat-enable	(Optional) For flow-disabled UDP ports, enables content-rule and source-group lookups for NAT. With this option, you can use Layer 3 (IP address) and Layer 4 (IP address and destination port) content rules and the sticky table (for example, **sticky-srcip**). However, without the benefit of a flow, the CSS cannot spoof the back-end connection, which is required to make Layer 5 content-based decisions.
nat-disable	(Optional) For flow-disabled UDP ports, the CSS does not perform content-rule and source-group lookups for NAT.

Usage Guidelines

By default, Domain Name Service (DNS) port 53 (TCP and UDP) and SIP port 5060 (UDP) are flow-enabled. You can change the flow states of the preconfigured ports, and you can configure any 16 unique TCP or UDP ports and their flow states.You can also set the port address translation (PAT) state for flow-disabled UDP ports only.

For more information, refer to the *Cisco Content Services Switch Content Load-Balancing Configuration Guide*.

Related Commands

(config) flow-state flow-disable timeout
(config) zero flow-state-counters
show flow-state-table

(config) flow-state flow-disable timeout

To set the wait time for any response from a server for a configured flow-disable port, use the **flow-state flow-disable timeout** command. Use the **no** form of the command to reset the default flow-disable timeout to 5 seconds.

flow-state flow-disable timeout *seconds*

no flow-state flow-disable timeout

Syntax Description

seconds	Time in seconds. Enter an interger from 5 to 20. The default value is 5.

Usage Guidelines

By default, the CSS times out a flow-disable (no flow) connection in 5 seconds if it does not receive a response from the server. In the case of DNS responses, they may take longer than 5 seconds causing the connection to fail. By using the **flow-state flow-disable timeout** command to set a longer wait time for server responses, these connections are less likely to fail.

Related Commands

(config) flow-state port_number
show flow-state-table

(config) flow permanent

To define a set of TCP or UDP ports that will have permanent connections and will not be reclaimed by the CSS when the flows are inactive, use the **flow permanent** command. By default, the CSS may reclaim TCP/UDP flows that have not received an ACK or content request after approximately 15 seconds. Use the **no** form of this command to disable a permanent connection by setting its port number to 0.

flow permanent [**port**[**1**|**2**|**3**|**4**|**5**|**6**|**7**|**8**|**9**|**10**]] *port_number*

no flow permanent [**port**[**1**|**2**|**3**|**4**|**5**|**6**|**7**|**8**|**9**|**10**]]

Syntax Description

port_number	Number of the port. Enter an integer from 0 to 65535. The default is 0, which disables a permanent connection on the port.

Usage Guidelines

Entering the **flow permanent** command disables Denial of Service protection and reclaiming of ports when there is asymmetrical routing on any flow with the specified transport layer port as a source or destination of a flow.

Do not configure the **flow permanent** command without enabling the **cmd-sched** command to periodically remove the permanent port and allow for cleanup. For details on using the **cmd-sched** command to configure the scheduled execution of any CLI commands, refer to the *Cisco Content Service Switch Administration Guide*.

(config) flow persist-span-ooo

To enable the reordering of persistent spanning packets, use the **flow persist-span-ooo** command. By default, the CSS disables the reordering of persistent spanning packets. Use the **no** form of this command to reset the default behavior of disabling the reordering of persistent spanning packets.

flow persist-span-ooo

no flow persist-span-ooo

(config) flow set-port-zero

To enable or disable the CSS to pass traffic using port of 0 using a TCP/UDP source and destination port of 0, use the **flow set-port-zero** command. By default, the CSS disables the passing of traffic that use TCP or UDP source or destination port 0. The CSS normally logs traffic with source or destination ports of 0 as a denial-of-service (DOS) attacks. If you enable traffic on port 0, the CSS does not log the flows as denial-of-service attacks.

flow set-port-zero enable | disable

Syntax Description

enable	Enables the passing of traffic that use a TCP/UDP source and destination port of 0
disable	Resets the CSS to its default behavior of not passing traffic using a TCP/UDP source and destination port of 0

(config) flow tcp-del-ack

By default, when an HTTP content request spans multiple packets, the CSS sends delayed TCP acknowledgements (ACKs) to the client at an interval of 200 milliseconds (ms) if the full HTTP content request is not received. To reenable TCP delayed ACKs for Layer 5 spanning packets, use the **flow tcp-del-ack** command. Use the **no** form of this command to send TCP ACKs immediately to a client upon receipt of each packet in an HTTP spanned content request.

flow tcp-del-ack

no flow tcp-del-ack

(config) flow tcp-mss

To configure the TCP maximum segment size (MSS), use the **flow tcp-mss** command. Use the **no** form of this command to reset the TCP maximum segment size to the default value of 1460 bytes.

flow tcp-mss *size*

Syntax Description

size	Maximum segment size (in bytes) from 1 to 1460. The default is 1460 bytes. Do not define a very small segment size. Smaller payloads may be less efficient due to increased overhead.

Usage Guidelines

The **flow tcp-mss** command applies only when the client is accessing a Layer 5 content rule. The CSS does not negotiate a TCP maximum segment size for Layer 3 or Layer 4 content rules. The MSS is the largest piece of TCP data that the CSS expects to receive from the other end. This command changes the MSS value in the TCP header options field of a SYN/ACK segment back to the client.

(config) flow tcp-reset-on-vip-unavailable

To configure a CSS to send a TCP RST (reset) to a client when a VIP is unavailable, use the **flow tcp-reset-on-vip-unavailable** command. Use the **no** form of this command to return the CSS behavior to the default of dropping the TCP packet when a VIP is unavailable.

flow tcp-reset-on-vip-unavailable

no flow tcp-reset-on-vip-unavailable

Usage Guidelines

The CSS sends the TCP reset only in response to a TCP packet that is destined for a VIP that the CSS is hosting and only if that VIP is unavailable.

Related Commands

show ip statistics

(config) ftp data-channel-timeout

To configure the wait time to initiate the FTP data channel on an active or passive FTP connection for CSS FTP content rules and source groups, use the **ftp data-channel-timeout** command. By default, the CSS waits 5 seconds to initiate the FTP data channel on an active or passive FTP connection for CSS FTP content rules and source groups. Use the **no** form of the command to reset the wait time to 5 seconds.

ftp data-channel-timeout *seconds*

no ftp data-channel-timeout

Syntax Description

seconds	The wait time in seconds. Enter a number from 5 to 120. The default value is 5.

(config) ftp non-standard-ports

To enable the CSS to handle FTP connections ona non-standard FTP control port, use the **ftp non-standard-ports** command. By default, this setting is disabled. Use the **no** form of this command to reset the default behavior of requiring the FTP connection to use the standard control port of 21.

ftp non-standard-ports

no ftp non-standard-ports

Usage Guidelines

When disabled, the CSS requires the FTP connection to use the standard FTP port. The CSS preserves and does not NAT the FTP data port when the FTP data connection is passed through the CSS.

When enabled with the **ftp non-standards-ports** command, the CSS allows the FTP control connection to use a non-standard port, not port 21. The CSS does not preserve the FTP data port when the FTP data connection is passed through the CSS.

When you use the **ftp non-standards-ports** command to allow the use of a non-standard FTP port and a content rule is using FTP, you must configure the **application ftp-control** command on the content rule.

Related Commands

(config-owner-content) application ftp-control

(config) ftp-record

To create a File Transfer Protocol (FTP) record file to use when accessing an FTP server from the CSS, use the **ftp-record** command. Use the **no** form of this command to delete an FTP record file from the CSS.

ftp-record *ftp_record ip_or_host username* [**"***password***"** |**des-password** *des_pwd*] {**base_directory**}

no ftp-record *ftp_record*

Syntax Description

ftp_record	Name for the FTP record file. Enter an unquoted text string with no spaces and a maximum length of 16 characters.
ip_or_host	IP address or host name of the FTP server you want to access. Enter an IP address in dotted-decimal notation (for example, 192.168.11.1) or a mnemonic host name (for example, myhost.mydomain.com).
username	Valid login username on the FTP server. Enter a case-sensitive unquoted text string with no spaces and a maximum length 16 characters.
"*password*"	Password for the valid login username on the FTP server. Enter a case-sensitive quoted text string with no spaces and a maximum length of 16 characters.
des-password *des_pwd*	Specifies the Data Encryption Standard (DES) encrypted password for the valid login username on the FTP server. Enter a case-sensitive unquoted text string with no spaces and a maximum length of 64 characters.
base_directory	(Optional) Base directory when using this record. Enter a case-sensitive unquoted text string with no spaces and a maximum length of 64 characters.

Usage Guidelines

The CSS FTP server supports only the active (normal) FTP mode of operation. It does not support the passive FTP mode of operation.

Related Commands

copy ftp
copy log
copy running-config
copy script
copy ssl
(config-boot) primary
(config-boot) secondary

(config) global-portmap

To control the global source-port translation (port mapping) for TCP flows on a CSS, use the **global-portmap** command. Use the **no** form of this command to reset the starting port number and the port range to their default values.

global-portmap base-port *number1* **range** *number2*

no global-portmap

Syntax Description

base-port *number1*	Starting port number for global port mapping on a CSS. Enter an integer from 2016 to 63456. The default is 2016. **Caution** Changing the value of the *number1* variable may cause port conflicts on existing flows.
range *number2*	The total number of ports in the port-map range that the CSS allocates to each of the 16 megamap banks in each SP. Each megamap bank in an SP can use the full range of configured ports. Because of the unique source address hash that the CSS uses to select a megamap bank in an SP, more than one SP can use the same port number without a tuple collision. **Caution** Changing the value of the *number2* variable may cause port conflicts on existing flows. Enter an integer from 2048 to 63488. The default is 63488. If you enter a value that is not a multiple of 32, the CSS rounds up the value to the next possible multiple of 32. If you enter a portmap range that exceeds the number of available ports, you get an error. To determine the number of available ports, subtract the starting port number you specify from 65504.

Usage Guidelines

The global portmapper in a CSS is called the megaportmapper. The megaportmapper database comprises 16 banks of portmap numbers (megamap banks) in each session processor (SP) with unique ranges. A CSS uses a source port hash algorithm to select a megamap bank.

Use the **global-portmap** command to control the global source-port translation (port mapping) for TCP flows on a CSS. This command is always enabled. Use this command to specify the source-port mapping range on:

- A Cisco 11500 series CSS when you configure a service that uses a nondefault destination port number. A CSS changes a TCP destination port number configured on a service in a content rule when a request hits the content rule and the CSS sends a packet to the selected server. The CSS uses the **global portmap** command parameters to translate the corresponding client source port number to distinguish it from other clients requesting the same service.
- A redundant Cisco 11500 series CSS peers in a session-level redundancy configuration. For information on session-level redundancy, refer to the *Cisco Content Services Switch Redundancy Configuration Guide*.
- Any CSS with back-end server remapping enabled (refer to the *Cisco Content Services Switch Content Load-Balancing Configuration Guide*).

Note When you configure a source group, the **portmap** command values take precedence over the **global-portmap** command. For details on configuring the **portmap** command in a source group, refer to the *Cisco Content Services Switch Content Load-Balancing Configuration Guide*. Note that the **portmap disable** command has no effect on TCP flows.

Related Commands

show global-portmap
(config-group) portmap

(config) group

To access group configuration mode and configure a group, use the **group** command. A group is a collection of local servers that initiate flows from within the local web farm. For example, after processing a group of real audio transmitters, they all appear on the same source IP address. The CSS lets you treat a group as a virtual server with its own source IP address.

Use the **no** form of this command to delete an existing group.

group *group_name*

no group *existing_group_name*

Syntax Description

group_name	Name of a new group you want to create or of an existing group. Enter an unquoted text string with no spaces and a maximum length of 31 characters. To see a list of existing group names, enter: `group ?`

Usage Guidelines

When you use the **group** command to access group mode, the prompt changes to (config-group [*name*]). For information about commands available in this mode, see the "Group Configuration Mode Commands" section.

Caution Before you use the **no group** command to delete an existing group, make sure you want to permanently delete the group. You cannot undo this command. If you want a prompt before the CSS performs a command, use the **no expert** command.

(config) gsdb

To start the global sticky database (GSDB) on a dedicated CSS 11150 with 256 MB of RAM when you are configuring GSLB with a GSDB or using DNS Sticky in a Network Proximity configuration, or specify a time-to-live (TTL) interval for the GSDB sticky domain entries, use the **gsdb** command. Use the **no** form of this command to disable the GSDB or reset the TTL interval for GSDB entries to 7200 seconds.

gsdb {**ttl** *seconds*}

no gsdb {**ttl**}

Syntax Description

ttl	(Optional) Specifies the time-to-live interval for the GSDB entries.
seconds	(Optional) Time-to-live interval in seconds. The value you enter determines the length of time that GSDB entries are valid. Enter a number from 300 to 1000000. The default value is 7200. Any new request from a D-proxy for a sticky domain that arrives before the timer expires resets the timer.

Usage Guidelines

Because the GSDB is dependent upon the presence of the PDB, you must configure the PDB prior to starting the GSDB.

You do not need to configure a GSDB to use the basic DNS Sticky feature in a global server load-balancing (GSLB) environment. However, a GSDB provides a more robust DNS Sticky and load-balancing configuration. For details on the types of DNS Sticky configurations, refer to the *Cisco Content Services Switch Global Server Load-Balancing Configuration Guide*.

Related Commands

gsdb zero
show gsdb

(config) gsdb zero

To reset the Sticky Lookups and Sticky Sets statistics that are displayed by the **show gsdb** command, use the **gsdb zero** command. This command applies only to a CSS 11150 with 256 MB of RAM that is configured as a global sticky database (GSDB).

The syntax for this global configuration mode command is:

gsdb zero

Related Commands

gsdb
show gsdb

(config) gsdb-interface

To create a primary or secondary interface to the GSDB on the CSS DNS server to communicate with a GSDB, or zero the GSDB interface statistics, use the **gsdb-interface** command. Use the **no** form of this command to remove a primary or secondary GSDB interface.

gsdb-interface [**primary** *ip_address*|**secondary** *ip_address*|**zero**]

no gsdb-interface [**primary**|**secondary**]

Syntax Description

primary	Specifies the primary interface for the GSDB. The CSS uses the primary GSDB for sticky requests.
secondary	Specifies the secondary interface for the GSDB. The CSS uses the secondary interface when the primary interface is unavailable.

ip_address	IP address of the GSDB. Enter the address in dotted-decimal notation (for example, 192.168.11.1). In a Network Proximity configuration, the IP address of the primary sticky interface is typically the same as the IP address of the PDB.
zero	Resets the GSDB interface statistics that are displayed by the **show gsdb-interface** command.

Usage Guidelines

The **gsdb-interface** command is part of the Enhanced feature set.

A GSDB responds with a zone index to sticky queries from CSS DNS servers. All GSDBs participating in a peer mesh share sticky TTL and sticky zone information over APP.

Related Commands

show gsdb-interface

(config) header-field-group

To access header-field-group configuration mode and configure a request header-field group, use the **header-field-group** command. A request header-field group contains a list of defined header-field entries used by the content rule lookup process. Each header-field group is given a unique name so different content rules can use them. A group can contain several header-field entries. Use the **no** form of this command to remove a header-field group.

header-field-group *group_name*

no header-field-group *group_name*

Syntax Description

group_name	Header-field group that you want to configure. You must define a unique name for each header-field group so different content rules can use the groups. Enter a text string with a maximum of 32 characters. To see an existing list of header-field groups, enter: `header-field-group ?`

Usage Guidelines

To access header-field-group configuration mode, use the **header-field-group** command from all configuration modes, except boot and RMON modes. The prompt changes to (config-header-field-group [*group_name*]). You can also use this command in header-field-group mode to access another group. For information about commands available in this mode, see the "Header-Field Group Configuration Mode Commands" section.

Note When there is more than one header-field entry in a group, each header-field entry must be successfully matched before the CSS uses the associated content rule.

Related Commands

show header-field-group
(config-owner-content) header-field-rule

(config) host

To manage entries in the Host table, use the **host** command. The Host table is the static mapping of mnemonic host names to IP address, analogous to the ARP table. Use the **no** form of this command to remove an existing host from the Host table.

host *host_name ip_address*

no host *host_name*

Syntax Description

host_name	Name of the host. Enter an unquoted text string with no spaces and a maximum length of 16 characters. To see a list of host names, enter: `show running-config global`
ip_address	IP address associated with the host name. Enter the IP address in dotted-decimal notation (for example, 192.168.11.1).

Usage Guidelines

To add a host to the Host table, the host name must not already exist. To change a current host's address, remove it and then add it again.

Related Commands

show running-config

(config) http-method parse

To configure the CSS to support all HTTP methods defined in RFC-2518 including RFC-2616 and configure user-defined methods, use the **http-method parse** command. Use the **no** form of this command to disable the parsing of RFC-2518 extension methods or remove a user-defined method.

http-method parse RFC2518-methods|user-defined-method *METHOD_NAME* {**uri** [**wildcard|authority|url**]}

no http-method parse RFC2518-methods|user-defined-method *METHOD_NAME*

Syntax Description

RFC2518-defined-method	Enables the CSS to support the extension methods defined in RFC2518.
user-defined-method	Allows you to configure a maximum of 16 user-defined methods
METHOD_NAME	The name for the method that performs the processing of the Request-URI field. Enter the name as an unquoted alphanumeric text string with a minimum of 3 characters and a maximum of 15 characters, including the hyphen (-) and underscore (_) characters. You must capitalize the alphabetic characters. You cannot use control, tspecial, and space characters. The tspecial characters include “(”, “)”, “<”, “>”, “@”, “,”, “;”, “:”, “\”, “"”, “/”, “[”, “]”, “?”, “=”, “{”, “}”, space (SP), and horizontal tab (HT) characters. The method name cannot conflict with currently supported HTTP methods defined in RFC-2616, extension methods defined in RFC-2518, or user-defined method.
uri	(Optional) Allows you to identify the resource in the Request-URI field that is applied to the method. By default, the method processes a resource in an absoluteURI or absolute path format.

authority	Indicates that the Request-URI field contains an authority format. The CONNECT method is the only method that uses the authority format. For example: `CONNECT` **`server.cisco.com:80`** `HTTP/1.1`

url	Indicates that the Request-URI field contains an absoluteURI or absolute path format. This keyword is the default resource. The absoluteURI format is required when the request is to a proxy. The proxy either forwards the request or services it from a valid cache, and then returns a response. For example:
	`GET` **`http://www.test3.org/pub/www/Test.html`** `HTTP/1.1`
	
	The absolute path identifies a resource on an origin server or gateway. The absolute path of the URI is transmitted as the Request-URI, and the network location of the URI (authority) is transmitted in a Host header field.
	`GET` **`/index.html`** `HTTP/1.1`
	`Host:www.test3.org`

wildcard	Indicates that the Request-URI field can contain a wildcard (*) character, an absolute URI, or an absolute path. The wildcard character indicates that the request does not apply to a particular resource, but to the server itself, and is only allowed when the method used does not necessarily apply to a resource. For example: `OPTIONS` **`*`** `HTTP/1.1` When you configure a user-defined method with a wildcard URI, you must configure a Layer 5 rule (url "/*") with a header-field group that contains a request line with the method name for the CSS to match user-defined method to the rule.

Usage Guidelines

By default, a Layer 5 content rule supports the HTTP CONNECT, GET, HEAD, POST, and PUT methods. Unless configured, the CSS recognizes and forwards the following HTTP methods directly to the destination server in a transparent caching environment, but does not load balance them: OPTIONS, TRACE, PROPFIND, PROPPATCH, MKCOL, MOVE, LOCK, UNLOCK, COPY, and DELETE.

When you enable the CSS to support all RFC-2518 methods, the CSS parses the Request-URI field in an attempt to match a Layer 5 rule. If the contents of the Request-URI field are not in a compliant format of an absolute URI or an absolute path, the CSS tries to match the field to the next best wildcard ("/*") rule. If the match fails, the CSS attempts to match the Layer 4 rule, and then the Layer 3 rule.

The CSS provides scripts for the configuration of RFC-2518 and custom methods required for Outlook Web Access (OWA). The **setup_owa_methods** script enables RFC-2518 methods and configures the POLL, SEARCH, SUBSCRIBE, BMOVE, BCOPY, BDELETE, and BPROPPATCH user-defined methods. The **remove_owa_methods** script disables the RFC-2518 methods and removes the OWA methods configured with the **setup_owa_methods** script.

Related Commands **show http-methods**

(config) http-method statistics clear

To clear the Hit Counter fields for the methods displayed through the **show http-methods** command, use the **http-method statistics clear** command.

http-method statistics clear

Related Commands **show http-methods**

(config) http-redirect-option

To configure the CSS to send specific TCP FIN and RST flags with HTTP 302 redirect messages, use the **http-redirect-option** command. By default, when the CSS sends an HTTP 302 redirect message, it sends a FIN flag on an initial connection and RST flags on subsequent requests in a persistent connection.

http-redirect-option [**fin-rst**|**fin-fin**|**rst-rst**]

Syntax Description

fin-rst	Sends a FIN flag for initial connections and an RST flag for persistent connection (default)
fin-fin	Always sends a TCP FIN flag
rst-rst	Always sends a TCP RST flag

Usage Guidelines

When the CSS sends packets to a client that contains a redirect message to a Microsoft IE browser, use the **http-redirect-option** command to select a behavior that is suitable for the browser.

Related Commands

show http-redirect-option

(config) idle timeout

To set the maximum amount of time that any Telnet, console, FTP, or web management session can be idle on the CSS before the CSS logs it out, use the **idle timeout** command. Use the **no** form of this command to set the idle timeout for the session connected to the CSS to the default of 0.

idle timeout {**web-mgmt**} *minutes*

no idle timeout {**web-mgmt**}

Syntax Description

web-mgmt	(Optional) Sets the maximum amount of idle time for active web management sessions. This option does not apply to the CVDM available in the CSS software release 8.10 or greater.
minutes	Maximum time in minutes. Enter a number from 0 to 65535. The default is 0.

Usage Guidelines

The **idle timeout** command without the **web-mgmt** option sets the global timeout for Telnet, console, SSH, and FTP sessions.

The **web-mgmt** option does not apply to the CVDM available in the CSS software release 8.10 or greater.

You can override the **idle timeout** command with the **terminal** command in SuperUser mode for Telnet, console, SSH, and FTP sessions.

(config) interface

To enter interface configuration mode and configure an interface, use the **interface** command.

interface *interface_name*

Syntax Description

interface_name	CSS interface that you want to configure. For a CSS 11501, enter the interface name in *interface-port* format (for example, e2). For a CSS 11503 or 11506, the interface format is *slot*/*port* (for example, 3/1). To see a list of valid interfaces for this CSS, enter: `interface ?`

Usage Guidelines

When you use the **interface** command to access this mode, the prompt changes to (config-if [*interface_name*]). For information about commands available in this mode, see the "Interface Configuration Mode Commands" section.

(config) ip

To enter global IP configuration commands, use the **ip** command. The options for this global configuration mode command are:

- **ip advanced-route-remap** - Remaps flows using the best available route
- **ip ecmp** - Sets the equal-cost multipath selection algorithm
- **ip firewall** - Configures an index that identifies a physical firewall
- **ip management no-icmp-redirect** - Configures the Ethernet management port to discard ICMP redirect packets
- **ip management route** - Configures a static route for the Ethernet management port
- **ip no-implicit-service** - Does not allow the CSS to start an implicit service for the next hop of static routes
- **ip opportunistic** - Configures opportunistic Layer-3 forwarding

- **ip record-route** - Enables processing of frames with a record-route option
- **ip redundancy** - Enables CSS-to-CSS redundancy
- **ip route** - Configures a static route
- **ip source-route** - Enables processing of source-routed frames
- **ip subnet-broadcast** - Enables forwarding of subnet broadcast addressed frames

For more information on these options and associated variables, see the following commands.

Related Commands **show ip config**
show ip summary

ip advanced-route-remap

To configure the CSS to remap flows using the best-available route, use the **ip advanced-route-remap** command. Use the **no** form of this command to disable the remapping of flows using the best-available route.

ip advanced-route-remap

no ip advanced-route-remap

Command Modes Global configuration mode

ip ecmp

To set the equal-cost multipath selection algorithm and the preferred reverse egress path, use the **ip ecmp** command. Use the **no** form of this command to reset the ingress path of a flow for its preferred reverse egress path.

ip ecmp [**address**|**no-prefer-ingress**|**roundrobin**]

no ip ecmp no-prefer-ingress

Syntax Description

address	Chooses among alternate paths based on IP addresses.
no-prefer-ingress	Does not prefer the ingress path of a flow for its reverse egress path. By default, the ingress path for a flow is its preferred egress path.
roundrobin	Alternates between equal paths in roundrobin fashion.

Command Modes

Global configuration mode

Usage Guidelines

The equal-cost multipath selection algorithm for non-TCP/UDP packets (for example, ICMP) is applied on a packet-by-packet basis. Multipath selection for TCP and UDP is performed on a per-flow basis and all packets for a particular flow take the same path.

ip firewall

To configure an index that identifies a physical firewall, use the **ip firewall** command. Use the **no** form of the **ip firewall** *index* command to delete a firewall index. Use the **no** form of the **ip firewall timeout** command to reset the firewall timeout to the default value of three seconds.

ip firewall [*index local_firewall_address remote_firewall_address remote_switch_address*|**timeout** *seconds*]

no ip firewall [*index*|**timeout**]

Syntax Description

index	Index number to identify the firewall. Enter a number from 1 to 254.
local_firewall_address	IP address of the firewall on a subnet connected to the CSS. Enter the IP address in dotted-decimal notation (for example, 192.168.11.1).
remote_firewall_address	IP address of the firewall on the remote subnet that connects to the remote switch. Enter the IP address in dotted-decimal notation (for example, 192.168.11.1).
remote_switch_address	IP address of the remote CSS. Enter the IP address in dotted-decimal notation (for example, 192.168.11.1).
timeout *seconds*	Number of seconds that the CSS waits to receive a keepalive message from the remote CSS before declaring the firewall to be unreachable. The timeout range is 3 to 16 seconds. The default is 3 seconds.

Command Modes

Global configuration mode

Usage Guidelines

You can configure indices for multiple parallel firewalls allowing for traffic load balancing. To avoid dropping packets, all connections in either direction between a pair of IP addresses cross the same firewall. If a failure occurs on one path, all traffic uses the remaining path.

A CSS must exist on each side of the firewall to control which firewall is selected for each flow. You must configure a firewall index identifier on the remote CSS with the same index number to the same physical firewall.

To configure the firewall route, use the **ip route** command. Firewalls cannot perform Network Address Translation (NAT). If your configuration requires NATing, you must configure a content rule or source group on the CSS to provide this function.

Caution When you delete a firewall index, all routes associated with that index are also deleted.

The two CSS switches at the endpoints of the firewall configuration must use the same firewall keepalive timeout value. Otherwise, routes on one CSS may not fail over simultaneously with those on the other CSS. This could permit asymmetric routing to occur across the firewalls.

Related Commands

ip route

ip management no-icmp-redirect

To configure the CSS to discard ICMP redirect packets on the Ethernet management port, use the **ip management no-icmp-redirect** command. By default, the Ethernet management port accepts all incoming ICMP redirect packets. Use the **no** form of this command to reset the default behavior of accepting ICMP redirect packets on the Ethernet management port.

ip management no-icmp-redirect

no ip management no-icmp-redirect

Command Modes

Global configuration mode

Usage Guidelines

If you do not configure static routes for the Ethernet management port, the CSS disregards any ICMP redirects. However, when you configure static routes for the Ethernet management port, the CSS incorporates the ICMP redirects as entries in the routing table.

To enhance security on the CSS when you configure static routes on the Ethernet management port, we strongly recommend that you configure the CSS Ethernet management port to discard ICMP redirects.

The Ethernet management port never transmits an ICMP redirect.

If you remove a static route when the Ethernet management port is configured to accept ICMP redirect packets, the CSS removes the router entry created by the ICMP redirect associated with the static route from the routing table.

Related Commands

show ip config

ip management route

The ability to configure static routes on the Ethernet management port provides access to the CSS from hosts on subnets that are different from the Ethernet management port subnet. To manage the CSS from a subnet that is different from the Ethernet management port, use the **ip management route** command. By default, this option is disabled. Use the **no** form of this command to disable a static route for the Ethernet Management port.

ip management route *ip_address1 subnet mask ip_address2*

no ip management route *ip_address1 subnet mask ip_address2*

Syntax Description

ip_address1	The destination network address. Enter the IP address in dotted-decimal notation (for example, 192.168.11.0).
subnet_mask	The IP subnet mask. Enter the mask as either: • A prefix length in CIDR bitcount notation (for example, /24). Do not enter a space to separate the IP address from the prefix length. • An IP address in dotted-decimal notation (for example, 255.255.255.0).
ip_address2	The next hop address for the route. Enter the IP address in dotted-decimal notation (for example, 172.16.6.1).

Command Modes

Global configuration mode

Usage Guidelines

You can configure a maximum of eight static routes for the Ethernet management port.

The CSS does not use an internal (implicit) service for the Ethernet management port to periodically poll the next hop address in a static route. The periodic polling of the next hop address with an ICMP echo (or ping) keepalive is performed only when you configure a static route for an Ethernet interface port.

The **rip redistribute static** and **ospf redistribute static** commands do not advertise static routes configured on the Ethernet management port. These two commands only advertise static routes configured on the Ethernet interface ports.

ip no-implicit-service

To stop the CSS from starting an implicit service for the next hop of static routes, use the **ip no-implicit-service** command. By default, this option is disabled. Use the **no** form of this command to reset the default setting.

ip no-implicit-service

no ip no-implicit-service

Command Modes

Global configuration mode

Usage Guidelines

By default, the CSS establishes an implicit (or internal) service for the gateway address when a static route is defined. The **ip no-implicit-service** command specifies that no implicit service is established to the next hop of the static route.

Note When you implement the **ip no-implicit-service** command, it does not affect any previously configured static routes. If you wish to stop the implicit service for a previously configured static route, you must delete and reconfigure that static route.

The purpose of the implicit service to the next hop of a static route is to monitor the availability of the next hop to forward data traffic. When the **ip no-implicit-service** command is in effect, traffic will be forwarded to the next hop even when it is unavailable. Because of the possibility of data loss if the next hop becomes unavailable, use of the **ip no-implicit-service** command is strongly discouraged.

ip opportunistic

To configure the opportunistic Layer 3 forwarding of packets, use the **ip opportunistic** command. Use the **no** form of this command to allow opportunistic Layer 3 forwarding for local destinations.

ip opportunistic [**all**|**disable**]

no ip opportunistic

Syntax Description

all	Allows opportunistic Layer 3 forwarding for all destinations; when the IP destination address matches any routing entry on the CSS. This mode is not recommended for a topology that includes multiple routers and the CSS does not know all the routes that the routers know.
disable	Disables opportunistic Layer 3 forwarding. Layer 3 forwarding only occurs for packets whose destination MAC address belongs to the CSS.

Command Modes

Global configuration mode

Usage Guidelines

Opportunistic Layer 3 forwarding allows the CSS to forward packets according to the IP destination address. The MAC destination address does not need to belong to the CSS. By default, the CSS allows this forwarding for local destinations when the IP destination address belongs to a node that resides on one of the subnets directly attached to the CSS and an ARP resolution is known for this node.

ip record-route

To enable the CSS to process frames with a record-route option, use the **ip record-route** command. Use the **no** form of this command to disable the processing of frames with a record-route option (the default behavior).

ip record-route

no ip record-route

Refer to the *Cisco Content Services Switch Routing and Bridging Configuration Guide* for more information about this command.

Command Modes Global configuration mode

ip redundancy

To enable CSS-to-CSS redundancy on two CSSs interfaced with a crossover cable, use the **ip redundancy** command. You can also use the **master** option to manually designate which CSS is the master. By default, redundancy is disabled on a CSS. Use the **no** form of the **ip redundancy** command to disable CSS-to-CSS redundancy. Use the **no** form of the **ip redundancy master** to unassign the CSS as the master CSS.

ip redundancy {**master**}

no ip redundancy {**master**}

Syntax Description

master	(Optional) Enables CSS-to-CSS redundancy on the CSS that you want to designate as the master CSS. Do not enter this command on both the master and backup CSSs. You can enter this command option on the CSS: • Whether it was initially booted as the master or the backup. If you enter this command on the backup CSS, it becomes the master and the other CSS automatically becomes the backup CSS. • When CSS-to-CSS redundancy is currently enabled.

Command Modes

Global configuration mode

Usage Guidelines

If you have no requirement to designate a specific CSS as the master, use the **ip redundancy** command with no keyword on each CSS. When you do not manually designate a master CSS, the CSSs negotiate to determine the master and backup. In this negotiation, the master CSS is the CSS that boots first. If both CSSs boot at the same time, the CSS with the higher IP address becomes the master. When the master CSS goes down, the backup CSS automatically becomes the master. When the former master CSS comes up again, it becomes the backup CSS.

To manually designate a CSS as the master CSS, enter the **master** option on it. You can enter this option on a negotiated master or backup. If you enter this option on a master, it remains the master. If you enter this option on the backup CSS, it becomes the master and the other CSS automatically becomes the backup.

Caution Do not enter the **ip redundancy master** command on both the master and backup CSSs. This can cause network problems.

Because the designated master CSS saves its configuration setting in the running-config, if it goes down and then comes up again, it regains its master status. For example, when the master CSS goes down, the backup CSS becomes master. When the former master CSS comes up again, it becomes the master again.

You cannot use the **ip redundancy master** command if you previously used the **(config-if) redundancy-phy** or **(config-service) type redundancy-up** command. Before you can use the **ip redundancy master** command, you must enter the **(config-if) no redundancy-phy** or **(config-service) no type** command.

The **no ip redundancy master** command does not disable CSS-to-CSS redundancy.

The CSS does not support simultaneous CSS-to-CSS redundancy and VIP redundancy configurations.

The CSS does not support a trace route of a redundant IP interface.

Related Commands

redundancy force-master
show redundancy
(config-if) redundancy-phy
(config-circuit) redundancy
(config-circuit-ip) redundancy-protocol

ip route

To configure a static route including routes for firewalls, use the **ip route** command. Use the **no** form of this command to remove a blackhole, static, or firewall route.

ip route *ip_address subnet_mask* [**blackhole**|*ip_address2* {*distance*|**originated-packets**}|**firewall** *index* {*distance*}]

no ip route *ip_address subnet_mask* [**blackhole**|*ip_address2* |**firewall** *index*]

Syntax Description

ip_address	Destination network address. Enter the IP address in dotted-decimal notation (for example, 192.168.11.1).
subnet_mask	IP subnet mask. Enter the mask as either: • A prefix length in CIDR bitcount notation (for example, /24). Do not enter a space to separate the IP address from the prefix length. • A subnet mask in dotted-decimal notation (for example, 255.255.255.0).
blackhole	Instructs the CSS to drop any packets addressed to the route.
ip_address2	Next hop address for a static route. Enter the IP address in dotted-decimal notation (for example, 192.168.11.1).
distance	(Optional) Administrative distance. Enter an integer from 1 to 254. A smaller number is preferable. The default value is 1.
firewall	Configures a firewall route.

index	Existing index number for the firewall route. For information on configuring a firewall index, see the **ip firewall** command.
originated-packets	(Optional) Instructs the CSS to use this route for flow and session packets going to and from the CSS (for example, a Telnet session to the CSS). Flows or session packets that go through the CSS (for example, between an attached server and a remote client) do not use this route.

Command Modes Global configuration mode

Usage Guidelines The CLI prevents you from configuring IP static routes that *are* firewall routes and IP static routes that *are not* firewall routes to identical destinations using identical administrative costs.

Note Ping responses and SNMP responses do not use the originated-response route. Ping *requests* sent from the CSS use the originated-response route. Ping *responses* sent from the CSS do not use the originated-response route.

ip source-route

To enable the processing of source-routed frames, use the **ip source-route** command. Use the **no** form of this command to disable the processing of source-routed frames (the default behavior).

ip source-route

no ip source-route

Refer to the *Cisco Content Services Switch Routing and Bridging Configuration Guide* for more information about this command.

Command Modes Global configuration mode

ip subnet-broadcast

To enable the forwarding of subnet broadcast addressed frames, use the **ip subnet-broadcast** command. Use the **no** form of this command to disable the forwarding of subnet broadcast addressed frames (the default behavior).

ip subnet-broadcast

no ip subnet-broadcast

Caution When the forwarding of the subnet broadcast is enabled, it can make the subnet susceptible to "smurf" attacks; an attacker sends an ICMP echo request frame using a subnet broadcast address as a destination and a forged address as the source. If the attack is successful, all the destination subnet hosts reply to the echo and flood the path back to the source. When the subnet broadcast forwarding is disabled, the original echo never reaches the hosts.

Command Modes Global configuration mode

(config) ip-fragment-enabled

To allow a CSS to flow-process UDP IP fragments, use the **udp-ip-fragment-enabled** command. By default, this feature is disabled. Use the **no** form of the command to reset the default behavior of the CSS to forwarding IP fragments.

Usage Guidelines The **ip-fragment-enabled** command has been deprecated (obsoleted). If you enter the **ip-fragment-enabled** command at the CLI or if your configuration already contains the **ip-fragment-enabled** command, the CSS automatically converts the command to the **udp-ip-fragment-enabled** command.

Related Commands **(config) udp-ip-fragment-enabled**
(config) tcp-ip-fragment-enabled

(config) ip-fragment max-assembled-size

To specify the maximum assembled size, use the **ip-fragment max-assembled-size** command. The maximum assembled size is the total length of an IP packet if all the IP fragments were assembled into the original packet. Assembled IP packets should be no larger than 64 KB.

As the CSS receives the IP fragments, it checks the fragments against the maximum assembled size value. If a fragment IP offset plus the IP payload (data) length is greater than the maximum assembled size, the CSS increments an error counter and discards the packet. Use the **no** form of this command to reset the maximum IP fragment assembled size to the default of 5120 bytes.

ip-fragment max-assembled-size *number*

no ip-fragment max-assembled-size

Syntax Description

number	Specifies the maximum size of an assembled packet in bytes. Enter an integer from 2048 to 65535. The default is 5120 bytes.

Related Commands

zero ip-fragment-stats
show ip-fragment-stats
(config) ip-fragment-enabled
(config) ip-fragment min-fragment-size

(config) ip-fragment min-fragment-size

To specify the smallest IP fragment payload based on your applications, use the **ip-fragment min-fragment-size** command. This command also provides protection against fragment attacks, which can consist of a chain of valid-looking, but very small, fragments. Use the **no** form of this command to reset the minimum IP fragment payload size to the default of 1024 bytes.

ip-fragment min-fragment-size *number*

no ip-fragment min-fragment-size

Syntax Description

number	Specifies the size of the smallest IP fragment payload that the CSS supports in bytes. Enter an integer from 64 to 1024. The default is 1024 bytes.

Related Commands

zero ip-fragment-stats
show ip-fragment-stats
(config) ip-fragment-enabled
(config) ip-fragment max-assembled-size

(config) keepalive

To access keepalive configuration mode and configure the properties for a global keepalive that you can apply to any service, use the **keepalive** command. Use the **no** form of this command to delete an existing keepalive.

keepalive *name*

no keepalive *existing_keepalive_name*

Syntax Description

name	Name of a new keepalive you want to create or of an existing keepalive. Enter an unquoted text string with no spaces and a maximum length of 31 characters. To see a list of existing keepalive names, enter: `keepalive ?`

Usage Guidelines

When you access keepalive mode, the prompt changes to (config-keepalive [*name*]). For information about commands available in this mode, see the "Keepalive Configuration Mode Commands" section.

Related Commands

show keepalive
(config-service) keepalive type named