MTU Adjustment Guide

Confidential

Intro

Radware uses traffic diversion technique in order to protect customers.
In order to avoid routing loops, a GRE tunnel is set up between Radware's designated Scrubbing Center and the Customer.
Due to the nature of the encapsulation, the GRE tunnel will add 24 Bytes of headers to the packet.

This document describes implications of using GRE tunnels and allows the customer to predict and to adjust the network parameters accordingly, while avoiding networking issues.

MTU issue explained

The addition of the 24 Bytes header, could cause interference when using encrypted protocol such as SSL.

HTTP servers that support SSL will usually set the DF bit (Don't Fragment) on packets. Setting this flag will not allow packets to be fragmented when reaching a network node that does not support the traversal of packets with the specific packet's size.

Instead, the router will drop the packet, and respond to the sender with an ICMP Fragmentation Needed (Type 3, Code 4). This process is called Path MTU Discovery (PMTUD).

PMTUD will, in some cases, fail when used in an internet environment since ICMP is usually blocked due to security reasons.

Schematic explanation of the issue

Looking at the following topology could further explain the issue at hand:

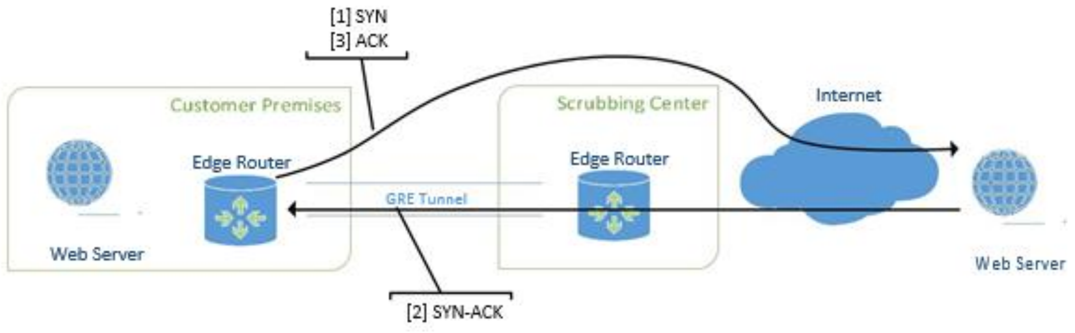

Customer premises: contains a webserver that needs to use an external SSL based service, and is routed via the edge router.

Scrubbing center: where traffic is diverted to, and attacks are currently being mitigated.

Web Server: is an example of an external webserver.

The Process:

1. The Web Server on the customer's premises sends a SYN packet with the server's defined MTU (generally, 1514 Bytes). The SYN packet exists the customer's network via it's default route as it usually does. Note that this packet usually does not traverse the Scrubbing Center as the default mode of operation of Radware's Scrubbing Centers is Asymmetric Ingress Only.

2. The external Webserver receives the SYN packet, checks that the requested packet size is supported, and sends a SYN-ACK back to the requesting Webserver on the Customers' premises.

3. The Webserver on the Customer's premises will then sends an ACK\PUSH-ACK in the same route that the first SYN packet was sent on, thus not traversing the Scrubbing Center.

When a response from the external Web Server reaches the scrubbing center and is bigger than the allowed size on the GRE Tunnel (Default 1360 Bytes) it will be fragmented by the router.

If the packet has a DF bit set, as customary in SSL connections, the packet will be dropped, and the session will not be established as needed.

Solution

To avoid this, Radware recommends to reduce the MSS value at the Customer's Premises Edge Router to the desired value of 1360.

Doing so, will restrict packets' sizes to 1360 Bytes, thus removing the need to fragment packets arriving at the tunnel.

One example of doing so would be using Cisco's mss-adjust feature on the outgoing interface to the internet:

```
ip tcp adjust-mss 1360
```

Setting this will cause the router to examine outgoing SYN packets that go through to the internet, reducing the value of the MSS to a lower value than the maximum value that the GRE tunnel could handle.

This will remove the need to fragment packets and drop non-fragmentable ones.