

INTERNET-DRAFT
<draft-wilson-wrec-wccp-v2-01.txt>

M Cieslak
D Forster
G Tiwana
R Wilson
Cisco Systems
03 Apr 2001
Expires Oct 2001

Web Cache Communication Protocol V2.0

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress".

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

1. Abstract

This document describes version 2.0 of the Web Cache Communication Protocol (WCCP). The WCCP V2.0 protocol specifies interactions between one or more routers and one or more web-caches. The purpose of the interaction is to establish and maintain the transparent redirection of selected types of traffic flowing through a group of routers. The selected traffic is redirected to a group of web-caches with the aim of optimising resource usage and lowering response times.

The protocol does not specify any interaction between the web-caches within a group or between a web-cache and a web-server.

2. Definitions

Assignment Method

The method by which redirected packets are distributed between

web-caches.

Designated Web-Cache.

The web-cache in a web-cache farm responsible for dictating to the router or routers how redirected traffic should be distributed between the members of the farm.

Forwarding Method

The method by which redirected packets are transported from router to web-cache.

Packet Return Method

The method by which packets redirected to a web-cache 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 web-cache.

Service Group

A group of one or more routers plus one or more web-caches 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 web-caches by a router or switch transparently to the end points of the traffic flow.

Usable Web-Cache.

From the viewpoint of a router a web-cache is considered a usable member of a Service Group when it has sent that web-cache a WCCP2_I_SEE_YOU message and has received in response a WCCP2_HERE_I_AM message with a valid "Receive ID".

Web-Cache Farm.

One or more web-caches associated with a router or routers.

3. Introduction

3.1 Protocol Overview

WCCP V2.0 defines mechanisms to allow one or more routers enabled for transparent redirection to discover, verify, and advertise connectivity to one or more web-caches.

Having established connectivity the routers and web-caches 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 web-caches and also the method used to transport traffic between router and cache.

A single web-cache within a Service Group is elected as the designated web-cache. It is the responsibility of the designated web-cache to provide routers with the data which determines how redirected traffic is distributed between the web-caches in the Service Group.

3.2 WCCP V2.0 enhancements

WCCP V2.0 supports the following enhancements to the WCCP V1.0 protocol.

- * Multi-Router Support.

WCCP V2.0 allows a farm of web-caches to be attached to more than one router.

- * Multicast Support.

WCCP V2.0 supports multicasting of protocol messages between web-caches and routers.

- * Improved Security.

WCCP V2.0 provides optional authentication of protocol packets received by web-caches and routers.

- * Support for redirection of non-HTTP traffic.

WCCP V2.0 supports the redirection of traffic other than HTTP traffic through the concept of Service Groups.

- * Packet return.

WCCP V2.0 allows a web-cache to decline to service a redirected packet and to return it to a router to be forwarded. The method by which packets are returned to a router is negotiable.

* Alternative Hashing.

WCCP V2.0 allows the designated web-cache to mark individual buckets in the Redirection Hash Table for a secondary hash. This allows the traffic being hashed to a particular bucket to be distributed across the members of a Service Group.

* Multiple Forwarding Methods

WCCP V2.0 allows individual web-caches to negotiate the method by which packets are forwarded to a web-cache from a router. Packets may now be forwarded unencapsulated using a Layer 2 destination address rewrite.

* Multiple Assignment Methods

WCCP V2.0 allows the designated web-cache to negotiate the method by which packets are distributed between the web-caches in a service group. Packets may now be assigned using a hashing scheme or a masking scheme.

* Command and Status Information

WCCP V2.0 includes a mechanism to allow a web-cache to pass a command to the routers in a Service Group. The same mechanism can be employed by the routers to pass status information to the web-caches in a Service Group.

4. Protocol Description

4.1 Joining a Service Group

A web-cache joins and maintains its membership of a Service Group by transmitting a WCCP2_HERE_I_AM message to each router in the Group at HERE_I_AM_T (10) second intervals. 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 web-cache by IP address. The Service Info component of the WCCP2_HERE_I_AM message identifies and describes the Service Group in which the web-cache wishes to participate.

A router responds to a WCCP2_HERE_I_AM message with a WCCP2_I_SEE_YOU message. 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 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 web-cache 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 web-caches to which the packet is addressed. A web-cache not

in the list should discard the WCCP2_I_SEE_YOU message.

4.2 Describing a Service Group

The Service Info component of a WCCP2_HERE_I_AM message describes the Service Group in which a web-cache 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 web-caches and do not require a description other than a Service ID.

In contrast Dynamic Services must be 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 web-cache to join the Service Group. A web-cache 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.

4.3 Establishing Two-Way Connectivity

WCCP V2.0 uses a "Receive ID" to verify two-way connectivity between a router and a web-cache. The Router Identity Info component of a WCCP2_I_SEE_YOU message contains a "Receive ID" field. This field is maintained separately for each Service Group and its value is incremented each time the router sends a WCCP2_I_SEE_YOU message to the Service Group.

The "Receive ID" sent by a router is reflected back by a web-cache in the Web-Cache View Info component of a WCCP2_HERE_I_AM message. A router checks the value of the "Receive ID" in each WCCP2_HERE_I_AM message received from a Service Group member. If the value does not match the "Receive ID" in the last WCCP2_I_SEE_YOU message sent to that member the message is discarded.

A router considers a web-cache to be a usable member of a Service Group only after it has sent that web-cache a WCCP2_I_SEE_YOU message and received a WCCP2_HERE_I_AM message with a valid "Receive ID" in response.

4.4 Negotiating the Forwarding Method

A web-cache and router may negotiate the method by which packets are forwarded to the web-cache by the router.

This negotiation is per web-cache, per Service Group. Thus web-caches participating in the same Service Group may negotiate different forwarding methods with the Service Group routers.

A router will advertise the supported forwarding 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.

A web-cache will inspect the forwarding method advertisement in the first WCCP2_I_SEE_YOU message received from a router for a particular Service Group. If the router does not advertise a method supported by the web-cache then the web-cache will abort its attempt to join the Service Group. Otherwise the web-cache will pick one method from those advertised by the router and specify that in the optional Capabilities Info component of its next WCCP2_HERE_I_AM message. Absence of a forwarding method advertisement in a WCCP2_HERE_I_AM message implies the cache is requesting the default GRE encapsulation method.

A router will inspect the forwarding method selected by a web-cache in the WCCP2_HERE_I_AM message received in response to a WCCP2_I_SEE_YOU message. If the selected method is not supported by the router the router will ignore the WCCP2_HERE_I_AM message. If the forwarding method is supported the router will accept the web-cache as usable and add it to the Service Group.

4.5 Negotiating the Assignment Method

A web-cache and router may negotiate the method by which packets are distributed between the web-caches in a Service Group.

The negotiation is per Service. Thus web-caches participating in several Service Groups may negotiate a different assignment method for each Service Group.

A router will 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.

A web-cache will inspect the assignment method advertisement in the first WCCP2_I_SEE_YOU message received from a router for the Service Group. If the router does not advertise a method supported by the

web-cache then the web-cache will abort its attempt to join the Service Group. Otherwise the web-cache will pick one method from those advertised by the router and specify that in the optional Capabilities Info component of its next WCCP2_HERE_I_AM message. Absence of an assignment method advertisement in a WCCP2_HERE_I_AM message implies the cache is requesting the default Hash assignment method.

A router will inspect the assignment method selected by a web-cache in the WCCP2_HERE_I_AM message received in response to a WCCP2_I_SEE_YOU message. If the selected method is not supported by the router the router will ignore the WCCP2_HERE_I_AM message. If the assignment method is supported the router will accept the web-cache as usable and add it to the Service Group.

4.5 Negotiating the Packet Return Method

A web-cache and router may negotiate the method by which packets are returned from a web-cache to a router for normal forwarding.

The negotiation is per Service. Thus web-caches participating in several Service Groups may negotiate a different packet return method for each Service Group.

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.

A web-cache will inspect the packet return method advertisement in the first WCCP2_I_SEE_YOU message received from a router for the Service Group. If the router does not advertise a method supported by the web-cache then the web-cache will abort its attempt to join the Service Group. Otherwise the web-cache will pick one method from those advertised by the router and specify that method in the optional Capabilities Info component of its next WCCP2_HERE_I_AM message. Absence of a packet return method advertisement in a WCCP2_HERE_I_AM message implies the cache is requesting the default GRE packet return method.

A router will inspect the packet return method selected by a web-cache in the WCCP2_HERE_I_AM message received in response to a WCCP2_I_SEE_YOU message. If the selected method is not supported by the router the router will ignore the WCCP2_HERE_I_AM message. If the packet return method is supported the router will accept the web-cache as usable and add it to the Service Group.

4.6 Advertising Views of the Service Group

Each router advertises its view of a Service Group via the Router View Info component in the WCCP2_I_SEE_YOU message it sends to web-caches. This component includes a list of the useable web-caches in the Service Group as seen by the router and a list of the routers in the Service Group as reported in WCCP2_HERE_I_AM messages from web-caches. A change number in the component is incremented if the Service Group membership has changed since the last WCCP2_I_SEE_YOU message sent by the router.

Each web-cache advertises its view of the Service Group via the Web Cache View Info component in the WCCP2_HERE_I_AM message it sends to routers in the Service Group. This component includes the list of routers that have sent the web-cache a WCCP2_I_SEE_YOU message and a list of web-caches learnt from the WCCP2_I_SEE_YOU messages. The Web Cache View Info component also includes a change number which is incremented each time Service Group membership information changes.

4.7 Security

WCCP V2.0 provides a security component in each protocol message to allow simple authentication. Two options are supported:

- * No Security (default)
- * MD5 password security

MD5 password security requires that each router and web-cache wishing to join a Service Group be configured with the Service Group password. Each WCCP protocol packet sent by a router or web-cache for that Service Group will contain in its security component the MD5 checksum of the WCCP protocol message (including the WCCP message header) and a Service Group password. Each web-cache or router in the Service Group will authenticate the security component in a received WCCP message immediately after validating the WCCP message header. Packets failing authentication will be discarded.

4.8 Distribution of Traffic Assignments

WCCP V2.0 allows the traffic assignment method to be negotiated. There are two types of information to be communicated depending on the assignment method:

- * Hash Tables
- * Mask/Value Sets

4.8.1 Hash Tables

When using hash assignment each router uses a 256-bucket Redirection Hash Table to distribute traffic for a Service Group across the member web-caches. It is the responsibility of the Service Group's designated web-cache to assign each router's Redirection Hash Table.

The designated web-cache uses a WCCP2_REDIRECT_ASSIGNMENT message to assign the routers' Redirection Hash Tables. This message is generated following a change in Service Group membership and is sent to the same set of addresses to which the web-cache sends WCCP2_HERE_I_AM messages. The designated web-cache will wait 1.5 HERE_I_AM_T seconds following a change before generating the message in order to allow the Service Group membership to stabilise.

The Redirection Hash Tables can be conveyed in either an Assignment Info Component or an Alternate Assignment Component within a WCCP2_REDIRECT_ASSIGNMENT. Both components contain an Assignment Key. This will be reflected back to the designated web-cache in subsequent WCCP2_I_SEE_YOU messages from the routers in the Service Group. A WCCP2_REDIRECT_ASSIGNMENT may be repeated after HERE_I_AM_T seconds if inspection of WCCP2_I_SEE_YOU messages indicates a router has not received an assignment.

A router will flush its Redirection Hash Table if a WCCP2_REDIRECT_ASSIGNMENT is not received within 5 HERE_I_AM_T seconds of a Service Group membership change. A router will flush its Redirection Hash Table if it receives a WCCP2_REDIRECT_ASSIGNMENT message in which it is not listed.

The designated web-cache lists the web-caches to which traffic should be distributed in either an Assignment Info Component or an Alternate Assignment Component within a WCCP2_REDIRECT_ASSIGNMENT message. Only those web-caches seen by every router in the Service Group are included.

4.8.2 Mask/Value Sets

When using mask assignment each router uses masks and a table of values to distribute traffic for a Service Group across the member web-caches. It is the responsibility of the Service Group's designated web-cache to assign each router's mask/value sets.

The designated web-cache uses the Alternate Assignment Component in a WCCP2_REDIRECT_ASSIGNMENT message to assign the routers' mask/value set. This message is generated following a change in Service Group membership and is sent to the same set of addresses to which the web-cache sends WCCP2_HERE_I_AM messages. The designated web-cache

will wait 1.5 HERE_I_AM_T seconds following a change before generating the message in order to allow the Service Group membership to stabilise.

The Alternate Assignment Info component of the WCCP2_REDIRECT_ASSIGNMENT contains an Assignment Key. This will be reflected back to the designated web-cache in subsequent WCCP2_I_SEE_YOU messages from the routers in the Service Group. A WCCP2_REDIRECT_ASSIGNMENT message may be repeated after HERE_I_AM_T seconds if inspection of WCCP2_I_SEE_YOU messages indicates a router has not received an assignment.

A router will flush its mask/value set if a WCCP2_REDIRECT_ASSIGNMENT is not received within 5 HERE_I_AM_T seconds of a Service Group membership change. A router will flush its mask/value set if it receives a WCCP2_REDIRECT_ASSIGNMENT in which it is not listed.

The designated web-cache lists the web-caches to which traffic should be distributed in the Alternate Assignment Info component of the WCCP2_REDIRECT_ASSIGNMENT message. Only those web-caches seen by every router in the Service Group are included.

4.9 Electing the Designated Web-cache

Election of the designated web-cache will take place once a Service Group membership has stabilised following a change. The designated web-cache must be receiving a WCCP2_I_SEE_YOU message from every router in the Service Group.

Election of the designated web-cache is not part of the WCCP protocol. However it is recommended that the web-cache with the lowest IP address is selected as designated web-cache for a Service Group.

4.10 Traffic Interception

A router will check packets passing through it against its set of Service Group descriptions. The Service Group descriptions are checked in priority order. A packet which matches a Service Group description is a candidate for redirection to a web-cache in the Service Group.

A router will not redirect a packet with a source IP address matching any web-cache in the Service Group.

4.11 Traffic Redirection

4.11.1 Redirection with Hash Assignment

Redirection with hash assignment is a two-stage process. In the first stage a primary key is formed from the packet (as defined by the Service Group description) and hashed to yield an index into the Redirection Hash Table.

If the Redirection Hash Table entry contains an unflagged web-cache index then the packet is redirected to that web-cache. If the bucket is unassigned the packet is forwarded normally. If the bucket is flagged as requiring a secondary hash then a secondary key is formed (as defined by the Service Group description) and hashed to yield an index into the Redirection Hash Table. If the secondary entry contains a web-cache index then the packet is directed to that web-cache. If the entry is unassigned the packet is forwarded normally.

4.11.2 Redirection with Mask Assignment

The first step in redirection using the mask assignment method is to perform a bitwise AND operation between the mask from the first mask/value set in the Service Group definition and the contents of the packet. The output of this operation is the set of fields in the packet which will be used for value matching. The selected fields from the packet are then compared against each entry in the list of values for that mask/value set. If a match is found the packet is redirected to the web-cache associated with the value entry. If no match is found the process is repeated for each mask/value set defined for the Service Group. If, after trying all of the mask/value sets defined for the Service Group, no match is found, the packet is forwarded normally.

Mask/value sets are processed in the order in which they are presented in the Alternate Assignment component. Value elements are compared in the order in which they appear in the mask/value set of which they are part.

4.12 Traffic Forwarding

WCCP allows the negotiation of the forwarding method between router and web-cache (See Negotiating the Forwarding Method). The currently defined forwarding methods are:

- * GRE Encapsulated
- * Unencapsulated with L2 rewrite

4.12.1 Forwarding with GRE Encapsulation

Redirected packets are encapsulated in a new IP packet with a GRE [\[1\]](#) header followed by a four-octet Redirect header.

The GRE encapsulation uses the simple four-octet GRE header with the two Flags and Version octets set to zero and a Protocol Type of 0x883E.

The Redirect header is as follows:

```
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|D|A| Reserved | Service ID | Alt Bucket | Pri Bucket |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---
```

- D Dynamic Service
 - 0: Well known service
 - 1: Dynamic service

- A Alternative bucket used
 - 0: Primary bucket used
 - 1: Alternative bucket used

Service ID

Service Group identifier

Alt Bucket

Alternative bucket index used to redirect the packet. Only valid for hash assignment.

Pri Bucket

Primary bucket index used to redirect the packet. Only valid for hash assignment.

4.12.2 Forwarding with L2 Rewrite

Redirected packets are not encapsulated. The router replaces the packet's destination MAC address with the MAC address of the target web-cache.

This forwarding method requires that the target web-cache be directly-connected to the router at Layer 2. A router will not allow a web-cache which is not directly attached to negotiate this forwarding method.

4.13 Packet Return

WCCP V2.0 allows a web-cache to decline a redirected packet and return it to a router for normal forwarding as specified by the packet's destination IP address. The method by which packets are returned from router to cache is a matter for negotiation (see Negotiating the Packet Return Method).

When a router receives a returned packet it must not attempt to redirect that packet back to a web-cache. Two methods are available to prevent any further redirection:

- * Interface Configuration
- * Encapsulation

The interface configuration method requires that a router is configured to inhibit redirection of packets arriving over interfaces connected to web-caches. Redirection may be disabled for all packets arriving on an interface or for packets where the source MAC address is that of a web-cache. This mechanism is efficient but is topology dependant and thus may not always be suitable. In this case the packet return method in use is L2.

The encapsulation method requires a web-cache to send returned packets to a router with encapsulation. Returned packets are encapsulated in a GRE packet [1] with a Protocol Type of 0x883E and contain the original Redirect Header or a null Redirect Header if none was present in the original redirected packet. The receiving router removes the GRE encapsulation from the packets and forwards them without attempting to redirect. The packet return method used in this case is GRE.

4.14 Querying Cache Time-Out

If a router does not receive a WCCP2_HERE_I_AM message from a Service Group member for $2.5 * \text{HERE_I_AM_T}$ seconds it will query the member by unicasting a WCCP2_REMOVAL_QUERY message to it. The target Service Group member should respond by sending a series of 3 identical WCCP2_HERE_I_AM messages, each separated by $\text{HERE_I_AM_T}/10$ seconds.

If a router does not receive a WCCP2_HERE_I_AM message from a Service Group member for $3 * \text{HERE_I_AM_T}$ seconds it will consider the member to be unusable and remove it from the Service Group. The web-cache will no longer appear in the Router View Info component of the WCCP2_I_SEE_YOU message.

The web-cache will be purged from the assignment data for the Service Group.

4.15 Command and Status Information

WCCP V2.0 includes a mechanism to allow web-caches to send commands to routers within a service group. The same mechanism can be used by the routers to provide status information to web-caches.

The mechanism is implemented by the Command Extension component. This component is included in the WCCP2_HERE_I_AM message from a web-cache passing commands to routers in a Service Group.

If a router needs to send status information to a web-cache it will include a command in the Command Extension component within its own WCCP2_I_SEE_YOU message. That command will indicate the type of status information being carried.

5. Protocol Messages

Each WCCP protocol message is carried in a UDP packet with a destination port of 2048. There are four WCCP V2.0 messages:

- * Here I AM
- * I See You
- * Redirect Assign
- * Removal Query

5.1 'Here I Am' Message

```
+-----+
|      WCCP Message Header      |
+-----+
|      Security Info Component  |
+-----+
|      Service Info Component   |
+-----+
| Web-Cache Identity Info Component |
+-----+
| Web-Cache View Info Component  |
+-----+
| Capability Info Component (optional) |
+-----+
| Command Extension Component (optional) |
+-----+
```

5.2 'I See You' Message

```
+-----+
|      WCCP Message Header      |
+-----+
|      Security Info Component   |
+-----+
|      Service Info Component    |
+-----+
|      Router Identity Info Component |
+-----+
|      Router View Info Component |
+-----+
|      Assignment Info Component  |
|              OR                 |
|      Assignment Map Component  |
+-----+
| Capability Info Component (optional) |
+-----+
| Command Extension Component (optional) |
+-----+
```


5.3 'Redirect Assign' Message

```
+-----+
|      WCCP Message Header      |
+-----+
|      Security Info Component   |
+-----+
|      Service Info Component    |
+-----+
|      Assignment Info Component  |
|              OR                 |
|      Alternate Assignment Component |
+-----+
```

5.4 'Removal Query' Message

```
+-----+
|      WCCP Message Header      |
+-----+
|      Security Info Component   |
+-----+
|      Service Info Component    |
+-----+
|      Router Query Info Component |
+-----+
```

5.5 WCCP Message Header



Type

- WCCP2_HERE_I_AM (10)
- WCCP2_I_SEE_YOU (11)
- WCCP2_REDIRECT_ASSIGN (12)
- WCCP2_REMOVAL_QUERY (13)

Version

0x200

Length

Length of the WCCP message not including the WCCP Message Header.

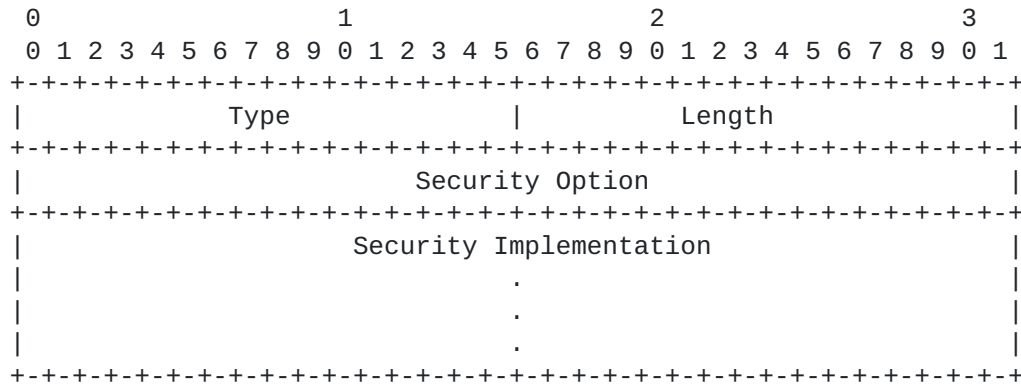
5.6 Message Components

Each WCCP message comprises a WCCP Message Header followed by a number of message components. The defined components are:

- * Security Info
- * Service Info
- * Router Identity Info
- * Web-Cache Identify Info
- * Router View Info
- * Web-Cache View Info
- * Assignment Info
- * Router Query Info
- * Capabilities Info
- * Alternate Assignment
- * Assignment Map
- * Command Extension

Components are padded to align on a four-octet boundary. Each component has a 4-octet header specifying the component type and length. Note that the length value does not include the 4-octet component header.

5.6.1 Security Info Component



Type

WCCP2_SECURITY_INFO (0)

Length

Length of the remainder of the component.

Security Option


WCCP2_NO_SECURITY (0)
WCCP2_MD5_SECURITY (1)

Security Implementation

If Security Option has the value WCCP2_NO_SECURITY then this field is not present. If Security Option has the value WCCP2_MD5_SECURITY this is a 16-octet field containing the MD5 checksum of the WCCP message and the Service Group password. The maximum password length is 8 octets.

Prior to calculating the MD5 checksum the password should be padded out to 8 octets with trailing zeros and the Security Implementation field of the Security Option set to zero. The MD5 checksum is calculated using the 8 octet padded password and the WCCP message (including the WCCP Message Header).

5.6.2 Service Info Component



Type

WCCP2_SERVICE_INFO (1)

Length

Length of the remainder of the component.

Service Type

WCCP2_SERVICE_STANDARD (0).

Service is a well known service and is described by the Service ID. All fields other than Service ID must be zero.

WCCP2_SERVICE_DYNAMIC (1).

Service is defined by the Protocol, Service Flags and Port fields.

Service ID

Service number. A number in the range 0-255. For well known services numbers in the range 0-50 are reserved. The numbers currently defined for well known services are:

0x00 HTTP

Priority

Service priority. The lowest priority is 0, the highest is 255. Packets for redirection are matched against Services in priority order, highest first. Well known services have a priority of 240.

Protocol

IP protocol identifier

Service Flags

0x0001 Source IP Hash
0x0002 Destination IP Hash
0x0004 Source Port Hash
0x0008 Destination Port Hash
0x0010 Ports Defined.
0x0020 Ports Source.
0x0100 Source IP Alternative Hash
0x0200 Destination IP Alternative Hash
0x0400 Source Port Alternative Hash
0x0800 Destination Port Alternative Hash

The primary hash flags (Source IP Hash, Destination IP Hash, Source Port Hash, Destination Port Hash) determine the key which will be hashed to yield the Redirection Hash Table primary bucket index. If only the Destination IP Hash flag is set then the packet destination IP address is used as the key. Otherwise if any of the primary hash flags are set then the key is constructed by XORing the appropriate fields from the packet with the key (which has an initial value of zero).

The key is hashed using the following algorithm:

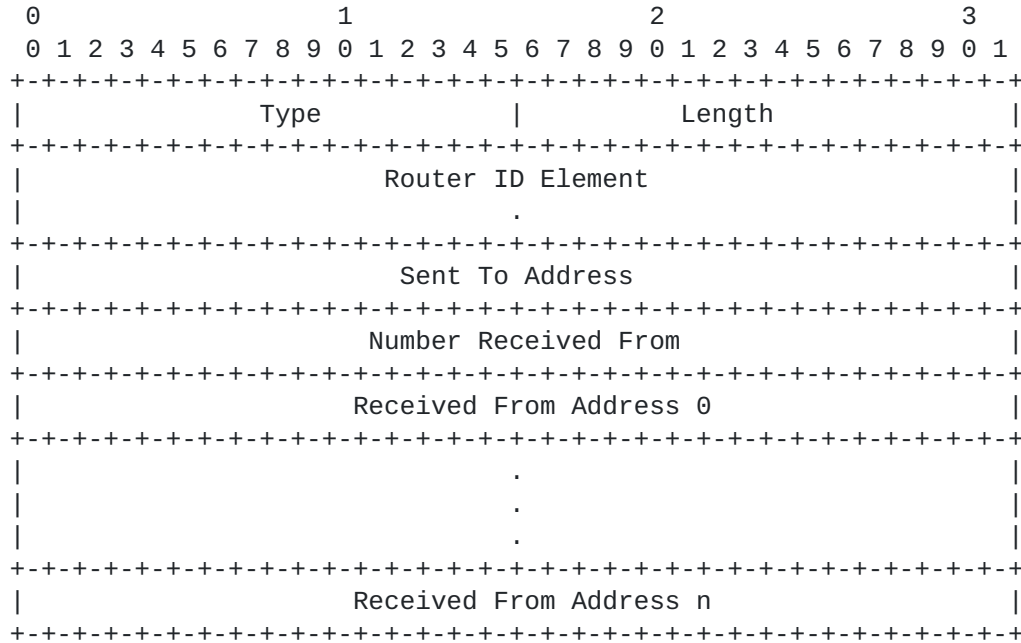
```
ulong hash = key;  
hash ^= hash >> 16;  
hash ^= hash >> 8;  
return(hash & 0xFF);
```

If alternative hashing has been enabled for the primary bucket (see Assignment Info Component) the alternate hash flags (Source IP Alternative Hash, Destination IP Alternative Hash, Source Port Alternative Hash, Destination Port Alternative Hash) determine the key which will be hashed to yield a secondary bucket index. The key is constructed by XORing the appropriate fields from the packet with a key (which has an initial value of zero).

Port 0-7

Zero terminated list of UDP or TCP port identifiers. Packets will be matched against this set of ports if the Ports Defined flag is set. If the Ports Source flag is set the port information refers to a source port, if clear the port information refers to a destination port.

5.6.3 Router Identity Info Component



Type

WCCP2_ROUTER_ID_INFO (2)

Length

Length of the remainder of the component.

Router ID Element

Element containing the router's identifying IP address and Receive ID. The IP address must be a valid, reachable address for the router.

Sent To Address

IP address to which the target web-cache sent the WCCP2_HERE_I_AM message. When this component is present in a unicast WCCP2_I_SEE_YOU message it will contain the IP address that the target web-cache used. When present in a multicast WCCP2_I_SEE_YOU message it will contain the Service Group multicast address.

Number Received From

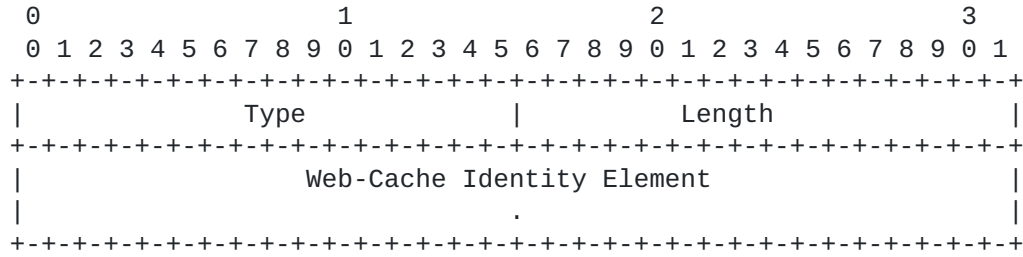
The number of web-caches to which this message is directed. When using multicast addressing it may be less than the number of caches which

actually see the message.

Received From Address 0-n

List of the IP addresses of web-caches to which this message is directed. When using multicast addressing it may be a subset of the caches which actually see the message.

5.6.4 Web-Cache Identity Info Component



Type

WCCP2_WC_ID_INFO (3)

Length

Length of the remainder of the component.

Web-Cache Identity Element

Element containing the web-cache IP address and Redirection Hash Table mapping.

5.6.5 Router View Info Component

This represents a router's view of the Service Group.

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
Type										Length																													
Member Change Number																																							
Assignment Key																																							
.																																							
Number of Routers																																							
Router 0																																							
.																																							
.																																							
.																																							
Router n																																							
Number of Web-Caches																																							
Web-Cache Identity Element 0																																							
.																																							
.																																							
.																																							
Web-Cache Identity Element n																																							
.																																							

Type

WCCP2_RTR_VIEW_INFO (4)

Length

Length of the remainder of the component.

Member Change Number

Incremented each time there is a change in Service Group membership.

Assignment Key

Assignment Key element received in the last WCCP2_REDIRECT_ASSIGNMENT message. Used by the designated web-cache to verify that an assignment has been executed.

Number of Routers

Number of routers in the Service Group

Router 0-n

IP addresses of routers in the Service Group. This list is constructed from routers reported by web-caches via WCCP2_HERE_I_AM messages. Note that a router does not include itself in the list unless it has also been reported via a WCCP2_HERE_I_AM message.

Number of Web-Caches

Number of useable web-caches in the Service Group

Web-Cache Identity Element 0-n

Identity elements of useable web-caches in Service Group. This list contains web-caches that have sent the router a WCCP2_HERE_I_AM message with a valid "Received ID".

5.6.6 Web Cache View Info Component

This represents a web-cache's view of the Service Group.

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
Type										Length																													
Change Number																																							
Number of Routers																																							
Router ID Element 0																																							
Router ID Element n																																							
Number of Web-Caches																																							
Web Cache address 0																																							
Web Cache address n																																							

Type

WCCP2_WC_VIEW_INFO (5)

Length

Length of the remainder of the component.

Change Number

Incremented each time there is a change in the view.

Number of Routers

Number of routers in the Service Group

Router ID Element 0-n

List of elements containing the identifying IP address for each router in the Service Group and the last "Received ID" from each.

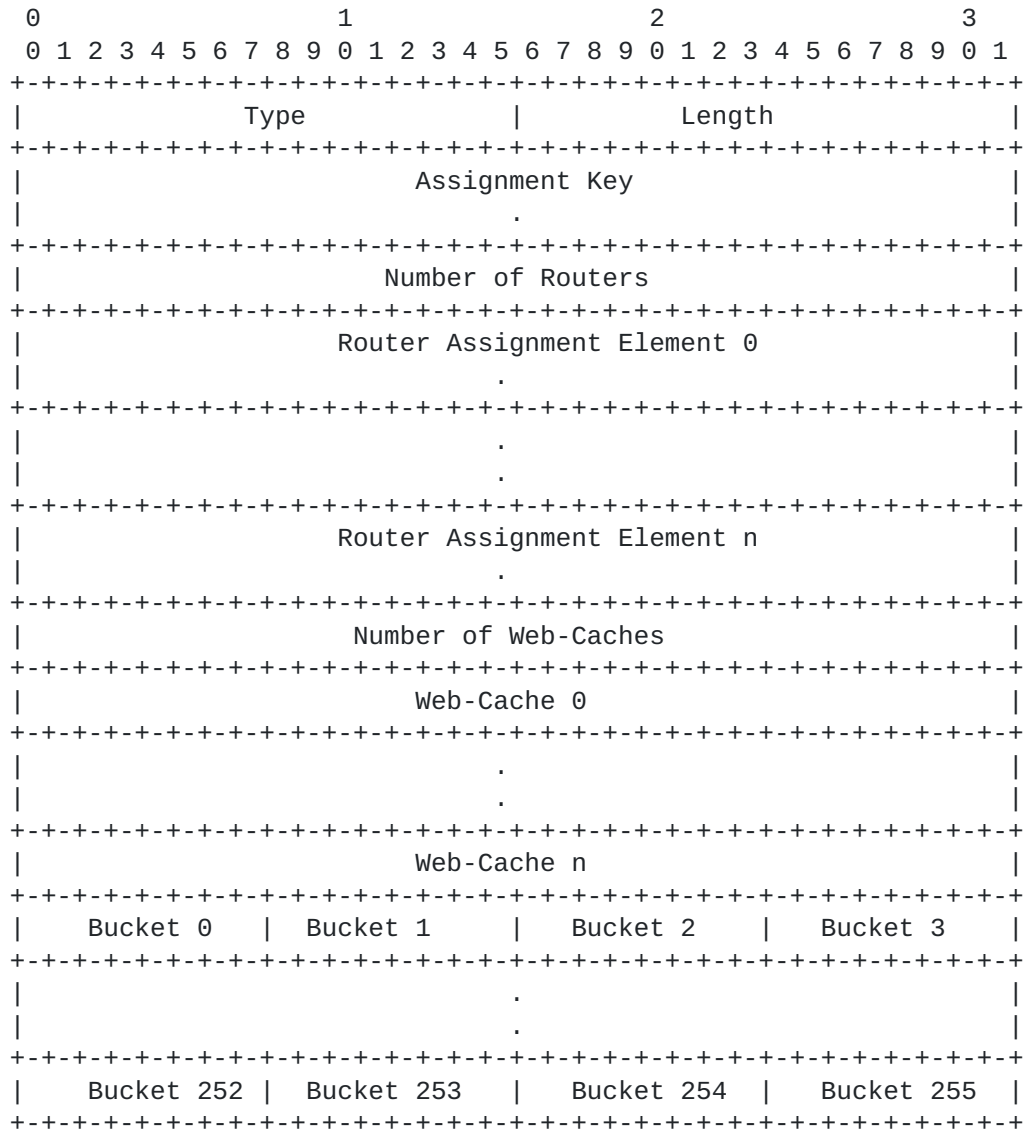
Number of Web-Caches

Number of web-caches in the Service Group

Web Cache address 0-n

List of web-cache IP addresses learnt from WCCP2_I_SEE_YOU messages.

5.6.7 Assignment Info Component



Type

WCCP2_REDIRECT_ASSIGNMENT (6)

Length

Length of the remainder of the component.

Assignment Key

The designated web-cache expects this element to be returned by a router in subsequent WCCP2_I_SEE_YOU messages.

Number of Routers

Number of routers reachable by the designated web-cache.

Router Assignment Element 0-n

Elements containing the router IP address, "Receive ID" and "Change Number" for each router.

Number of Web-Caches

Number of useable web-caches in the Service Group seen by all routers.

Web Cache 0-n

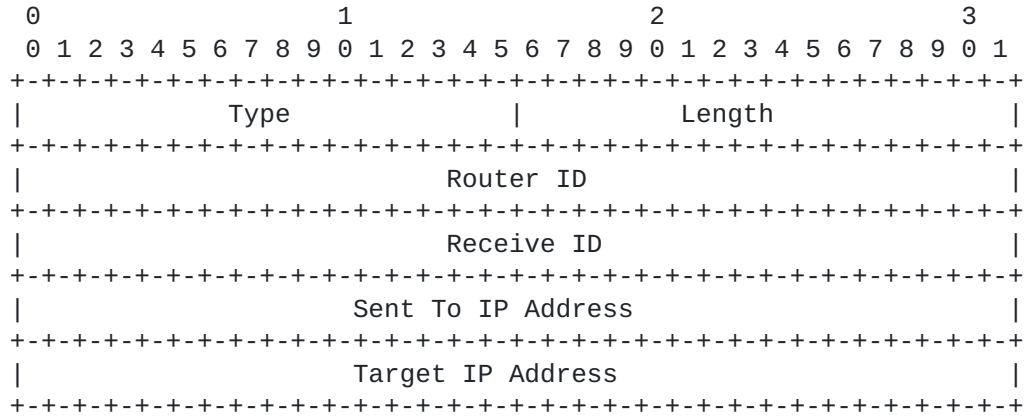
List of the IP addresses of useable web-caches in Service Group. The position of a web-cache identifier in this list is the web-cache index. The first entry in the list has an index of zero.

Bucket 0-255

Contents of the Redirection Hash Table. The content of each bucket is a web-cache index value in the range 0-31. If set the A flag indicates that alternative hashing should be used for this web-cache. The value 0xFF indicates no web-cache has been assigned to the bucket.

```
0 1 2 3 4 5 6 7
+--+--+--+--+--+--+
| Index      |A|
+--+--+--+--+--+--+
```

5.6.8 Router Query Info Component



Type

WCCP2_QUERY_INFO (7)

Length

Length of the remainder of the component.

Router ID

Router IP address. The same address advertised in a WCCP2_I_SEE_YOU message.

Receive ID

Receive ID expected by the router.

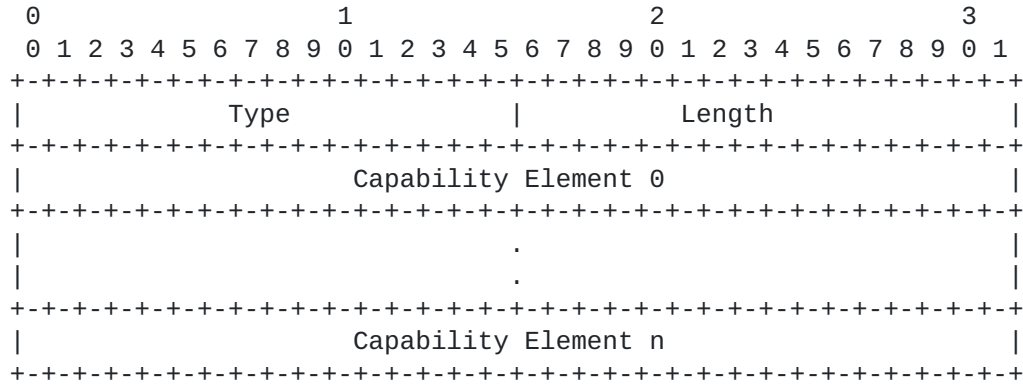
Sent To IP Address

IP address to which the web-cache sent its last WCCP2_HERE_I_AM message. This will not be the Router ID if the web-cache is multicasting its WCCP2_HERE_I_AM messages.

Target IP Address

IP address of web-cache being queried.

5.6.9 Capabilities Info Component



Type

WCCP2_CAPABILITY_INFO (8)

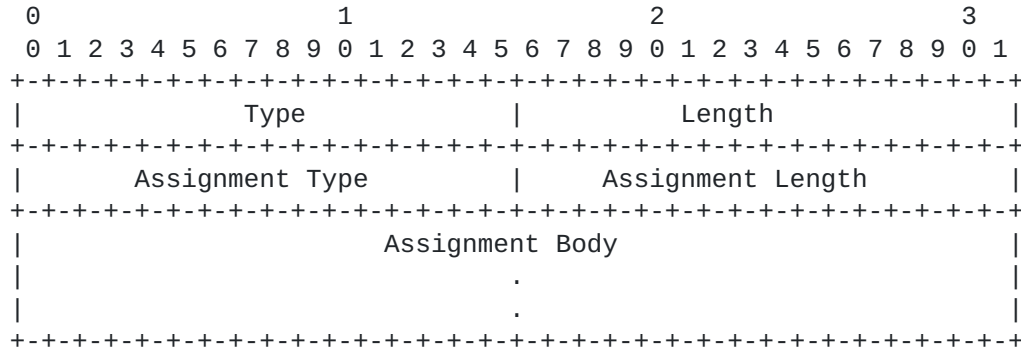
Length

Length of the remainder of the component.

Capability Element

Element in Type-Length-Value format (TLV) describing a router or web-cache capability.

5.6.10 Alternate Assignment Component



Type

WCCP2_ALT_ASSIGNMENT (13)

Length

Length of the remainder of the component.

Assignment Type

Currently defined values:

WCCP2_HASH_ASSIGNMENT (0x00)

WCCP2_MASK_ASSIGNMENT (0x01)

Assignment Length

Length of Assignment Body

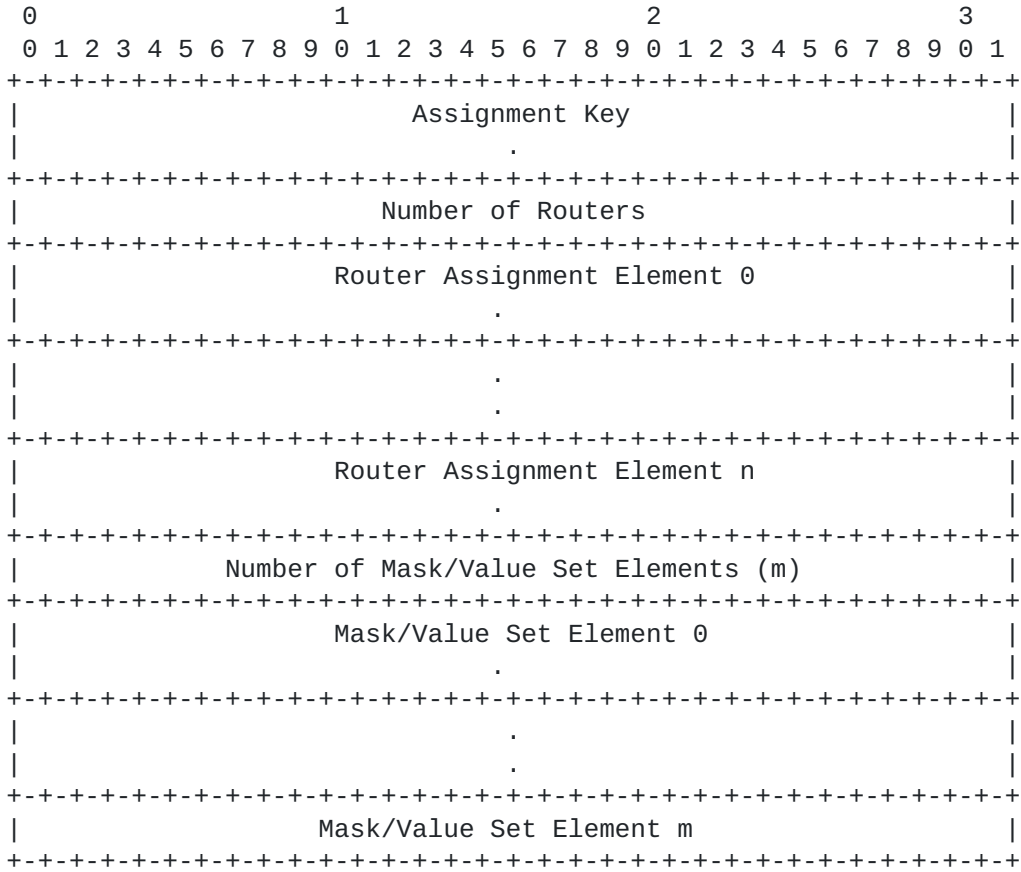
Assignment Body

The format of Assignment Body depends upon the value of Assignment Type.

Assignment Type = WCCP2_HASH_ASSIGNMENT

In this case the body of the message is identical to the Assignment Info Component with the Type and Length fields omitted.

Assignment Type = WCCP2_MASK_ASSIGNMENT



Assignment Key

The designated web-cache expects this element to be returned by a router in subsequent WCCP2_I_SEE_YOU messages.

Number of Routers

Number of routers reachable by the designated web-cache.

Router Assignment Element 0-n

Element containing the router IP address, Receive ID and Change Number for each router.

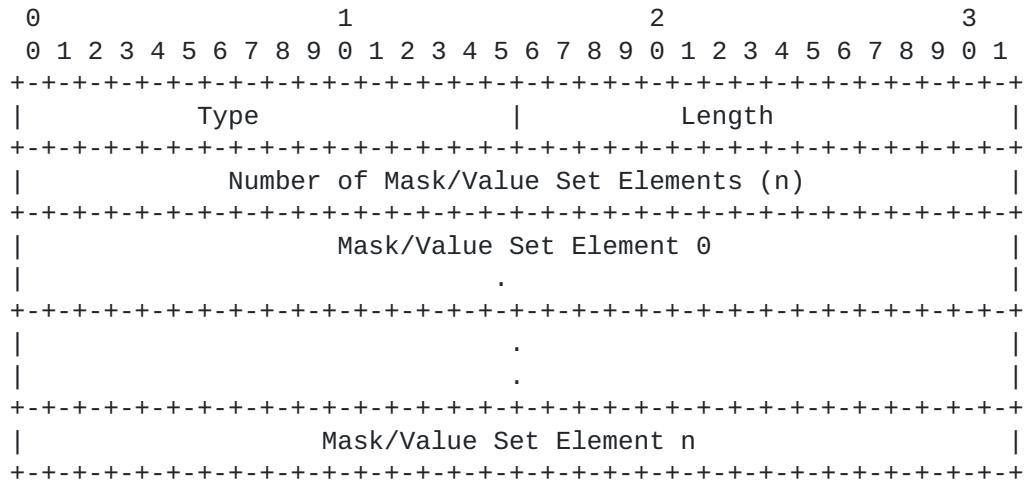
Number of Mask/Value Set Elements (m)

Number of Mask/Value Set elements in this message

Mask/Value Set Element 0-m

A list of the Mask/Value Element Sets for the Service Group

5.6.11 Assignment Map Component



Type

WCCP2_ASSIGN_MAP (14)

Length

Length of the remainder of the component.

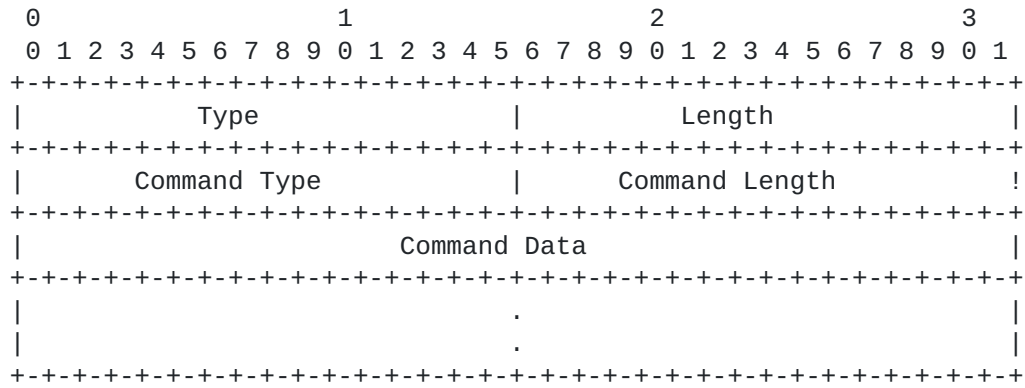
Number of Mask/Value Set Elements (n)

Number of Mask/Value Set elements in the message

Mask/Value Set Element 0-n

A list of the Mask/Value Element Sets for the Service Group

5.6.12 Command Extension Component



Type

WCCP2_COMMAND_EXTENSION (15)

Length

Length of the remainder of the component.

Command Type

The command specifier.

Command Length

The length of the Command Data field of this command


The defined Command Types are:

Command Type: WCCP2_COMMAND_TYPE_SHUTDOWN (01)
 Command Length: 4
 Command Data: Web-cache IP address
 Description: This command is used by a web-cache to indicate to the routers in a Service Group that it is shutting down and should no longer receive any redirected traffic.

Command Type: WCCP2_COMMAND_TYPE_SHUTDOWN_RESPONSE (02)
 Command Length: 4
 Command Data: Web-cache IP address.
 Description: This command is used by a router to acknowledge receipt of a SHUTDOWN command received from the web-cache identified by the IP address in the Command Data field.

5.7 Information Elements

5.7.1 Router ID Element



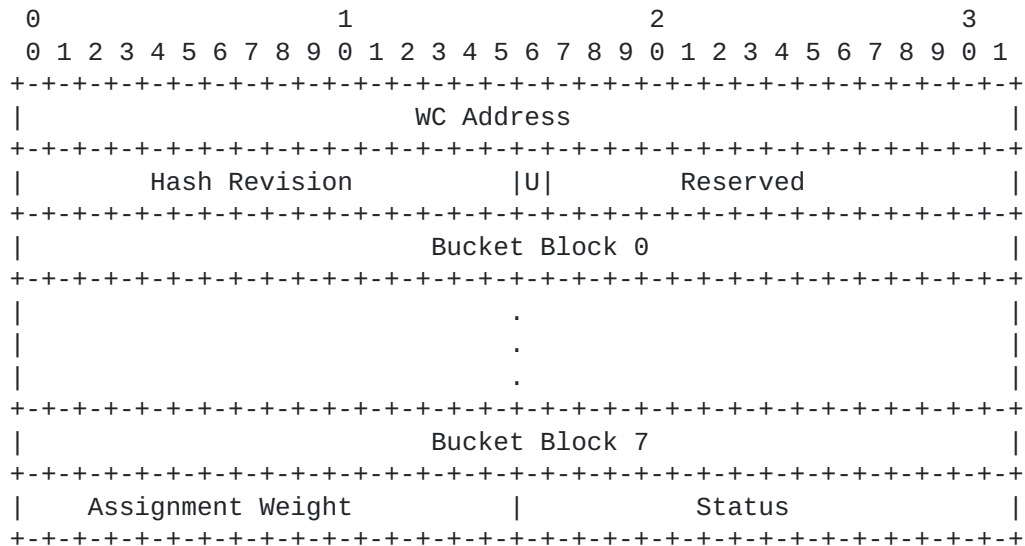
Router ID

Router's identifying IP address. This must be a valid IP address by which the router is reachable.

Receive ID

Defined per Service Group. Incremented each time the router sends a WCCP protocol message including a Router Identity Info component. Will never be zero.

5.7.2 Web-Cache Identity Element



WC Address

Web-Cache IP address

Hash Revision

0x00

U

If set indicates that the web cache does not have an assignment in the Redirection Hash Table and that Bucket Block data is historical. Historical data may be used by the designated web-cache to re-assign the same bucket set to a web-cache that left and subsequently rejoined a Service Group.

Bucket Block 0-7

256-bit vector. A set bit indicates the corresponding Redirection Hash Table bucket is assigned to this web-cache.

Assignment Weight


Hash weight. May be used to indicate to the designated web-cache how new assignments should be made.

Status

Hash status. May be used to indicate to the designated web-cache how new assignments should be made.

5.7.3 Assignment Key Element

This element identifies a particular assignment.



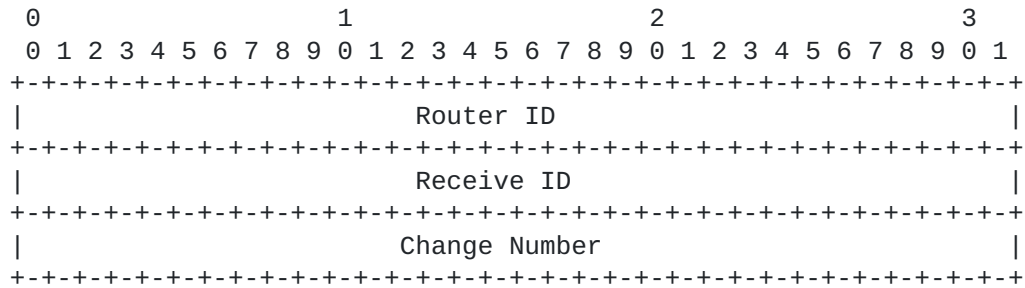
Key IP Address

Designated web-cache IP address

Key Change Number

Incremented if a change has occurred.

5.7.4 Router Assignment Element



Router ID

Router's identifying IP address. It must be a valid address by which the router is reachable.


Receive ID

Last Receive ID received from the router identified by Router ID. A router will ignore an assignment if Receive ID is invalid.

Change Number

Last Member Change Number received from the router identified by Router ID. A router will ignore an assignment if Change Number is invalid.

5.7.5 Capability Element



Type

Currently defined types are:

WCCP2_FORWARDING_METHOD	0x01
WCCP2_ASSIGNMENT_METHOD	0x02
WCCP2_PACKET_RETURN_METHOD	0x03

Length

Length of Capability element Value

Value

The length and format of the value field is dependant on the capability type.

Type = WCCP2_FORWARDING_METHOD

A 32-bit bitmask indicating supported/selected forwarding methods. Currently defined values are:

WCCP2_FORWARDING_METHOD_GRE 0x00000001
WCCP2_FORWARDING_METHOD_L2 0x00000002

Type = WCCP2_ASSIGNMENT_METHOD

A 32-bit bitmask indicating supported/selected assignment methods. Currently defined values are:

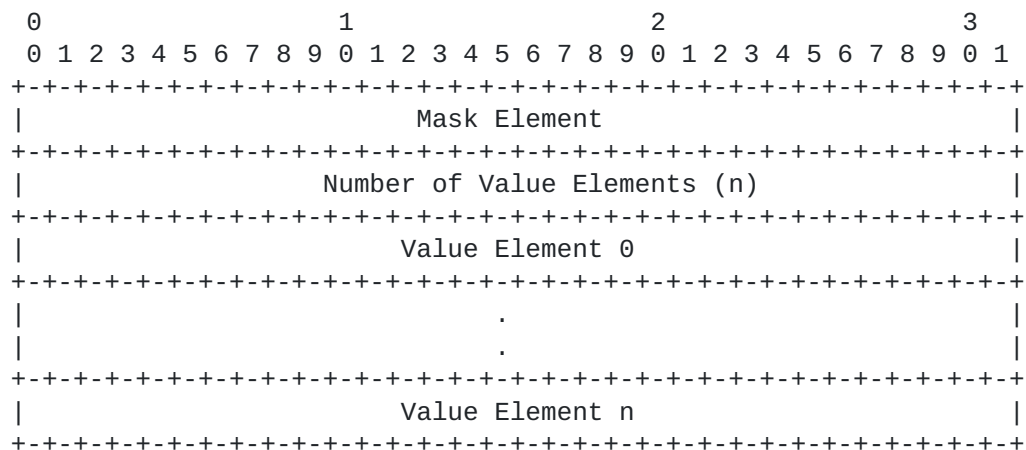
WCCP2_ASSIGNMENT_METHOD_HASH 0x00000001
WCCP2_ASSIGNMENT_METHOD_MASK 0x00000002

Type = WCCP2_PACKET_RETURN_METHOD

A 32-bit bitmask indicating supported/selected packet return methods. Currently defined values are:

WCCP2_PACKET_RETURN_METHOD_GRE 0x00000001
WCCP2_PACKET_RETURN_METHOD_L2 0x00000002

5.7.6 Mask/Value Set Element



Mask Element

Mask element for this set.

Number of Value Elements (n)

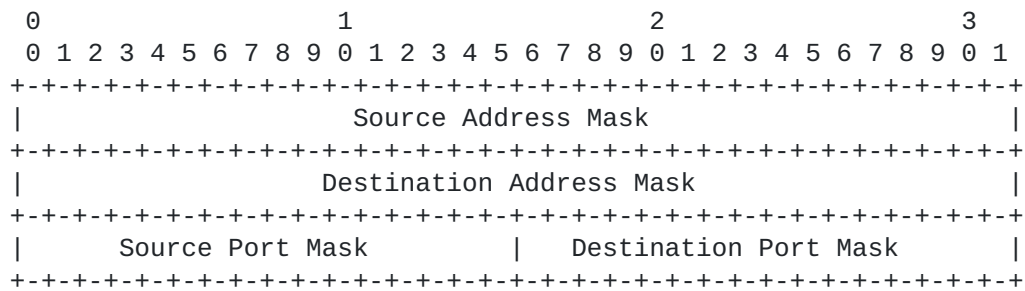
The number of value elements in this set.

Value Element 0-n

The list of value elements for this set.

5.7.7 Mask Element

Note that in all of the mask fields of this element a zero means "Don't care".



Source Address Mask

The 32 bit mask to be applied to the source IP address of the packet.

Destination Address Mask

The 32 bit mask to be applied to the destination IP address of the packet.

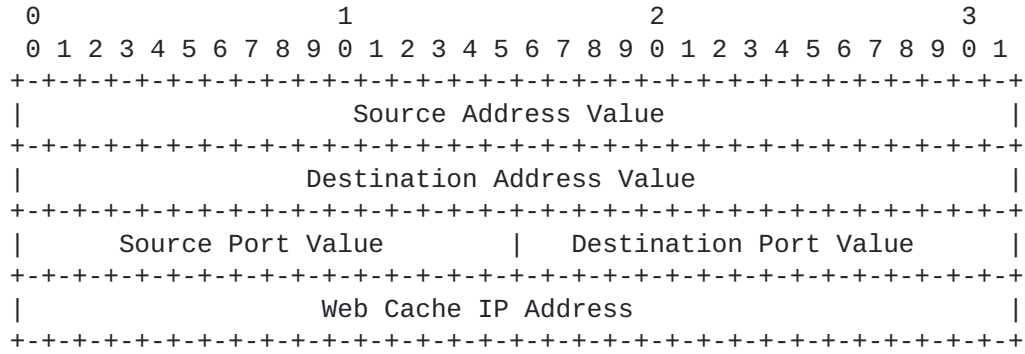
Source Port Mask

The 16 bit mask to be applied to the TCP/UDP source port field of the packet.

Destination Port Mask

The 16 bit mask to be applied to the TCP/UDP destination port field of the packet.

5.7.8 Value Element



Source Address Value

The value to match against the source IP address of the packet after masking.

Destination Address Value

The value to match against the destination IP address of the packet after masking.

Source Port Value

The value to match against the TCP/UDP source port number of the packet after masking.

Destination Port Value

The value to match against the TCP/UDP destination port number of the packet after masking.

Web-cache IP address

The IP address of the web-cache to which packets matching this value element should be sent.

6. Security Considerations

WCCP V2 provides a mechanism for message authentication. It is described in [section 4.7](#) of this document. The authentication mechanism relies on a password known to all routers and web-caches in a Service Group. The password is part of the Service Group configuration and is used to compute message checksums which can be verified by other members of the group. Should the password become known to a host attempting to disrupt the operation of a Service Group it would be possible for that host to spoof WCCP messages and appear as either a router or web-cache in the Service Group.

To pose as a router in a Service Group a host would advertise its presence to the members of the group in I_SEE_YOU messages. If accepted as part of the Service Group the host would receive the configuration for the group in a HERE_I_AM message from the designated web-cache. This situation would not pose any threat to the operation of the Service Group because the host would not be performing any packet redirection and all packets would flow normally.

To pose as a web-cache within a Service Group a host would advertise its presence in HERE_I_AM messages. Acceptance of the host as part of the Service Group would be decided by the designated cache and may be subject to additional security checks not specified by WCCP. Should the host become part of the Service Group it would be assigned a proportion of the traffic redirected by the routers in the Service Group. Assuming that the host drops any redirected packets the net effect to clients would be that some attempts to retrieve content via the Service Group routers would fail.

7. References

[1] Hanks, Li, Farinacci & Traina, "Generic Routing Encapsulation (GRE)", [RFC 1701](#), October 1994

8. Authors' Addresses

Martin Cieslak
Cisco Systems
170 Tasman Drive
San Jose, CA 95143

David Forster
Cisco Systems
170 Tasman Drive
San Jose, CA 95143

Gurumukh Tiwana
Cisco Systems
170 Tasman Drive
San Jose, CA 95143

Rob Wilson
Cisco Systems
170 Tasman Drive
San Jose, CA 95143

email: robewils@cisco.com

Expires January 2001