WSA Behavior on Path MTU Discovery with Use of WCCP

Document ID: 118843

Contributed by Ivo Sabev, Cisco TAC Engineer.
Mar 17, 2015

Contents

Introduction

This document describes a problem encountered where the router drops packets when your configuration includes both Web Cache Communication Protocol (WCCP) and path Maximum Transmission Unit (MTU) discovery, and it provides a solution to the problem.

Background Information

Pre−phase

When looked at separately, many features are excellent to handle a specific problem. Sometimes though, if you combine two or three techniques, it produces some awkward behavior and you must introduce another feature or workaround in order to make it work properly. For example, use spanning tree and Open Shortest Path First (OSPF) and Layer 2 (L2) convergence takes longer (20s) than OSPF (1s if minimum dead interval is used), but replace spanning tree with Multiple Spanning−Tree (MST) and it functions properly again.

The same interoperability behavior has been observed between WCCP and path MTU discovery; many think that it is the Generic Routing Encapsulation (GRE) header problem. However, this document explains the real cause.

How Path MTU Dicovery and WCCP Work Separately

Path MTU Discovery

Each line has its limit on how large a packet can be. If you send a larger packet than is supported, then it is dropped. One of the roles of the L3 devices (routers) on the way is to take care and chop large packets from one of the lines to the other one in order to make sure that end−to−end communication is transparent to each line's capabilities.

Sometimes though, end hosts are configured in such a way that their packets cannot be chopped (for example,

encrypted files, voice calls). This information is communicated via the Don't Fragment (DF) bit inside the IP header. Routers drop packets like these, but the router tries to report to the end host via Internet Control Message Protocol (ICMP) message (type 3−Destination unreachable, code 4 − fragmentation needed, but DF bit set). This way, the host knows to send smaller packets in the future.

This is the heart of path MTU discovery. You can send large packets with the DF bit set in order to see whether they make it towards the end or if you receive an ICMP report as previously described. Once you determine the maximum workable packet size, use it for any further communications. Refer to RFC 1191 for more information.

The Web Security Appliance (WSA) employs path MTU discovery by default. Thus, all its generated packets have the DF bit set by the default configuration.

WCCP

If you need to impose security into your network on the web traffic without others' knowledge, you run their traffic via a proxy that is not visible. WCCP is the protocol that is used to communicate between the device that intercepts (router/firewall) and the web cache engine/proxy, which is WSA in this case.

This diagram illustrates how traffic flows in this scenario:

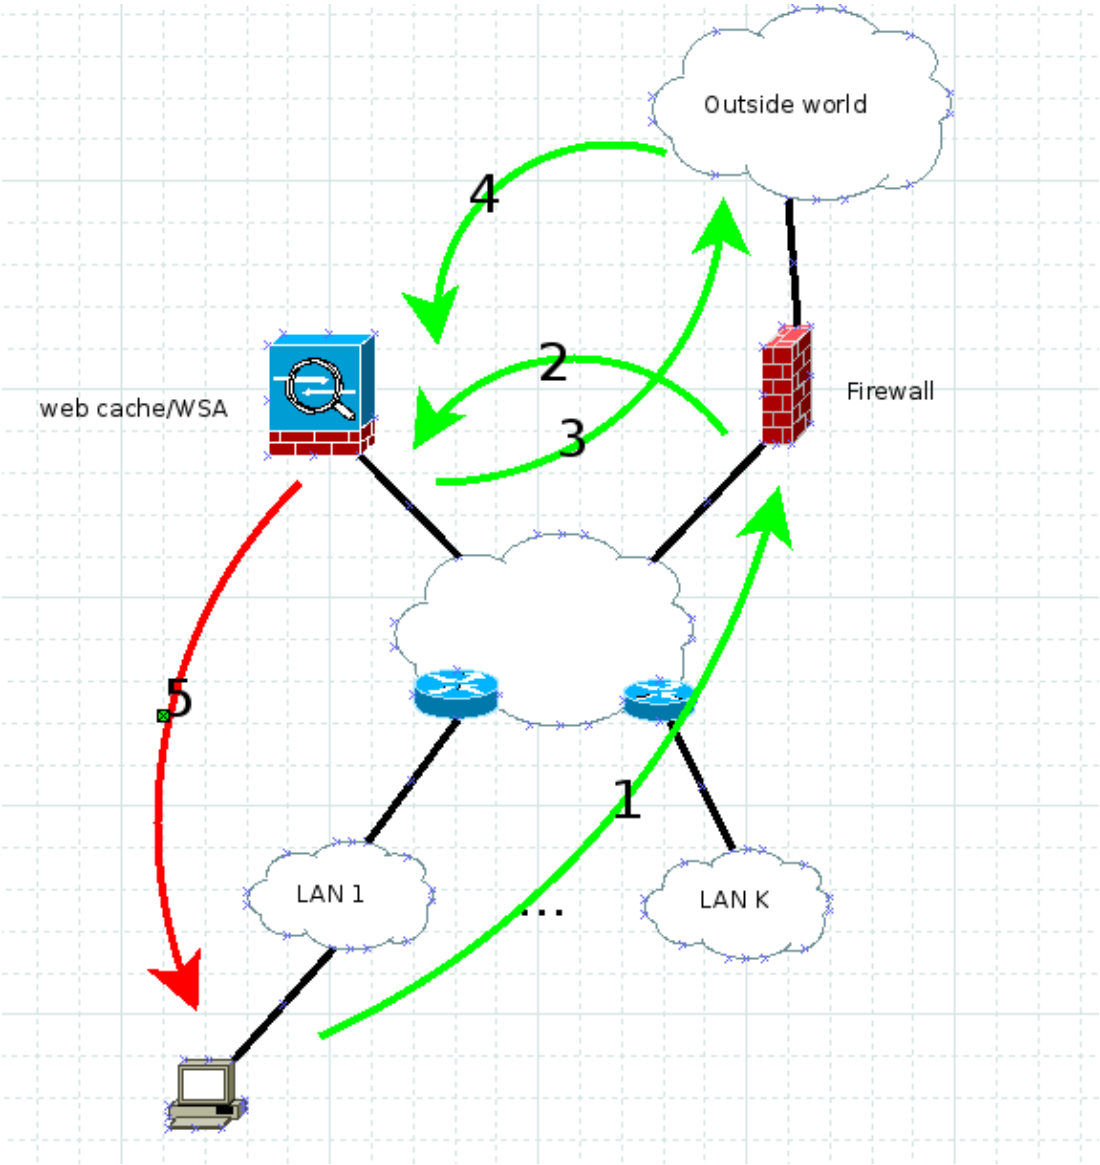

It works like this:

1. Client sends HTTP GET with the IP source, its IP address (client IP address), and the destination server IP address.

2. The firewall or router intercepts the HTTP GET and forwards it via WCCP GRE or pure L2 to web cache/WSA. The source is still the client IP address and the destination is still the web server IP address.

3. The WSA inspects the request and, if it is legitimate, mirrors it towards the web server. Here the destination IP address is the web server IP address and the source IP address might be the WSA or the client, based on whether you enabled client IP address spoofing. For this example, it does not matter because the return traffic in both cases has to hit the WSA.

4. The return traffic is inspected at the WSA.

5. The WSA sends the response to the client with the source IP address, ALWAYS the web server IP address (so the client does not get suspicious), and the destiantion client IP address.

Problem

What happens if one of the routers from the diagram has to fragment the traffic? The WSA puts the DF bit on packet number 5, but it has to be fragmented. The router drops it and tells the sender that fragmentation is needed but the DF bit is set (ICMP type 3 code 4). After all, RFC 1191 has to work now and the sender must lower its packet size.

With WCCP, the source IP address is the web server IP address, so this ICMP never goes to the WSA; rather, it tries to go to the real web server (remember, this router on the bottom is not aware of WCCP). This is how WCCP and path MTU discovery together sometimes break your network design.

Solution

There are four ways to solve this problem:

- Discover the real MTU and then use ***etherconfig*** on the WSA to lower the interface's MTU. Remember that the TCP header is 60, IP is 20, and when you use ICMP, that adds 8 bytes to the IP header.

- Disable path MTU discovery (***pathmtudiscovery*** CLI WSA command). This results in TCP MSS of 536, which might cause a performance problem.

- Change the network so there is no L3 fragmentation between the WSA and clients.

- Use the ***ip tcp mss−adjust 1360*** (or other calculated number) command on each Cisco router on the way on the relevant interfaces.

Additional Notes

While this problem was under investigation, it was discovered that if you set the proxy explicitly into the client for a couple of minutes and then remove it, the issue is resolved for the next four to five hours. This is due to the fact that, in explicit mode, the path MTU discovery mechanism between the WSA and the client works. Once the WSA discovers the path MTU, it stores it along with the discovered TCP MSS onto the

internal table for reference. Apparently this table is refreshed every four to five hours, which renders the solution to not work again after so much time.

Updated: Mar 17, 2015 Document ID: 118843