ASA: Inbound Access to NAT Addresses Fails After Upgrade to 8.4(3)

Contents

Introduction

This document provides information about NAT addresses that fail after upgrading the Adaptive Security Appliance (ASA) to version 8.4(3). This document also provides a resolution to this issue.

Prerequisites

Requirements

Readers of this document should have knowledge of these topics.

- Basic understanding of the concept of Address Resolution Protocol (ARP) and proxy ARP

Components Used

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

- Any Cisco ASA 5500 Series Adaptive Security Appliance
- Adaptive Security Appliance version 8.4(3) or later

Conventions

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

Symptoms

Starting with ASA version 8.4(3), the ASA does not respond to ARP requests received on an interface, for IP addresses that are not a part of that interface's IP subnet. Before version 8.4(3), the ASA would respond to ARP requests that were not in the IP subnet of the ASA's interface.

This change can manifest itself immediately after upgrading the ASA to version 8.4(3). In some cases, Internet users cannot connect to the global address of a translated server through the ASA.

This message is displayed if this situation is encountered, and 'debug arp' is enabled on the ASA's CLI:

```
arp-in: Arp packet received from 192.168.10.1 which is in different subnet
than the connected interface 192.168.11.1/255.255.255.0
```

The root cause of this issue is not a bug. See the information below to learn more about potential causes and solutions to the issue.

Conditions / Environment

In order to encounter this situation, the ASA must receive an ARP request for an IP address that matches a global address in a configured NAT translation. The global IP address must reside in an IP subnet that is different from the IP subnet configured on the ASA's interface.

Cause / Problem Description

In order to understand the full ramifications of this issue, it is important to get a complete understanding of how this issue can appear and the best way to mitigate the problem.

These are some instances where this situation can be encountered:

Upstream device has IP routes configured with no next-hop IP address

This is probably the most common cause of this situation. It is due to a non-optimal configuration of an upstream device. It is preferred to configure IP routes such that the next hop of the IP route is an IP address in the same subnet as that interface's address:

```
ip route 10.1.2.0 255.255.255.0 192.168.1.2
```

However, sometimes network administrators configure an interface instead of an IP address as the next hop:

```
ip route 10.1.2.0 255.255.255.0 FastEthernet0/1
```

This causes the router to route traffic destined to the 10.1.2.0/24 network to the FastEthernet0/1 interface, and send an ARP request for the destination IP address in the IP packet. It is assumed that some device will respond to the ARP request, and the router then forwards the packet to the MAC address that was resolved due to the ARP process. The benefits of this type of configuration is that it is very easy to configure and administer. The administrator does not have to explicitly configure a next hop IP address for the route, and they assume that an adjacent device will have proxy-ARP enabled and will respond to the ARP request if it is capable of routing the packets to the destination IP address.

However, there are serious problems with this type of IP route configuration:

- By sending an ARP request to determine the next hop for IP traffic, the router is exposed to problems caused by other devices that might incorrectly respond to that ARP request. The result is traffic can be black-holed when sent to an incorrect device.
- The route will cause the device to send an ARP request for every unique destination address in the packets that match the route. This can cause a large amount of ARP traffic on the subnet and negatively affect performance as well as the memory space required to hold a potentially large amount of ARP entries.
- Because ARP table space is a memory bound resource, an excessive number of entries can negatively impact the router's performance and stabilty.

Therefore, the best practice is to configure all routes with explicit IP next hop addresses and not use routes that have an interface name by itself to identify the outgoing interface. If the interface is needed to tie the route to the egress interface for failover, enter both the egress interface name and the next hop in the static route.

Given the administrative implications for some Cisco customers, an Enhancement Request has been opened in order to make the new secure behavior configurable: Cisco bug ID CSCty95468 (registered customers only) (ENH: Add Command to Allow ARP Cache Entries from Non-Connected Subnets).

Mismatched IP subnet masks on adjacent devices

Mismatched IP subnet masks configured on the ASA's interface and the adjacent device's interface can cause a similar situation. If the adjacent device had a subnet mask that was a supernet (255.255.240.0) of the ASA's interface IP subnet mask (255.255.255.0), the adjacent device will ARP for IP addresses that are not in the ASAs interface IP subnet. Ensure that the subnet masks are correct.

Transparent Mode Implications

Another side effect of this change is the inability to learn MAC addresses from non-directly-connected subnets in Transparent mode. This affects communication in these scenarios:

- The transparent ASA does not have a management IP address configured or the configuration is incorrect.
- The transparent ASA is using secondary subnets on the same segment.

There is no workaround for this issue in Transparent mode other than the downgrade. However, this Enhancement Request has been opened in order to make ASA interoperate with secondary subnets in Transparent mode: Cisco bug ID CSCty49855 (registered customers only) (ENH: Support Non Directly Connected Hosts in MAC Discovery Mechanism).

Resolution

The solution to this problem (in the case that the IP address in question is not in the same layer-3 subnet as the ASA's interface IP) is to make the changes necessary to ensure that devices adjacent to the ASA route traffic directly to the ASA's interface IP address as the next hop device, instead of relying on a device to proxy-ARP on behalf of the IP address.

Related Information

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