CISCO PUBLIC WHITE PAPER

CISCO WAAS DEPLOYMENT USING WEB CACHE COMMUNICATION PROTOCOL VERSION 2 (WCCPV2)

Cisco Wide Area Application Services (WAAS) relies on network interception to be integrated into the network as well as to receive packets from flows that should be optimized. WCCPv2 gives IT organizations a mechanism to transparently intercept and redirect this network traffic to a nearby WAAS device for application acceleration and WAN optimization. This document provides a basic understanding of how WCCPv2 works and how to configure and manage a WCCPv2 deployment for WAAS.

OVERVIEW

WCCPv2 is a protocol that allows devices such as the Cisco WAE to join service groups with network devices (such as routers, switches, or firewalls) for the purposes of injecting itself into the flow of application traffic to optimize or otherwise manipulate flows. When configured

properly, the WCCPv2 process on the network device will examine traffic to identify flows related to applications that match the criteria defined in the configured service groups. When this traffic is identified, the network device will then redirect the traffic to one of the registered service group devices, such as a Cisco WAE, using either a Generic Route Encapsulation tunnel (GRE) or through frame header rewriting, called Layer 2 redirection (L2-redirect). Once the WAE has received packets, it can then apply a function to the flow, such as local response handling, message suppression, or compression.

Figure 1 Basic WCCP Functionality

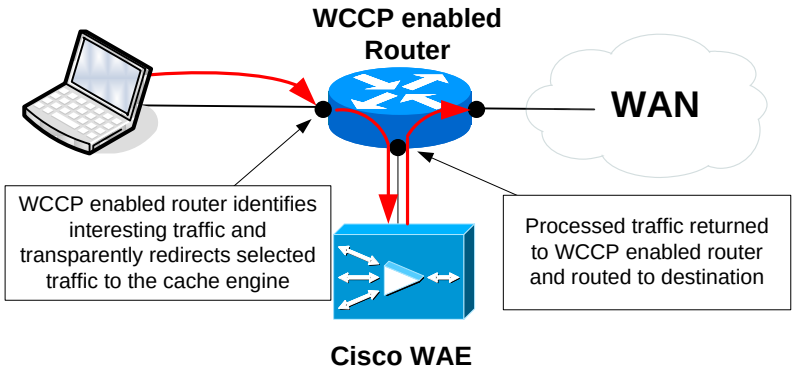

SpeedBump: Note that, throughout this paper, the term "cache engine" is used in the WCCP protocol definition to refer to any device to which it redirects traffic. All WAAS devices including WAE, WAVE, and NME-WAE network modules can be "cache engines" receiving WCCP redirected packets.

WCCP PROTOCOL BASICS

WCCPv2 defines mechanisms to allow one or more routers enabled for transparent redirection to discover, verify, and advertise connectivity to one or more cache engines. Having established connectivity, the routers and WAAS devices form Service Groups to handle the redirection of traffic whose characteristics are part of the Service Group definition.

The protocol provides the means to negotiate the specific method used for load distribution among WAAS devices and also the method used to transport traffic between router and cache. A single WAAS device within a Service Group is elected as the *designated* cache engine. It is the responsibility of the designated cache engine to provide routers in the Service Group with the data which determines how redirected traffic is distributed between the WAAS devices in the Service Group.

Service Groups

A ***Service Group*** is a group of one or more routers plus one or more WAAS devices working together in the redirection of traffic whose characteristics are part of the Service Group definition. A WAAS device joins and maintains its membership in a Service Group by transmitting a WCCP2_HERE_I_AM (HIA) message (Figure 3) to each router in the Group at 2 second intervals[1]. This may be by unicast to each router or multicast to the configured Service Group multicast address. The Web Cache Info component in the WCCP2_HERE_I_AM message identifies the WAAS device by IP address and the Service Info component identifies and describes the Service Group in which the WAAS device wishes to participate.

A router responds to a WCCP2_HERE_I_AM message with a WCCP2_I_SEE_YOU (ISU)message (Figure 2). If the WCCP2_HERE_I_AM message was unicast then the router will respond immediately with a unicast WCCP2_I_SEE_YOU message. If the

[1] The WCCP V2.0 Protocol Specification defines a default HERE_I_AM_T interval of 10 seconds. Cisco WAAS uses 2 seconds as the HERE_I_AM_T interval.

WCCP2_HERE_I_AM message was multicast, the router will respond via the scheduled multicast WCCP2_I_SEE_YOU message for the Service Group.

A router responds to multicast cache engine members of a Service Group using a multicast WCCP2_I_SEE_YOU message transmitted at 9 second intervals with a 10% jitter. The Router Identity component in a WCCP2_I_SEE_YOU message includes a list of the WAAS devices to which the packet is addressed.

The Service Info component of a WCCP2_HERE_I_AM message describes the Service Group in which a WAAS device wishes to participate. A Service Group is identified by Service Type and Service ID. There are two types of Service Group:

- Well Known Services
- Dynamic Services.

Well Known Services are known by both routers and cache engines and do not require a description other than a Service ID. Service IDs can range from 0 to 255 with the 0-50 range reserved for Well Know Services. Currently, web-cache is the only defined Well Known Service. In contrast Dynamic Services must be fully described to a router. A router may be configured to participate in a particular Dynamic Service Group, identified by Service ID, without any knowledge of the characteristics of the traffic associated with the Service Group. The traffic description is communicated to the router in the WCCP2_HERE_I_AM message of the first cache engine to join the Service Group. A cache engine describes a Dynamic Service using the Protocol, Service Flags and Port fields of the Service Info component. Once a Dynamic Service has been defined a router will discard any subsequent WCCP2_HERE_I_AM message which contains a conflicting description. A router will also discard a WCCP2_HERE_I_AM message which describes a Service Group for which the router has not been configured.

Table 1. Cisco WCCP Service Groups

Service Name	Service Number	Protocol	Port	Priority
web-cache	0	tcp	80	240
dns	53	udp	53	202
ftp-native	60	tcp		200
tcp-promiscuous	61	tcp	*	34
tcp-promiscuous	62	tcp	*	34
https-cache	70	tcp	443	231
rtsp	80	tcp	554	200
wmt	81	tcp	1755	201
mmsu	82	udp	1755	201
rtspu	83	udp	5005	201
cifs-cache	89	tcp	139, 445	224
custom	90-97			220-227
custom-web-cache	98	tcp	80	230
reverse-proxy	99	tcp	80	235

A router considers a WAAS device to be a usable member of a Service Group only after it has sent that WAAS device a WCCP2_I_SEE_YOU message and received a WCCP2_HERE_I_AM message with a valid "Receive ID" in response. The "Receive ID" is sent by a router in the WCCP2_I_SEE_YOU message and is reflected back by a WAAS device in the Cache engine View Info component of a WCCP2_HERE_I_AM message.

Figure 2. WCCPv2 I See You packet

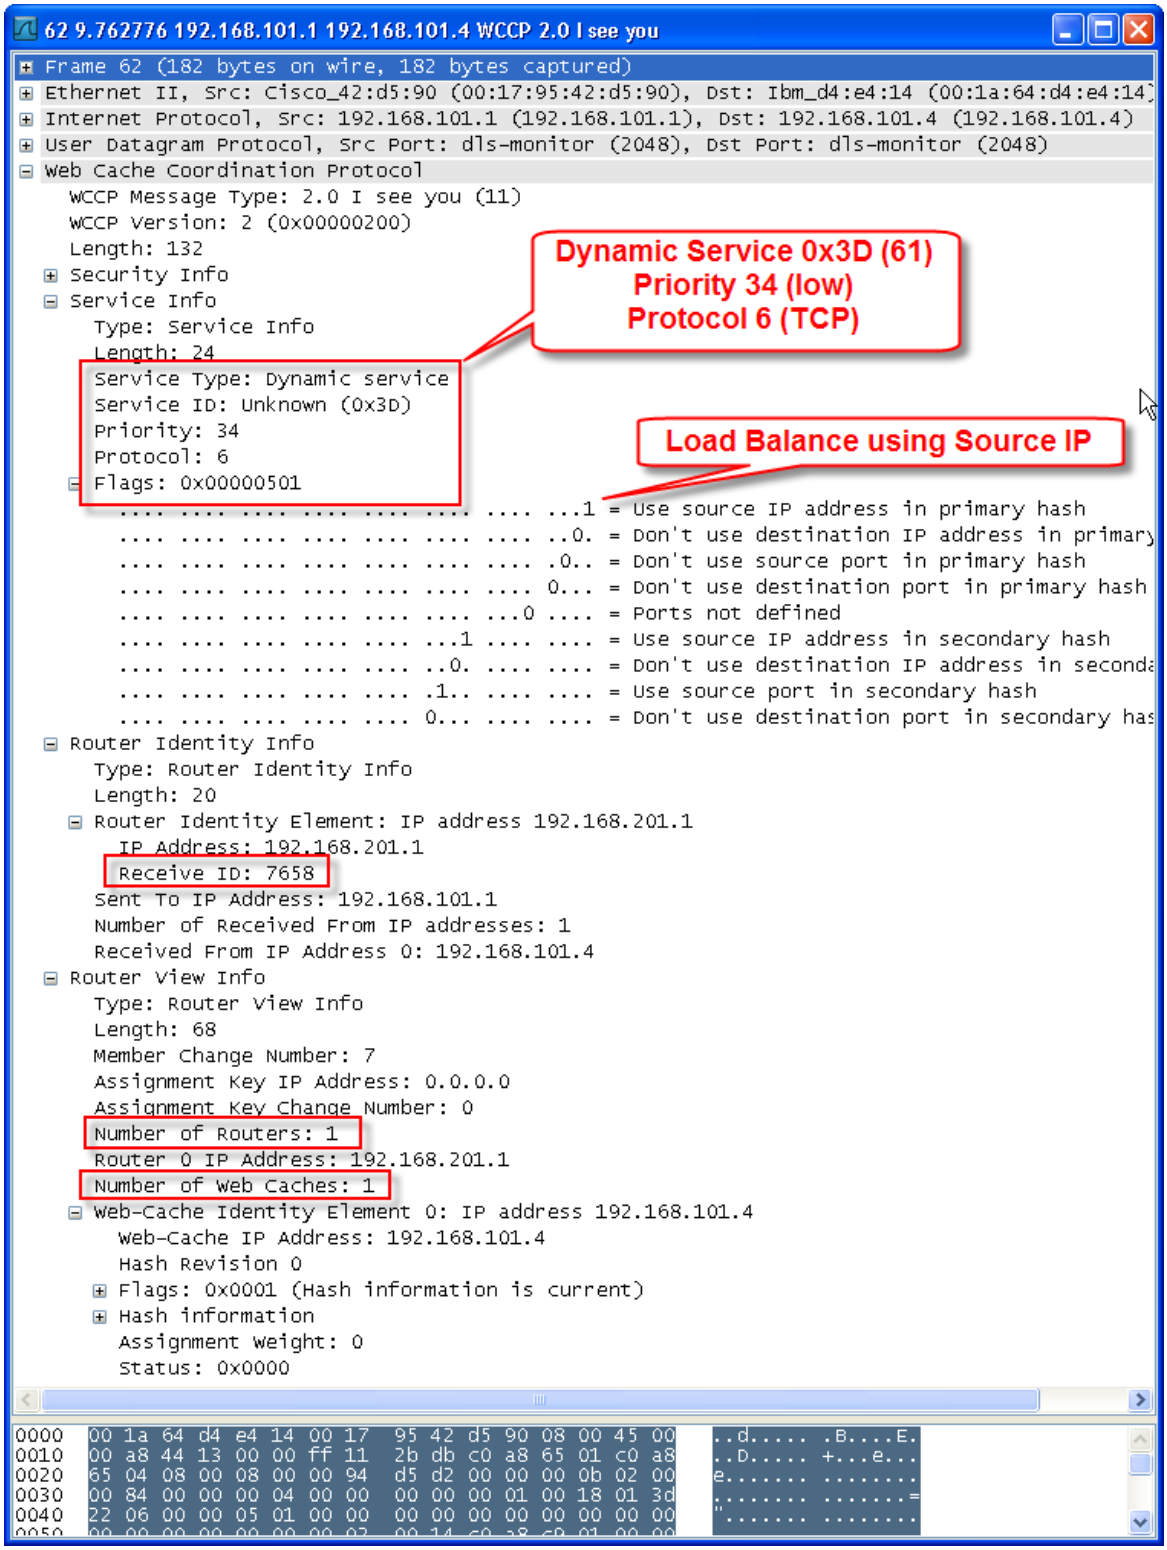

Figure 3. WCCPv2 Here I Am packet

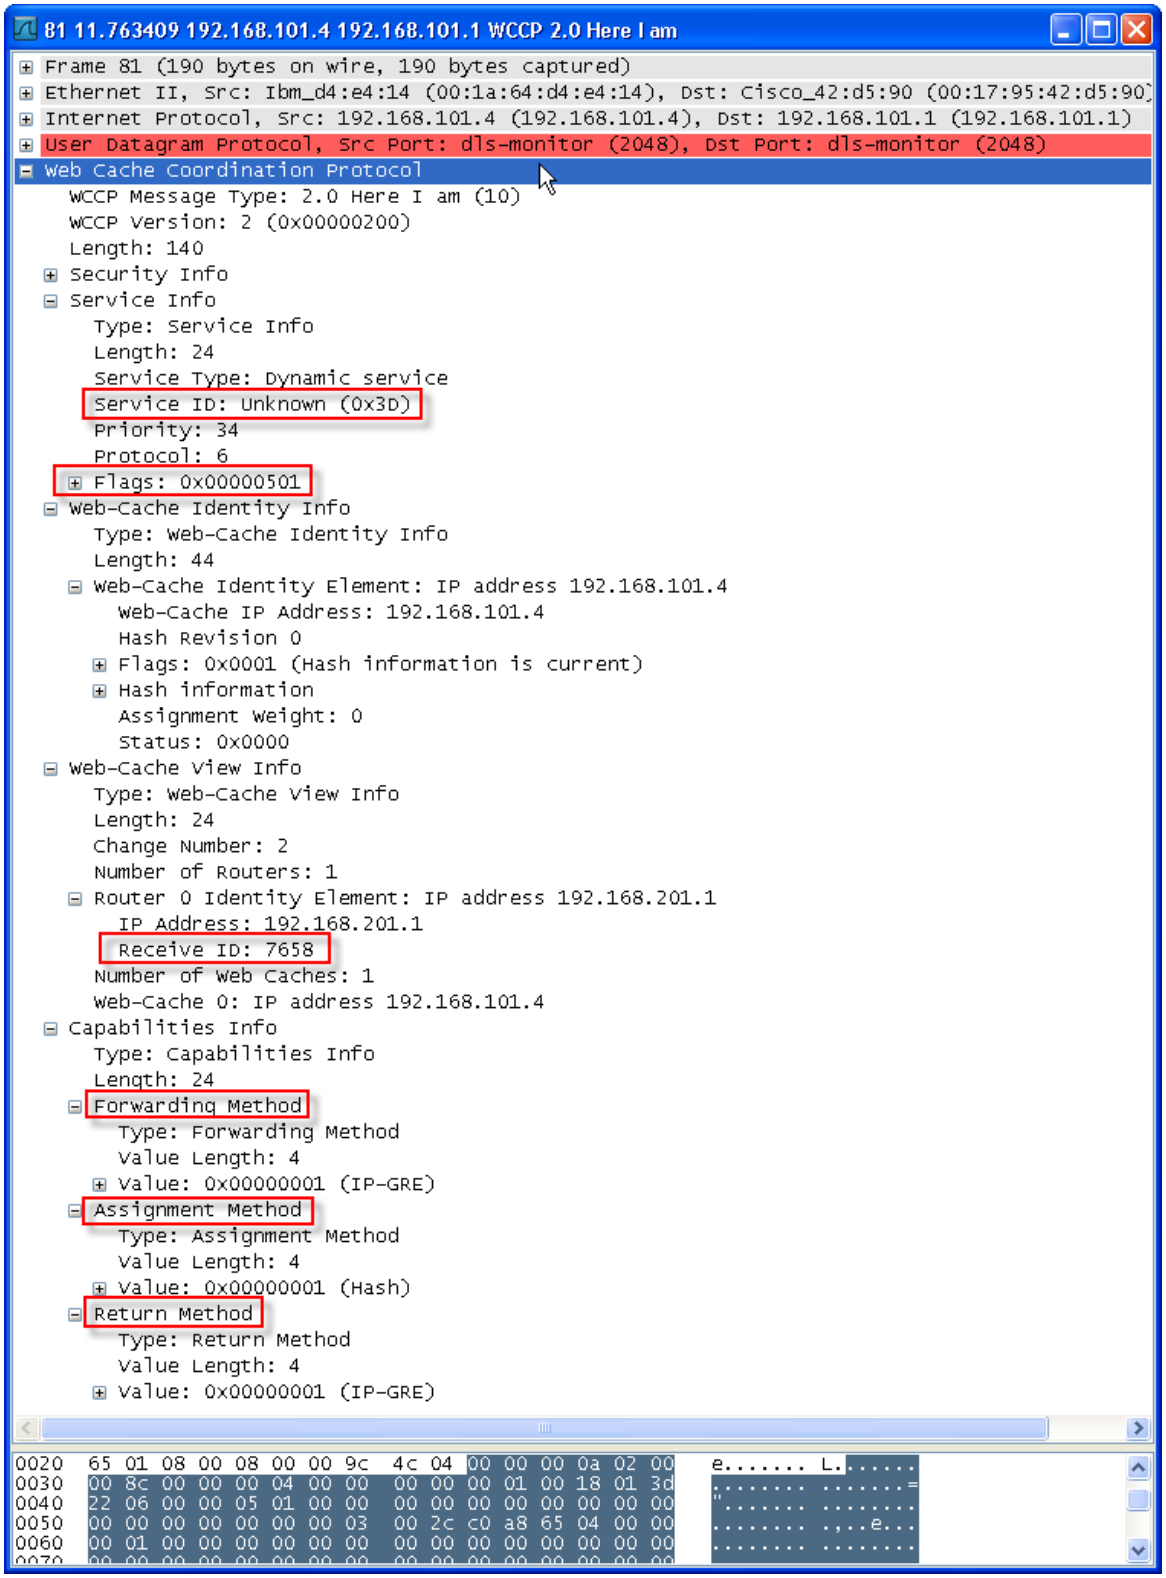

Redirect Method

The *Redirect Method,* also known as the Forwarding Method, is the method by which redirected packets are transported from router to WAAS device. This method is negotiated between the router and the WAAS device. While the WCCPv2 protocol allows different cache engines to use different redirect methods, Cisco WAAS and IOS-WCCP prevent this from happening. Each Cisco WAAS device belonging to a service group will use the same redirect method.

A router will advertise the supported redirect methods for a Service Group using the optional Capabilities Info component of the WCCP2_I_SEE_YOU message. The absence of such an advertisement implies the router supports the default GRE encapsulation method only.

Cisco WAAS supports two different Forwarding Methods:

1. WCCP GRE
2. WCCP Layer 2 (L2)

Best Practice: Use WCCP GRE encapsulation when working with routers because most routers do not support L2 redirection. Use WCCP L2 Redirection as the packet forwarding method when working with switches or Cisco 7600 series routers, because redirection can be performed in hardware. This approach minimizes the CPU workload on the Cisco WAE and the switch, and can improve overall performance.

WCCP GRE

WCCP GRE, also known as Layer 3 Generic Routing Encapsulation (GRE), allows packets to reach the WAAS device even if there are other routers in the path between the forwarding router and the WAAS device. The connection between the router and the WAAS device is also known as a *GRE Tunnel*. Packet redirection is handled entirely by the router software. GRE encapsulates the selected datagram with the GRE header containing the routing information to the selected WAAS device. The WAAS device de-encapsulates the datagram, evaluates the payload using the static bypass rules and WAAS Policy specification, and either accepts or rejects the packet. If the packet is accepted for optimization, standard TCP connection setup occurs between the client and the WAAS device and between the WAAS device and the destination server. If the packet is rejected because of a static bypass rule, it is re-encapsulated and returned to the router. The router understands that the WAE is not interested in this connection and forwards the packet to its original destination. All other packets, pass-through or optimized, are returned to the router using the configured packet egress method.

	IP	Port	
Source	1.1.1.1	5432	Payload
Destination	5.5.5.5	80	

Original Packet sent from IP 1.1.1.1 to IP 5.5.5.5

	IP	IP	Port	
Source	2.2.2.2	1.1.1.1	5432	Payload
Destination	3.3.3.3	5.5.5.5	80	

GRE Encapsulated Packet redirected from router at IP 2.2.2.2 to WAAS device at 3.3.3.3. Original packet persevered.

When using GRE encapsulation for WCCP redirection, the router uses the router ID IP address as its source IP address. The router ID IP address is the highest loopback address on the router, or if the loopback interface is not configured, the router ID IP address is the highest address of the physical interfaces. The router ID IP address is used as the source address for packets redirected from the router to the Cisco WAAS device, and as a result it is also used as the destination address for traffic from the Cisco WAE to the router.

Speedbump: **Be sure that a route exists from the Cisco WAAS device to the router. This can be guaranteed by configuring a static route on the Cisco WAAS device to the router ID IP address. The router ID can be identified with the command "show wccp routers" on the Cisco WAAS device. If the WCCP server group contains multiple routers, a static route should be added to each of these routers' router ID. The command to configure such static routes is:**

```
WAE(config)# ip route router-ID  netmask gateway-ip
```

Best Practice: **To ensure that the static route between the WAAS device and the router is always valid regardless of the state of any physical interface on the router, we recommend configuring a loopback interface as the router ID on all routers in the service group.**

WCCP L2

WCCP L2 (Layer-2) redirection takes advantage of internal switching hardware that either partially or fully implements the WCCP traffic interception and redirection functions at Layer 2. Redirection occurs by overwriting the original MAC header of the IP packet with the MAC header of one of the WAAS devices in the Service Group. With L2 Redirection, the first redirected traffic packet is handled by either the router software or router hardware, depending on the platform and/or software version. The rest of the traffic may be handled by the router hardware on supported routers and switches making L2 redirection more efficient than Layer 3 GRE. Using L2 Redirection as a forwarding method allows direct forwarding to the WAAS device without further lookup. Layer-2 redirection requires that WAAS devices be directly connected to an interface on each WCCP router. Unless multicast IP addresses are used, WCCP configuration of the WAAS device must reference the directly connected interface IP address of the WCCP router and not a loopback IP address or any other IP address configured on the WCCP router.

	MAC	IP	Port	
Source	01.01.01.01.01.01	1.1.1.1	5432	Payload
Destination	05.05.05.05.05.05	5.5.5.5	80	

Original Packet sent from IP 1.1.1.1 to IP 5.5.5.5

	MAC	IP	Port	
Source	02.02.02.02.02.02	1.1.1.1	5432	Payload
Destination	03.03.03.03.03.03	5.5.5.5	80	

L2 Rewrite redirects packet from router at MAC 02.02.02.02.02.02 to WAAS device at MAC 03.03.03.03.03.03. Original packet persevered.

Assignment Method

The *Assignment Method* is the method by which redirected packets are distributed between the WAAS devices in a Service Group effectively providing load balancing among the WAAS devices. This method is negotiated between a router and all cache engines on a per

Service Group basis. Cache engines participating in multiple Service Group may have different assignment methods for each Service Group but all cache engines within a single Service Group will use the same Assignment Method. A router may advertise the supported assignment methods for a Service Group using the optional Capabilities Info component of the WCCP2_I_SEE_YOU message. The absence of such an advertisement implies the router supports the default Hash assignment method only.

There are two types of assignment methods:

1. Hash Table Assignment
2. Mask/Value Sets Assignment

Hash Assignment

The default Assignment Method uses Hash Tables to load balance and select a particular WAAS device from those registered in the Service Group. With Hash Assignment, each router in the Service Group uses a 256-bucket Redirection Hash Table to distribute traffic for a Service Group across the member WAAS devices. The hash key may be based on any combination of the source and destination IP and port of the packet. For WAAS, load-balancing hashing is based on a source IP address (default), a destination IP address, or both.

Mask/Value Assignment

When using mask assignment, each router uses masks and a table of values to distribute traffic for a Service Group across the member WAAS devices. It is the responsibility of the Service Group's designated cache engine to assign each router's mask/value sets. For WAAS, the default mask value is 0x1741 and is applied to the source IP address for service 61 and the destination IP address for service 62. The Mask Value can be specified with a maximum of 7 bits and like the hash key, can be configured to cover both the source as well as the destination address space.

Best Practice: Do not use the default mask when using Mask Assignment. See Choosing a Mask below.

Packet Return Method

The *Packet Return Method* is the method by which packets ***not*** selected for optimization (i.e. packets which satisfy a static bypass rule or packets which do not satisfy *any* policy classifier) are returned to a router/switch for normal forwarding. The Packet Return Method is not required to match the Forwarding method. The return method is negotiated between a router and all WAAS devices on a per Service Group basis. While the WCCPv2 protocol allows different cache engines to use different packet return methods, Cisco WAAS and IOS-WCCP prevent this from happening. A router will advertise the supported packet return methods for a Service Group using the optional Capabilities Info component of the WCCP2_I_SEE_YOU message. The absence of such an advertisement implies the router supports the default GRE packet return method only.

Cisco WAAS supports two different Packet Return Methods:

1. WCCP GRE (default)
2. WCCP L2

WCCP GRE

WCCP GRE return, also known as Layer 3 GRE return, performs the reverse of GRE Encapsulation. Returning packets are encapsulated with a GRE header specifying the original WCCP Service Group router as the destination. Upon receipt, the router strips off the GRE information and routes the resulting packet normally. However, these packets sent to the original router are processed on the router at the software level potentially increasing the CPU load on certain routers.

WCCP L2

WCCP L2, also known as Layer-2 Rewrite Return, rewrites the destination MAC address to the MAC address of the WCCP Service Group router. Upon receipt, the router routes the packet normally. Layer-2 Rewrite requires that the WAAS devices in the Service Group be directly connected to the router at Layer 2.

Packet Egress Method

Beginning with Cisco WAAS Release 4.0.13, WAAS provides an alternate packet return method for packets satisfying a policy classifier including both optimized packets and packets specified by policy action as pass-through. This method is called the *Packet Egress Method* and can be defined separately from the Packet Return Method used for bypassed packets.

Cisco WAAS supports three different Packet Egress Methods:

1. IP Forwarding (default)
2. WCCP Negotiated Return
3. Generic GRE Return

IP Forwarding

IP Forwarding is the default Packet Egress Method and sends optimized packets to the configured default gateway of the WAAS device. With the IP forwarding egress method WAAS devices cannot be placed on the same VLAN or subnet as the clients and servers, and it does not ensure that packets are returned to the original intercepting router.

WCCP Negotiated Return

WCCPv2 is capable of negotiating the redirect method and the return method for intercepted connections. Cisco WAAS supports both WCCP L2 Rewrite and WCCP GRE as negotiated Packet Return Methods. If WCCP negotiates a WCCP Layer 2 Rewrite return, as Cisco WAAS does not support L2 Rewrite return as a Packet Egress method, the WAAS device defaults to using IP forwarding as the egress method.

If WCCP GRE return is negotiated as the packet egress method, the behavior of packet egress return depends on the intercept method. For GRE intercept, all packets are egressed to the designated router's routerID. The routerID must be reachable, and preferably via the shortest path, which may require static route configuration. For L2 intercept, packets are egressed to the L2 adjacent address of the router (except for flow protection, where packets may be egressed to a router's routerID), and Cisco WAAS requires (but does not enforce) that the WAAS device's routerlist should contain the router's L2 adjacent addresses.

Best Practice: To ensure that the WCCP HIA/ISU exchange verifies the data path of egress packets with GRE return, we recommend configuring a loopback interface on the router as the routerID, as well as the WAAS device's routerlist IP address.

Generic GRE Return

Like WCCP GRE Return, generic GRE Return also performs the reverse of GRE Encapsulation. The generic GRE egress method is supported only when the WCCP GRE is specified as the Redirect Method. If the interception method is set to WCCP Layer 2 and you configure generic GRE return as the egress method, the WAAS device will default to IP forwarding as the egress method as WCCP L2 Redirect method and generic GRE return are not compatible.

With generic GRE return, after optimization is performed, the packet will be returned to the same router from which it was originally redirected, preserving the original packet flow path. The generic GRE egress method returns packets to the intercepting router by using a GRE tunnel that must be manually configured on the router[2]. Unlike the WCCP GRE Return method, generic GRE Return was designed specifically to allow packets to be processed in hardware on platforms like the Cisco 7600 series router or the Catalyst 6000 series router with Sup32 or Sup720, increasing the overall performance on the router and eliminating the risk of CPU overload.

WCCP CONFIGURATION

Configuring WCCPv2 for transparent packet redirection requires configuration on both the forwarding routing/switching device and the receiving Cisco WAAS device.

WAAS Service Groups

WAAS supports a single dynamic WCCP service called tcp-promiscuous. Tcp-promiscuous defined on a WAAS device requires that two dynamic WCCP services be enabled on the forwarding routing/switching device, service groups 61 and 62. On the Cisco WAAS device, configuration requires that WCCPv2 be enabled and that routers in the service group be identified.

When configuring WCCP on a WAAS device, the WCCP service needs to be enabled and associated with one or more WCCP enabled routers. The minimal CLI required to accomplish this and to show the current WCCP status is:

```
WAE# show run | inc wccp
wccp router-list 1 192.168.5.1
wccp tcp-promiscuous router-list-num 1
wccp version 2
WAE#
WAE#show wccp stat
WCCP version 2 is enabled and currently active
WAE#
WAE# show wccp routers

Router Information for Service: TCP Promiscuous 61
        Routers Seeing this Wide Area Engine(1)
                Router Id        Sent To        Recv ID         AssKeyIP    AssKeyCN
MemberCN
                192.168.200.1   192.168.5.1     003F869F        10.86.46.74     1       12
        Routers not Seeing this Wide Area Engine
                -NONE-
        Routers Notified of from other WAE's
                -NONE-
        Multicast Addresses Configured
                -NONE-

Router Information for Service: TCP Promiscuous 62
        Routers Seeing this Wide Area Engine(1)
                Router Id        Sent To        Recv ID         AssKeyIP    AssKeyCN
MemberCN
                192.168.200.1   192.168.5.1     003F856D        10.86.46.74     1       12
```

[2] See "Configuring a GRE Tunnel Interface on a Router" in the Cisco WAAS Configuration Guide for detail on configuring the GRE tunnel interface on the router.

```
        Routers not Seeing this Wide Area Engine
                -NONE-
        Routers Notified of from other WAE's
                -NONE-
        Multicast Addresses Configured
                -NONE-
WAE#
```

On the router, WCCP must also be enabled for each service group (61 and 62) and redirection assigned to all necessary interfaces. The minimal CLI required to accomplish this and to show the current WCCP status is:

```
Rtr#show run | beg ip wccp
ip wccp 61
ip wccp 62
!
ip cef
!
interface GigabitEthernet0/1
 no ip address
 duplex full
 speed auto
!
interface GigabitEthernet0/1.2
 description WAN side
 encapsulation dot1Q 2
 ip address 192.168.2.2 255.255.255.0
 ip wccp 62 redirect in
!
interface GigabitEthernet0/1.3
 description Internal LAN side
 encapsulation dot1Q 3
 ip address 192.168.200.1 255.255.255.0
 ip wccp 61 redirect in
!
… ! many lines deleted
Rtr#
Rtr# show ip wccp
Global WCCP information:
    Router information:
        Router Identifier:                   192.168.200.1
        Protocol Version:                    2.0

    Service Identifier: 61
        Number of Service Group Clients:     1
        Number of Service Group Routers:     1
        Total Packets s/w Redirected:        3417705
          Process:                           3
          Fast:                              0
          CEF:                               3417702
        Service mode:                        Open
        Service access-list:                 -none-
        Total Packets Dropped Closed:        0
        Redirect access-list:                -none-
        Total Packets Denied Redirect:       0
        Total Packets Unassigned:            55
```

```
        Group access-list:                    -none-
        Total Messages Denied to Group:       0
        Total Authentication failures:        0
        Total Bypassed Packets Received:      0

    Service Identifier: 62
        Number of Service Group Clients:      1
        Number of Service Group Routers:      1
        Total Packets s/w Redirected:         4926706
          Process:                            28
          Fast:                               0
          CEF:                                4926678
        Service mode:                         Open
        Service access-list:                  -none-
        Total Packets Dropped Closed:         0
        Redirect access-list:                 -none-
        Total Packets Denied Redirect:        0
        Total Packets Unassigned:             51
        Group access-list:                    -none-
        Total Messages Denied to Group:       0
        Total Authentication failures:        0
        Total Bypassed Packets Received:      0

Rtr#
Rtr#
```

WCCP Router Lists and Router-IDs

WAAS and WCCP support up to 32 routers defined as part of a WCCP service group. Up to 7 router ids can be specified on a single "wccp router-list" global config command line. Multiple "wccp router-list" commands using the same router-list number can be entered to extend the list up to a maximum of 32 routers. The router-list is assigned to the tcp-promiscuous service groups with the "wccp tcp-promiscuous router-list-num" command. Regardless of the router IP address entered in the "wccp router-list" command, the router ID listed in the "show wccp routers" EXEC command will be the address of the routers first loopback interface or highest active physical interface.

WAAS Assignment Methods

Hash Assignment

Load balancing within a WCCP service group uses Hash Assignment by default. A hash function is a mathematical function applied to a 32-bit integer (an IP address) and results in an integer value with uniform distribution over a specified range of values. The WCCP hashing algorithm is applied by the router/switch to either the packet source IP address (for Service Group 61) or to the packet destination IP address (for Service Group 62) and results in a value in the range of zero thru 255 representing a total of 256 "address buckets". The packet is then forwarded to the WAAS device assigned to service that address bucket. When a WAAS device joins a WCCP service group, it is allocated a portion of the address buckets. If there is only one WAAS device, then all buckets are assigned to that one WAAS device. When multiple WAAS devices join the service group, the address buckets are divided among them. By default, each WAAS device is allocated an equal proportion of address buckets. An unequal division can be accomplished specifying a "weight" factor. See Load Balancing Weight Assignment for more information.

```
WAAS-Core# show run | inc wccp
wccp router-list 1 192.168.101.1
wccp tcp-promiscuous router-list-num 1
```

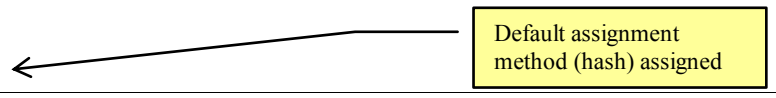

```
wccp version 2
WAAS-Core# show wccp flows tcp-promiscuous summary

Flow summary for service: TCP Promiscuous 61

Total Buckets

OURS = 128

  0- 59: .......... .......... .......... .......... .......... ..........
 60-119: .......... .......... .......... .......... .......... ..........
120-179: ........OO OOOOOOOOOO OOOOOOOOOO OOOOOOOOOO OOOOOOOOOO OOOOOOOOOO
180-239: OOOOOOOOOO OOOOOOOOOO OOOOOOOOOO OOOOOOOOOO OOOOOOOOOO OOOOOOOOOO
240-255: OOOOOOOOOO OOOOOO

 (… more )
```

256 hash buckets available – half assigned to (O = Owned) this cache engine

Mask Assignment

Mask assignment is an alternative method for distributing the WCCP traffic load among members of a WCCP service group. With this method, WCCP concatenates the source IP address, destination IP address, source port number, and destination port number into a 96-bit value (12 bytes). Each CE provides a 96-bit selection mask and 96-bit comparison value to the router. WCCP performs a bit-wise AND of the mask and packet data, then selects the CE whose value matches the result.

Figure 4 WCCP Mask Assignment

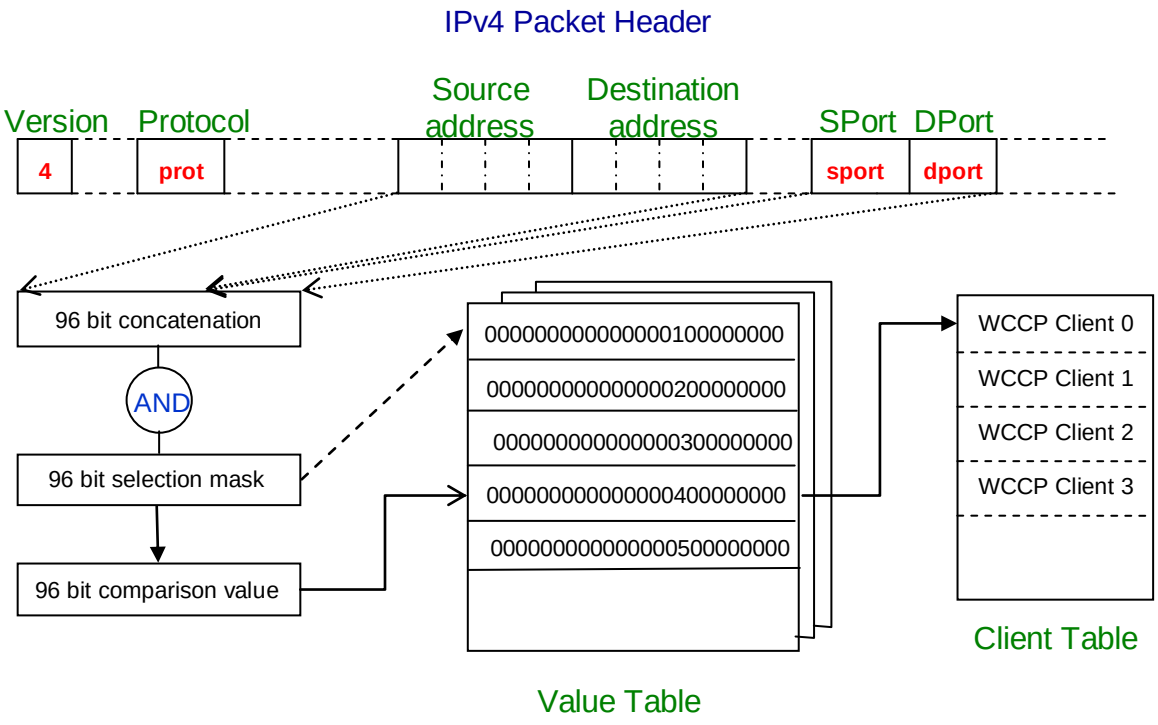

For WAAS, a single mask value is assigned by the WAAS device to the tcp-promiscuous WCCP service using the config command "wccp tcp-promiscuous mask" command. The mask value can be specified with up to 7 bits total and these bits can be assigned to either the source IP address or to the destination IP address. While the WAAS Central Manager GUI and device CLI allows a mask assignment to include both source and destination IP address as defined and allowed by the WCCP specification, redirection in WAAS is typically based on either the source or destination IP address alone. The number of bits specified determines the number of address buckets created for the assignment pool.

Bits	Buckets (2^n)	Maximum WAAS devices
1	2	2

2	4	4
3	8	8
4	16	16
5	32	32
6	64	32
7	128	32

When assigning a mask value to the tcp-promiscuous service, that mask value is actually assigned to the WCCP service group 61. The specified mask assignment is reversed for the WCCP service group 62. For instance, with the CLI config command "`wccp tcp-promiscuous mask src-ip-mask 0x7`" command, the 96 bit mask assigned to the WCCP service group 61 would specify a source IP address mask of 0x7 and service group 62 would specify a destination IP address mask of 0x7. The following CLI listing illustrates this concept using a src-ip-mask of 0x7:

```
WAAS-Core# show run | inc wccp
wccp router-list 1 192.168.101.1
wccp tcp-promiscuous mask src-ip-mask 0x7 dst-ip-mask 0x0
wccp tcp-promiscuous router-list-num 1 weight 25 mask-assign
wccp version 2
WAAS-Core# show wccp masks tcp-promiscuous


Service-id is : 61
        Mask  SrcAddr    DstAddr    SrcPort DstPort
        ----  -------    -------    ------- -------
        0000: 0x00000007 0x00000000 0x0000  0x0000

        Value SrcAddr    DstAddr    SrcPort DstPort CE-IP
        ----- -------    -------    ------- ------- -----
        0000: 0x00000000 0x00000000 0x0000  0x0000  0x00000000 (Not assigned)
        0001: 0x00000001 0x00000000 0x0000  0x0000  0x00000000 (Not assigned)
        0002: 0x00000002 0x00000000 0x0000  0x0000  0x00000000 (Not assigned)
        0003: 0x00000003 0x00000000 0x0000  0x0000  0x00000000 (Not assigned)
        0004: 0x00000004 0x00000000 0x0000  0x0000  0x00000000 (Not assigned)
        0005: 0x00000005 0x00000000 0x0000  0x0000  0x00000000 (Not assigned)
        0006: 0x00000006 0x00000000 0x0000  0x0000  0xC0A86503 (192.168.101.3)
        0007: 0x00000007 0x00000000 0x0000  0x0000  0xC0A86503 (192.168.101.3)


Service-id is : 62
        Mask  SrcAddr    DstAddr    SrcPort DstPort
        ----  -------    -------    ------- -------
        0000: 0x00000000 0x00000007 0x0000  0x0000

        Value SrcAddr    DstAddr    SrcPort DstPort CE-IP
        ----- -------    -------    ------- ------- -----
        0000: 0x00000000 0x00000000 0x0000  0x0000  0x00000000 (Not assigned)
        0001: 0x00000000 0x00000001 0x0000  0x0000  0x00000000 (Not assigned)
        0002: 0x00000000 0x00000002 0x0000  0x0000  0x00000000 (Not assigned)
```

```
        0003: 0x00000000 0x00000003 0x0000  0x0000  0x00000000 (Not assigned)
        0004: 0x00000000 0x00000004 0x0000  0x0000  0x00000000 (Not assigned)
        0005: 0x00000000 0x00000005 0x0000  0x0000  0x00000000 (Not assigned)
        0006: 0x00000000 0x00000006 0x0000  0x0000  0xC0A86503 (192.168.101.3)
        0007: 0x00000000 0x00000007 0x0000  0x0000  0xC0A86503 (192.168.101.3)
WAAS-Core#
```

Note in the above listing that, in addition to specifying mask assignment, this WAAS device has been assigned a weight of 25 and no other WAAS device has joined or is available for this service group. This results in only two of the available 8 address buckets being assigned to a WAAS device. *Any IP address whose mask result falls into one of the other 6 buckets (Not assigned) will not be redirected and will be routed normally.* This may be intended behavior or not, depending on the capacity of the WAAS device and the failover requirements. See the section "Load Balancing Weight Assignment" for more information on load balancing and failover scenarios.

Choosing a Mask

Choosing the correct mask for WCCP Mask Assignment can be a challenging task and there are no easy answers on how to construct a mask. An in-depth knowledge of the network structure and possible IP subnet addresses that flow thru the router/switch is required. An understanding of decimal to hexadecimal conversion and logical binary operations is also needed to test the defined mask against a possible IP address.

Minimally, the number of bits used in the mask must provide enough buckets to be apportioned to each WAAS device assigned to the service group, taking into account the load balancing weight assigned to each device (see the next section "Load Balancing Weight Assignment"). A 1 bit mask can support only 2 WAAS devices (2^1 = 2) while a 5 bit (or more) mask can support 32 WAAS devices (2^5 = 32), the maximum number allowed in a service group.

An IP address consists of 4 tuples when specified in dot notation. When written in hexadecimal notation, each tuple consists of two hexadecimal digits in the range of 0 thru F. For example, the IP address 192.168.101.3 can be written hexadecimal dot notation as C0.A8.65.03 and manipulated as a single 4 byte entity of 0xC0A86503. This example address is typical of one assigned by DHCP in a branch office. The WCCP mask is also typically specified in hexadecimal with leading zeros omitted.

The default source IP address mask, if one is not explicitly defined, is 0x1741 (the default destination mask is 0x0). While this is a default, it is not a recommended mask assignment and will be unlikely to provide adequate load balancing in most cases. Note that if there is only one available WAAS device, the mask used is irrelevant and the default mask can be used.

	Hex	Binary
IP Address 192.168.101.3	C0.A8.65.03	11000000.10101000.01100101.00000011
Default Mask 0x1741	00.00.17.41	00000000.00000000.00010111.01000001
Binary AND		00000000.00000000.00000101.00000001

In a branch office, source IP addresses are typically assigned sequentially by DHCP. This results in addresses that vary only in the least significant bits. Mask distribution can be maximized here by specifying a WCCP mask in the range of 0x1 thru 0x7F depending on the number of WAAS devices.

In a regional data center or network hub, many lower level subnets are switched and routed. The mask assigned here should take into account the subnet specification so that entire subnets, and thus branch offices/branch WAAS devices, are redirected and peered with one or a small number of WAAS devices. This logical dispersion assumes that employees within a single subnet typically share work responsibilities and are more likely to take advantage of the redundancy elimination of a common DRE cache. Typically, in a large enterprise network address allocation scheme, these subnets vary in the third tuple of the IP address quadruple (/24 address mask) so the WCCP mask can be specified between 0xF00 thru 0x7F00.

Following in this logic, the mask used for a main enterprise Data Center would seek to redirect and consolidate based on regional address groupings typically specified in the second tuple of the IP address quadruple (/16 address mask). Data Center WCCP masks are typically in the range of 0xF0000 and 0x7F0000. For a large scale environment where core DRE size becomes a concern, higher order octets in the IP address would be used to load balance entire sites rather than hosts into the WAAS core to minimize the number of peers.

Load Balancing Weight Assignment

For deployments where the WAAS devices belonging to a service group are not the same hardware model, an even distribution of redirected packets may overwhelm the smaller WAAS devices in the service group. In this case, a load balancing "weight" can be assigned to each device in the service group to allocate address buckets according to the maximum number of TCP connections a device can handle (see the WAAS Sizing Guidelines for TCP connection maximums).

For hash assignment, the assigned weight factors for all WAAS devices in the service group can be displayed with the "show wccp wide-area-engine" status command:

```
WAAS-Core1# config
WAAS-Core1(config)# wccp tcp-promiscuous router-list-num 8 weight 25
WCCP configuration for TCP Promiscuous service 61 and 62 succeeded. WCCP configuration for
TCP Promiscuous succeeded.Please remember to configure WCCP service 61 and 62 on the
corresponding router.
WAAS-Core1(config)# end
WAAS-Core1# show wccp wide-area-engine

Wide Area Engine List for Service: TCP Promiscuous 61

Last Received Assignment Key IP address: 192.168.101.3
Last Received Assignment Key Change Number: 3
Last WAE Change Number: 3
Asssignment Made Flag = FALSE

        IP address = 192.168.101.3      Lead WAE = YES  Weight = 25
        Routers seeing this Wide Area Engine(1)
                192.168.201.1

        IP address = 192.168.101.4      Lead WAE = NO   Weight = 75
        Routers seeing this Wide Area Engine(1)
                192.168.201.1

Wide Area Engine List for Service: TCP Promiscuous 62

Last Received Assignment Key IP address: 192.168.101.3
Last Received Assignment Key Change Number: 3
Last WAE Change Number: 3
Asssignment Made Flag = FALSE

        IP address = 192.168.101.3      Lead WAE = YES  Weight = 25
        Routers seeing this Wide Area Engine(1)
```

25% of traffic allocated to this WAE

75% of traffic allocated to this WAE

```
                192.168.201.1

        IP address = 192.168.101.4      Lead WAE = NO   Weight = 75
        Routers seeing this Wide Area Engine(1)
                192.168.201.1
WAAS-Core1#
```

For mask assignment, the weight value for the "show wccp wide-area-engine" status command displays zero (a CDET has been filed on this - CSCta11645). However, the weight distribution can be inferred from the "show wccp flows tcp-promiscuous summary" command:

```
WAAS-Core#sho run | include wccp
wccp router-list 1 192.168.101.1
wccp tcp-promiscuous mask src-ip-mask 0xf dst-ip-mask 0x0
wccp tcp-promiscuous router-list-num 1 weight 250 mask-assign
wccp version 2
WAAS-Core#show wccp wide-area-engine

Wide Area Engine List for Service: TCP Promiscuous 61

Last Received Assignment Key IP address: 192.168.101.3
Last Received Assignment Key Change Number: 1
Last WAE Change Number: 3
Asssignment Made Flag = FALSE

        IP address = 192.168.101.3      Lead WAE = YES  Weight = 0
        Routers seeing this Wide Area Engine(1)
                192.168.201.1

        IP address = 192.168.101.4      Lead WAE = NO   Weight = 0
        Routers seeing this Wide Area Engine(1)
                192.168.201.1

Wide Area Engine List for Service: TCP Promiscuous 62

Last Received Assignment Key IP address: 192.168.101.3
Last Received Assignment Key Change Number: 1
Last WAE Change Number: 3
Asssignment Made Flag = FALSE

        IP address = 192.168.101.3      Lead WAE = YES  Weight = 0
        Routers seeing this Wide Area Engine(1)
                192.168.201.1

        IP address = 192.168.101.4      Lead WAE = NO   Weight = 0
        Routers seeing this Wide Area Engine(1)
                192.168.201.1
WAAS-Core# show wccp flows tcp-promiscuous summary

Flow summary for service: TCP Promiscuous 61

Total Buckets
OURS = 4
  0- 15: .......... ..OOOO
BYP  = 0
```

4 bit mask used results in 16 buckets

25% of buckets (250 out of a total of 1000) allocated to this WAE equals 4 of the 16 buckets

16 buckets allocated (4 bit mask) of which 4 are "**O**urs" (250 or 25% weight)

```
  0- 15: .......... ......
AWAY = 12
  0- 15: AAAAAAAAAA AA....
IN   = 0
  0- 15: .......... ......
LAST_IN   = 0
  0- 15: .......... ......
Flow summary for service: TCP Promiscuous 62

Total Buckets
OURS = 4
  0- 15: .......... ..OOOO
BYP  = 0
  0- 15: .......... ......
AWAY = 12
  0- 15: AAAAAAAAAA AA....
IN   = 0
  0- 15: .......... ......
LAST_IN   = 0
  0- 15: .......... ......
WAAS-Core#
```

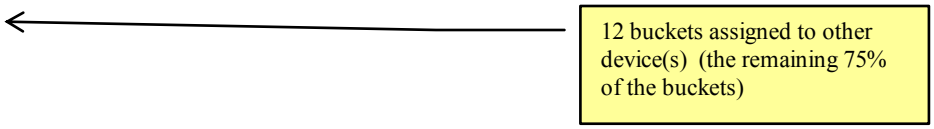

The CLI listing above shows a fully subscribed service group where all available buckets are assigned to a specific WAAS device. As noted in a previous CLI sample on page 17, any bucket that is not assigned to a specific WAAS device will be switched or routed without redirection and subsequent processing. This may be desired behavior as when two different model WAAS devices are deployed within the same WCCP service group, and the smaller device or devices are not capable of handling the entire traffic load. Should a larger device fail, traffic that would normally be redirect to it will, in effect, be bypassed.

This specific behavior is the result of using weight factors for all WAAS devices that sum to less than or equal to 100. To avoid this scenario and guarantee failover coverage, the sum of all weight factors must exceed 100 *at all times*. Should a device within a device group fail and the sum of the weights of the remaining devices be less than 100, there can be buckets for which there is no assigned WAAS device.

Best Practice: When using WCCP weight factors, to guarantee complete WCCP failover coverage, use weight factors for individual devices that are greater than 100.

When the sum of all weight factors is greater than 100, the specific percentage of buckets assigned to a specific WAAS device is the weight assigned to that WAAS device divided by the total weight and *rounded up*. Rounding up guarantees that each WAAS device will be assigned at least one bucket.

Best Practice: When using WCCP weight factors, use the "show wccp flow tcp-promiscuous summary" CLI command to verify that all available buckets have been assigned to a WAAS device. You will need to execute this command on all WAAS devices in the service group to see the buckets assigned to each.

Speed Bump: When using WCCP weight factors with values less than 100, verify that the sum of all weight factors is greater than or equal to 100. Understand that should one or more WAAS devices fail, some TCP traffic will not be redirected and thus will not be optimized.

WAAS Egress Method

Network Path Affinity

Multiple routers or multiple network paths pose challenges in Cisco WAAS deployments when using IP forwarding for Cisco WAAS egress. With multiple equal cost paths or redundant router designs, the Cisco WAE might send egress packets to a router other than the one that was used for interception and redirection. Figure 5 depicts a scenario where all WCCP defaults are taken such that traffic is intercepted on one router (1) and redirected to the WAAS device, but is returned to a different router (2) because it is specified as the WAAS device's default gateway.

Figure 5 Network Path Affinity Issue

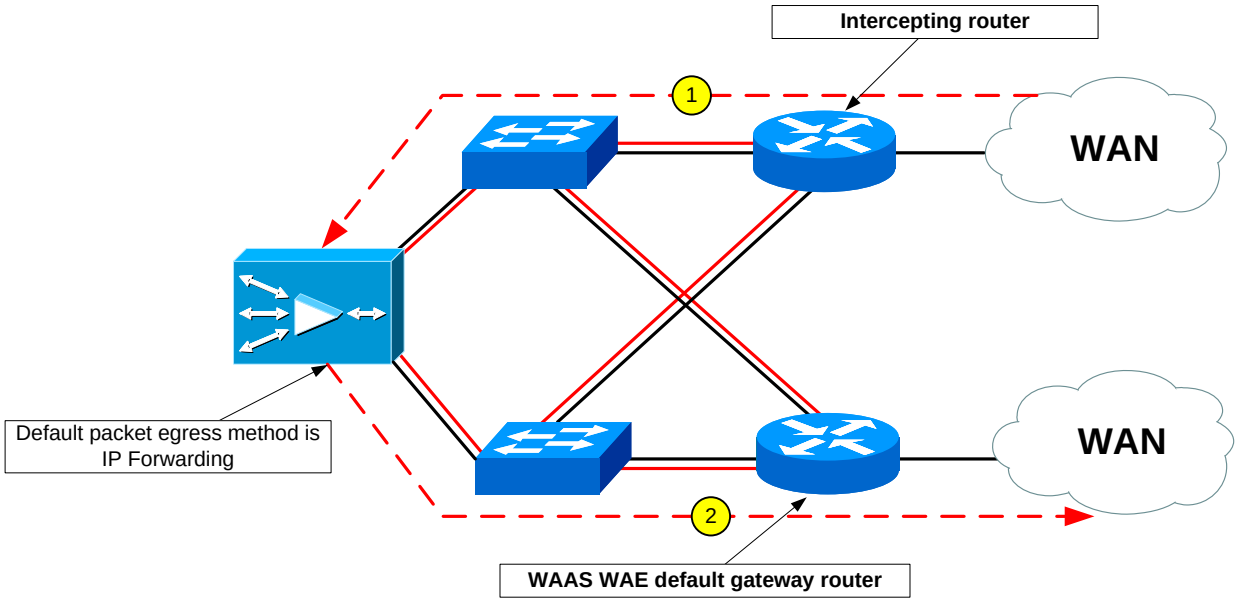

If the Cisco WAAS WAE is not L2 adjacent to the intercepting router, or if the intercepting router is not the same as the Cisco WAE default gateway, then by default you cannot ensure that packets are returned to the intercepting router. This also means that, in Gateway Load Balancing Protocol (GLBP) router situations or HSRP or VRRP redundancy situations, the Cisco WAE cannot reliably send packets back to the same router that intercepted the connection. In these cases, it is important to specify Generic GRE Return as the packet egress method.

Limiting WAAS Redirected Traffic

In general, it is desirable to limit the amount of traffic processed by a WAAS device to that traffic which will benefit from the optimizations offered by WAAS. There are three methods of excluding undesirable traffic from being processed by WAAS, Router Redirect Lists, WAAS Redirect Lists, and WAAS Policy Bypass.

Router Redirect Lists

Cisco WAAS tcp-promiscuous redirection selection on a WCCPv2 enabled router/switch is based on the protocol specified in the IP packet header. All TCP packets processed by a properly configured WCCP enabled interface will, by default, be redirected to the appropriate WAAS device for processing. Further limitation of the selection process can be configured using a WCCPv2 redirect list consisting of a standard or extended IOS access list associated with a specific WCCPv2 service group. Once applied to a service, traffic passing through an interface with redirection applied must match not only the default protocol specified for the service group (i.e. TCP), but also one of the permit statements defined by the access list. Failure to match a permit statement will result in the packet being routed normally. The following router IOS commands create an extended access list that when applied as a WCCPv2 service group, redirects only TCP traffic to or from the Class A network 10.0.0.0. Note that the redirect list must be applied in both directions.

```
…
ip wccp 61 redirect-list 100
ip wccp 62 redirect-list 100
!
ip cef
…
!
access-list 100 permit ip 10.0.0.0 0.255.255.255 any
access-list 100 permit ip any 10.0.0.0 0.255.255.255
access-list 100 deny   ip any any
!
…
```

Speed Bump: Avoid specifying tcp ports as part of an ACL. This can cause severe issues if the network configuration results in any packet fragmentation.

Speed Bump: Router redirect list ACLs must be applied to *both* service groups, 61 and 62, in order to maintain packet flow consistency. Failure to do so will cause auto-discovery to fail and the connections will be bypassed.

WAAS Static Bypass Lists

In addition to router redirect lists, static bypass lists can be defined for a device or a device group and used on a WAAS device to accomplish a similar function. When using inline interception, WAAS Static Bypass Lists are the only option and can be used in place of router redirect lists. However, whenever possible, router redirect lists are recommended because they more efficiently move packets through the router and avoid unnecessary processing.

The global config command to establish a WAAS bypass list is "bypass static". This WAAS command is limited to selecting an individual client and/or destination server. Wildcards are not permitted.

WAAS Policy Pass-through

The last resort bypass mechanism is the pass-through action built into an Application Policy definition. After identifying "uninteresting" traffic with a defined Application Classifier, that traffic can be bypassed and returned for routing via the WCCP packet return method without having any optimization processing performed. Bypassing large volumes of traffic in this manner incurs significant overhead both on the router and WAAS device.

Best Practice: Apply WCCPv2 redirect list ACLs on the router/switch wherever possible to prevent unnecessary processing and packet routing on the router/switch.

SIMPLE DEPLOYMENT SCENARIO

The simplest WAAS WCCP deployment scenario involves a single Cisco WAAS device configured with a single router. Figure 6 shows a remote office WAAS WAE peerred with a single Data Center Cisco WAAS WAE.

Figure 6 - Simple "One-Arm" Deployment

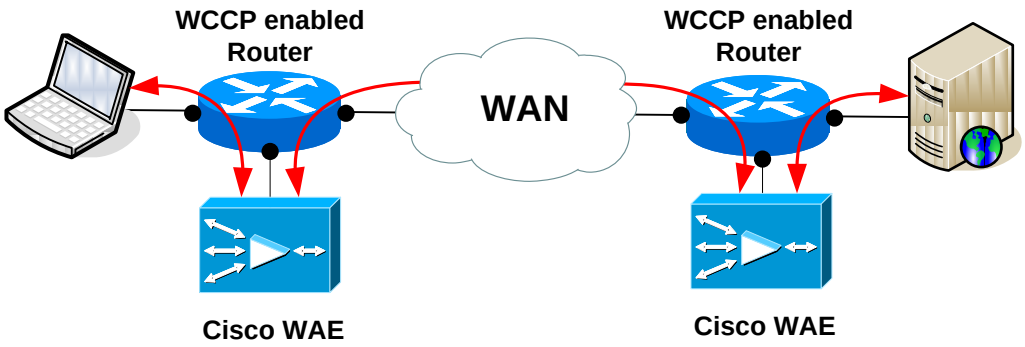

Figure 7 shows the WAAS Central Manager GUI where a WAAS device has been selected and **WCCP Settings** has been selected from the **Configure** menu on the left. To identify the router/switches that will be part of this service group, click the **New Router List** button.

Figure 7. Enabling WCCP Interception

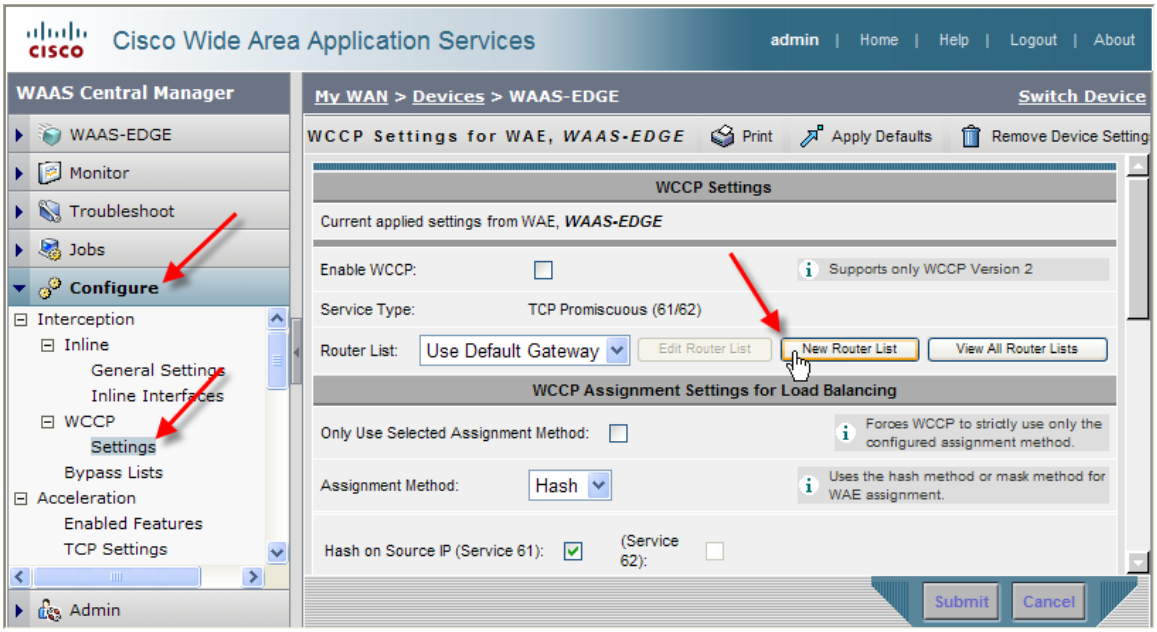

Figure 8 shows the GUI for creating a service group router list. In this case, we are creating router-list-number 1 with two routers defined, 192.168.5.1 and 192.168.6.1. Press **Add Router>>>>** to add the entered IP address as a router in this list and then **Submit** to create the router list.

Figure 8 Creating a WCCP Router List

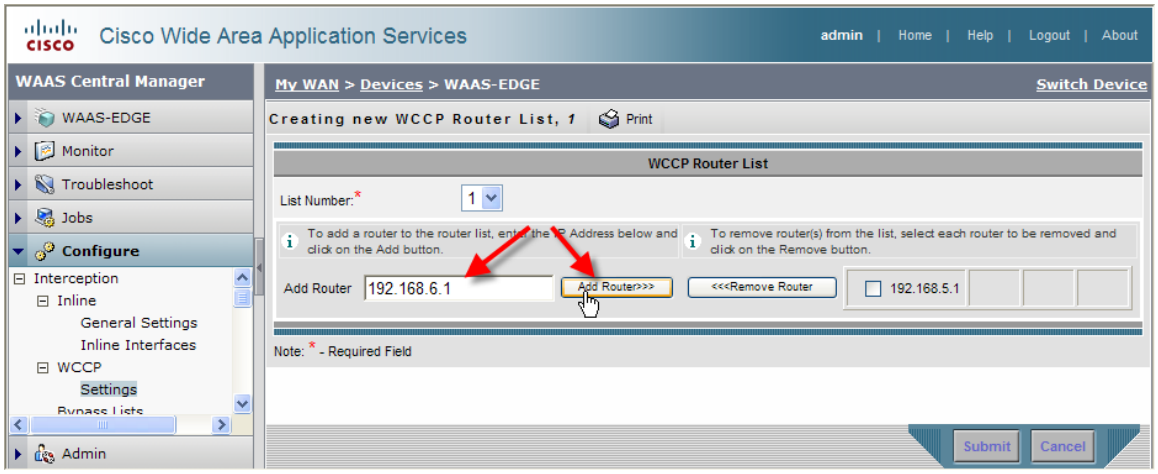

After creating the router list, the WAAS GUI will return to the WCCP Settings display as shown in Figure 9. Check the **Enable WCCP** checkbox and select **Router List** 1 from the drop down menu. Click **Submit** to complete the WCCP service group setup process.

Figure 9 Enable WCCPv2 and select Router List

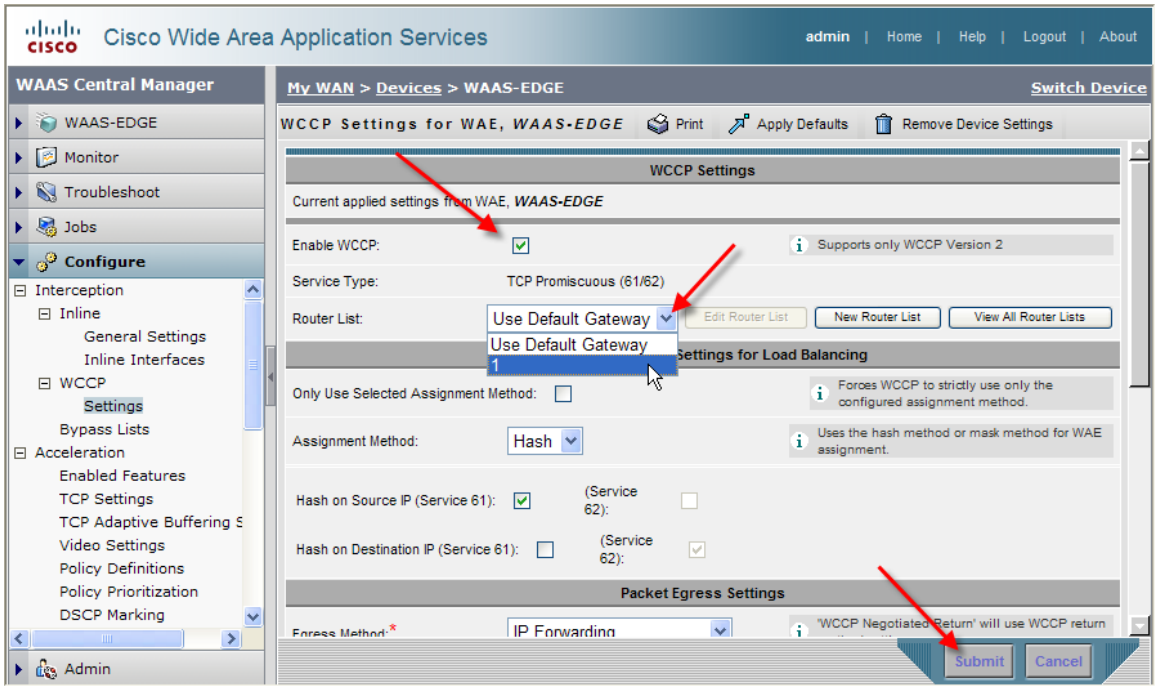

The following WAAS CLI commands accomplish the same results.

```
!
wccp router-list 1 192.168.5.1 192.168.6.1
wccp tcp-promiscuous router-list-num 1
wccp version 2
!
```

Note that additional options available for the WCCP Settings GUI and the "wccp tcp-promiscuous router-list-num" CLI command provide for selecting forwarding, assignment and return methods, and other advanced parameters. These options will be discussed in the section titled "Advanced Deployment Scenarios".

Once WCCPv2 has been enabled, the WAAS CLI command "show wccp services detail" can provide the configured details for the tcp-promiscuous services. The following shows the detail associated with the default settings assigned to the tcp-promiscuous services:

```
WAE# show wccp services detail

Service Details for TCP Promiscuous 61 Service
        Service Enabled                    : Yes
        Service Priority                   : 34
        Service Protocol                   : 6
        Service Flags (in Hex)             : 501
        Service Ports                      :      0     0     0     0
                                           :      0     0     0     0
        Security Enabled for Service       : No
        Multicast Enabled for Service      : No
        Weight for this Web-CE             : 0
        Negotiated forwarding method       : GRE
        Negotiated assignment method       : HASH
        Negotiated return method           : GRE
        Received Values:
        Source IP mask (in Hex)            : 0
        Destination IP mask (in Hex)       : 0
        Source Port mask (in Hex)          : 0
        Destination Port mask (in Hex)     : 0
        Calculated Values:
        Source IP mask (in Hex)            : 1741
        Destination IP mask (in Hex)       : 0
        Source Port mask (in Hex)          : 0
        Destination Port mask (in Hex)     : 0

Service Details for TCP Promiscuous 62 Service
        Service Enabled                    : Yes
        Service Priority                   : 34
        Service Protocol                   : 6
        Application                        : Unknown
        Service Flags (in Hex)             : 502
        Service Ports                      :      0     0     0     0
                                           :      0     0     0     0
        Security Enabled for Service       : No
        Multicast Enabled for Service      : No
        Weight for this Web-CE             : 0
        Negotiated forwarding method       : GRE
        Negotiated assignment method       : HASH
        Negotiated return method           : GRE
        Received Values:
        Source IP mask (in Hex)            : 0
        Destination IP mask (in Hex)       : 0
        Source Port mask (in Hex)          : 0
        Destination Port mask (in Hex)     : 0
        Calculated Values:
        Source IP mask (in Hex)            : 0
        Destination IP mask (in Hex)       : 1741
        Source Port mask (in Hex)          : 0
```

```
        Destination Port mask (in Hex)      : 0
```

Note that both services are identical with the exception of the assignment method hash key. Service 61 uses the source IP as the hash key while service 62 uses the destination IP as the hash key. Service 61 intercepted traffic will be distributed across all available WAAS devices based on the source of the traffic, and service 62 traffic is distributed based on the destination. Additional configuration beyond the default settings may be required depending on the deployment scenario. See Advanced Deployment Scenarios in this document for more information.

On each router/switch configured to be part of a WCCPv2 service group, WCCPv2 must also be enabled for each service group and redirection assigned to participating interfaces. The following CLI log shows the minimal configuration required to set up the WAAS WCCP service groups and assign redirection to the router interfaces. Cisco Express Forwarding (cef) is enabled whenever possible to increase the overall performance of the routing within the device. Selection of the interface for service group placement is discussed in the next section.

```
…
ip wccp 61                                      ! enable wccp service group 61
ip wccp 62                                      ! enable wccp service group 62
!
!
ip cef                                          ! enable Cisco Express Forwarding
!
…
!
interface GigabitEthernet0/1
 no ip address
 duplex full
 speed auto
!
interface GigabitEthernet0/1.2
 description LAN client side
 encapsulation dot1Q 2
 ip address 192.168.200.1 255.255.255.0
 ip wccp 61 redirect in                         ! apply service group 61 redirection on
!                                               !   inbound traffic arriving on this
!                                               !   interface
interface GigabitEthernet0/1.3
 description WAN server side
 encapsulation dot1Q 3
 ip address 192.168.2.2 255.255.255.0
 ip wccp 62 redirect in                         ! apply service group 66 redirection on
!                                               !   inbound traffic arriving on this
!                                               !   interface
…
```

Best Practice: While CEF on the router is not required, it is strongly recommended for improved performance. WCCP can use IP CEF if CEF is enabled on the router.

Service Group Placement

Cisco WAAS auto-discovery and the establishment of an optimized connection requires that request and response packets be redirected and processed by the same WAAS device. In a multiple WAAS device deployment, the two service groups with opposing assignment methods required for the routing/switching devices allow response packets to be redirected to the same WAAS device that handled the

request packet, regardless of on which interface, or even on which routing/switch device participating in the service group the packets arrive. A request packet from a client with source IP x.x.x.x will be redirected to the same WAAS devices as the corresponding response packet back from the server with a destination IP of the client at x.x.x.x. However, care must be taken to ensure that the correct service is placed on *each and every interface of each router/switch* included in the service group.

WCCP packet redirection can be processed on any interface in either direction, both inbound (ingress) and outbound (egress). While multiple service groups can be defined on an interface in either direction, only one service group in each direction will be chosen for redirection based on the defined service group priority (see Appendix A -Cisco WCCP Service Groups for assigned service group priorities). *Packets returning from redirection are not candidates for further redirection on the same interface.* For instance, if both web-cache (service group 0 – protocol TCP, port 80, priority 240) and tcp-promiscuous service 61 (service group 61 – protocol TCP, port any, priority 34) are defined for inbound redirection on the same interface, all web traffic (TCP port 80) would be redirected to the web-cache service group cache engines because the web-cache service group has a higher priority, while all other TCP traffic would be redirected to the tcp-promiscuous WAAS devices. Packets returning from the web-cache service group cache engines will not be redirected again to the tcp-promiscuous WAAS device.

Best Practice: Typically, placing service group redirection on the in-bound interface wherever possible will prevent unnecessary processing on the router/switch. There are exceptions where out-bound redirection offers advantages that simplify configuration and actually may reduce load on the router such as the case where routing may be done between subnets off the router and the WAN interface may not be involved.

Best Practice: When WCCP GRE return is used as egress method, the behavior depends on the intercept method. For GRE intercept, all packets are egressed to the routerID. The routerID must be reachable, and preferably via the shortest path, which may require static route configuration. For L2 intercept, packets are egressed to the L2 adjacent address of the router (except for flow protection, where packets may get egressed to a routerID), and we require (but don't enforce) that the WAAS device's router list should contain the router's L2 adjacent addresses. With generic GRE return, we don't have the above issues with routerID reachability. Also, to ensure that the WCCP HIA/ISU exchange verifies the data path of egress packets with GRE return, we recommend configuring a loopback interface as routerID, as well as the WAAS device's routerlist IP address.

In most cases, Service Group 61 should be placed on the inbound interface closest to the client while Service Group 62 should be placed on the inbound interface closest to the requested server. Figure 10 depicts the simplest and recommended configuration for WAAS service group placement.

Figure 10. WCCP Recommended Service Group Placement

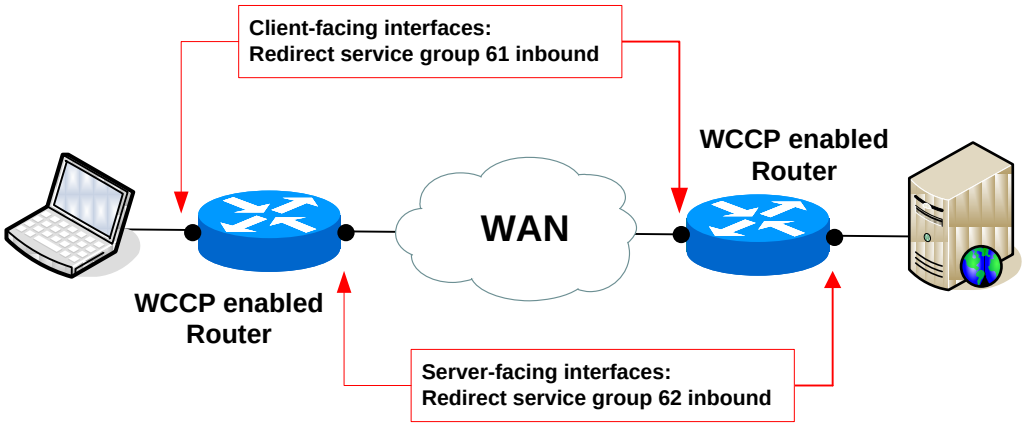

Alternative placement is possible when needed such as when deploying multiple WCCP enabled functions on a single device. An example would be when deploying a caching and streaming service such as ACNS along with WAAS[3]. Both optimization services require redirection of TCP packets and only one can be placed on an interface in any one direction. In cases such as this, one of the service groups must be placed on an outbound interface.

There are three options for service group redirection:

1. ip wccp redirect in
2. ip wccp redirect out
3. ip wccp exclude in

To enable redirection on a router interface, one of the first two options, redirect in or redirect out, should be specified as part of the router interface configuration. The third option, exclude in, is required on the router interface to which the WAAS devices are connected, and only if the wccp redirect out method is used on any of the other interfaces. The reason for this is that packets returning from a WAAS device are routed normally and if redirection were not specifically excluded, they would be redirect again on egress from the router causing a routing loop. Figure 11 diagrams the packet flow when using egress redirection and not excluding returning packets from further redirection.

1. Packets arrive from the source and are routed completely through the router arriving on the egress interface.
2. The egress interface is configured to redirect traffic to the service group WAAS devices.
3. The WAAS device processes the packets and returns them to the router for further routing.
4. Because WAAS is fully transparent, there is no way to distinguish a packet returning from the WAAS device from a packet arriving from the source, the egress interface again redirect the packets to the WAAS device causing a routing loop where the packet will eventually be dropped.

[3] See ACNS/WAAS Co-deployment using Transparent WCCP Redirection

Figure 11 WCCP Black Hole

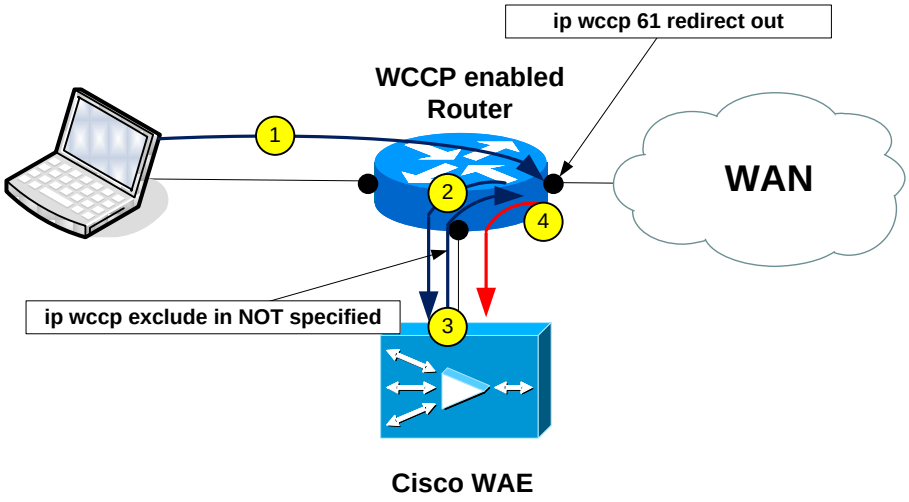

This is known as the WCCP "Black Hole" syndrome where all TCP traffic just "disappears." This can be particularly annoying when using a tcp service such as telnet to configure the router, as all tcp connections to the router will be lost to the black hole and the only way to recover the device will be with a console connection.

Best Practice: When using a one-arm or two-arm WCCP deployment, always place an "ip wccp exclude in" interface configuration statement on the router's WAAS subnet interface.

Speed Bump: Before placing any "ip wccp redirect out" CLI command on a router interface, always check to make sure you have already placed an "ip wccp exclude in" on the interface through which the WAAS WAEs will return traffic to the router. Failure to heed this may result in total loss of connectivity with the router!

ADVANCED DEPLOYMENT SCENARIOS

Advanced deployment scenarios can involved a myriad of configurations involving multiple router/switches and multiple WAAS devices belonging to a single WCCP service group. It can also involve various alternatives and combinations of alternative redirection, assignment, packet return and packet egress methods. A common data center deployment involves multiple redundant switches redirecting traffic to multiple WAAS WAEs supporting high throughput and high availability. Typically this kind of deployment would involve two or more high performance switches such as the Cisco Catalyst 6500 switches redirecting traffic to a farm of high-end WAAS WAE devices as depicted in Figure 12.

Figure 12 High Throughput and High Availability Data Center Deployment

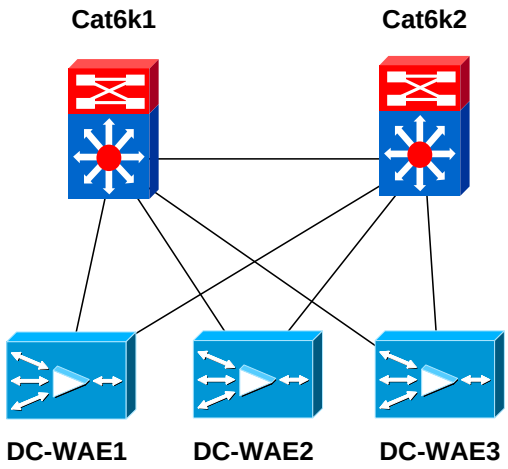

An excellent source of a validated advanced deployment scenario has been developed as part of the Cisco Safe Harbor Initiative. The document, titled "WCCP Network Integration Validated Design and Best Practices", can be accessed internally from the Safe Harbor Results Documentation site under the subheading of WCCP. A customer facing version of the document is also available for download from that page, or it can be requested from your Cisco representative by asking for document EDCS-773004 Phase 1 Catalyst 6500: Ingress Layer 2 Redirection and Mask Assignment.

The Safe Harbor document is an initiative to provide recommended WCCP configuration and operational best practices for multiple platforms while identifying potential unknown issues. This project uses an end to end architecture based verification which consists of simulated common Application, Data-center, Campus, Branch and WAN environments. WCCP stability is verified while undergoing commonly used operational tasks like configuration modifications, hardware, software upgrades, failover, and other negative scenarios under moderate system level stress. The system level stress is a representative mix of applications that invoke most of the component features and functions. This document identifies issues and provides recommendations for those deploying products that leverage WCCP including but not limited to Wide Area Application Services (WAAS), Application and Content Networking System (ACNS), Ironport Web Security Appliance (WSA), and other Cisco or third party WCCP client devices.

L2 Forwarding and Return

While WCCP GRE forwarding is the default method of redirecting packets to a WAAS device, L2 forwarding is a more efficient method and should be preferred whenever supported by the switch/router. Because packet switching is done at Layer 2 by rewriting the destination MAC address, the WAAS devices must be directly connected to the switch/router. Likewise, WCCP GRE is the default packet return method while L2 Return is more efficient. Both L2 Forwarding and L2 Return can be configured by assigning the features using the "wccp tcp-promiscuous router-list-num" global config CLI command.

```
WAAS-Core(config)# wccp tcp-promiscuous router-list-num 1 l2-redirect l2-return
WCCP configuration for TCP Promiscuous service 61 and 62 succeeded. WCCP configuration for
TCP Promiscuous succeeded.Please remember to configure WCCP service 61 and 62 on the
corresponding router.
WAAS-Core(config)# exit
WAAS-Core#sho run | inc wccp
wccp router-list 1 192.168.101.1
wccp tcp-promiscuous mask src-ip-mask 0x7f dst-ip-mask 0x0
wccp tcp-promiscuous router-list-num 1 l2-redirect l2-return
wccp version 2
```

```
WAAS-Core#
```

Specifying l2-redirect as the forwarding and/or return method has no effect if the switch/router does not support that method.

SpeedBump	**Layer 2 redirection requires that content engines be directly connected to an interface on each WCCP router. Unless multicast IP addresses are used, WCCP configuration of the content engine must reference the IP address of each WCCP router that is directly connected. The configuration should not reference a loopback IP address or an IP address that is configured on the WCCP router.**

Note that while L2 return is supported and preferred by WAAS as the packet return method for packets not selected for optimization, WAAS does not support L2 rewrite as the packet egress method for pass-through packets and optimized packets. If the egress packet return method configured is WCCP Negotiated Return, and WCCP negotiates a WCCP L2 Rewrite return method, the WAAS device will default to IP Forwarding as the egress method.

Avoiding WCCP-induced Routing Loops

WAAS WCCP deployments involving multiple routing paths between a remote client and a server are prone to routing loops as traffic is redirected to the WAAS WAEs. Figure 13 shows a typical scenario involving multiple Catalyst 6500 switches doing ingress WCCP redirection (properly) on the WAN and LAN interfaces. A routing loop can be introduced when redirection and return are not from and to the same routing device. There are a few simple solutions to avoid the problem. The easiest method is to provide an additional (preferred) path directly between the switches involved on which WCCP redirection does not occur. Another less preferable option is to perform both ingress and egress redirection (service groups 61 and 62) on the WAN interface. However, some platforms do not support egress redirection. See the Platform Variations section for platform limitations.

Figure 13 WCCP-induced Routing Loop

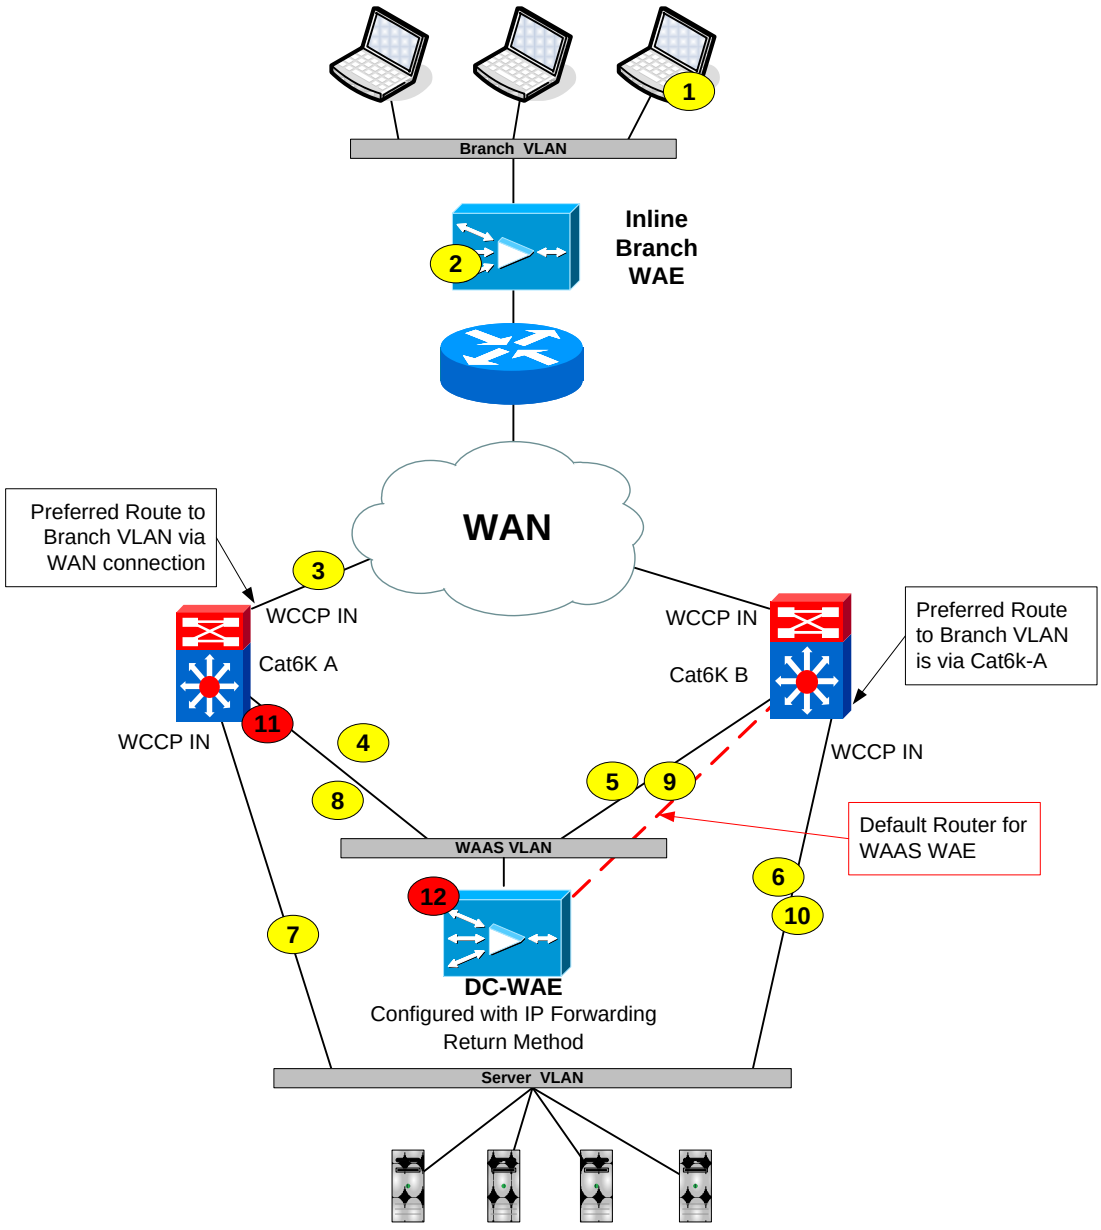

Step	Element	Function
1	Client Laptop	Client initiates connection to server in a data center.
2	Branch WAE	WAAS WAE marks TCP Syn packet with TCP Options.
3	WAN	TCP packet routed to either Cat6k-A or Cat6k-B (assumes equal cost routing).
4	Cat6k-A	WCCP running on the Cat6k-A redirects TCP traffic incoming from the WAN to the WAAS device.
5	DC WAE	The DC WAE is configured with IP Forwarding as the WCCP packet return method and returns the processed packet to its default router, Cat6k-B.
6	Cat6k-B	Traffic arriving from the WAAS VLAN is routed without redirection to the requested server.
7	Server	The server responds via its defined default gateway, Cat6k-A.
8	Cat6k-A	WCCP running on the Cat6k-A redirects TCP traffic incoming from the Server LAN to the WAAS device.
9	DC WAE	The DC WAE marks TCP Syn packet with TCP Options. The WAE is configured with IP Forwarding as the WCCP packet return method and returns the processed packet to its default router, Cat6k-B.
10	Cat6k-B	The Cat6k does a route lookup for the preferred route to the Client VLAN and routes the packet to the Cat6k-A router via the Server VLAN.
11	Cat6k-A (loop)	WCCP running on the Cat6k-A redirects TCP traffic incoming from the Server LAN to the WAAS device.
12	DC WAE (loop)	The DC WAE recognizes the TCP Option marking as its own indicating a routing loop. The packet is dropped.

Best Practice: When sharing a WAAS cluster in a WCCP service group among multiple routers/switches with the potential for asymmetric routing, always provide direct paths between the switches/routers whose interfaces do not participate in WCCP redirection.

PLATFORM VARIATIONS

Various WCCP supported router and switches are listed below. Each device has its own features and capabilities. Most can be used with Cisco WAAS and exceptions are noted. As experience is gained with each device and IOS software release, this information is subject to change. Please consult the internal WAASipedia and Engineering WCCP web pages for the latest information.

IOS Version

WCCP, like any software, was developed over time, adapted to various hardware environments, and refined in capabilities. Different features and hardware support are available in various release and trains of IOS. In general, for Mainline and T train IOS releases, IOS release 12.4(20)T or later is recommended. Prior version may work with a specific device and feature set requirement. Please consult the internal WAASipedia and Engineering WCCP web pages for the latest information.

Cisco Integrated Services Router

Consisting of the 1800, 2800 and 3800 series routers, the ISRs are the workhorse of the remote office in large enterprises and the foundation for networking in small businesses. The addition of a Cisco WAAS Network Module, NME-WAE, provides a "go fast" mode for remote offices within the enterprise. WCCP provides a simple and convenient way to redirect traffic on an ISR to the WAAS WAE.

The minimum IOS version required for the ISR routers is 12.4(11)T3.

Redirection Methods

Cisco ISRs support GRE redirection only.

Packet Return

Cisco ISRs support GRE Packet Return only.

Packet Egress

Cisco ISRs support GRE Packet Egress only.

Interface Assignment

While both ingress and egress redirection is supported, ingress redirection is recommended where possible to minimize the CPU load on the switch. Cisco Express Forwarding is always recommended to improve routing performance within the router. Where egress WCCP redirection is required, CEFs is strongly recommended. Redirection ACLs are supported and recommended to limit the redirection processing required on the switch.

Cisco ASR 1000 Series Aggregation Services Routers

The minimum IOS XE version required for the ASR 1000 routers is 2.2.

Assignment Method

For the ISO XE 2.2 release, the ASR 1000 supports only mask assignment. Hash assignment may be available in future releases.

Redirection Methods

The Cisco ASR 1000 Series supports both L2 redirection method and GRE redirect in Cisco IOS Software XE Release 2.2.

Packet Return

The Cisco ASR 1000 Series supports both L2 return method and GRE return in Cisco IOS Software XE Release 2.2.

Packet Egress

The Cisco ASR 1000 Series supports IP Forwarding, WCCP GRE and generic GRE egress return in Cisco IOS Software XE Release 2.2.

Interface Assignment

The Cisco ASR 1000 Series supports only ingress redirection. Full extended redirection ACLs are supported and recommended to limit the redirection processing required on the switch.

Cisco ASA 5500 Series

As of release version 8.2, the Cisco ASA 5500 series Adaptive Security Appliance supports WCCPv2 for web-caching, but cannot be use with Cisco WAAS. The current implementation does not support WCCP services that preserve the original client IP address.

Cisco Catalyst 3550

The Catalyst 3550 supports WCCPv2 but is not recommended for use with Cisco WAAS.

Cisco Catalyst 3560/3750

The Catalyst 3560 and 3750 minimum IOS version is 12.2(46)SE.

Redirection Methods

The Catalyst 3560 and 3750 switches support only L2 redirection. The 3560 and 3750 do not support GRE redirection. WCCP L2 redirection requires that the WAAS WAEs be directly connected to the Catalyst switch.

Packet Return

The Catalyst 3650 supports WCCP L2 return and WCCP GRE. Since these devices support only L2 redirection requiring that the WAAS WAEs be directly connected to the switch, WCCP L2 packet return is the recommended return method.

Packet Egress

Since Layer 2 rewrite is supported as the packet return method, Packet Egress should also be implemented as L2 return.

Interface Assignment

The Catalyst 3650 supports only ingress redirection. Redirection ACLs are supported and recommended to limit the redirection processing required on the switch.

Additional Notes

- For the ACLs used in the router redirect-list, the deny keyword is not supported to explicitly deny certain traffic. Instead, use permit statement for all permitted traffic, and deny all will be implicitly applied at the end.
- While the 3560 and 3750 do not support GRE redirection, other GRE traffic passing through the switch will break because this traffic is incorrectly redirected to WAAS. Non E-series 3560 and 3750 have a hardware related issue where GRE over such a switch does not work (see CSCee62674). For the 4500, this problem will be fixed in IOS release 12.2(52)SG .

Cisco Catalyst 4500/4900

The Catalyst 4500 and 4900 Sup3 and Sup6 do not support WCCP. For other supervisors, the Catalyst 4500 and 4900 family minimum IOS version is 12.2(31)SG.

Assignment Method

The Catalyst 4500 and 4900 family supports mask assignment only. For optimum performance, use mask assignment whenever possible.

Redirection Methods

The Catalyst 4500 and 4900 switches support only L2 redirection. GRE redirection is not supported. WCCP L2 redirection requires that the WAAS WAEs be directly connected to the Catalyst switch.

Packet Return

The Catalyst 4500 and 4900 switches support only WCCP L2 return.

Packet Egress

Since Layer 2 rewrite is the only supported return method, Packet Egress should also be implemented as L2 return.

Interface Assignment

The Catalyst 4500 and 4900 support ingress redirection only. Redirection ACLs are not supported.

Additional Notes

- The Catalyst 4500 and 4900 family of switches do not support router redirect ACLs. Use WAAS Static Bypass Lists instead.
- Layer 2 redirection requires that WAAS devices be directly connected to the router and should be on separate IP subnetworks.
- The Cat4k does not allow PBR and WCCP to be configured on the same interface.
- WCCP is not compatible with VRF-lite enabled on the same Cat4k interface.
- See Configuring WCCPv2 on a Cisco 4500.

Cisco Catalyst 6500

The Cisco Catalyst 6500 family LAN switches can be configured to use WCCPv2 for network interception and redirection. WCCPv2 supported for the Catalyst 6500 models requires the following minimum software versions:

- Supervisor 1A or older – Not supported for use with Cisco WAAS
- Supervisor 2 – IOS version 12.1(27b)E or newer
- Supervisor 32 – IOS version 12.2(18)SXF14 or newer, hybrid requires CatOS 8.5 or newer
- Supervisor 720 – IOS version 12.2(18)SXF14 or newer

Assignment Method

The Catalyst 6500 family supports both hash and mask assignment. For optimum performance, use mask assignment whenever possible.

Redirection Methods

The Catalyst 6500 supports L2 redirection in hardware and GRE redirection software[4]. Wherever possible, L2 redirection in hardware is recommended. This requires that the WAAS WAEs be directly connected to the Catalyst switch.

Packet Return

The Catalyst 6500 supports WCCP L2 return in hardware and WCCP GRE return in software. Wherever possible, WCCP L2 return in hardware is recommended.

Packet Egress

The Catalyst 6500 supports IP Forwarding, WCCP Negotiated Return and Generic GRE return. IP Forwarding is typically not recommended as there is no guarantee of network path affinity in a multiple Catalyst 6500 deployment. WCCP Negotiated return is recommended over Generic GRE return only if the selected Packet Return method is WCCP L2. Generic GRE return is recommended in a multiple Catalyst 6500 deployment.

Additional Notes

- Avoid redirecting traffic on egress.
- Use Mask Assignment, L2 Redirect and Redirect in where possible.
- An excellent treatment of "WCCP on 6500" (circa 2007) can be found on the Advanced Services Campus VT.
- See Configuring Web Cache Services using WCCP on a Cisco 6500.

Cisco Catalyst 7600

The Catalyst 7600 minimum IOS version is 12.2(18)SXD1.

Assignment Method

The Catalyst 7600 family supports both hash and mask assignment. For optimum performance, use mask assignment whenever possible.

Redirection Methods

The Catalyst 7600 supports L2 redirection in hardware and GRE redirection software. Wherever possible, L2 redirection in hardware is recommended. This requires that the WAAS WAEs be directly connected to the Catalyst switch.

Packet Return

The Catalyst 7600 supports WCCP L2 return in hardware and WCCP GRE return in software. Wherever possible, WCCP L2 return in hardware is recommended.

Packet Egress

The Catalyst 7600 supports IP Forwarding, WCCP Negotiated Return and Generic GRE return. IP Forwarding is typically not recommended as returning packets put an unnecessary burden on the Catalyst CPU and there is no guarantee of network path affinity in a multiple Catalyst 7600 deployment. WCCP Negotiated return is recommended over Generic GRE return only if the selected

[4] GRE redirection in hardware is supported on the Catalyst 6500 with the SUP720 Policy Feature Card 3 (PFC3).

Packet Return method is WCCP L2. Generic GRE return is recommended in a multiple Catalyst 7600 deployment where WCCP L2 return is not possible.

Interface Assignment

While both ingress and egress redirection is supported, ingress redirection preferred and recommended.

Additional Notes

- Use Mask Assignment, L2 Redirect and Redirect in where possible.

Cisco Nexus 7000

The Nexus 7000 minimum NXOS version is 4.2.1.

Assignment Method

The Nexus 7000 supports only mask assignment.

Redirection Methods

The Nexus 7000 supports L2 redirection in hardware. GRE redirection is not supported in the current release. This requires that the WAAS WAEs be directly connected to the Nexus switch.

Packet Return

The Nexus 7000 supports WCCP L2 return in hardware. WCCP GRE return is not supported.

Packet Egress

Since the Nexus 7000 does not support GRE return (WCCP or generic), IP Forwarding is the only available egress return method.

Interface Assignment

While both ingress and egress redirection is supported, ingress redirection preferred and recommended.

Additional Notes

- Redirect ACL lists are supported.
- While the WCCPv2 protocol specification does not have support for VRF, the Nexus 7000 does provide support for VRF:
    - User will be provided with the flexibility of enabling service groups on a per VRF basis.
    - The service group definitions will be local to the VRF it was defined in.
    - Service groups defined within a VRF will be logically separated from each other. This means that the service group numbers could be reused across VRF.
    - If no VRF was specified in the service group configuration, it will be associated with the default or global VRF (i.e. no VRF).
    - When WCCP redirect is enabled on an interface, the VRF association can be derived from the interface itself (single interface can belong to only one VRF). To be consistent with IOS implementation, the VRF id will be required when WCCP redirect is enabled on an interface.

TROUBLESHOOTING

An excellent WCCP Troubleshooting Guide is available on the IP Technologies Engineering UK web page.

TOP 10 COMMON MISTAKES WHEN DEPLOYING WAAS WITH WCCP

1. Making major WCCP changes on the router/switch without first disabling the WCCP on all WAAS devices in the service group.

   Before making any WCCP global service group configuration changes on a service group router, first deregister the WAAS devices from the service group by issuing a "no wccp version 2" global config command one each device or by un-checking the WCCP Version 2 check box on the Central Manager device or device group configuration panel. When the change is complete, re-register the WAAS devices with the service group by enabling WCCP with the "wccp version 2" global config command. This is not necessary when making router interface configuration changes or router redirect-list changes.

   When configuring WCCP on a WAAS device, configure the WCCP egress method, define the WCCP router list, and enable the service groups with the "wccp tcp-promiscuous" command, in that order, before enabling WCCP with the "wccp version 2" config command. When making changes to the WAAS device WCCP configuration, disable WCCP first before making the change.

2. Routing on the host subnet in the branch using 61 in.

3. Forgetting to turn on the egress-method when deploying on the host subnet causing a loop.

   See Service Group Placement above.

4. Using default WCCP settings on a Catalyst 6500 switch.

   See Cisco Catalyst 6500 Platform Variations above.

5. Using the default WCCP mask for packet forwarding load balancing.

   See Choosing a Mask above.

6. No router redirect-list ACLs to minimize unnecessary routing and packet processing.

   See Limiting WAAS Redirected Traffic above.

7. Incorrect ACL usage on specific platforms (Catalyst 3750, Catalyst 4500 etc.)

   See Platform Variations above.

8. Using egress redirection on a Catalyst 6500 switch.

   See Cisco Catalyst 6500 Platform Variations above.

9. Interaction between active gateway and routing path (host subnet routing below).

   See Avoiding WCCP-induced Routing Loops above.

10. WCCP peering over the WAN with GRE return.

    WCCP peering over the WAN with GRE return can occur when deploying 2 routers with 2 WAAS devices, with the WAAS devices directly attached to both routers and configured with GRE return. In this case, the WAAS device to router peering could be either over the client LAN (good) or the WAN (very bad).

GLOSSARY

Assignment Method	The method by which redirected packets are distributed between cache engines.
Designated Cache Engine	The cache engine in a WAAS device farm responsible for dictating to the router or routers how redirected traffic should be distributed between the members of the farm.
Egress redirection	The router examines traffic as it is leaving an interface toward an external network. Generally, the traffic has already entered an interface prior to reaching the interface where egress redirection is configured, and a load has already been placed on the router resources.
Forwarding Method	The method by which redirected packets are transported from router to WAAS device.
Ingress redirection	The router examines traffic as it enters an interface from an external network. Generally, ingress redirection is less resource intensive than egress redirection.
One-arm deployment	The simplest form of off path WCCP deployment where one of the Cisco WAE's Gigabit Ethernet interfaces is attached directly to the router or to the dedicated Cisco WAE VLAN on the switch and the second interface is not used.
Packet Return Method	The method by which packets redirected to a WAAS device are returned to a router for normal forwarding.
Redirection Hash Table	A 256-bucket hash table maintained by the router or routers. This table maps the hash index derived from a packet to be redirected to the IP address of a destination WAAS device.
Service Group	A group of one or more routers plus one or more WAAS devices working together in the redirection of traffic whose characteristics are part of the Service Group definition.
Transparent Redirection	Transparent redirection is a technique used to deploy caching without the need for reconfiguration of clients or servers. It involves the interception and redirection of traffic to one or more WAAS devices by a router or switch transparently to the end points of the traffic flow.
Two-arm deployment	An alternative to the one-arm off path deployment mode uses the second interface (the external interface on the Cisco NME-WAE). The first interface is connected directly to the router or to the switch in the Cisco WAE VLAN, and the second interface connects to the switch in the VLAN with the users or the servers. With this deployment mode, all traffic from the Cisco WAE must traverse the router with the exception of traffic going to any node adjacent to the Cisco WAE's second interface (users or servers that are in the same VLAN). The two-arm off path mode can provide a slight improvement in performance and less overall router CPU workload.
Usable Cache Engine	From the viewpoint of a router, a cache engine is considered a usable member of a Service Group when it

	has sent that cache engine a WCCP2_I_SEE_YOU message and has received in response a WCCP2_HERE_I_AM message with a valid "Receive ID".
Cache Engine Farm	One or more cache engines associated with a router or routers.

APPENDIX A – CISCO WCCP SERVICE GROUPS

Service Name	Service Number	Protocol	Port	Priority
web-cache	0	tcp	80	240
dns	53	udp	53	202
ftp-native	60	tcp		200
tcp-promiscuous	61	tcp	*	34
tcp-promiscuous	62	tcp	*	34
https-cache	70	tcp	443	231
rtsp	80	tcp	554	200
wmt	81	tcp	1755	201
mmsu	82	udp	1755	201
rtspu	83	udp	5005	201
cifs-cache	89	tcp	139, 445	224
custom	90			220
custom	91			221
custom	92			222
custom	93			223
custom	94			224
custom	95			225
custom	96			226
custom	97			227
custom-web-cache	98	tcp	80	230
reverse-proxy	99	tcp	80	235

DO NOT DELETE ANYTHING BELOW THIS LINE (INCLUDING THE CARRIAGE RETURNS). REMOVE THIS & THE ABOVE TEXT PRIOR TO PUBLISHING OR PRINTING TO PDF.

Corporate Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Web site at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Printed in the USA