UDP Traffic through ASA Fails after Primary ISP Link Comes Back Online in a Dual ISP Setup

Contents

Introduction
Before You Begin
Requirements
Components Used
Conventions
Problem
Solution
Related Information

Introduction

If an Adaptive Security Appliance (ASA) has two egress interfaces per destination subnet and the preferred route to a destination is removed from the routing table for some time, User Datagram Protocol (UDP) connections can fail when the preferred route gets re-added to the routing table. TCP connections might also be affected by the problem, but since TCP detects packet loss, these connections are torn down automatically by the endpoints, and re-built using the more optimal routes after the routes change.

This problem can also be seen if a routing protocol is used and a topology change triggers a change in the routing table on the ASA.

Before You Begin

Requirements

In order to encounter this problem, the ASA's routing table must change. This is common with dual ISP links in a redundant fashion or when the ASA is learning routes via an IGP (OSPF, EIGRP, RIP).

This issue occurs when the primary ISP link comes back online or the said IGP sees a reconvergence due to which a less preferred route that was being used by the ASA is replaced with the preferred lower-metric-route. You would then see long-lived connections, such as UDP SIP registrations, GRE, etc, failing once the primary or preferred route is re-installed into the ASA's routing table.

Components Used

The information in this document is based on these hardware and software versions:

- Any Cisco ASA 5500 Series Adaptive Security Appliance
- ASA versions 8.2(5), 8.3(2)12, 8.4(1)1, 8.5(1) and later

Conventions

For more information on document conventions, refer to Cisco Technical Tips Conventions.

Problem

If a routing table entry is removed from the ASA's routing table and there are no routes out of an interface to reach a destination, connections built through the firewall with that foreign destination will be deleted by the ASA. This occurs so that the connections can be built again using a different interface with routing entries for the destination present.

However, if more specific routes are added back to the table, the connections will not be updated to use the new, more specific routes, and will continue to use the less-optimal interface.

For example, consider that the firewall has two interfaces that face the Internet - "outside" and "backup" - and these two routes exist in the ASA's configuration:

```
route outside 0.0.0.0 0.0.0.0 10.1.1.1 1 track 1
route backup 0.0.0.0 0.0.0.0 172.16.1.1 254
```

If both the outside and backup interfaces are "up", then connections built outbound through the firewall will use the outside interface, as it has the preferred metric of 1. If the outside interface is shut down (or the SLA monitoring function that is tracking the route encounters a loss of connectivity to the tracked IP), connections using the outside interface would be torn down and re-built using the backup interface, as the backup interface is the only interface with a route to the destination.

The problem occurs when the outside interface is brought back up or the tracked route becomes the favored route again. The routing table is updated to prefer the original route, but existing connections continue to exist on the ASA and traverse the backup interface and are NOT deleted and recreated on the outside interface with the more-preferred metric. This is because the backup default route still exists in the ASA's interface-specific routing table. The connection continues to use the interface with the less preferred route until the connection is deleted; in the case of UDP, this might be indefinite.

This situation can cause problems with long-lived connections, such as external SIP registrations or other UDP connections.

Solution

In order to address this specific problem, a new feature was added to the ASA that will cause connections to be torn down and rebuilt on a new interface if a more preferred route to the destination is added to the routing table. In order to activate the feature (it is disabled by default), set a non-zero timeout to the **timeout floating-conn** command. This timeout (specified in HH:MM:SS) specifies the time the ASA waits before it tears down the connection once a more preferred route is added back to the routing table:

This is a CLI example of enabling the feature. With this CLI, if a packet is received on an existing connection for which there is now a different, more preferred route to the destination, the connection will be torn down 1 minute later (and rebuilt using the new, more preferred route):

```
ASA# config terminal ASA(config)# timeout floating-conn 0:01:00 ASA(config)# end ASA#
show run timeout timeout conn 1:00:00 half-closed 0:10:00 udp 0:50:00 icmp 0:00:02
timeout sunrpc 0:10:00 h323 0:05:00 h225 1:00:00 mgcp 0:05:00 mgcp-pat 0:05:00
timeout sip 0:30:00 sip_media 0:02:00 sip-invite 0:03:00 sip-disconnect 0:02:00
timeout sip-provisional-media 0:02:00 uauth 0:01:00 absolute timeout tcp-proxy-
reassembly 0:01:00 timeout xlate 0:01:00 timeout pat-xlate 0:00:30 timeout floating-
conn 0:01:00 ASA#
```

This feature is added to the ASA platform in versions 8.2(5), 8.3(2)12, 8.4(1)1, and 8.5(1), including later versions of ASA software.

If you run a version of ASA code that does not implement this feature, a workaround to the issue would be to manually flush the UDP connections that continue to take the less preferred route despite a better route being made available via a **clear local-host** <*IP*> or **clear-conn** <*IP*> .

The command reference lists this new feature under the **timeout** section.

Related Information

- **Technical Support & Documentation - Cisco Systems**