corero

Corero Network Security

Remote Workers and the Rise of OpenVPN Amplification DDoS Attacks

AUTHORS

Huy Nguyen
Andrei Tsarou

corero

Contents

corero

Overview

Since the widespread lockdowns resulting from the COVID-19 pandemic, millions of people worldwide have begun working from home and many of them are using virtual private networks (VPNs) to connect to their corporate office networks.

Corero has observed an increase in the number of Distributed Denial of Service (DDoS) attacks and targets across our customer base since the latter part of Q1 2020, which we believe correlates with the increase in remote working.

OpenVPN is a popular open source application, allowing companies or individuals to extend their private network in a secure and reliable manner. However, proof-of-concept source code for a Denial of Service attack exploiting an OpenVPN reflection/amplification vulnerability was posted on the Internet as far back as 2017 but has pretty much laid dormant until recently. This added another new weapon to the cybercriminal's arsenal.

In October 2019, a more significant reflection/amplification vulnerability was found in SoftEther – a derivative version of OpenVPN. The damaging impact of that vulnerability has become more apparent now, during the COVID-19 pandemic, due to the increased number of remote workers which appears to be driving the continued increase in the deployment of OpenVPN servers.

How OpenVPN is used for DDoS attack

In this section we are going to focus on the message exchange between an OpenVPN client and server during a connection handshake, and the configuration settings on server side that make the software vulnerable to being leveraged for DDoS reflection attacks.

FIGURE 1

OpenVPN Session establishment

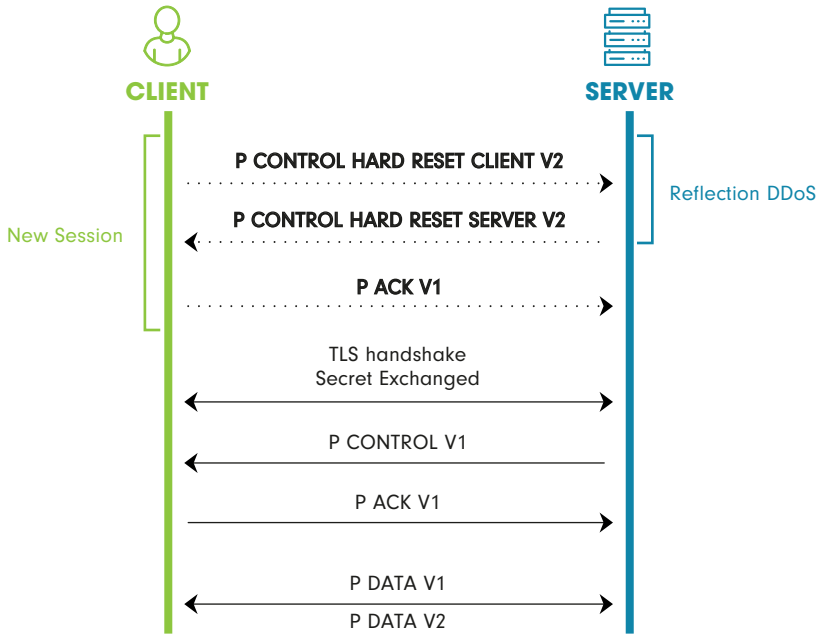

In modern versions of OpenVPN, when the client wants to initiate a new session, it sends a message containing P_CONTROL_HARD_RESET_CLIENT_V2 (Opcode 0x07) to the server. The server then replies to the client with the message P_CONTROL_HARD_RESET_SERVER_V2 (Opcode 0x08). The server then expects to receive P_ACK_V1 (Opcode 0x05) to complete the handshake.

(Refer here for more information about message types and opcode)

FIGURE 2

Reflection with OpenVPN

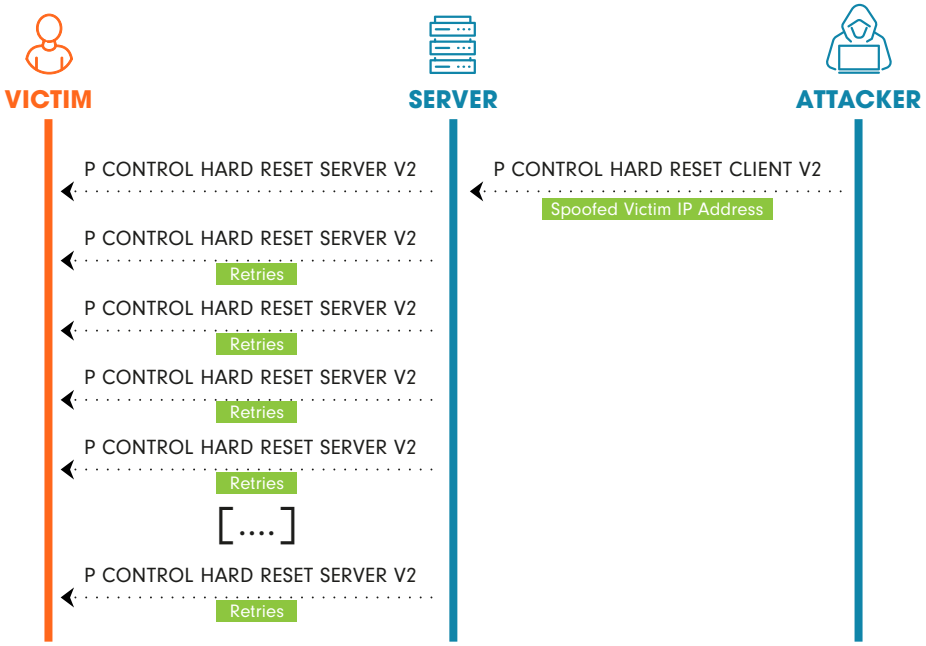

In the case where the OpenVPN server is not receiving the P_ACK_V1, the server assumes the packet got lost, and will resend the P_CONTROL_HARD_RESET_SERVER_V2 at an interval defined by the reliable_send_timeout parameter. With OpenVPN servers using the default configuration, the exponential increase timeout would be used, and it could be resent up to 5 times before reaching the 60 seconds limit of the tls_process timeout.

However, in some distributions of OpenVPN — for instance, the SoftEther version prior to release 5.01.9672 — the P_CONTROL_HARD_RESET_SERVER_V2 will be resent every "OPENVPN_CONTROL_PACKET_RESEND_INTERVAL" (0.5 seconds by default), until OPENVPN_NEW_SESSION_DEADLINE_TIMEOUT is reached (30 seconds by default). The resend timeout and deadline timeout may vary on the servers deployed in production. Assuming default values are used, the resulting amplification factor is 60.

Here is an example of the server response, parsed by Wireshark, from a UDP payload.

FIGURE 3

OpenVPN server response during session establishment handshake

```
▾ OpenVPN Protocol
  ▾ Type: 0x40 [opcode/key_id]
      0100 0... = Opcode: P_CONTROL_HARD_RESET_SERVER_V2 (0x08)
      .... .000 = Key ID: 0
    Session ID: 11068696068725122376
    Message Packet-ID Array Length: 0
    Message Packet-ID: 0
```

As evidence that this vulnerability is now being used by attackers, we have examined ICMP server unreachable responses in the wild. In the case where the server being targeted for reflection or amplification is unreachable, or does not exist, (for example due to a firewall blocking the OpenVPN port), the victim often receives a so-called "failed reflector" packet. This is the original packet, which the attacker sent to the server, encapsulated inside the request section an ICMP "destination unreachable" response packet. This effectively proves the attacker created a reflection attack using legitimate servers as reflectors, rather than just directly generating bogus OpenVPN server replies with random server IPs.

Here is an example from a failed reflection attempt: The original P_CONTROL_HARD_RESET_CLIENT_V2 request did not make it to the server and was returned encapsulated in an ICMP Type 3 Destination Unreachable packet. Because the attacker was spoofing the victim's address, as part of the reflection attempt, this packet is forwarded to the victim allowing us to collect evidence of the attack.

FIGURE 4

ICMP Type 3 failed reflector response

```
▾ Internet Control Message Protocol, Opcode: P_CONTROL_HARD_RESET_CLIENT_V2, Key ID: 0
    Type 3: (Destination unreachable)
    Code 3: (Port unreachable)
    Checksum: 0xae2c [correct]
    [Checksum Status: Good]
    Unused: 00000000
                                          victim IP as source
  ▸ Internet Protocol Version 4, Src: 201.23.34.234, 127.67.276.363 2x7
  ▾ User Datagram Protocol, Src Port:53499, Dst Port: 1194
      Source Port: 53499
      Destination Port: 1194
      Length: 22
      [Checksum: [missing]]
      [Checksum Status: Not present]
      [Strem index: 0]
  ▾ OpenVPN Protocol
    ▾ Type: 0x38 [opcode/key_id]
        0011 1... = Opcode: P_CONTROL_HARD_RESET_CLIENT_V2 (0X07)
        .... .000 = key ID: 0
      Session ID: 7647933796043154430
      Message Packet-ID Array Length: 0
      Message Packet-ID: 0
```

Potential Risk

A simple search for OpenVPN default port 1194 on shodan.io shows how many potential reflectors are out there. In March 2020, the result returned approximately 827K servers (which is still growing by approximately 10K new servers a week as of August 2020), more than enough to launch a very powerful volumetric DDoS attack.

While many popular reflection attacks (Including DNS amplification, NTP reflection, Memcached and reflective CLDAP) generate large, often fragmented, packets, the size of the replies the victim gets from OpenVPN reflectors is relatively small – usually 60 to 72 Bytes. However, the amplification factor multiplied by number of available reflectors can generate enormous packet rates, which could easily result in a Denial of Service condition for many applications, even those that are hosted on a public cloud.

So far, Corero researchers have observed OpenVPN reflection attacks reaching 30Gbps.

FIGURE 5

Shodan search for port 1194

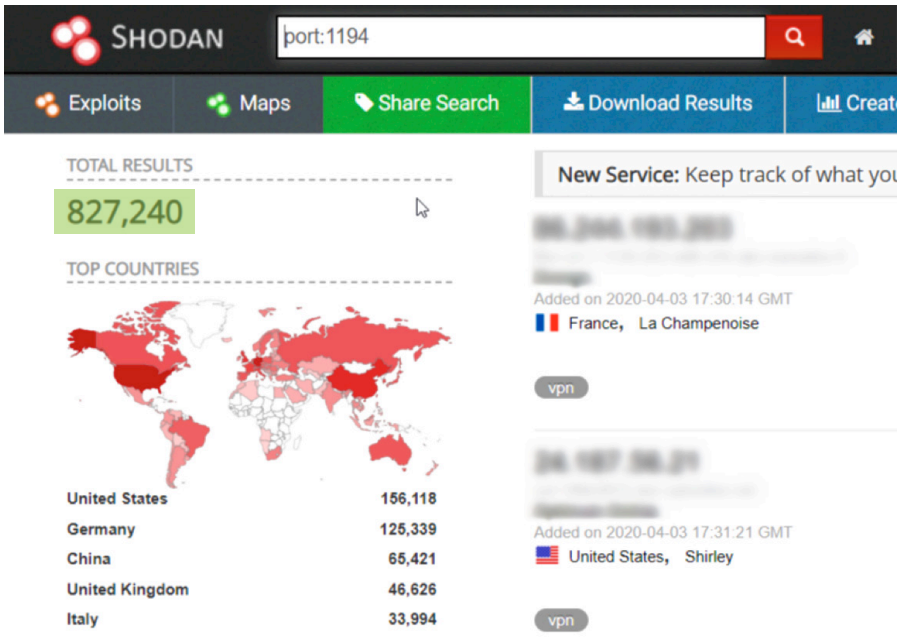

Additional damage may also result from false-positives during the mitigation of an OpenVPN reflection DDoS attack, if legitimate traffic to an OpenVPN port is blocked, for example by a crude DDoS defense system that is rate-limiting traffic after it reaches a certain threshold.

The Corero SmartWall avoids such problems, by selectively blocking just the malicious reflected OpenVPN traffic while allowing legitimate OpenVPN traffic through.

The Analysis

The chart below represents the percentage of total number of OpenVPN reflection attacks exceeding 10Kpps or 10Mbps per month for the period from January 1, 2019 to June 30, 2020. We see a clear elevation in the number of attacks since the beginning of the year.

FIGURE 6

Number of OpenVPN reflection attacks between January 1, 2019 and June 30, 2020

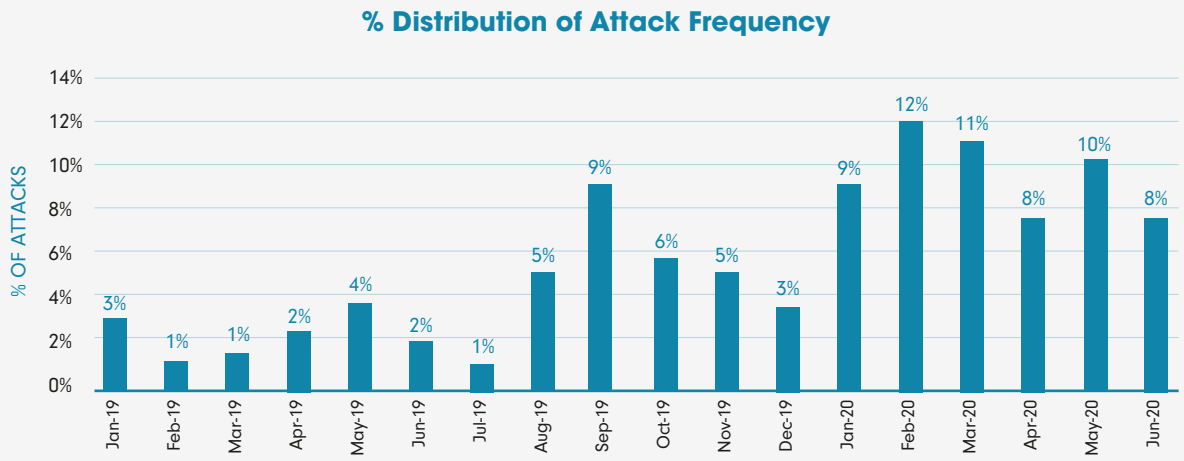

From Corero's observations, the increase in use of OpenVPN reflection tends to be for single vector rather than multi-vector attacks.

FIGURE 7

Number of single vector OpenVPN reflection attacks (blue) vs multi-vector attack (green) OpenVPN reflection attacks, between January 1, 2019 and June 30, 2020

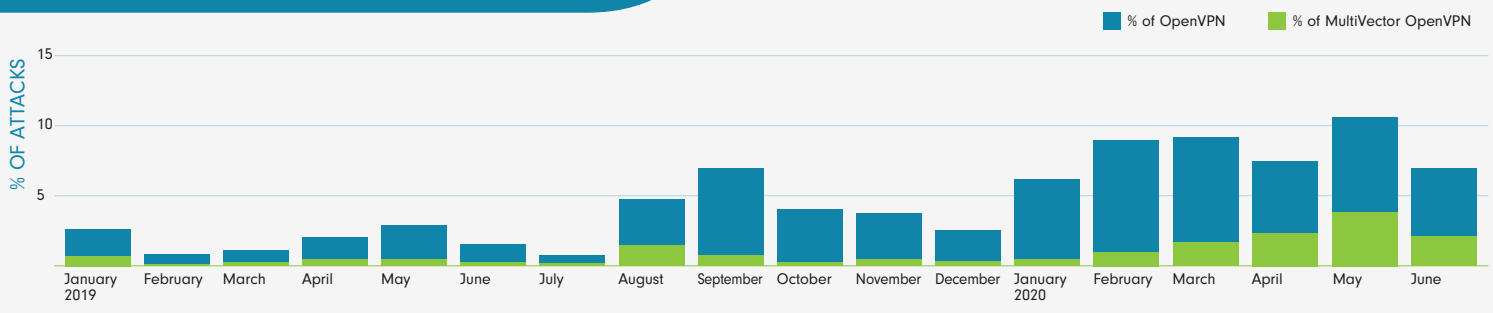

Analysis of reflected packets indicates that the same OpenVPN server often has more than one active session targeting the victim. We suspect this may cause collateral damage on OpenVPN servers that are used as reflectors, as they become heavily loaded, especially if the attacker uses the same server to attack multiple targets at the same time. This can result in indirect collateral damage to the legitimate users of OpenVPN servers that are being used for reflection DDoS attacks.

Mitigation

To deliver its industry leading DDoS protection, Corero developed a patented, proprietary, heuristic-based detection and mitigation mechanism called Smart-Rules, in addition to the surgically accurate, exact-match Flex-Rules. The Smart-Rules continuously inspect a broad range of packets and their associated attributes, looking for those which exhibit specific traits, or indicators, which identify them as potentially being part of a DDoS attack. When repeated packets are seen with the same suspicious characteristics, this enables them to be accurately identified as part of a DDoS attack and automatically blocked, even if that specific packet type has never been seen before. This allows Corero solutions to detect and mitigate the majority of attacks automatically and surgically, without affecting legitimate traffic, based on traffic behavior or payload pattern.

Reflective OpenVPN attacks normally originate from source port 1194. Based on analysis across the Corero customer base, the vast majority of observed reflected packets have UDP length 22 Bytes and contain the same Remote Session ID hex value 6a22eb445adb63fe. We also see the same value used as the Session ID in failed reflectors returned inside ICMP "destination unreachable" packets. This suggests that an attack tool is being used to stimulate these reflectors and generate the attacks.

FIGURE 123

Request packet with Session ID

```
▾ OpenVPN Protocol
  ▾ Type: 0x38 [opcode/key_id]
        0011 1... = Opcode: P_CONTROL_HARD_RESET_CLIENT_V2 (0x07)
        .... .000 = Key ID: 0
    Session ID: 7647933796043154430
    Message Packet-ID Array Length: 0
    Message Packet-ID: 0
```

```
0000  bc ca b5 ef 03 47 44 85   00 de 4d dd 08 00 45 00    .....GD. ..M...E.
0010  00 2a f9 dc 00 00 80 11   aa 51 0a 00 00 8d 3e 63    .*...... .....›c
0020  4d 4d e3 75 04 aa 00 16   08 51 38 6a 22 eb 44 5a    MM.u.... .Q8j".DZ
0030  db 63 fe 00 00 00 00 00                             .c......
```

FIGURE 123

Server response packet with Remote Session ID

```
▾ OpenVPN Protocol
  ▾ Type: 0x40 [opcode/key_id]
        0100 0... = Opcode: P_CONTROL_HARD_RESET_SERVER_V2 (0x08)
        .... .000 = Key ID: 0
    Session ID: 4846841803552123662
    Message Packet-ID Array Length: 1
  ▾ Packet-ID Array
      Message Packet-ID Array Element: 0
    Remote Session ID: 7647933796043154430
    Message Packet-ID: 0
```

```
0000  44 85 00 de 4d dd bc ca   b5 ef 03 47 08 00 45 00    D...M... ...G..E.
0010  00 36 17 28 00 00 71 11   9c 52 3e 63 4d 4d 0a 00    .6.(..q. .R›cMM..
0020  00 8d 04 aa e3 75 00 22   ad 9c 40 43 43 70 f0 cf    .....u." ..@CCp..
0030  3c eb 0e 01 00 00 00 00   6a 22 eb 44 5a db 63 fe    ‹....... j".DZ.c.
0040  00 00 00 00                                         ....
```

Looking at the two packets above, we can see the attacker was sending out requests with Opcode 0x38 and the bogus Session ID 6a22eb445adb63fe to the reflectors. In this case, the first response packet from the reflectors would contain the Remote Session ID, which related to the original request with same ID value 6a22eb445adb63fe. An attacker would need to send this request to as many reflectors as possible, to launch an effective reflection DDoS attack on a victim.

How an OpenVPN Amplification attack is blocked by Corero Smart-Rules

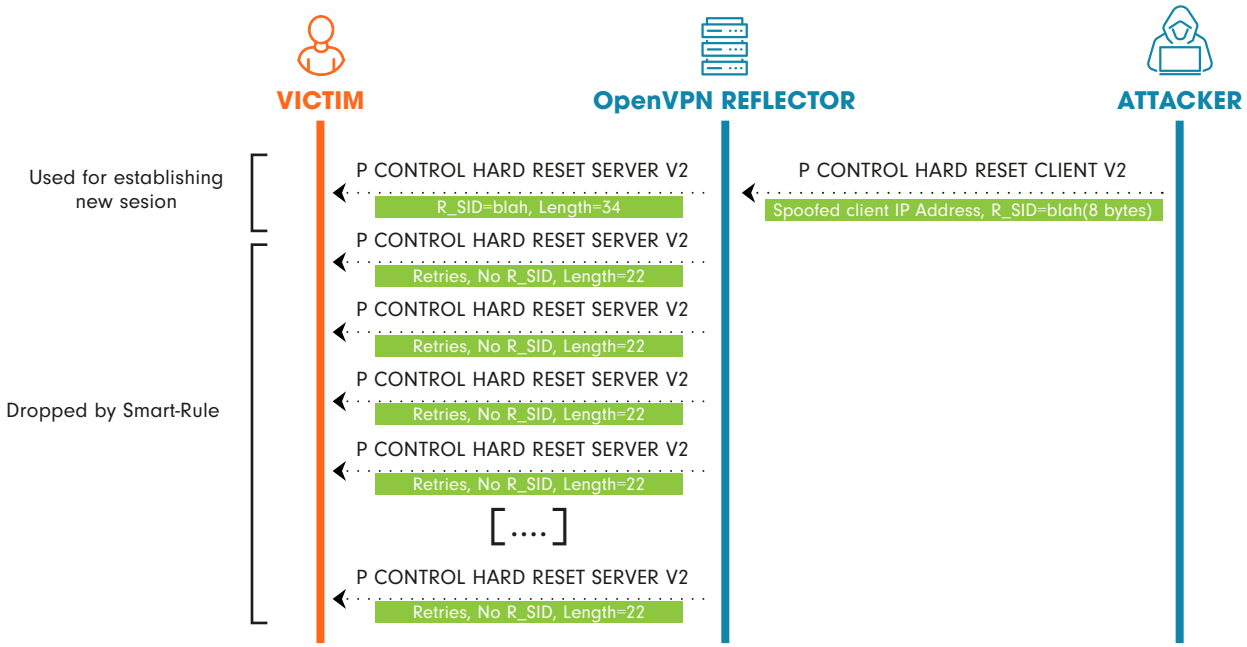

Fortunately, because the retry packets in a single reflection/amplification exhibit the same unique patterns, all the packets from the retries would trigger anomaly thresholds and be blocked automatically by Smart-Rules, while the first packets in the sequence of larger size would still go through and allow legitimate OpenVPN traffic to successfully establish a new session. This means that, during attack time, all retry packets from reflectors and around 98% of bad traffic (59 out of 60 packets) will be blocked, automatically.

How can a Corero Flex-Rule be used to block additional attack traffic?

A Flex-Rules can be used to further increase accuracy, by mitigating the additional 2% of leakage traffic, without impacting good traffic. With some attack samples, identified using Corero's SecureWatch Analytics tool, you will see a pattern with the same Remote Session ID across many source IP addresses; this situation should not happen in normal OpenVPN traffic. By using a Flex-Rule to block all packets with that same Remote Session ID, most of the bad traffic would be mitigated and legitimate users could still establish a new session without any problem.

The Importance of Comprehensive Visibility into DDoS Attacks

In addition to instant real-time mitigation by SmartWall, Corero offers SecureWatch® Analytics, a powerful web-GUI security analytics application that delivers comprehensive and easy-to-read security dashboards for traffic visibility, attack analysis, reporting and alerting. This analytics portal gives Hosting Providers, Service Providers and Enterprises a window into DDoS attacks targeting their Internet-facing services. The real-time security engineered dashboards provide industry leading visibility into an organization's network and security activity for rapid response in combating these threats. Additionally, SecureWatch Analytics supports archived security event data to enable historical forensic analysis and compliance reporting.

As remote workers continue to increase the use of VPNs during and after the COVID-19 pandemic, cybersecurity teams in across all types of organizations should be vigilant and deploy modern-day DDoS mitigation protection to guard against DDoS attacks.

For over a decade, Corero has been providing state-of-the-art, highly effective, real-time automatic DDoS protection solutions for enterprise, hosting and service provider customers around the world. Our SmartWall® DDoS mitigation solutions protect on-premise, cloud, virtual and hybrid environments.

For more on Corero's diverse deployment models, click here. If you'd like to learn more, please contact us.

US HEADQUARTERS

Corero Network Security Inc.
293 Boston Post Road West, Suite 310
Marlborough, MA 01752
Tel: +1 978 212 1500
Email: info@corero.com

EMEA HEADQUARTERS

Corero Network Security (UK) Ltd.
St Mary's Court, The Broadway,
Amersham, Buckinghamshire, HP7 0UT, UK
Tel: +44 (0) 1494 590404
Email: info_uk@corero.com