WCCP on Catalyst 6500 Platforms with High CPU Utilization Configuration Guide

Document ID: 116134

Contributed by Rahul Parameswaran and Dixon Ho, Cisco TAC
Engineers.

Jul 16, 2013

Contents

Introduction

This document describes how to use the Web Cache Communication Protocol (WCCP) on the Cisco Catalyst 6500 Series platform.

WCCP was originally designed as a method to intercept web traffic (HTTP) and forward it to a local cache device, where it could be served to a client from a local location and conserve expensive WAN bandwidth.

From a network user perspective, WCCP is transparent because it is used at the network level, without any special configuration by the user, in order to identify and redirect the web traffic that traverses a Layer 3 (L3) device to a local cache device. Although WCCP was originally designed for web traffic, the transparent method of redirection has become a very useful mechanism to address other problems with high−volume content over low−volume links. For this reason, additional protocol support was added to later WCCP versions. These additional technologies include protocols such as HTTP, HTTPS, FTP, video streaming, and file caching technologies, such as the Common Internet File System (CIFS). These technologies support newer Cisco solutions and platforms, such as Wide Area File Services (WAFS), Wide Area Application Services (WAAS), Application Oriented Networking (AONs), and enhanced capabilities of the Application and Content Networking Software (ACNS).

WCCP adoption is increasing as enterprises implement the latest productivity tools such as video−based communications and training, as well as live and on−demand video. Efforts to control costs, such as consolidated data centers, create the need for WCCP to support additional, extremely high bandwidth services.

Because of the importance of WCCP with today's content rich networks, some platforms, such as the Catalyst 6500, have implemented hardware−assisted performance with WCCP so that the CPU load required for the protocol is reduced. This document describes how to deploy WCCP on the Catalyst 6500 in order to maximize hardware utilization and reduce CPU load.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- WCCP
- Cisco Catalyst 6500 Series switches
- Cisco IOS® Software

Components Used

The information in this document is based on these software and hardware versions:

- Cisco Catalyst 6500 Series Switches
- Cisco IOS Software Release 12.1(50)SY or later with Supervisor Engine 2T (Policy Feature Card 4)
- Cisco IOS Software Release 12.2(18)SXD1 or later with Supervisor Engine 720 (Policy Feature Card 3)

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

WCCP Functional Overview

The functionality that is generically referred to as WCCP actually involves three components:

1. Functionality implemented only at the router or switch

2. Functionality implemented only at the traffic−processing entity, such as a web cache
3. WCCP implemented on both sides

This document examines these three WCCP operational characteristics:

1. The ***assignment*** method determines which WCCP traffic and which WCCP device are chosen for the destination of the traffic.The WCCP protocol supports multiple devices clustered together.
2. The ***redirection*** method forwards the traffic flow to the WCCP appliance when there is a need to change from the normal network path that the traffic was traversing.
3. The ***return*** method determines how the traffic is sent back to the router from the WCCP appliance if the traffic could not be serviced. In cases where the WCCP appliance does service the request, the packets are returned directly to the requestor.

WCCP and Catalyst 6500

The Catalyst 6500 Supervisor Engine 2, Supervisor Engine 32, and Supervisor Engine 720 support these WCCP features and methods:

- WCCP Version 2 (WCCPv2)
- Hash−based assignment method
- Mask−based assignment method (hardware−accelerated)
- L3 generic routing encapsulation (GRE) redirection method (hardware−accelerated on Supervisor Engine 32 and Supervisor Engine 720)
- Layer 2 (L2) redirection method (hardware−accelerated on Supervisor Engine 2, Supervisor Engine 32, and Supervisor Engine 720)
- GRE return method
- L2 return method with Cisco IOS Software Release 12.2SXH (hardware−accelerated on Supervisor Engine 2, Supervisor Engine 32, and Supervisor Engine 720)

For more details on these features, refer to Configuring WCCP in the Cisco 6500 Series Cisco IOS Software Configuration Guide, 12.2SX.

WCCP Assignment Method

WCCP assignment determines which traffic the WCCP protocol redirects and which WCCP entity receives redirected traffic.

When WCCP is configured on an interface of a router and on a WCCP entity, the redirecting device (the Catalyst 6500) needs to know what traffic is to be redirected and where the traffic is to be sent. The WCCP entities within a service group cluster all communicate through the WCCP protocol with the Catalyst 6500; however, one WCCP appliance is selected to represent the cluster in order to control how the cluster operates (by the assignment method, redirection method, and so forth). The WCCP appliance and router may negotiate the method by which packets are distributed between the web caches in a service group. A service group is defined as a many−to−many relationship between up to 32 routers and 32 WCCP entities. The negotiation is by service group; thus, web caches that participate in several service groups may negotiate a different assignment method for each service group. The currently available WCCP services are:

WCCP Service	***Protocol***
web cache	HTTP
53	DNS Cache
60	FTP
61	WAAS − forward

62	WAAS − reverse
70	HTTPS
80	Real−Time Streaming Protocol (RTSP)
81	Microsoft Media Server (MMS) over UDP (MMSU)
82	MMS over TCP (MMST)
83	RTSP using UDP (RTSPU)
89	CIFS−Cache WAAS
98	Custom Web Cache
99	Reverse proxy
90−97	User−configurable *

* User−configurable services are implemented in the Catalyst 6500 with an interface level command applied to an inbound or outbound direction. The implications of the choice of inbound or outbound are discussed later, but inbound is the preferred method because a redirect list can be coupled with WCCP for maximum hardware performance.

Once configured for WCCP, a router advertises the supported assignment methods for a service group in the WCCP communications messages. The absence of such a message implies that the Catalyst 6500 supports the default hash−based assignment method only.

Once the negotiation between the Catalyst 6500 and the WCCP appliance is complete, the WCCP designated entity, through WCCP, informs the Catalyst 6500 what traffic is to be redirected and to which WCCP entity(s) the traffic is assigned. As an example, the WCCP entity may inform the Catalyst 6500 to redirect all web traffic from a particular subnet to cache engines 1 − 4 in the service group when there are more than four WCCP devices available.

There are two assignment methods available for WCCP:

1. ***Hash−based assignment*** (default) utilizes a software−based hash algorithm along with directives from the WCCP−designated appliance in order to determine which WCCP appliance receives traffic. In a Supervisor Engine 32 or Supervisor Engine 720 platform, NetFlow hardware resources are used in order to apply some level of hardware assist.
2. ***Mask−based assignment*** utilizes the hardware capabilities of the Catalyst 6500, specifically the access control list (ACL) ternary content addressable memory (TCAM), in order to assign traffic to WCCP entities. This is the preferred method.

All devices within a WCCP service group must use the same assignment method. Assignment methods are configured on the WCCP entity and learned by the Catalyst 6500. Refer to WCCP Design Recommendations for more details.

Hash−Based Assignment Method Detail

The hash−based assignment mechanism relies on an algorithm performed in software. In order to leverage the hash algorithm, the first packet in a particular flow is sent from the hardware path to the software path where the hash is performed.

The software performs an XOR hash of various components of the flow and comes up with a hash that splits the traffic flows to the various WCCP entities. The hash mechanism determines how the traffic is distributed among the available WCCP entities.

The hash result is programmed into the hardware NetFlow table where subsequent packets in that flow are forwarded. Regardless of the fields available for hashing by WCCP, the full five−tuple is used. This means

NetFlow is put into interface, full−flow mode when WCCP is enabled. This has implications on other features that may require NetFlow resources. See the WCCP Defects section for more details.

A common question about WCCP on the Catalyst 6500 is, "Why does the CPU utilization increase when I enable WCCP?" When hash−based assignments are in use, the software−based processing of the initial packet in each flow places a burden on the CPU and is most often the cause of increased utilization. With the currently available Policy Feature Card 3 (PFC3) forwarding hardware, if WCCP is configured as an egress feature or if hash−based assignment is in use (ingress or egress), some level of software processing is always required.

The use of the hash−based assignment method impacts these features:

- ***NetFlow table*** − The number of entries supported by the PFC is limited, and the flow mask changes to interface full−flow for the entire NetFlow table.
- ***CPU utilization*** − There is an increase in CPU utilization as the first packet in each flow is software switched.
- ***Performance*** − The rate at which traffic is sent to the CPU for lookup is limited so that the CPU is protected.
- ***NetFlow features*** − Other features that use NetFlow resources might be impacted if the NetFlow resources are consumed by WCCP.

The limitations and implications caused by the hash−based assignment requirement for software processing are applicable to both ingress and egress traffic. Impact on the CPU can be exacerbated if the network is undergoing atypical traffic patterns, such as a Denial of Service (DoS) attack. In a typical attack or worm outbreak, every packet sent by a host is to a new destination or port, which causes every packet to be processed in software. Since WCCP redirected traffic is explicitly being sent to the CPU for first−packet processing, there are limited methods of protection. The use of 'deny' ACL entries on the interface can limit what is sent to the CPU; however, there are no rate−limiters or other protections against these types of attacks.

Mask−Based Assignment Method Detail

Mask−based assignment is handled differently dependent upon whether it is configured on ingress or on egress.

With ingress mask−based assignment, the mask is programmed into the ACL TCAM before packet forwarding, so the NetFlow table and software processing are not needed. The WCCP entity chooses a number of hash−buckets and assigns an address mask and WCCP appliance to each bucket. Once the assignments are complete, the supervisor programs one TCAM entry and one hardware adjacency for each bucket and redirects packets that match the address mask to the associated WCCP appliance by means of an L2 rewrite.

If WCCP is configured as an ingress feature, it may use an ACL redirect−adjacency entry in the hardware ACL table. Once WCCP matches the entry, it uses an appropriate adjacency in order to perform either an L2 rewrite or GRE encapsulation. Thus, when mask assignment is used on ingress, both L2 rewrite (Supervisor Engine 2, Supervisor Engine 32, and Supervisor Engine 720) and GRE encapsulation (Supervisor Engine 32 and Supervisor Engine 720 only) are performed in hardware.

If WCCP is configured as an egress feature, ACL redirect−adjacencies are not supported in hardware because the packets in the flow have already been routed by the system. The first packet of a flow is sent to software for processing. Once the proper redirect−adjacency is determined, it is programmed into the NetFlow hardware (instead of ACL TCAM), where the entry points to an adjacency that performs either an L2 rewrite or GRE encapsulation. Subsequent packets in the flow are redirected in hardware by the NetFlow hardware.

Note: If WCCP is configured as an egress feature, mask assignment requires software processing, which negates any benefit of the mask−based assignment method.

Of the two mask−based options, only the ingress mask−based assignment enables full hardware−based forwarding for initial and subsequent packets. Any other option, such as the use of hash−based assignment or egress processing, causes software switching of the initial packet and hardware−NetFlow switched forwarding of subsequent packets.

WCCP Redirection Method

The WCCP entity, not the Catalyst 6500, dictates the hash tables and mask/value sets to the Catalyst 6500, so configuration of the redirect method is done on that appliance, and not on the Catalyst 6500 switch. The Catalyst 6500 determines the best redirect method available, based on the WCCP communications with the WCCP entity/group. This negotiation determines how redirected traffic is forwarded to the appliance. There are two redirection options: L3 (GRE) and L2 (MAC address rewrites).

With WCCPv1, the only option is L3 redirection, also known as GRE encapsulation. With L3 redirection, each WCCP redirected packet is encapsulated in a GRE header marked with a protocol type 0x883E followed by a four−octet WCCP redirect header, which is subsequently sent to the WCCP appliance (such as a cache engine).

With the introduction of WCCPv2, L2 redirection, also known as accelerated WCCP redirection, was added in order to take advantage of hardware switching platforms such as the Catalyst 6500. When WCCP uses L2 redirection, the WCCP appliance and Catalyst 6500 must be L2 adjacent (within the same L2 VLAN). Redirected L2 traffic does not use GRE encapsulation; instead, the MAC destination address is rewritten by the Catalyst 6500 to that of the L2−connected WCCP entity and forwarded through normal hardware switching.

Note: The method of forwarding to the WCCP device may not be the same method that the WCCP device uses in order to send traffic back to the Catalyst 6500. WCCP is used in order to negotiate a forward and return method that both devices support. See WCCP Return Method.

L3 (GRE) Forwarding Method

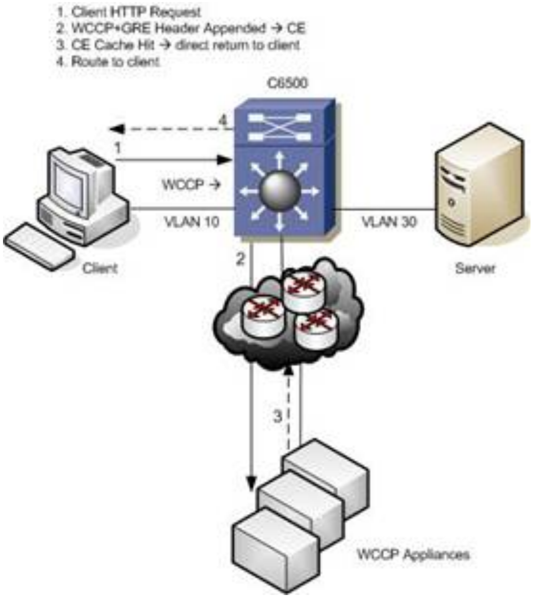

WCCP L3 operation involves the use of GRE as an encapsulation method. Redirected packets are encapsulated in a GRE header with a protocol type of 0x883e, along with a 4−byte WCCP redirection header that includes a service ID and hash bucket matched (WCCPv2 only). The use of GRE enables the WCCP client to be separated from the Catalyst 6500 by multiple L3 (routed) hops.

In this scenario, the options available for WCCP redirection include:

1. Ingress − L3 (GRE) redirection + hash assignment; this requires software processing.
2. Ingress − L3 (GRE) redirection + mask assignment; this requires full hardware processing and is available only on the Supervisor Engine 32 or Supervisor Engine 720.
3. Egress − L3 (GRE) redirection + hash assignment; this requires software processing.
4. Egress − L3 (GRE) redirection + mask assignment; this requires software processing.

Ingress − L3 (GRE) Redirection + Hash Assignment

On the Supervisor Engine 2, each GRE packet is sent to the multilayer switch feature card (MSFC) for processing. Since GRE encapsulation is not supported in hardware, the MSFC must apply both GRE and WCCP headers, which forces software switching for all traffic.

With the hash−based assignment method, the Supervisor Engine 32 and Supervisor Engine 720 forward the first packet of every flow in software so that a NetFlow table entry is established. The packet is then encapsulated in GRE (the encapsulation and forwarding of the initial packet is done in software) and forwarded to the WCCP appliance.

The establishment of the NetFlow entry impacts CPU utilization, but subsequent packet forwarding is done in hardware for Supervisor Engine 720 and Supervisor Engine 32. Traffic patterns, especially the number of unique flows, dictate how much the CPU is utilized. If the NetFlow resources of the Catalyst 6500 are consumed, then all traffic is forwarded in software.

The NetFlow resources of the supervisor PFC differ across various platforms. Currently, the largest NetFlow resources are available on the PFC−3BXL on the Supervisor Engine 720 platform.

Ingress − L3 (GRE) Redirection + Mask Assignment

On the Supervisor Engine 2, each GRE packet is sent to the MSFC for processing. Since GRE encapsulation is not supported in hardware, the MSFC must apply both GRE and WCCP headers, which forces software switching for all traffic.

With the mask−based assignment method, the Supervisor Engine 32 and Supervisor Engine 720 forward the initial and subsequent packets in hardware, because GRE is supported natively, and the mask assignment uses the ACL TCAM hardware for forwarding.

Egress − L3 (GRE) Redirection + Hash Assignment

On the Supervisor Engine 2, each packet is sent to the MSFC for processing. Since GRE encapsulation is not supported in hardware, the MSFC must apply both GRE and WCCP headers, which forces software switching for all traffic.

With the hash−based assignment method with the Supervisor Engine 32 and Supervisor Engine 720, the Catalyst 6500 forwards the initial packet of every flow in software so that the NetFlow table entry is established. The packet is then encapsulated in GRE and forwarded to the WCCP entity.

The establishment of the NetFlow entry impacts CPU utilization, but subsequent packet forwarding is done in hardware. Traffic patterns, especially the number of unique flows, dictate how much the CPU is utilized. If the NetFlow resources of the Catalyst 6500 are consumed, then all traffic is forwarded in software.

The NetFlow resources of the supervisor PFC differ across the various platforms. Currently, the largest NetFlow resources are available on the PFC−3BXL on the Supervisor Engine 720 platform.

Egress − L3 (GRE) Redirection + Mask Assignment

On the Supervisor Engine 2, each packet is sent to the MSFC for processing. Since GRE encapsulation is not supported in hardware, the MSFC must apply both GRE and WCCP headers, which forces software switching for all traffic.

With the mask−based assignment method with the Supervisor Engine 32 and Supervisor Engine 720, the first packet of every flow is software switched so that the NetFlow table entry is established. None of the supervisors support egress ACL adjacency programming, which forces this software processing and uses NetFlow resources (instead of hardware ACL TCAM) for the initial packet in each flow. The packet is then encapsulated in GRE and forwarded to the WCCP appliance.

The establishment of the NetFlow entry impacts CPU utilization, but subsequent packet forwarding is done in hardware. Traffic patterns, especially the number of unique flows, dictate how much the CPU is utilized. If the NetFlow resources of the Catalyst 6500 are consumed, then all traffic is forwarded in software.

The NetFlow resources of the supervisor PFC differ across the various platforms. Currently, the largest NetFlow resources are available on the PFC−3BXL on the Supervisor Engine 720 platform.

L2 Forwarding Method

With L2 forwarding, the WCCP entities (ACNS, WAFS, WAAS, and so forth) within a service group are part of the same subnet and are L2 adjacent to the Catalyst 6500. This enables high throughput, low latency redirection of traffic. The ingress interface (where WCCP is configured) and the interface where the WCCP appliance(s) are located must be on different VLANs.

Note: With L2 redirection, the packet is rewritten with the source MAC set to the router and the destination MAC set to the cache engine. The only disadvantage of this redirection method is that the cache engine must be ***L2 reachable*** by the Catalyst 6500 and must reside on a ***different L3 interface*** than the configured ingress WCCP interface.

Note: The method of forwarding to the WCCP device may not be the same method that the WCCP device uses in order to send traffic back to the Catalyst 6500. WCCP is used in order to negotiate a forward and return method that both devices support. See WCCP Return Method.

The options available for WCCP redirection in this scenario include:

- Ingress − L2 redirection + hash assignment; this requires software processing.
- Ingress − L2 redirection + mask assignment this requires full hardware processing and is recommended.
- Egress − L2 redirection + hash assignment; this requires software processing.
- Egress − L2 redirection + mask assignment; this requires software processing.

Ingress − L2 + Hash Assignment

When configured on ingress with L2 + hash assignment, WCCP traffic sends the first packet in every flow to be software switched, which creates a NetFlow entry in the hardware NetFlow table.

Note: The NetFlow flow−mask is set to interface full−flow mode, which might impact other NetFlow features configured on the switch.

Since WCCP is a stateless mechanism, the information is not maintained in software; rather, it is maintained in hardware as entries in the NetFlow table. Subsequent traffic in the flow is forwarded in hardware as long as a NetFlow table entry exists.

The establishment of the NetFlow entry impacts CPU utilization, but subsequent packet forwarding is done in hardware. Traffic patterns, especially the number of unique flows, dictates how much the CPU is utilized. If the NetFlow resources of the Catalyst 6500 are consumed, then all traffic is forwarded in software.

The NetFlow resources of the supervisor PFC differ across the various platforms. Currently, the largest NetFlow resources are available on the PFC−3BXL on the Supervisor Engine 720 platform.

Ingress − L2 + Mask Assignment

When configured on ingress, L2 + mask assignment is the most efficient WCCP method supported on the Catalyst 6500. All traffic is hardware switched, including the initial packet in each flow. No software redirection is required; initial and subsequent packet forwarding are provided by hardware.

The hardware ACL TCAM resources of the Catalyst 6500 are used in order to preprogram the hardware entries before any WCCP packets are received.

In order to use this method and utilize full hardware switching, the WCCP entity must also support L2 redirect and the mask−based assignment method. Configuration of this method is completed on the WCCP entity, and the Catalyst 6500 negotiates the best method during the WCCP initial communications with the WCCP entity/group.

Egress − L2 + Hash Assignment

With egress L2 + hash assignment, WCCP traffic sends the first packet in every flow to be software switched, which creates a NetFlow entry in the hardware NetFlow table.

Note: The NetFlow flow−mask is set to interface full−flow mode, which might impact other NetFlow features configured on the switch.

Additionally, when configured in the egress direction, an additional forwarding information base (FIB) lookup is required on the first packet of the flow in order to determine the adjacency associated with the CE, which requires packet recirculation within the Catalyst 6500. Subsequent packets are NetFlow switched in hardware.

The establishment of the NetFlow entry impacts CPU utilization, but subsequent packet forwarding is done in hardware. Traffic patterns, especially the number of unique flows, dictates how much the CPU is utilized. If the NetFlow resources of the Catalyst 6500 are consumed, then all traffic is forwarded in software.

The NetFlow resources of the supervisor PFC differ across the various platforms. Currently, the largest NetFlow resources are available on the PFC−3BXL on the Supervisor Engine 720 platform.

Egress − L2 + Mask Assignment

When configured in the egress direction, L2 + mask assignment switches the first packet in each flow in software, just like the L2 + hash assignment case. Subsequent packets are NetFlow switched in hardware.

Note: The NetFlow flow−mask is set to interface full−flow mode, which might impact other NetFlow features configured on the switch.

The PFC2 and PFC3 do not support egress ACL adjacency programming, which forces software processing for the initial packet in each flow; subsequent packets in the flow are forwarded in hardware.

The establishment of the NetFlow entry impacts CPU utilization, but subsequent packet forwarding is done in hardware. Traffic patterns, especially the number of unique flows, dictates how much the CPU is utilized. If the NetFlow resources of the Catalyst 6500 are consumed, then all traffic is forwarded in software.

The NetFlow resources of the supervisor PFC differ across the various platforms. Currently, the largest NetFlow resources are available on the PFC−3BXL on the Supervisor Engine 720 platform.

WCCP Return Method

WCCP Entity Can Service the Request

When WCCP is used to intercept traffic and the WCCP entity performs a complete operation on those packets, the packets are then ready to be sent back to client from the WCCP device. This normally processed traffic, which is destined back to the client on the network, does not require any special encapsulation when sent out of the WCCP device back toward the client.

Because the WCCP interception has resulted in the client request being successfully serviced (file from cache, split stream from cache, file from WAAS), it can be sent back to the network as normal traffic with the destination address in the packets being the original requestor. This traffic can be normally L3/switched by the Catalyst 6500 if it is in the network path from the WCCP appliance to the client; with an L2 attached WCCP appliance, traffic is in the network path. No encapsulation is needed in order to send it back from the WCCP device to the Catalyst 6500 because the destination is now the original client instead of a server on the internet or intranet. The network then handles this like any other IP traffic flow and uses hardware forwarding in the Catalyst 6500 in order to return the requested traffic to the client.

WCCP Entity Cannot Service the Request

In some cases where the WCCP entity cannot perform the requested operation, the WCCP appliance may need to send the traffic back to the Catalyst 6500 and preserve the original destination of the packets. Forwarding this traffic from the WCCP entity without encapsulation may result in traffic loops. In order to hide an unsuccessful service attempt from the client and send the packets to the original destination to be serviced, the packets must remain unchanged, be placed back into their original forwarding path, and forwarded without WCCP interception to the original destination.

In the WCCP return method,WCCP can be used to encapsulate these packets, send them back to the device that intercepted them in the first place, strip off any encapsulation, and place them back into the forwarding path from which they were intercepted. These packets need to be sent normally as if they were never intercepted by WCCP.

Examples of these cases include:

- Cache overloaded traffic, where traffic is being bypassed
- Rules on cache devices that deny the WCCP entity from servicing the traffic
- Bypassed traffic that is seeking a service not available on the device

At this time, this return method can be done only with GRE encapsulation and is not yet supported in any Catalyst 6500 hardware. If large amounts of traffic are sent back to the Catalyst 6500 with this method, high CPU utilization might occur because this traffic is processed in software. In Cisco IOS Software Release 12.1(18)SXH, there is an L2 return method supported by the Catalyst 6500 hardware.

GRE Return

In Cisco IOS software releases earlier than 12.2(18)SXH, the only return method supported for the Catalyst 6500 is GRE encapsulation. In addition to the GRE header appended to the return traffic, a WCCP header is also appended. Although GRE is natively supported in the Supervisor Engine 32 and Supervisor Engine 720 hardware, this additional header results in the GRE not being hardware assisted. Note that both the Catalyst 6500 and the WCCP appliance must support and negotiate the L2 return method.

No GRE hardware support exists in the Supervisor Engine 2 for any GRE, ingress, egress, or WCCP return. In order to process this type of GRE de−encapsulation, the Cisco IOS software programs the WCCP GRE tunnel receive−adjacency on the WCCP enabled interface to point to the route processor (RP), which results in software processing of return traffic.

Use of redirect lists at the Catalyst 6500 in order to avoid traffic that may need to be returned through GRE is an effective method to reduce the software processing requirements of traffic that would be sent back from the WCCP entity. This is much more effective than having traffic denied on the WCCP entity and forcing it to be GRE encapsulated and sent back to the Catalyst 6500.

Remember that the WCCP service group is scalable. If excess traffic is bypassed due to load, then this traffic is sent back, which creates CPU load on the Catalyst 6500. Proper scaling or even overbuilding of the WCCP service group is the only method to avoid this situation.

L2 Return Enhancement

In 12.2(18)SXH, an option allows the WCCP entity to rewrite the L2 MAC address rather than encapsulate return traffic. This L2 return enhancement (Cisco bug ID CSCuk59825) enables hardware processing of returned traffic when WCCP is configured to use ingress redirection with mask assignment.

Summary of WCCP Options

When implemented on the Catalyst 6500, WCCP offers many configuration options, as shown in this table. Note that the WCCP appliance negotiates these options and controls which options are used by the Catalyst 6500. Configuration is done on the WCCP appliance side of the WCCP connection.

Redirect Method	***Assignment Method***	***Ingress/ Egress***	***Switching Result***
L2	Hash	Ingress	Software processing
L2 (Recommended)	***Mask***	***Ingress***	***Full hardware processing with ACL TCAM***
L2	Hash	Egress	Software processing
L2	Mask	Egress	Software processing
GRE	Hash	Ingress	Software processing
GRE (PFC3 or newer)	***Mask***	***Ingress***	***Full hardware processing with NetFlow full−flow***
GRE	Hash	Egress	Software processing
GRE	Mask	Egress	Software processing

From a hardware perspective, all egress WCCP configurations require software processing and impact CPU utilization. Software processing is also required on ingress when the hash−based assignment method is used and results in the same potential impact on CPU utilization.

The recommended method of WCCP deployment on the Catalyst 6500 is L2 redirection with mask assignment and, when available, L2 return.

Tip: Only one option ensures high performance, full hardware−based forwarding: the ***ingress−based L2 redirection with mask assignment*** on Supervisor Engine 2, Supervisor Engine 32, and Supervisor Engine 720.

WCCP Design Recommendations

Use these configuration recommendations so you can determine the best method of WCCP deployment for your situation.

Summary

Design the network such that WCCP ingress can be used as a redirection method. A good design method is to have a caching switchblock as part of a hierarchical, L3 distribution network; this ensures that WCCP serviced traffic can be identified at a few major ingress ports.

- Use 12.2(18)SXF7 or newer Cisco IOS software.
- Identify and redirect traffic on ingress interfaces.
- Use WCCP appliances that support the L2 redirect method.
- Use WCCP appliances that support the mask−based assignment method.
- If possible, use a redirect list at the Catalyst 6500 for traffic that cannot be serviced by the WCCP appliance.
- Tweak NetFlow timers with egress or any hash−based assignment configurations with Supervisor Engine 720.
- WCCP appliances should have a dedicated L3 environment and SVI/VLAN interface.

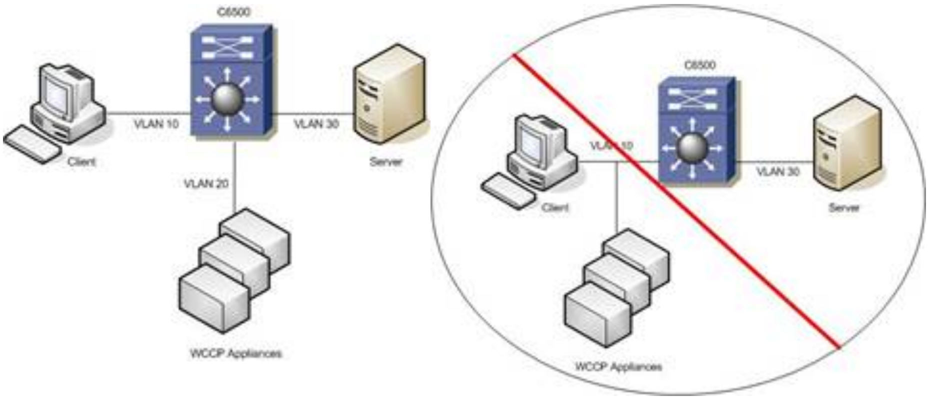

In addition, Cisco Advanced Service recommends these design considerations:

- In an environment that is not fully L2 redirect with mask assignment, the Supervisor Engine 720 does not perform better than the Supervisor Engine 2 platform. Do not upgrade hardware and expect better performance in this situation.
- For large, central site deployments with high volume traffic requirements, consider a design with policy−based routing (PBR) and Cisco Content Switching Model (CSM)/Application Control Engine (ACE) for traffic interception and forwarding.
- WCCPv2 functions as expected under perfect circumstances. However, under some circumstances, CPU utilization on the router can reach high levels that render the device unusable and that require reload:
    - WCCPv2 falls back from ingress hardware−accelerated L2 mask−based assignment mode into any other mode that requires the CPU on the MSFC.
    - WCCP is incorrectly configured (for example, egress instead of ingress or L2 hash instead of mask assignment).

- Any high volume WCCP deployment with Catalyst 6500 (caching, streaming, WAAS, or other 'high traffic rate' scenarios) should be performance tested with live traffic before full production deployment.

Additional Recommendations

Use a redirect list at the switch in order to avoid packets that are sent back to the Catalyst 6500. If any rules of the cache devices can be moved to the Catalyst 6500 as a redirect list, this may provide better hardware performance.

NetFlow resources on the Supervisor Engine 720 platform can be exhausted quickly if you use any method other than ingress−L2 mask assignment. Supervisor Engine 720 does not provide better performance than Supervisor Engine 2 with any other method.

In cases where the Supervisor Engine 720 or Supervisor Engine 32 must be used in a non−optimal design, consider use of the ***mls ip netflow creation software−mode fast*** command so that the NetFlow processing of the WCCP initial packet can be speeded up. This removes the enhancements added to Supervisor Engine 32 and Supervisor Engine 720 NetFlow and provides performance equal to that of the Supervisor Engine 2 NetFlow hardware.

The configuration for a Cisco content engine (CE) for mask assignment is:

- ***wccp version 2***
- ***wccp router list*** *number ip−address*
- ***wccp*** *service* ***router−list−num*** *num* ***mask−assign***

Use these commands in order to review NetFlow utilization and determine if WCCP is using up NetFlow entries and is using software processing:

- ***show mls netflow table−contention detailed***
- ***show mls netflow ip sw−installed***
- ***show mls netflow aging***
- ***show mls netflow ip dynamic count***
- ***show mls netflow ip count***
- ***show mls ip***
- ***show fm netflow counters***

If you experience WCCP software issues because the NetFlow resources are being consumed, these commands might aggressively clear out existing entries and create room for new entries. (This does not help if there are simply more entries than there is NetFlow space.)

- ***mls aging fast time 3 threshold 1***
- ***mls aging long 64***
- ***mls aging normal 32***
- ***mls ip netflow creation software−mode fast*** (This disables some Supervisor Engine 720 NetFlow changes that were not in Supervisor Engine 2.)

To avoid packet drops, the WCCP entities must redirect traffic out of an interface that is not the interface that the WCCP is configured on. Catalyst 6500 WCCP drops packets in this situation when all conditions are met:

- Both clients and WCCP entity are on the same L3 interface.
- WCCP redirection policy is applied on this interface.
- GRE redirection method is used.
- Redirecting packet fragmentation is required.

- Supervisor Engine 720 is used.

This situation is caused by protection mechanisms built in to the Catalyst 6500; the Cisco IOS software has checks that prevent the packet from entering and leaving the same Cisco IOS software virtual interface where it could potentially loop and cause undesired behavior. Move WCCP appliances to their own dedicated L3 environment in order to prevent this.

User−based rate limiting (UBRL) and WCCP do not work simultaneously on an interface because of flowmasks. There is one flowmask for each interface for each unicast feature. WCCP requires full−flow, and UBRL uses src−only or dst−only.

WCCP support was added for Supervisor Engine 2 and L2 return in 12.2(18)SXF5. This was not in Supervisor Engine 720 until 12.2(18)SXH in April/May 2007.

Only WCCP L2 PFC redirection is supported with Cisco IOS software server load balancing (SLB); other WCCP configurations are not compatible, and GRE does not work. The WCCP accelerated command only applies to Supervisor Engine 2/MSFC2. Its purpose is to force the router to negotiate mask assignment and L2 redirect, which means that all WCCP redirection is done in hardware. The Supervisor Engine 32 and Supervisor Engine 720 negotiate this without the need for this command.

Solutions

Note: Use the Command Lookup Tool (registered customers only) in order to obtain more information on the commands used in this section.

ACNS

For standard transparent caching redirection, recall that the WCCP entity provides the WCCP router with the supported methods and may need to be configured to do so. For Cisco ACNS, this example configuration requests the optimized L2−redirect and mask−based assignment methods:

1. Ensure you have WCCPv2 for the Catalyst 6500 optimizations:

   ```
   ContentEngine(config)# wccp version 2
   ```
2. Configure a router list that defines the Catalyst 6500s to use:

   ```
   ContentEngine(config)# wccp router−list 1 172.16.16.1
   ```
3. Configure the service to use the optimization methods:

   ```
   ContentEngine(config)# wccp service router−list−num 1 l2−redirect mask assign
   ```

From the router side, the Catalyst 6500 design should ensure that the WCCP devices are on a dedicated L3 interface that is not in the current traffic path (ingress or egress). For hardware performance, traffic flows should be captured inbound to the Catalyst 6500, even if this requires configuration of more interfaces than if a single egress interface was chosen. An ideal design would aggregate all traffic before reaching this device, and only a few interfaces would require the WCCP ingress configuration.

The WCCP configuration on the Catalyst 6500 should be:

1. Configure WCCPv2:

```
6500Switch# ip wccp version {1 | 2}
```

2. Configure a WCCP service group.

```
6500Switch (config)# ip wccp service [accelerated] redirect−list access−list
```

Use the accelerated command only for Supervisor Engine 2 platforms with 12.1E Cisco IOS software.

The redirect list is used in order to identify the traffic that should be selected or not selected for redirection. Recall that this ACL can be performed in hardware, which is a much more efficient way to prevent redirection for traffic that cannot be serviced by the WCCP device. Traffic that is sent to the device and cannot be serviced there must be returned to this Catalyst 6500 in order to be put back into the original traffic path, which requires additional processing. The WCCP access lists are standard or extended access−lists.

This example shows that any requests from 10.1.1.1 to 12.1.1.1 bypass the cache and that all other requests are redirected.

```
6500Switch(config)# ip wccp service redirect−list 120
6500Switch(config)# access−list 120 deny tcp host 10.1.1.1 any
6500Switch(config)# access−list 120 deny tcp any host 12.1.1.1
6500Switch(config)# access−list 120 permit ip any any
```

Configure the ingress WCCP method on each ingress interface that receives the traffic to be redirected:

```
Router(config−if)# ip wccp service redirect in
```

This completes the configuration on the WCCP appliance and the switch, so traffic redirection should be occurring at this point.

The final WCCP configurations of the devices look like this.

Device	***Configuration***
WCCP appliance	wccp version 2 wccp router−list 1 *router−ip−addresses* wccp *service* router−list−num 1 l2−redirect mask assign
WCCP router: global	ip wccp version 2 ip wccp *service* redirect−list 120 access−list 120 deny tcp ... access−list 120 deny udp ... access−list 120 permit ip any any
WCCP router: each ingress interface	ip wccp redirect *service* in

To check this configuration, enter this command:

```
Show ip wccp service detail
```

For additional WCCP configuration options, such as group addressing using multicast or additional WCCP security, refer to Configuring Web Cache Services using WCCP.

WCCP IOS Show and Debug Commands

- ***show ip wccp*** *service−number* − provides the 'Total Packets Redirected' count. This count is the number of flows, or sessions, that are redirected.
- ***show ip wccp*** *service−number* ***detail*** − provides the 'Packets Redirected' count. The count is the number of flows, or sessions, that are redirected.

- ***show ip wccp web−cache detail*** − provides an indication of how many flows, rather than packets, are using L2 redirection.
- ***clear ip wccp*** − resets the counter for the 'Packets Redirected' information.
- ***show ip wccp*** *service−number* ***view*** − displays the WCCP devices that are part of the service group.
- ***show ip wccp*** *service−number* ***service*** − displays the hash, ports, and WCCP priority of the service. Higher priority services are hit first when multiple services are configured on the interface.
- ***debug ip wccp events*** − troubleshoots the state of the WCCP protocol.
- ***debug ip wccp packets*** − reviews the communications between the WCCP packet processing entities.

When you use WCCP and hardware forwarding, some counters may not display as expected:

- If the WCCP ACL is processed completely in hardware, WCCP counters may not display accurate packet counts.
- If the WCCP is using hash−based assignment and NetFlow hardware resources, the WCCP counters may reflect the number of flows instead of the number of packets.

NetFlow Commands

When you have WCCP configurations that require the use of NetFlow hardware resources, use these multilayer switching (MLS) and Fabric Manager (FM) commands so you can review the status of the NetFlow resources:

- ***show mls netflow table−contention detailed***
- ***show mls netflow ip sw−installed***
- ***show mls netflow aging***
- ***show mls netflow ip dynamic count***
- ***show mls netflow ip count***
- ***show mls ip***
- ***show fm netflow counters***

WCCP Defects

This table of Cisco bug IDs and resolutions supports the general recommendation to use Cisco IOS Software Release 12.2(18)SXF7 or later for the best support of WCCP.

Cisco Bug ID	***Resolved in Cisco IOS Software Release***	***Details***
CSCsd20327	12.2(18)SXF7	WCCP for service 90 goes up and down, and causes a loss of WCCP service. This problem occurs when services 81, 82, and 90 are configured. The packet traces indicate that the router might respond to 'Here_I_Am' messages from the cache with 'I_See_You' messages that contain an incorrect destination IP address.
CSCsa77785	12.2(18)SXF6	A reload might occur when you use WCCP L2 redirection and mask assignment mode with a host−based standard ACL as a WCCP redirect ACL.
CSCse69713	12.2(18)SXF6	When all cache engines in a WCCP service group are lost, traffic is processed in software instead of switched in hardware.
CSCsd28870	12.2(18)SXF5	In a WCCP redirect ACL list, ACEs that are configured with the log keyword are not programmed into the TCAM table.
CSCsb61021	12.2(18)SXF5	High CPU utilization might occur on a Supervisor Engine 720 or a Supervisor Engine 32 when the IP spoofing feature is configured on a

		cache engine and when WCCP redirection is configured in the egress direction. IP−spoofed packets from the cache engine, with a destination of either the client or the server, are switched in software instead of hardware. As a workaround, use the ***ip wccp*** *service* ***redirect in*** command for both the inbound and the outbound interfaces.
CSCsb21972	12.2(18)SXF2	With both WCCP and NDE configured, you might see numerous tracebacks caused by alignment errors, and CPU utilization might be unacceptably high.
CSCeh85087	12.2(18)SXF	When there is a user−configured 'deny ip any any' in the WCCP redirect ACL and when many WCCP service groups are being serviced, the traffic associated with some service groups is not redirected to CE routers.
CSCeh56916	12.2(18)SXF	When a WCCP service is enabled, when mask assignment is configured as the assignment method, and when five or more caches are in the service group, protocol messages sent to the cache may overflow and cause memory corruption and reload.
CSCsb18740	12.2(18)SXF and SXE6	In GRE−based forwarding mode, WCCP unnecessarily uses a software cache that increases MSFC CPU utilization.
CSCsb26773	12.2(18)SXF	An inbound ACL might cause WCCP redirection to fail with the loss of all redirected traffic.
CSCsa90830	12.2(18)SXE2	WCCP ingress−redirected traffic uses the NetFlow table for hardware switching when the cache engine is configured for GRE forwarding with the mask assignment mode. When the NetFlow table is full, WCCP ingress redirection fails.
CSCec55429	12.2(18)SXE	The WCCP service group list is scanned in the order in which service groups are created, rather than by priority. If multiple dynamic WCCP services are defined, traffic that matches the selection criteria for more than one service group is not redirected to the service group with the highest priority.
CSCuk50878	12.2(18)SXE	In a release where Cisco bug ID CSCec55429 is resolved, after a number of WCC 'cache lost' and 'cache found' events have occurred for all the caches in a service group, these events may occur: • Spurious memory accesses might occur. • The addition and deletion of WCCP services might fail. • The ***show ip wccp*** command displays the WCCP service, but the output of the ***show ip wccp*** *service_number* command does not show the WCCP service.
CSCsa67611	12.2(18)SXE	Incoming Multiprotocol Label Switching (MPLS) packets that exit on a non−MPLS interface (tag to IP path) on which an output feature is configured (for example, egress ACL or egress WCCP) might not have the output features applied. This problem occurs because the output ACL lookup is bypassed.
CSCeh13292	12.2(18)SXD4	Configuration of WCCPv2 on a Supervisor Engine 720 causes high CPU utilization.
CSCeb28941	12.2(18)SXD1	Network Address Translation (NAT) does not work with WCCP configured.
CSCed92290	12.2(17d)SXB2	WCCP−redirected packets that have no next−hop Address Resolution Protocol (ARP) cache entry are process switched to generate an ARP

		request. Because of the WCCP redirection, however, no ARP request is sent, the ARP cache is never populated for the next hop, and subsequent WCCP−redirected packets continue to be process switched.
CSCuk59825	12.2(17d)SXF5 −Sup2 Whitney1.0 for Sup720	This Cisco IOS software release added hardware support for L2 return traffic. The WCCP Request for Comment (RFC) specifies L2 return as an optional capability for negotiation between router and cache. Until now, WCCP on Cisco IOS software has not allowed the negotiation of this capability because the required hardware support is absent. That support is now available, so the negotiation of L2 return in the WCCP protocol exchange between router and cache can be enabled.

Updated: Jul 16, 2013 Document ID: 116134