

Configuring Web Cache Services By Using WCCP

This chapter describes how to configure your Catalyst 3550 switch to redirect traffic to cache engines (web caches such as the Cisco Cache Engine 550) by using the Web Cache Communication Protocol (WCCP). WCCP is a Cisco-developed content-routing technology that you can use to integrate cache engines into your network infrastructure. The cache engines transparently store frequently accessed content and then fulfill successive requests for the same content, eliminating repetitive transmissions of identical content from web servers. Cache engines accelerate content delivery and ensure maximum scalability and availability of content. In a service-provider network, you can deploy the WCCP and cache engine solution at the points of presence (POPs). In an enterprise network, you can deploy the WCCP and cache engine solution at the regional site and the small branch office.

To use this feature, you must have the IP services image (formerly known as the enhanced multilayer image [EMI]) installed on your switch.

Note

For complete syntax and usage information for the commands used in this chapter, see the "WCCP Router Configuration Commands" section in the "System Management Commands" part of the Cisco IOS Configuration Fundamentals Command Reference, Release 12.2.

This chapter consists of these sections:

- Understanding WCCP, page 34-1
- Configuring WCCP, page 34-4
- Monitoring and Maintaining WCCP, page 34-8

Understanding WCCP

The WCCP and Cisco cache engines (or other caches running WCCP) localize web-traffic patterns in the network, enabling content requests to be fulfilled locally.

WCCP enables supported Cisco routers and switches to transparently redirect content requests. With transparent redirection, users do not have to configure their browsers to use a web proxy. Instead, they can use the target URL to request content, and their requests are automatically redirected to a cache engine. The word *transparent* means that the end user does not know that a requested file (such as a web page) came from the cache engine instead of from the originally specified server.

Γ

When a cache engine receives a request, it attempts to service it from its own local cache. If the requested information is not present, the cache engine sends a separate request to the end server to retrieve the requested information. After receiving the requested information, the cache engine forwards it to the requesting client and also caches it to fulfill future requests.

This software release supports only WCCP version 2 (WCCPv2). Only a subset of WCCPv2 features are supported. For more information, see the "Unsupported WCCPv2 Features" section on page 34-4.

With WCCPv2, multiple routers or switches can service the cache-engine cluster (a series of cache engines); however, in this release, only one Catalyst 3550 switch can service the cluster, as shown Figure 34-1. Content is not duplicated on the cache engines.

Figure 34-1 Cisco Cache Engine and WCCPv2 Network Configuration



WCCP Message Exchange

This sequence of events describes the WCCP message exchange:

- 1. The cache engines send their IP addresses to the WCCP-enabled switch by using WCCP, signaling their presence through a *Here I am* message. The switch and cache engines communicate to each other through a control channel based on UDP port 2048.
- 2. The WCCP-enabled switch uses the cache engine IP information to create a cluster view (a list of caches in the cluster). This view is sent through an *I see you* message to each cache engine in the cluster, essentially making all the cache engines aware of each other. A stable view is established after the membership of the cluster remains the same for a certain amount of time.
- **3.** When a stable view is established, the cache engine in the cluster with the lowest IP address is elected as the designated cache engine.

L

WCCP Negotiation

In the exchange of WCCP protocol messages, the designated cache engine and the WCCP-enabled switch negotiate these items:

- Forwarding method (the method by which the switch forwards packets to the cache engine). The switch rewrites the Layer 2 header by replacing the packet's destination MAC address with the target cache engine's MAC address. It then forwards the packet to the cache engine. This forwarding method requires the target cache engine to be directly connected to the switch at Layer 2.
- Assignment method (the method by which packets are distributed among the cache engines in the cluster). The switch uses some of the least-significant bits of the destination IP address to determine which cache engine receives the redirected packet. The number of bits used is based on the number of cache engines. If the number of cache engines is equal to a power of 2 (for example, 1, 2, 4 and so forth), the switch evenly distributes (load balances) the traffic among the cache engines.

The switch does not support the mask assignment method described in the WCCP V2.0 Internet Draft.

- Packet-return method (the method by which packets are returned from the cache engine to the switch for normal forwarding). These are the typical reasons why a cache engine rejects packets and initiates the packet-return feature:
 - The cache engine is overloaded and has no room to service the packets.
 - The cache engine receives an error message (such as a protocol or authentication error) from the web server and implements the dynamic client bypass feature. The bypass enables clients to bypass the cache engines and to connect directly to the web server.

The cache engine returns a packet to the WCCP-enabled switch to forward to the web server as if the cache engine is not present. The cache engine does not intercept the reconnection attempt. In this way, the cache engine effectively cancels the redirection of a packet to the cache engine and creates a bypass flow. The switch receives the returned packet through a generic-route encapsulation (GRE) tunnel. The switch CPU uses Cisco express forwarding (CEF) to send these packets to the target web server. When the server responds with the requested information, the switch uses the normal Layer 3 forwarding to return the information to the requesting client.

MD5 Security

WCCPv2 provides an optional security component in each protocol message to enable the switch to use MD5 authentication on messages between the switch and the cache engine. Messages that do not authenticate (when authentication of the switch is enabled) are discarded by the switch. You enable the security feature by using the **ip wccp web-cache password** global configuration command. The password string is combined with the MD5 value to create security for the connection between the switch and the cache engine. You must configure the same password on each cache engine.

Packet Redirection

After WCCP is configured on the switch, the switch forwards all HTTP TCP port 80 packets received from clients to the cache engines. However, these packets are not redirected:

- Packets originating from the cache engine and targeted to the web server.
- Packets originating from the cache engine and targeted to the client.
- Packets returned or rejected by the cache engine. These packets are sent to the web server.

Unsupported WCCPv2 Features

These WCCPv2 features are not supported in this software release:

• WCCP service numbers, which are configured by using the **ip wccp** [*service-number*] global and interface configuration commands. These commands are not supported.

This software release supports caching only for TCP port 80.

• Packet redirection on an outbound interface, which is configured by using the **ip wccp redirect out** interface configuration command. This command is not supported.

This software release supports packet redirection only on an inbound interface.

• The connection of multiple Catalyst 3550 switches to multiple cache engines.

This software release supports the connection of only one switch to multiple cache engines.

- WCCP multicasting. The **ip wccp web-cache group-address** and **ip wccp web-cache group listen** global configuration commands are not supported.
- WCCP access lists. The **ip wccp web-cache redirect-list** and **ip wccp web-cache group-list** global configuration commands are not supported.
- Statistics for WCCP-related counters. Statistics for counters are not provided; they appear as zeros in the **show ip wccp web-cache view** privileged EXEC command output.

Configuring WCCP

These sections describe how to configure WCCP on your switch:

- Default WCCP Configuration, page 34-4
- WCCP Configuration Guidelines, page 34-5
- Enabling the Web Cache Service, Setting the Password, and Redirecting Traffic Received From a Client, page 34-5 (required)

Default WCCP Configuration

Table 34-1 shows the default WCCP configuration.

Table 34-1 Default WCCP Configuration

Feature	Default Setting
WCCP enable state.	WCCP services are disabled.
Protocol version.	WCCPv2.
Redirecting traffic received on an interface.	Disabled.

L

WCCP Configuration Guidelines

Before configuring WCCP on your switch, make sure to follow these configuration guidelines:

- Do not configure the cache engine for GRE because the switch does not support traffic forwarding by using GRE. For more information, see the documentation that shipped with the cache engines.
- Make a direct Layer 2 connection from the cache engines to the switch so that the switch can perform Layer 2 rewrites for WCCP redirection. The Cisco Cache Engines require the use of a Fast Ethernet interface for a direct connection. You also can connect the switch to the cache engine by using a 10/100/1000 port if the connection is a direct Layer 2 connection.
- Connect up to 32 cache engines to a single Catalyst 3550 switch.
- Connect only one Catalyst 3550 switch to multiple cache engines. Do not connect multiple Catalyst 3550 switches to multiple cache engines.
- Configure the switch interfaces that are connected to the web clients, the cache engines, and the web server as Layer 3 interfaces (routed ports and switch virtual interfaces [SVIs]). For HTTP packet redirection to work, the servers, cache engines, and clients must be on different subnets.
- Do not configure the clients, cache engines, or web servers on the same switch interface.
- Do not configure the switch with both WCCP and multiple VPN routing/forwarding (multi-VRF) instances in customer edge (CE) devices.
- Modify the Switch Database Management (SDM) template to enable the switch to support 144-bit Layer 3 TCAM by using the sdm prefer extended-match, sdm prefer access extended-match, or sdm prefer routing extended-match global configuration command. For more information on the SDM templates, see the "Optimizing System Resources for User-Selected Features" section on page 6-26.
- Do not configure WCCP and policy-based routing (PBR) on the same switch interface.

Enabling the Web Cache Service, Setting the Password, and Redirecting Traffic Received From a Client

MD5 password security requires that the switch and cache engines be configured with the same password. Each cache engine or switch authenticates the security component in a received WCCP packet immediately after validating the WCCP message header. Packets failing authentication are discarded.

For WCCP packet redirection to operate, you must configure the switch interface connected to the client to redirect inbound HTTP packets.

This procedure shows how to configure these features on routed ports. To configure these features on SVIs, see the configuration examples that follow the procedure.

Beginning in privileged EXEC mode, follow these steps to enable the web cache service, to set a password, to configure routed interfaces, and to redirect inbound packets received from a client to the cache engine. This procedure is required.



Before configuring WCCP commands, configure the SDM template, and reboot the switch. For more information, see the "Optimizing System Resources for User-Selected Features" section on page 6-26.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip wccp web-cache [password encryption-number password]	Enable the web cache service on your switch. By default, this feature is disabled.
		(Optional) For [password <i>encryption-number password</i>], specify an encryption number. The range is 0 to 7. Use 0 for not encrypted, and use 7 for proprietary. Specify a password name up to seven characters in length. The switch combines the password with the MD5 authentication value to create security for the connection between the switch and the cache engine. By default, no password is configured, and no authentication is performed.
		You must configure the same password on each cache engine.
		When authentication is enabled, the switch discards messages that are not authenticated.
Step 3	interface interface-id	Specify the interface connected to the cache engine or the web server, and enter interface configuration mode.
Step 4	no switchport	Enter Layer 3 mode.
Step 5	ip address ip-address subnet-mask	Configure the IP address and subnet mask.
Step 6	no shutdown	Enable the interface.
Step 7	exit	Return to global configuration mode. Repeat Steps 3 through 7 for each cache engine and web server.
Step 8	interface interface-id	Specify the interface connected to the client, and enter interface configuration mode.
Step 9	no switchport	Enter Layer 3 mode.
Step 10	ip address ip-address subnet-mask	Configure the IP address and subnet mask.
Step 11	no shutdown	Enable the interface.
Step 12	ip wccp web-cache redirect in	Redirect packets received from the client to the cache engine.
Step 13	exit	Return to global configuration mode. Repeat Steps 8 through 13 for each client.
Step 14	end	Return to privileged EXEC mode.
Step 15	show ip wccp web-cache	Verify your entries.
	and	
	show running-config	
Step 16	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the web cache service, use the **no ip wccp web-cache** global configuration command. To disable inbound packet redirection, use the **no ip wccp web-cache redirect in** interface configuration command.

This example shows how to configure routed interfaces and to enable the web cache service. Fast Ethernet port 1 is connected to the cache engine, is configured as a routed port with an IP address of 172.20.10.30, and is re-enabled. Gigabit Ethernet port 1 is connected through the Internet to the web server, is configured as a routed port with an IP address of 175.20.20.10, and is re-enabled. Fast Ethernet

ports 2 to 5 are connected to the clients and are configured as routed ports with IP addresses 175.20.30.20, 175.20.40.30, 175.20.50.40, and 175.20.60.50. The switch redirects HTTP packets received from the client interfaces to the cache engine.

```
Switch# configure terminal
Switch(config) # ip wccp web-cache
Switch(config) # interface fastethernet0/1
Switch(config-if) # no switchport
Switch(config-if) # ip address 172.20.10.30 255.255.255.0
Switch(config-if) # no shutdown
Switch(config-if) # exit
Switch(config) # interface gigabitethernet0/1
Switch(config-if) # no switchport
Switch(config-if) # ip address 175.20.20.10 255.255.255.0
Switch(config-if) # no shutdown
Switch(config-if) # exit
Switch(config) # interface fastethernet0/2
Switch(config-if) # no switchport
Switch(config-if)# ip address 175.20.30.20 255.255.255.0
Switch(config-if) # no shutdown
Switch(config-if) # ip wccp web-cache redirect in
Switch(config-if) # exit
Switch(config) # interface fastethernet0/3
Switch(config-if) # no switchport
Switch(config-if) # ip address 175.20.40.30 255.255.255.0
Switch(config-if) # no shutdown
Switch(config-if) # ip wccp web-cache redirect in
Switch(config-if)# exit
Switch(config) # interface fastethernet0/4
Switch(config-if) # no switchport
Switch(config-if)# ip address 175.20.50.40 255.255.255.0
Switch(config-if)# no shutdown
Switch(config-if) # ip wccp web-cache redirect in
Switch(config-if) # exit
Switch(config) # interface fastethernet0/5
Switch(config-if) # no switchport
Switch(config-if)# ip address 175.20.60.50 255.255.255.0
Switch(config-if) # no shutdown
Switch(config-if) # ip wccp web-cache redirect in
Switch(config-if)# exit
```

This example shows how to configure SVIs and how to enable the web cache service. VLAN 299 is created and configured with an IP address of 175.20.20.10. Gigabit Ethernet port 1 is connected through the Internet to the web server and is configured as an access port in VLAN 299. VLAN 300 is created and configured with an IP address of 172.20.10.30. Fast Ethernet port 1 is connected to the cache engine and is configured as an access port in VLAN 301. The switch redirects HTTP packets received from the client interfaces to the cache engine.

```
Switch# configure terminal
Switch(config)# ip wccp web-cache
Switch(config)# vlan 299
Switch(config-vlan)# exit
Switch(config)# interface vlan 299
Switch(config-if)# ip address 175.20.20.10 255.255.255.0
Switch(config-if)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan 299
Switch(config)# vlan 300
Switch(config-vlan)# exit
```

```
Switch(config)# interface vlan 300
Switch(config-if)# ip address 172.20.10.30 255.255.255.0
Switch(config-if)# exit
Switch(config)# interface fastethernet0/1
Switch(config-if) # switchport mode access
Switch(config-if) # switchport access vlan 300
Switch(config-if)# exit
Switch(config) # vlan 301
Switch(config-vlan)# exit
Switch(config)# interface vlan 301
Switch(config-if) # ip address 175.20.30.20 255.255.255.0
Switch(config-if)# ip wccp web-cache redirect in
Switch(config-if)# exit
Switch(config)# interface range fastethernet0/2 - 5
Switch(config-if-range)# switchport mode access
Switch(config-if-range)# switchport access vlan 301
Switch(config-if-range)# exit
```

Monitoring and Maintaining WCCP

To monitor and maintain WCCP, use one or more of the privileged EXEC commands in Table 34-2:

Command	Purpose
clear ip wccp web-cache	Removes statistics for the web-cache service.
show ip wccp web-cache	Displays global information related to WCCP.
show ip wccp web-cache detail	Displays information for the switch and all cache engines in the WCCP cluster.
show ip interface	Displays status about any IP WCCP redirection commands that are configured on an interface; for example, <i>Web Cache Redirect is enabled / disabled.</i>
show ip wccp web-cache view	Displays which other members have or have not been detected.

Table 34-2 Commands for Monitoring and Maintaining WCCP