Configuring Web Cache Services By Using WCCP

This chapter describes how to configure your Catalyst 3560 switch to redirect traffic to wide-area application engines (such as the Cisco Cache Engine 550) by using the Web Cache Communication Protocol (WCCP). This software release supports only WCCP version 2 (WCCPv2).

WCCP is a Cisco-developed content-routing technology that you can use to integrate wide-area application engines—referred to as *application engines*—into your network infrastructure. The application engines transparently store frequently accessed content and then fulfill successive requests for the same content, eliminating repetitive transmissions of identical content from web servers. Application engines accelerate content delivery and ensure maximum scalability and availability of content. In a service-provider network, you can deploy the WCCP and application engine solution at the points of presence (POPs). In an enterprise network, you can deploy the WCCP and application engine solution at the regional site and the small branch office.

To use this feature, the switch must be running the IP services image (formerly known as the enhanced multilayer image [EMI]).

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 Cisco IOS Configuration Fundamentals Command Reference, Release 12.2.* Documentation > Cisco IOS Software 12.2 Mainline Command References

This chapter consists of these sections:

Understanding WCCP

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

can use the target URL to request content, and their requests are automatically redirected to an application engine. The word   means that the end user does not know that a requested file (such as a web page) came from the application engine instead of from the originally specified server.

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

With WCCP, the application-engine cluster (a series of application engines) can service multiple routers or switches, as shown Figure 43-1.

Figure 43-1   *Cisco Cache Engine and WCCP Network Configuration*

WCCP Message Exchange

1.
                              Here I am
        communicate to each other through a control channel based on UDP port 2048.

2. The WCCP-enabled switch uses the application engine IP information to create a cluster view (a list of application engines in the cluster). This view is sent through an   message to each application engine in the cluster, essentially making all the application 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.

WCCP Negotiation

-

-

- switch for normal forwarding). These are the typical reasons why an application engine rejects packets and starts the packet-return feature:

    -

    -

(GRE), the switch receives the returned packet through a GRE tunnel that is configured in the application engine. The switch CPU uses Cisco express forwarding to send these packets to the target web server. If the return method is Layer 2 rewrite, the packets are forwarded in hardware to the target web server. When the server responds with the requested information, the switch uses normal Layer 3 forwarding to return the information to the requesting client.

MD5 Security

and the application engine. You must configure the same password on each application engine.

Packet Redirection and Service Groups

service group

Service groups are configured to map to a protocol and Layer 4 port numbers and are established and maintained independently. WCCP allows dynamic service groups, where the classification criteria are provided dynamically by a participating application engine.

You can configure up to 8 service groups on a switch or switch stack and up to 32 cache engines per service group. WCCP maintains the priority of the service group in the group definition. WCCP uses the priority to configure the service groups in the switch hardware. For example, if service group 1 has a priority of 100 and looks for destination port 80, and service group 2 has a priority of 50 and looks for source port 80, the incoming packet with source and destination port 80 is forwarded by using service group 1 because it has the higher priority.

WCCP supports a cluster of application engines for every service group. Redirected traffic can be sent to any one of the application engines. The switch supports the mask assignment method of load balancing the traffic among the application engines in the cluster for a service group.

After WCCP is configured on the switch, the switch forwards all service group packets received from clients to the application engines. However, these packets are not redirected:

Packets originating from the application engine and targeted to the web server.

Packets originating from the application engine and targeted to the client.

Packets returned or rejected by the application engine. These packets are sent to the web server.

You can configure a single multicast address per service group for sending and receiving protocol messages. When there is a single multicast address, the application engine sends a notification to one address, which provides coverage for all routers in the service group, for example, 225.0.0.0. If you add and remove routers dynamically, using a single multicast address provides easier configuration because you do not need to specifically enter the addresses of all devices in the WCCP network.

You can use a router group list to validate the protocol packets received from the application engine. Packets matching the address in the group list are processed, packets not matching the group list address are dropped.

To disable caching for specific clients, servers, or client/server pairs, you can use a WCCP redirect access control list (ACL). Packets that do not match the redirect ACL bypass the cache and are forwarded normally.

Before WCCP packets are redirected, the switch examines ACLs associated with all inbound features configured on the interface and permits or denies packet forwarding based on how the packet matches the entries in the ACL.

Only **permit** ACL entries are supported in WCCP redirect lists.

When packets are redirected, the output ACLs associated with the redirected interface are applied to the packets. Any ACLs associated with the original port are not applied unless you specifically configure the required output ACLs on the redirected interfaces.

Unsupported WCCP Features

-
-
-
- There is no SNMP support for WCCP.

Configuring WCCP

-
-
-

Default WCCP Configuration

Table 43-1   Default WCCP Configuration

Feature	Default Setting

WCCP Configuration Guidelines

-

- the web server as Layer 3 interfaces (routed ports and switch virtual interfaces [SVIs]). For WCCP packet redirection to work, the servers, application engines, and clients must be on different subnets.

- Use only nonreserved multicast addresses when configuring a single multicast address for each application engine.

- WCCP entries and PBR entries use the same TCAM region. WCCP is supported only on the templates that support PBR: access, routing, and dual IPv4/v6 routing.

- When TCAM entries are not available to add WCCP entries, packets are not redirected and are forwarded by using the standard routing tables.

- The number of available policy-based routing (PBR) labels are reduced as more interfaces are enabled for WCCP ingress redirection. For every interface that supports service groups, one label is consumed. The WCCP labels are taken from the PBR labels. You need to monitor and manage the labels that are available between PBR and WCCP. When labels are not available, the switch cannot add service groups. However, if another interface has the same sequence of service groups, a new label is not needed, and the group can be added to the interface.

- The routing maximum transmission unit (MTU) size configured on the stack member switches should be larger than the client MTU size. The MAC-layer MTU size configured on ports connected to application engines should take into account the GRE tunnel header bytes.

- You cannot configure WCCP and VPN routing/forwarding (VRF) on the same switch interface.

- You cannot configure WCCP and PBR on the same switch interface.

•

Enabling the Web Cache Service

Beginning in privileged EXEC mode, follow these steps to enable the web cache service, to set a multicast group address or group list, to configure routed interfaces, to redirect inbound packets received from a client to the application engine, enable an interface to listen for a multicast address, and to set a password. This procedure is required.

Before configuring WCCP commands, configure the SDM template, and reboot the switch. For more information, see Chapter 8, "Configuring SDM Templates."

	Command	Purpose
Step 1	**configure terminal**	
Step 2	**ip wccp** {**web-cache**	*service-number*} [*groupaddress*] [*access-list*] [*access-list*] [*encryption-number password*	*groupaddress* *access-list* *access-list* *encryption-number password* 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 application engine. By default, no password is configured, and no authentication is performed. You must configure the same password on each application engine. When authentication is enabled, the switch discards messages that are not authenticated.
		Specify the interface connected to the application engine or the web server, and enter interface configuration mode.
Step 4		
Step 5	*ip-address subnet-mask*	

	Command	Purpose
Step 6		
Step 7	exit	
Step 8		
Step 9		
Step 10		
Step 11		
Step 12		
Step 13		
Step 14		
Step 15		
Step 16		
Step 17	copy running-config startup-config	

```
Switch# configure terminal
Switch(config)# ip wccp web-cache 80 group-address 224.1.1.100 redirect list 12
            access-list 12 permit host 10.1.1.1
            interface gigabitethernet0/1
Switch(config-if)# no switchport
Switch(config-if)# ip address 172.20.10.30 255.255.255.0
              no shutdown
              ip wccp web-cache group-listen
              exit
            interface gigabitethernet0/2
              no switchport
              ip address 175.20.20.10 255.255.255.0
              no shutdown
              exit
            interface gigabitethernet0/3
              no switchport
```

```
interface gigabitethernet0/6
   no switchport
   ip address 175.20.60.50 255.255.255.0
   no shutdown
   ip wccp web-cache redirect in
   exit
```

ACL entries are being used in the redirect-list;   entries are unsupported.

```
access-list 15 permit host 171.69.198.102
access-list 15 permit host 171.69.198.104
access-list 15 permit host 171.69.198.106
vlan 299
```
Switch(config-vlan)#
Switch(config)#
Switch(config-if)#
Switch(config-if)#
Switch(config)#
Switch(config-if)#
Switch(config-if)#
Switch(config)#
Switch(config-vlan)#
Switch(config)#
Switch(config-if)#
Switch(config-if)#
Switch(config)#
Switch(config-if)#
Switch(config-if)#
Switch(config-if)#
Switch(config)#
Switch(config-vlan)#
Switch(config)#
Switch(config-if)#

```
Switch(config-if-range)#
Switch(config-if-range)#
Switch(config-if-range)#
```

Monitoring and Maintaining WCCP

Table 43-2   Commands for Monitoring and Maintaining WCCP

	Web Cache Redirect is enabled / disabled.
view	