

Web Cache Coordination Protocol V1.0

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>

Cisco has a patent pending that may relate to WCCP V1.0. If any patents issue to Cisco or its subsidiaries with claims that are necessary for practicing WCCP V1.0, then any party will be able to obtain a license from Cisco to use any such patent claims under openly specified, reasonable, non-discriminatory terms to implement WCCP V1.0. No license is required for nonprofit institutions.

Abstract

This draft documents the Web Cache Coordination Protocol (WCCP) V1.0. This protocol is used (a) to associate a single router with one or more web-caches for the purposes of transparent redirection of HTTP traffic, and (b) to allow one of the web-caches to dictate how the router distributes transparently-redirection traffic across the associated web-caches.

This draft describes the interactions between a router and one or more web-caches. It does not describe the interactions between a group of associated web-caches or those between a web-cache and a web-server.

Definitions

Transparent Redirection.

Transparent redirection is a technique used to deploy web-caching without the need for reconfiguration of web-clients. It involves the interception and redirection of HTTP traffic to one or more web-caches by a router or switch, transparently to the web-client.

Web-Cache Farm.

One or more web-caches associated with a router.

Designated Web-Cache.

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

Redirection Hash Table.

A 256-bucket hash table maintained by the router. This table maps the IP destination address of a packet for redirection to the IP address of a web-cache in the farm

Description of Protocol

WCCP has two main functions. The first is to allow a router enabled for transparent redirection to discover, verify, and advertise connectivity to one or more web-caches.

The second function is to allow one of the web-caches, the designated web-cache, to dictate how the router distributes redirected traffic across the web-cache farm.

It is recommended that the web-cache with the lowest IP address is elected as designated web-cache for a farm.

Discovery

WCCP V1.0 allows a single router to be associated with one or more web-caches. A group of web-caches associated with a router is referred to as a web-cache farm. A web-cache may be directly attached to the router or some hops distant.

Joining a web-cache farm

A web-cache joins a web-cache farm by periodically unicasting a

WCCP_HERE_I_AM packet to the router associated with the farm at intervals of HERE_I_AM_T (10) seconds. The source IP address of the WCCP_HERE_I_AM uniquely identifies the web-cache. The router unicasts a WCCP_I_SEE_YOU packet back to the web-cache in response to each WCCP_HERE_I_AM it receives.

Verifying connectivity

The Received_ID fields in the WCCP_HERE_I_AM and WCCP_I_SEE_YOU packets are used to verify two-way connectivity between the router and web-cache. The router increments the value of the Received_ID field each time it sends a WCCP_I_SEE_YOU to a web-cache and expects to receive the same value back in the Received_ID field of the next WCCP_HERE_I_AM from that cache. WCCP_HERE_I_AM packets containing an invalid Received_ID are ignored.

The Received_ID in the initial WCCP_HERE_I_AM sent from a web-cache is ignored. The router will only consider a web-cache to be reachable when it has received a subsequent WCCP_HERE_I_AM with a correct Received_ID. Note that a useable web-cache is merely reachable; the router will not redirect traffic to a newly-acquired useable web-cache until instructed to do so in a WCCP_ASSIGN_BUCKETS packet from the designated web-cache.

Advertising connectivity

The router includes a list of the web-caches it considers to be usable in each WCCP_I_SEE_YOU packet it transmits. Each entry in the list includes the IP address of the web-cache and indicates which buckets in the Redirection Hash Table are currently assigned to that web-cache. This information is provided for the benefit of the designated web-cache.

A Change ID in the WCCP_I_SEE_YOU packet is incremented whenever the web-cache list changes or the bucket allocation for an entry in the list is modified.

Timing-out a web-cache

If the router does not receive a valid WCCP_HERE_I_AM for 3 * HERE_I_AM_T seconds it will no longer consider a web-cache to be useable. In this case the web-cache is no longer advertised in the WCCP_I_SEE_YOU packet and all buckets previously assigned to the web-cache in the router's Redirection Hash Table are marked as unassigned.

Assignment

The router associated with a web-cache farm distributes redirected traffic by destination IP address across the members of the farm as directed by the designated web-cache via the WCCP_ASSIGN_BUCKETS packet.

How the designated web-cache arrives at the traffic distribution described by the WCCP_ASSIGN_BUCKETS packet is outside the scope of this draft.

Since the router has no knowledge of the designated web-cache election process it will accept a WCCP_ASSIGN_BUCKETS packet from any member of the web-cache farm.

The value of the Received_ID in the WCCP_ASSIGN_BUCKETS packet must match that in the last WCCP_I_SEE_YOU sent to the designated web-cache. If the Received_ID is not valid the router will ignore the WCCP_ASSIGN_BUCKETS packet.

On receipt of a valid WCCP_ASSIGN_BUCKETS packet the router will set its Redirection Hash Table from information contained in the packet. This information comprises a list of web-caches followed by a 256-bucket hash table. The position of a web-cache in the list is its index number, the index number of the first entry being zero. Each bucket in the hash table may contain the value 0xFF, indicating no web-cache has been assigned to that bucket, or the index number of a web-cache.

The router does not generate a packet in response to the WCCP_ASSIGN_BUCKET. However the change in the Redirection Hash Table will be reflected in subsequent WCCP_I_SEE_YOU packets generated by the router.

Packet Redirection

Detection

The router detects HTTP packets (TCP packets with a destination port number of 80) and redirects them to a web-cache in the web-cache farm.

The destination IP address of a candidate packet is hashed to yield an index into the 256-bucket Redirection Hash Table. The indexed bucket indicates to which web-cache the packet should be redirected. If the bucket in the Redirection Hash Table is unassigned the packet cannot be redirected and should be forwarded normally.

Encapsulation

Each redirected packet is encapsulated in a GRE packet[1]. The encapsulation uses the base four-octet GRE header with the two Flags and version octets set to zero and a Protocol Type of 0x883E.

An encapsulated packet may be fragmented if it exceeds the output interface's MTU.

Returned packets

The router must ensure that HTTP traffic passing through it from members of the web-cache farm en-route to a web-server is not redirected.

The router will not redirect any packet with a source address belonging to a member of the web-cache farm.

Format of Protocol Packets

This section defines the format of the WCCP packets.

Each WCCP protocol packet is carried in a UDP packet with a destination port of 2048

Here I Am

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								
Type																																							
Protocol Version																																							
Hash Revision																																							
Hash Information (0)																																							
.																																							
.																																							
.																																							
Hash Information (7)																																							
Reserved																																							
Received ID																																							

Type

WCCP_HERE_I_AM (7)

Protocol Version

4

Hash Revision

0

Hash Information

A 256-element bit-vector. A set bit indicates that the corresponding bucket in the Redirection Hash Table is assigned to this web-cache. Normally the value of the Hash Information present in the last

WCCP_I_SEE_YOU message received by this cache. In the initial WCCP_HERE_I_AM sent to the router it may be zero or the value assigned to the cache in a previous membership of this web-cache farm. This information may be used by the designated web-cache to re-assign buckets to the cache.

U

Normally the value of the U flag present in the last WCCP_I_SEE_YOU message received by this cache. Set in first WCCP_HERE_I_AM to indicate that Hash Information is historical.

Received ID

The value of the Received ID present in the last WCCP_I_SEE_YOU received by this web-cache.

I See You

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Type                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Protocol Version                         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Change Number                           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Received ID                             |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Number of WCs                           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Web-Cache List Entry(0)                   |
|                                     .                                         |
|                                     .                                         |
|                                     .                                         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Web-Cache List Entry(n)                   |
|                                     .                                         |
|                                     .                                         |
|                                     .                                         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type

WCCP_I_SEE_YOU (8)

Protocol Version

4

Change Number

Incremented if a Web-Cache List Entry has been added, removed or its hash information has been modified since the last WCCP_I_SEE_YOU sent by the router.

Received ID

Incremented each time the router generates a WCCP_I_SEE_YOU. Will never be zero.

Number of WCs

Number of Web-Cache List Entry elements in the packet.

Web-Cache List Entry

The Web-Cache List Entry describes a Web-Cache by IP Address and lists the redirection hash table entries assigned to it.

[illegible]

IP Address

Web-cache IP Address

Hash Revision

0

Hash Information

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

U

If set indicates web-cache is not assigned in the Redirection Hash Table and that the web-cache hash information is historical. This information may be used by the designated web-cache to reassign buckets to a web-cache which has rejoined the farm.

Assign Bucket

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Type                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Received ID                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Number of Web Caches                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Web-Cache 0 IP Address                   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     .                                         |
|                                     .                                         |
|                                     .                                         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Web-Cache n IP Address                   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Bucket 0   |   Bucket 1   |   Bucket 2   |   Bucket 3   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     .                                         |
|                                     .                                         |
|                                     .                                         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Bucket 252 |   Bucket 253 |   Bucket 254 |   Bucket 255 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type

WCCP_ASSIGN_BUCKET (9)

Received ID

Value of Received ID in last WCCP_I_SEE_YOU received from router.

Number of Web Caches

Number of Web Caches to which redirect traffic can be sent.

Web-Cache IP Address, 0-n

IP Addresses of Web-Caches to which redirect traffic can be sent. The position of a Web-Cache's IP Address in this list is the Web-Cache's index number. The first entry in the list has an index number of zero.

Bucket 0-255

These 256 buckets represent the redirection hash table. The value of each bucket may be 0xFF (Unassigned) or a Web-Cache index number (0-31).

References

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

Authors' Addresses

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

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

Phone: +44-208-7568967
Email: dforster@cisco.com

Expires January 2001