ASA Has High CPU Usage Due to a Traffic Loop When VPN Clients Disconnect

Contents

Introduction

This document describes a common issue that occurs when VPN clients disconnect from a Cisco Adaptive Security Appliance (ASA) that runs as a remote access VPN headend. This document also describes the situation where a traffic loop occurs when VPN users disconnect from an ASA firewall. This document does not cover how to configure or set up remote access to the VPN, only the specific situation that arises from certain common routing configurations.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Remote Access VPN configuration on the ASA
- Basic Layer 3 routing concepts

Components Used

The information in this document is based on an ASA Model 5520 that runs ASA code Version 9.1(1).

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Related Products

This document can be used with these hardware and software versions:

- Any ASA model
- Any ASA code version

Background Information

When a user connects to the ASA as a remote access VPN concentrator, the ASA installs a host-based route in the ASA routing table that routes traffic to that VPN client out of the outside interface (towards the Internet). When that user disconnects, the route is removed from the table, and the packets on the inside network (destined to that disconnected VPN user) might be looped between the ASA and an internal routing device.

Another problem is that directed (network) broadcast packets (generated by the removal of the VPN clients) might be forwarded by the ASA as a unicast frame towards the internal network. This might forward it back to the ASA, which causes the packet to be looped until the Time to Live (TTL) expires.

This document explains these issues and shows what configuration techniques can be used in order to prevent the problem.

Problem: Packets Destined for a Disconnected VPN Client Loop Inside Internal Network

When a remote access VPN user disconnects from an ASA firewall, the packets still present on the internal network (destined for those disconnected users) and the assigned IP VPN address might become looped within the internal network. These packet loops might cause the CPU usage on the ASA to increase until the loop stops either due to the IP TTL value in the IP packet header decrementing to 0, or the user reconnects and the IP address is re-assigned to a VPN client.

In order to understand this scenario better, consider this topology:

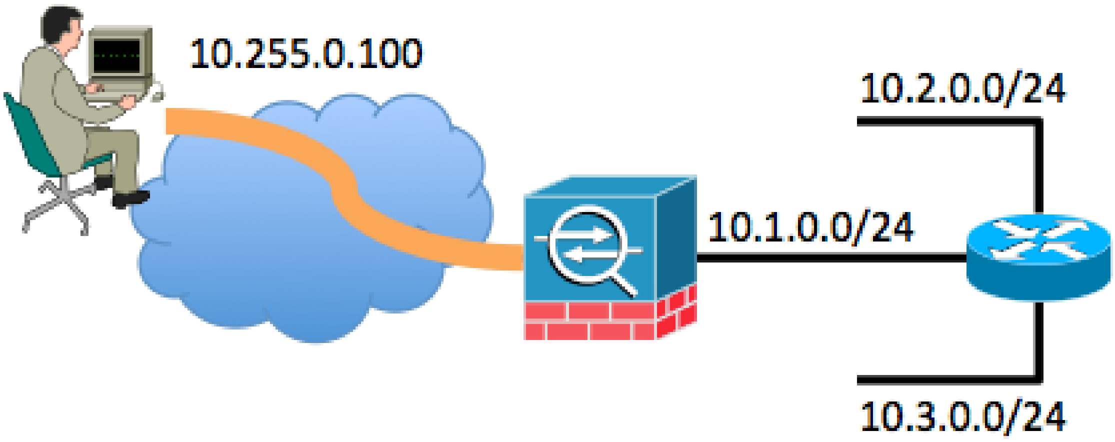

In this example, the remote access client has been assigned the IP address of 10.255.0.100. The ASA in this example is connected to the same inside network segment along with a router. The router has two additional Layer 3 network segments connected to it. The relevant interface (routing) and VPN configurations of the ASA and router are shown in the examples.

ASA configuration highlights are shown in this example:

```
interface GigabitEthernet0/0
nameif outside
security-level 0
ip address 198.51.100.100 255.255.255.0
!
interface GigabitEthernet0/1
nameif inside
security-level 100
ip address 10.1.0.1 255.255.255.0
!
same-security-traffic permit intra-interface
!
ip local pool VPNpool 10.255.0.1-10.255.0.255
!
route outside 0.0.0.0 0.0.0.0 198.51.100.1
route inside 10.0.0.0 255.0.0.0 10.1.0.2
```

Router configuration highlights are shown in this example:

```
interface FastEthernet0
description connected to the inside interface of the ASA G0/1
ip address 10.1.0.2 255.255.255.0
!
interface FastEthernet1
 description connected to network segment
 ip address 10.2.0.1 255.255.255.0
!
interface FastEthernet2
 description connected to other network segment
 ip address 10.3.0.1 255.255.255.0
!
ip route 0.0.0.0 0.0.0.0 10.1.0.1
```

The routing table of the router connected to the inside of the ASA simply has a default route pointed to the ASA inside interface of 10.1.0.1.

While the user is connected via VPN to the ASA, the ASA routing table shows as follows:

```
ASA# show route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default, U - per-user static route, o - ODR
P - periodic downloaded static route
Gateway of last resort is 198.51.100.1 to network 0.0.0.0
S 10.255.0.100 255.255.255.255 [1/0] via 198.51.100.1, outside
S 10.0.0.0 255.0.0.0 [1/0] via 10.1.0.2, inside
C 198.51.100.0 255.255.255.0 is directly connected, outside
C 10.1.0.0 255.255.255.0 is directly connected, inside
S* 0.0.0.0 0.0.0.0 [1/0] via 198.51.100.1, outside
```

The problem occurs when the remote access VPN user disconnects from the VPN. At this point, the host-based route is removed from the ASA routing table. If a host inside the network attempts to send traffic to the VPN client, that traffic is routed to the ASA inside interface by the router. This series of steps occurs:

1. The packet destined to 10.255.0.100 arrives on the inside interface of the ASA.

2. Standard ACL checks are performed.

3. The ASA routing table is checked in order to determine the egress interface for this traffic.

4. The destination of the packet matches the broad 10.0.0.0/8 route that points back out of the inside interface toward the router.

5. The ASA verifies if hair pinning traffic is allowed - it searches for **same-security permit intra-interface** and finds that it is allowed.

6. A connection is built to and from the inside interface and the packet is sent back to the router as a next hop.

7. The router receives a packet destined to 10.255.0.100 on the interface that faces the ASA. The router checks its routing table for a suitable next hop. The router finds that the next hop would be the ASA inside interface, and the packet is sent to the ASA.
8. Return to Step 1.

An example is shown here:

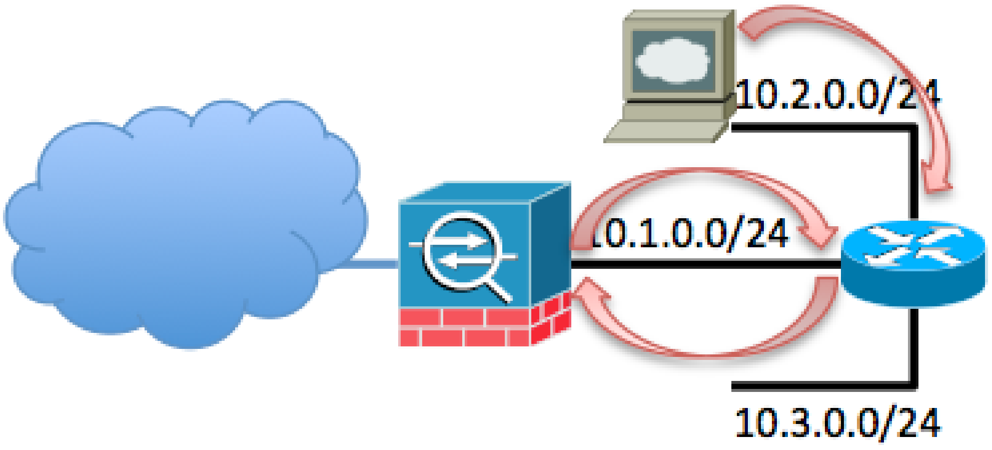

This loop occurs until the TTL of this packet decrements to 0. Note that the ASA Firewall **does not** decrement the TTL value by default when it processes a packet. The router decrements the TTL as it routes the packet. This prevents the occurrence of this loop indefinitely, but this loop does increase the traffic load on the ASA and causes the CPU usage to spike.

Problem: Directed (network) Broadcast Packets Generated by VPN Clients are Looped on an Inside Network

This issue is similar to the first problem.. If a VPN client generates a directed broadcast packet to its assigned IP subnet (10.255.0.255 in the previous example), then that packet might be forwarded as a unicast frame by the ASA to the inside router. The inside router might then forward it back to the ASA, which causes the packet to loop until the TTL expires.

This series of events occur:

1. The VPN client machine generates a packet destined to the network broadcast address 10.255.0.255, and the packet arrives at the ASA.
2. The ASA treats this packet as a unicast frame (due to the routing table) and forwards it to the inside router.

3. The inside router, which also treats the packet as a unicast frame, decrements the TTL of the packet and forwards it back to the ASA.

4. The process repeats until the TTL of the packet is reduced to 0.

Solutions to the Problem

There are several potential solutions to this issue. Depending on the network topology and the specific situation, one solution might be easier to implement than another.

Solution 1- Static Route for Null0 Interface (ASA Version 9.2.1 and Later)

When you send traffic to a **Null0** interface, it causes the packets destined to the specified network to be dropped. This feature is useful when you configure Remotely Triggered Black Hole (RTBH) for Border Gateway Protocol (BGP). In this situation, if you configure a route to Null0 for the remote access client subnet, it forces the ASA to drop traffic destined to hosts in that subnet if a more specific route (provided by Reverse Route Injection) is not present.

```
route Null0 10.255.0.0 255.255.255.0
```

Solution 2 - Use a Different IP Pool for VPN Clients

This solution is to assign the remote VPN users an IP address that does not overlap with any internal network subnet. This would would prevent the ASA from forwarding packets destined to that VPN subnet back to the inside router if the VPN user was not connected.

Solution 3 - Make the ASA Routing Table More Specific for Internal Routes

This solution is to ensure the routing table of the ASA does not have any very broad routes that overlap with the VPN IP pool. For this specific network example, remove the 10.0.0.0/8 route from the ASA and configure more specific static routes for the subnets that reside off of the inside interface. Dependent upon the number of subnets and the network topology, this might be a large number of static routes and it might not be possible.

Solution 4 - Add a More Specific Route for the VPN Subnet Back Out of the Outside Interface

This solution is more complicated that the others that are described in this document. Cisco recommends that you attempt to use the other solutions first due to the situation that is described in the Note later in this section. This solution is to prevent the ASA from forwarding IP packets sourced from the VPN IP subnet back to the internal router; you can do this if you add a more specific route for the VPN subnet out of the outside interface. Since this IP subnet is reserved for outside VPN users, packets with a source IP address from this VPN IP subnet should never arrive inbound on the ASA inside interface. The easiest way to achieve this is to add a route for the remote access VPN IP Pool out of the outside interface with a next hop IP address of the upstream ISP router.

In this network topology example, that route would look like this:

```
route outside 10.255.0.0 255.255.255.0 198.51.100.1
```

In addition to this route, add the **ip verify reverse-path inside** command in order to cause the ASA to drop any packets received inbound on the inside interface sourced from the VPN IP subnet due to the more preferred route that exists on the outside interface:

```
ip verify reverse-path inside
```

After these commands are implemeted, the ASA routing table looks similar to this when the user is connected:

```
ASA# show route

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default, U - per-user static route, o - ODR
P - periodic downloaded static route

Gateway of last resort is 198.51.100.1 to network 0.0.0.0

S 10.255.0.100 255.255.255.255 [1/0] via 198.51.100.1, outside
S 10.0.0.0 255.0.0.0 [1/0] via 10.1.0.2, inside
S 10.255.0.0 255.255.255.0 [1/0] via 198.51.100.1, outside
C 198.51.100.0 255.255.255.0 is directly connected, outside
C 10.1.0.0 255.255.255.0 is directly connected, inside
S* 0.0.0.0 0.0.0.0 [1/0] via 198.51.100.1, outside
```

When the VPN client is connected, the host-based route to that VPN IP address is present in the table and is preferred. When the VPN client disconnects, traffic sourced from that client IP address that arrives on the inside interface is checked against the routing table and dropped due to the **ip verify reverse-path inside** command.

If the VPN client generates a directed network broadcast to the VPN IP subnet, then that packet is forwarded to the inside router and forwarded by the router back to the ASA, where it is dropped due to the **ip verify reverse-path inside** command.

> **Note**: After this solution is implemented, if the **same-security permit intra-interface** command is present in the configuration and the access policies permit it, traffic sourced from a VPN user destined to an IP address in the VPN IP pool for a user that is not connected might be routed back out of the outside interface in clear-text. This is a rare situation and can be mitigated with the use of vpn-filters within the VPN policy. This situation only occurs if the **same-security permit intra-interface** command is present in the configuration of the ASA.
>
> Likewise, if internal hosts generate traffic destined to an IP address in the VPN pool and that IP address is not assigned to a remote VPN user, that traffic might egress the outside of the ASA in clear-text.