

# Group Configuration Mode Commands

Group configuration mode allows you to configure a group. 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.

To access group configuration mode, use the **group** command from any mode except ACL, boot, and header-field-group configuration modes. The prompt changes to (config-group [name]). You can also use this command from group mode to access another group. For information about commands available in this mode, see the following commands.

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 ?**

---

## (config-group) active

To activate the specified group, use the **active** command.

**active**

---

**Related Commands** (config-group) suspend

## (config-group) add destination service

To add a destination service to a source group, use the **add destination** command.

**add destination service** *service\_name*

---

### Syntax Description

<i>service_name</i>	Name of the service to add to the group. Enter an unquoted text string. To see a list of services, enter: <b>show service ?</b>
---------------------	--

---



---

### Usage Guidelines

You can configure a maximum of 64 services per source group.

You cannot use a service with the same name in other source groups or the source service list within the same source group. You can use services with duplicate addresses among destination services since the actual service is chosen through content rule selection.

If the group is active and the same service is hit through a content rule, ACL preferred service, or sorry service, the source group is used to NAT the source address.

The service must be active and added to a content rule to perform destination address NATing for the source group.

**Note**

Adding a destination service to a group will not allow that specific service flows to be NATed by the group when initiated flows are from the service. The destination service applies group membership based on rule and service match. To ensure service-initiated connections are NATed, you must additionally configure an ACL match criteria or additional service names with duplicate addresses, and then add those services to a source group. The source group used could be the current group with the destination service or any other group.

If your topology consists of a CSS 11800 using ECMP to the servers and server port NAT configured on the services, to ensure the correct processing of packets either:

- Enable Service Remapping with the **persistence reset remap** command.
- Create source groups for the services in the content rule with the **add destination service** command.

**Related Commands**

**show group**  
**show service**  
**(config-group) remove destination service**

## (config-group) add service

To add a source service to a source group, use the **add service** command.

**add service** *service\_name*

**Syntax Description**

*service\_name*

Name of the service to add to the group. Enter an unquoted text string. To see a list of services, enter:

**show service ?**

---

**Usage Guidelines**

You can configure a maximum of 64 services per source group.

You cannot use a service with:

- The same name in other source groups or the destination service list within the same source group
- The same address as a source service on another source group

If the service matches the client, the source group is used.

Before you can add a service, you must suspend the group.

The services configured under a source group must be active to perform NATing through the group.

---

**Related Commands**

**show group**  
**show service**  
**(config-group) remove service**

## **(config-group) flow-timeout-multiplier**

To specify the number of seconds for which an idle flow can exist before the CSS tears it down, use the **flow-timeout-multiplier** command. Use the **no** form of this command to restore the default timeout for the port type.

**flow-timeout-multiplier** *timeout-multiplier*

**no flow-timeout-multiplier**

---

**Syntax Description***timeout-multiplier*

Value that the CSS multiplies by 16 to calculate the flow timeout in seconds. Enter an integer from 0 to 65533. The default value depends on the port type (see the **show flow-timeout default** command). This default value applies only to flows that are created under the specified source group.

A value of zero (no timeout) instructs the CSS to never tear down the flow, resulting in a permanent flow and lost resources. This is equivalent to entering the global configuration **flow permanent port** command.

---

---

**Usage Guidelines**

We do not recommend that you set the **flow-timeout multiplier** command to 0 for UDP flows on Layer 3 and Layer 4 content rules. If the value is set to 0, the CSS does not clean up the resources for the UDP flows.

Use the **flow-timeout-multiplier** command to configure flow inactivity timeout values for TCP and UDP flows on a per-rule and per-source group basis. Note that this timeout value is *not* the frequency with which a CSS reclaims flow resources, but the time period that must elapse for an idle flow before the CSS cleans up the flow.

If you configure a source group with destination services for client source NATing, you need to configure the **flow-timeout multiplier** command only on the content rule. The CSS sets the same flow timeout value for flows in both directions. If you configure different timeout values on the content rule and on the source group, the CSS uses the timeout value configured on the content rule for both flows.

To set up and keep track of flows, a CSS uses data structures called flow control blocks (FCBs). For optimal performance, the CSS reuses FCBs that are no longer needed by flows. Flow resource reclamation involves removing FCBs from the TCP and UDP lists.

Normally, flow cleanup occurs at a rate that is directly related to the overall number of flows that are currently active on a CSS. The fewer the number of active flows there are on a CSS, the less frequently the CSS reclaims FCBs. A CSS also cleans up long-lived TCP flows that have received a FIN or a RST, or whose timeout values have been met.

The CSS uses the following precedence when reclaiming flow resources:

1. If a flow matches on a content rule, the CSS checks for a user-configured timeout value and uses that value if one exists.
2. If the flow matches on a source group, the CSS checks for a user-configured timeout and uses that value if one exists.
3. If you have configured a permanent port using the global configuration **flow permanent port** command, the CSS sets the flow timeout value to 0, which means that the flow should never time out.
4. If none of the previous conditions are met, then the CSS uses the default timeout value for the protocol type.

---

**Related Commands**    **show flow-timeout**  
**(config) flow permanent**

## (config-group) ip address

To specify the source IP address for the group, use the **ip address** command. This address is substituted for the source address in flows originating from one of the group's sources. This command's function is identical to the **(config-group) vip address** command.

**ip address** *ip\_address*

---

<b>Syntax Description</b>	<i>ip_address</i>	IP address for the group. Enter an IP address in dotted-decimal notation (for example, 192.168.11.1).
---------------------------	-------------------	---

---



---

**Related Commands**    **show group**

## (config-group) no

To negate a command or set it to its default, use the **no** command. For information on general **no** commands you can use in this mode, see the general **no** command. The following options are available in group mode.

Syntax Description		
	<b>no acl</b> <i>index</i>	Deletes an ACL
	<b>no flow-timeout-multiplier</b>	Restores the default flow timeout for the port type
	<b>no portmap base-port</b>	Resets the starting SFP port number to its default value
	<b>no portmap number-of-ports</b>	Resets the number of ports per SFP to its default value
	<b>no redundancy-l4-stateless</b>	Disables stateless redundancy failover
	<b>no redundant-index</b>	Disables redundancy on the source group

## (config-group) portmap

To enable or disable the NATing of source IP addresses and source ports for a configured source group, or define the source port translation of flows from the services configured in a source group, use the **portmap** command. Use the **no** form of this command to reset the starting SFP port number to its default value of 2016 or number of ports to its default value of 63488.

**portmap** [**base-port** *base\_number*|**disable**|**enable**|**number-of-ports** *number*|**vip-address-range** *number*]

**no portmap** [**base-port**|**number-of-ports**|**vip-address-range**]

Syntax Description		
	<b>base-port</b> <i>base_number</i>	Defines the base port (starting port number) for the CSS. Enter a base number from 2016 to 63456. The default is 2016.

---

<b>disable</b>	<p>Instructs the CSS to perform Network Address Translation (NAT) only on the source IP addresses and not on the source ports of UDP traffic hitting a particular source group. This option does not affect TCP flows.</p> <p>For applications with high-numbered assigned ports (for example, SIP and WAP), we recommend that you preserve those port numbers by configuring destination services in source groups. Destination services cause the CSS to NAT the client source ports, but not the destination ports.</p> <p><b>Note</b> If you disable flows for a UDP port using the flow-state table and configure the <b>portmap disable</b> command in a source group, traffic for that port that matches on the source group does not successfully traverse the CSS.</p> <p>The CSS maintains but ignores any <b>base-port</b> or <b>number-of ports</b> (see the options above) values configured in the source group. If you later reenables port mapping for that source group, any configured <b>base-port</b> or <b>number-of ports</b> values will take effect. The default behavior for a configured source group is to NAT both the source IP address and the source port for port numbers greater than 1023.</p>
<b>enable</b>	<p>Restores the default CSS behavior of NATing source IP addresses and source ports for a configured source group.</p>

---



---

<b>number-of-ports</b> <i>number</i>	<p>Defines the total number of ports in the portmap range for the entire CSS. The CSS allocates the total number of configured ports proportionally among all the session processors in the CSS chassis. The allocation is based on the session processor relative weight value. To display the relative weight value of a session processor, enter the <b>show chassis session-processors</b> command as described in the <i>Cisco Content Services Switch Administration Guide</i>.</p> <p>The more modules you add to the CSS chassis, the fewer session processing each module performs and the fewer ports the CSS assigns to each module. To display the number of ports that the CSS allocates to each module, enter the <b>show group portmap</b> command.</p> <p>Enter a number from 2048 to 63488. The default is 63488. This default value should be fine for most applications. 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.</p>
---	---

---

---

**vip-address-range** *number* Specifies a VIP address range for port mapping. Use this option to increase the number of available ports for port mapping.

For each additional VIP address that you configure, the CSS creates a new port mapper to manage the available ports for that VIP. When the CSS performs PAT, the source group roundrobin among all the configured port mappers and the selected port mapper chooses the next eligible port for a given VIP.

The *number* variable indicates a range of VIP addresses starting with the address specified by the group configuration mode **vip address** command. Enter an integer from 1 to 255. The default is 1.

With a VIP range of 255, the maximum number of eligible ports on an SCM in a fully populated CSS 11506 chassis is 63240. For other SPs or chassis configurations, the number of ports is greater.

If you observe no-portmap errors, increase the *number* variable to make additional source ports available for port mapping.

Note that configuring a VIP address range for port mapping is different from a Virtual Web Hosting (VWH) configuration where you configure a VIP address range on a source group, not the port mapper. In a VWH configuration, there is only one port mapper available. For complete details, refer to the *Cisco Content Services Switch Content Load-Balancing Guide*.

---

---

### Usage Guidelines

Before you can change the port mapping, you must suspend the group.

The services configured under the source group must be active to perform source address NATing through the group.

## (config-group) redundancy-l4-stateless

To enable the Stateless Redundancy Failover feature for a source group on a redundant CSS, use the **redundancy-l4-stateless** command. The CSS can set up a connection for a mid-stream TCP flow, allowing TCP traffic to continue when a failure occurs at the load-balancing CSS. By default, the CSS rejects TCP sessions that do not begin with a TCP/SYN frame. Use the **no** form of this command to reset the default behavior of the CSS.

**redundancy-l4-stateless**

**no redundancy-l4-stateless**

---

### Command Modes

Group configuration mode

---

### Usage Guidelines

The Stateless Redundancy Failover feature has specific environment and configuration requirements. The environment requirements are as follows:

- Layer 3 and Layer 4 content rules with a VIP address. This feature is not supported in Layer 5 configurations.
- Source IP address load balance method only.
- CSS-to-CSS identical server and content rule configuration including:
  - Content VIP address.
  - Content balance method.
  - Failover method.
  - Service IP address, number, and order. The CSS orders services alphabetically. Apply identical service names on the master and backup CSSs.
- Visibility of identical servers to keepalive traffic from CSS to CSS. This ensures that the redistribution of the balance method does not occur in a failover event.

Redundant routes in a high availability topology surrounding the CSS are supported. However, the topology must not balance packets in a TCP/IP socket connection across more than one Ethernet port on the CSS.

IP and VIP redundant configurations are supported. The configuration requirement for each server farm is synchronization across all CSSs of:

- Membership and IP addresses of the server farms.
- Content rule VIP address. Each CSS must share the content VIP address that is used as a balance point for the server farm.
- Source group VIP address. Define each CSS with a source group VIP address as the content VIP address to NAT source addresses for packets returning from the server. In case of a failover, the source group handles connection setups for TCP/IP retransmissions that arrive at the CSS from a server. All servers on the farm must be a member of the source group.

Do not configure source groups for outbound traffic from the servers because the backup CSS does not know which ports were mapped by the source group on the master CSS. This restriction also applies to active FTP because the server initiates the data connection.

For more detailed information on Stateless Redundancy Failover, refer to the *Cisco Content Services Switch Redundancy Configuration Guide*.

---

**Related Commands**

**show redundancy**  
**(config) ip redundancy**  
**(config) group**  
**(config) interface**  
**(config) service**  
**(config-owner) content**  
**(config-owner-content) redundancy-l4-stateless**

## (config-group) **redundant-index**

To configure the global content index for a redundant source group, use the **redundant-index** command. A CSS uses the global content index to keep track of redundant content rules and associated flow state information. Use the **no** form of this command to disable redundancy on the source group.

**redundant-index** *number*

**no redundant-index**

---

### Syntax Description

<i>number</i>	Redundant index for the source group. Enter a unique integer from 0 to 32767, where a value of 0 disables ASR for a source group. The default is 0, but it does not appear in the running-config even if you configure it explicitly.
---------------	---

---

---

### Usage Guidelines

If you enter the **no redundant-index** command on an active redundant source group on live redundancy peers, the command automatically suspends the source group. Flows already mapped by a CSS are not affected. However, if a failover occurs during the life of an active flow that matches on such a suspended source group, the backup CSS cannot map the flow because it cannot find the source group with the same global index as that on the original master.



#### Note

For implicit or explicit Layer 5 rules, where there is delayed binding, binding is not complete until the CSS processes the SYN/ACK from the server. This means that, if a failover occurs in the middle of a spanned content request, the master CSS will not receive the SYN/ACK from the server and the flow will not be replicated on the backup CSS. No data is lost and users can simply refresh their browsers to restart the connection.

---

For information on redundant indexes and configuring Adaptive Session Redundancy (ASR) on Cisco 11500 series CSS peers, including requirements and restrictions that apply to both CSS peers in an ASR configuration, refer to the *Cisco Content Services Switch Redundancy Configuration Guide*.

**Related Commands** (config-group) vip address  
 (config-owner-content) redundant-index  
 (config-service) redundant-index

## (config-group) remove destination service

To remove a previously configured destination service from a source group, use the **removedestination service** command.

**remove destination service** *service\_name*

Syntax Description	<i>service_name</i>
	Name of an existing service you want to remove from the group. Enter a case-sensitive unquoted text string. To see a list of services for this group, enter:  <b>show group</b>

**Related Commands** show group  
 show service  
 (config-group) add destination service

## (config-group) remove service

To remove a previously configure a source service from a source group, use the **remove service** command.

**remove service** *service\_name*

Syntax Description	<i>service_name</i>
	Name of an existing service you want to remove from the group. Enter a case-sensitive unquoted text string. To see a list of services for this group, enter:  <b>show group</b>

---

**Usage Guidelines** Before you can remove a service, you must suspend the group.

---

**Related Commands** `show group`  
`show service`  
`(config-group) add service`

## (config-group) suspend

To suspend the specified group, use the **suspend** command. The group and its attributes remain the same but it no longer has an effect on flow creation.

**suspend**

---

**Usage Guidelines** To reactivate the group, use the **(config-group) active** command.

---

**Related Commands** `show group`  
`(config-group) active`

## (config-group) vip address

To specify the source virtual IP address or a range of IP addresses for the group, use the **vip address** command. The address is substituted for the source address in flows originating from one of the group's sources. This command's function is identical to the **(config-group) ip address** command. Use the **no** form of this command to remove the VIP address for the group.

**vip address** *ip\_or\_host* {**range number**}

**no vip address**

<b>Syntax Description</b>	<i>ip_or_host</i>	IP address or name for the group. Enter the address in either dotted-decimal IP notation (for example, 192.168.11.1) or mnemonic host-name format (for example, myhost.mydomain.com).
	<b>range</b> <i>number</i>	(Optional) Defines the range of IP addresses for the group. Enter a <i>number</i> from 1 to 65535. The default is 1. The <i>ip_or_host</i> variable is the first address in the range.

**Usage Guidelines** Before you can change the address to 0 or use the **no vip address** command, you must suspend the group.

**Related Commands** `show group`

## (config-group) zero all

To clear the statistics displayed by the **show group** command for the group in the current mode, use the **zero all** command.

**zero all**

**Related Commands** `show group`