New VPNFilter malware targets at least 500K networking devices worldwide

blog.talosintelligence.com/2018/05/VPNFilter.html

Intro

For several months, Talos has been working with public- and private-sector threat intelligence partners and law enforcement in researching an advanced, likely state-sponsored or state-affiliated actor's widespread use of a sophisticated modular malware system we call "VPNFilter." We have not completed our research, but recent events have convinced us that the correct way forward is to now share our findings so that affected parties can take the appropriate action to defend themselves. In particular, the code of this

malware overlaps with versions of the BlackEnergy malware — which was responsible for multiple large-scale attacks that targeted devices in Ukraine. While this isn't definitive by any means, we have also observed VPNFilter, a potentially destructive malware, actively infecting Ukrainian hosts at an alarming rate, utilizing a command and control (C2) infrastructure dedicated to that country. Weighing these factors together, we felt it was best to publish our findings so far prior to completing our research. Publishing early means that we don't yet have all the answers — we may not even have all the questions — so this blog represents our findings as of today, and we will update our findings as we continue our investigation.

Both the scale and the capability of this operation are concerning. Working with our partners, we estimate the number of infected devices to be at least 500,000 in at least 54 countries. The known devices affected by VPNFilter are Linksys, MikroTik, NETGEAR and TP-Link networking equipment in the small and home office (SOHO) space, as well at QNAP network-attached storage (NAS) devices. No other vendors, including Cisco, have been observed as infected by VPNFilter, but our research continues. The behavior of this malware on networking equipment is particularly concerning, as components of the VPNFilter malware allows for theft of website credentials and monitoring of Modbus SCADA protocols. Lastly, the malware has a destructive capability that can render an infected device unusable, which can be triggered on individual victim machines or en masse, and has the potential of cutting off internet access for hundreds of thousands of victims worldwide.

The type of devices targeted by this actor are difficult to defend. They are frequently on the perimeter of the network, with no intrusion protection system (IPS) in place, and typically do not have an available host-based protection system such as an anti-virus (AV) package. We are unsure of the particular exploit used in any given case, but most devices targeted, particularly in older versions, have known public exploits or default credentials that make compromise relatively straightforward. All of this has contributed to the quiet growth of this threat since at least 2016.

This post provides the technical findings you would normally see in a Talos blog. In addition, we will detail some thoughts on the tradecraft behind this threat, using our findings and the background of our analysts, to discuss the possible thought process and decisions made by the actor. We will also discuss how to defend against this threat and how to handle a device that may be infected. Finally, we will share the IOCs that we have observed to this point, although we are confident there are more that we have not seen.
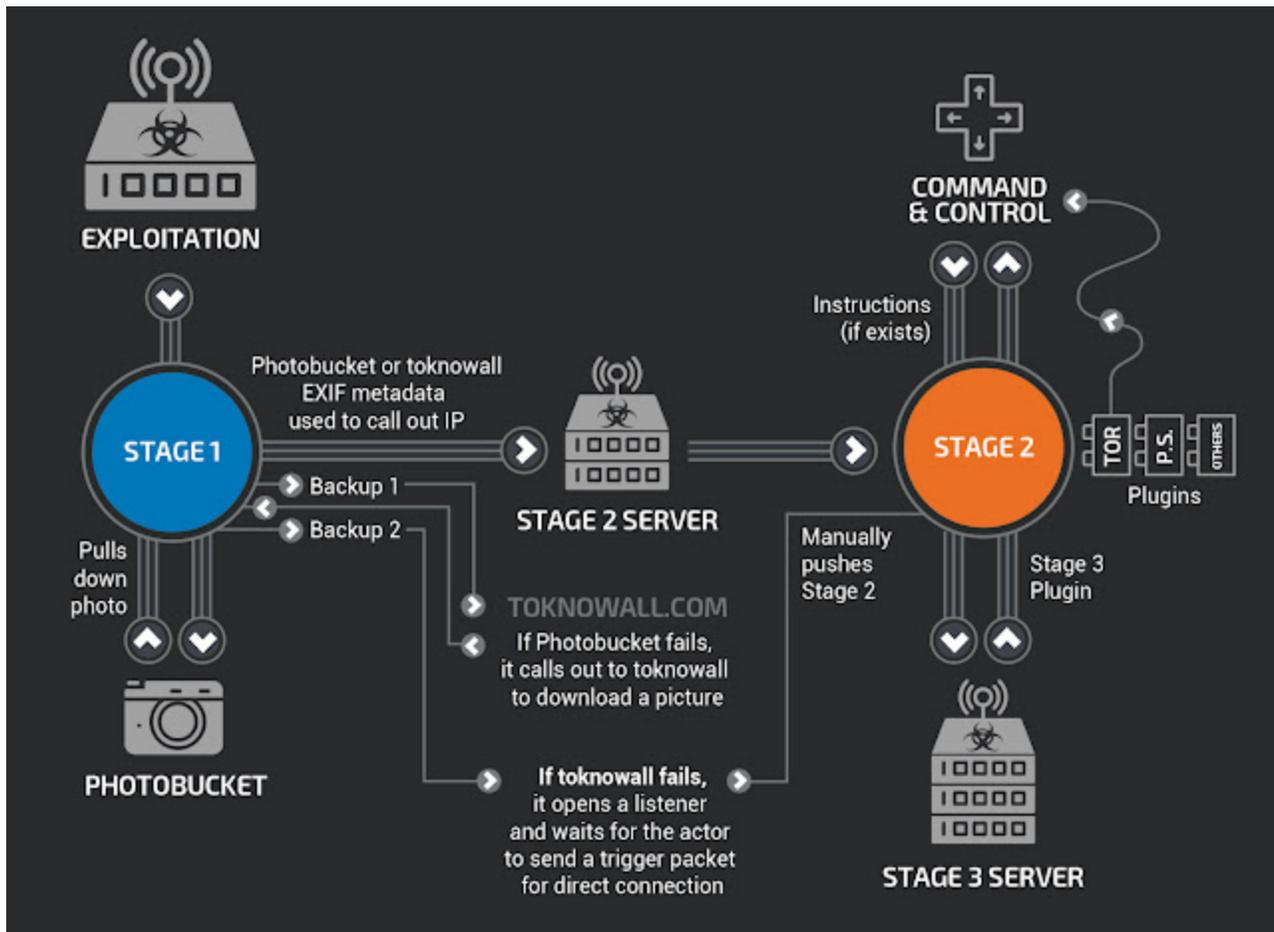
Brief technical breakdown

The VPNFilter malware is a multi-stage, modular platform with versatile capabilities to support both intelligence-collection and destructive cyber attack operations.

The stage 1 malware persists through a reboot, which sets it apart from most other malware that targets internet-of-things devices because malware normally does not survive a reboot of the device. The main purpose of stage 1 is to gain a persistent foothold and enable the deployment of the stage 2 malware. Stage 1 utilizes multiple redundant command and control (C2) mechanisms to discover the IP address of the current stage 2 deployment server, making this malware extremely robust and capable of dealing with unpredictable C2 infrastructure changes.

The stage 2 malware, which does not persist through a reboot, possesses capabilities that we have come to expect in a workhorse intelligence-collection platform, such as file collection, command execution, data exfiltration and device management. However, some versions of stage 2 also possess a self-destruct capability that overwrites a critical portion of the device's firmware and reboots the device, rendering it unusable. Based on the actor's demonstrated knowledge of these devices, and the existing capability in some stage 2 versions, we assess with high confidence that the actor could deploy this self-destruct command to most devices that it controls, regardless of whether the command is built into the stage 2 malware.

In addition, there are multiple stage 3 modules that serve as plugins for the stage 2 malware. These plugins provide stage 2 with additional functionality. As of this writing, we are aware of two plugin modules: a packet sniffer for collecting traffic that passes through the device, including theft of website credentials and monitoring of Modbus SCADA protocols, and a communications module that allows stage 2 to communicate over Tor. We assess with high confidence that several other plugin modules exist, but we have yet to discover them.

Diagram labels: EXPLOITATION; STAGE 1; Photobucket or toknowall EXIF metadata used to call out IP; Backup 1; Backup 2; Pulls down photo; PHOTOBUCKET; STAGE 2 SERVER; TOKNOWALL.COM; If Photobucket fails, it calls out to toknowall to download a picture; If toknowall fails, it opens a listener and waits for the actor to send a trigger packet for direct connection; COMMAND & CONTROL; Instructions (if exists); STAGE 2; TOR; P.S.; OTHERS; Plugins; Manually pushes Stage 2; Stage 3 Plugin; STAGE 3 SERVER

Tradecraft discussion

We assess with high confidence that this malware is used to create an expansive, hard-to-attribute infrastructure that can be used to serve multiple operational needs of the threat actor. Since the affected devices are legitimately owned by businesses or individuals, malicious activity conducted from infected devices could be mistakenly attributed to those who were actually victims of the actor. The capabilities built into the various stages and plugins of the malware are extremely versatile and would enable the actor to take advantage of devices in multiple ways.

Advanced threat actors, including nation-states, will try to make attribution of their cyber activities extremely difficult, unless it is in their interest for it to be openly known that they conducted a specific act. To this end, advanced threat actors use multiple techniques, including co-opting infrastructure owned by someone else to conduct their operations. The actor could easily use devices infected with this malware as hop points before connecting to their final victim in order to obfuscate their true point of origin.

The malware can also be leveraged to collect data that flows through the device. This could be for straightforward data-collection purposes, or to assess the potential value of the network that the device serves. If the network was deemed as having information of potential interest to the threat actor, they may choose to continue collecting content that passes

through the device or to propagate into the connected network for data collection. At the time of this posting, we have not been able to acquire a third-stage plugin that would enable further exploitation of the network served by the device. However, we have seen indications that it does exist, and we assess that it is highly likely that such an advanced actor would naturally include that capability in malware that is this modular.

Finally, this malware could be used to conduct a large-scale destructive attack by using the "kill" command, which would render some or all of the physical devices unusable. This command is present in many of the stage 2 samples we've observed, but could also be triggered by utilizing the "exec" command available in all stage 2 samples. In most cases, this action is unrecoverable by most victims, requiring technical capabilities, know-how, or tools that no consumer should be expected to have. We are deeply concerned about this capability, and it is one of the driving reasons we have been quietly researching this threat over the past few months.

Observed activities of concern

As we have researched this threat, we have put into place monitoring and scanning to gain an understanding of the scope of this threat and the behaviors of infected devices. Our analysis has shown that this is a global, broadly deployed threat that is actively seeking to increase its footprint. While our research continues, we have also observed activity potentially associated with this actor that indicates possible data exfiltration activity.

In early May, we observed infected devices conducting TCP scans on ports 23, 80, 2000 and 8080. These ports are indicative of scanning for additional Mikrotik and QNAP NAS devices, which can be found using these ports. These scans targeted devices in more than 100 countries.

We also used our telemetry to discover potentially infected devices globally. We evaluated their collective behavior to try and identify additional features of the C2 infrastructure. Many of these victim IPs appeared to demonstrate behavior that strongly indicated data exfiltration.

Finally, on May 8, we observed a sharp spike in VPNFilter infection activity. Almost all of the newly acquired victims were located in Ukraine. Also of note, a majority of Ukrainian infections shared a separate stage 2 C2 infrastructure from the rest of the world, on IP 46.151.209[.]33. By this point, we were aware of the code overlap between BlackEnergy and VPNFilter and that the timing of previous attacks in Ukraine suggested that an attack could be imminent. Given each of these factors, and in consultation with our partners, we immediately began the process to go public before completing our research.

As we continued to move forward with the public disclosure, we observed another substantial increase in newly acquired VPNFilter victims focused in Ukraine on May 17. This continued to drive our decision to publish our research as soon as possible.
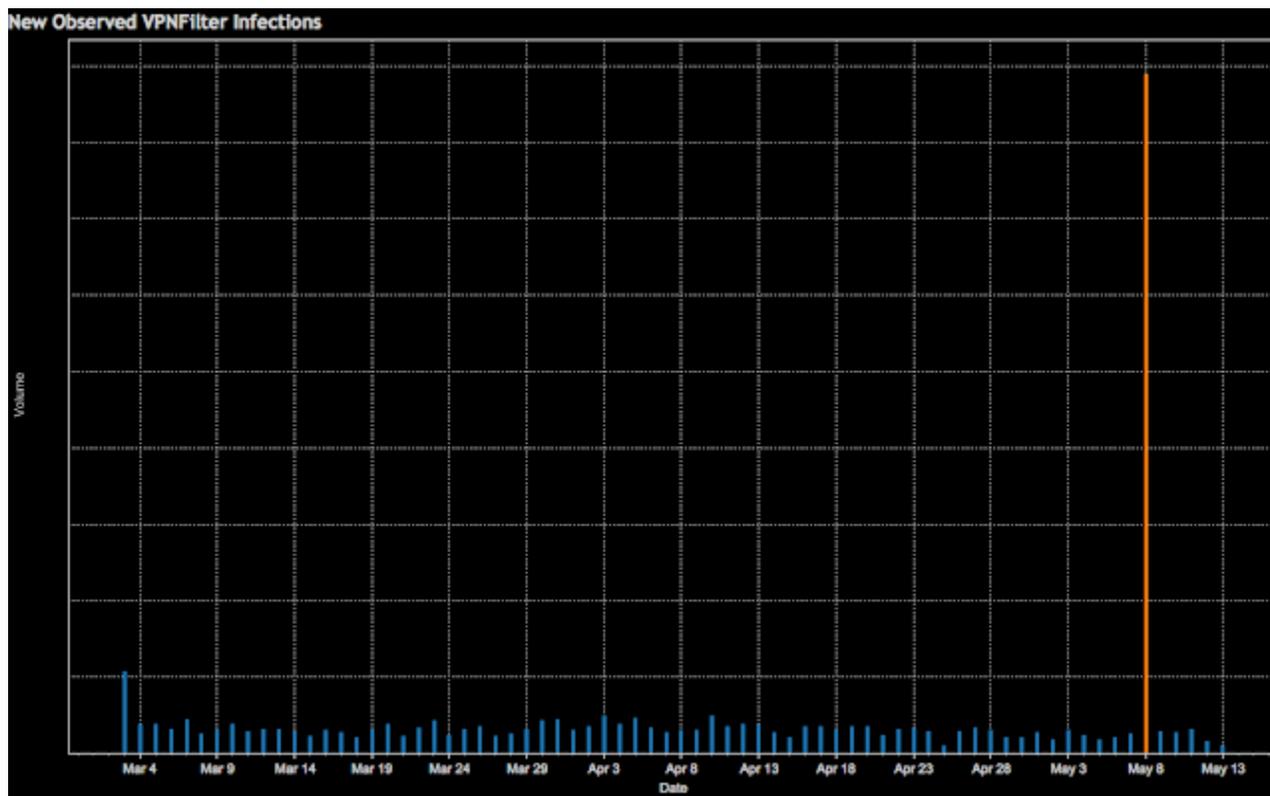
Diagram 1. New observed VPNFilter infections over time

Defending against this threat

Defending against this threat is extremely difficult due to the nature of the affected devices. The majority of them are connected directly to the internet, with no security devices or services between them and the potential attackers. This challenge is augmented by the fact that most of the affected devices have publicly known vulnerabilities which are not convenient for the average user to patch. Additionally, most have no built-in anti-malware capabilities. These three facts together make this threat extremely hard to counter, resulting in extremely limited opportunities to interdict malware, remove vulnerabilities, or block threats.

Despite these challenges, Talos has released protections for this threat from multiple angles, to try to take advantage of the limited options that exist. We developed and deployed more than 100 Snort signatures for the publicly known vulnerabilities for the devices that are associated with this threat. These rules have been deployed in the public Snort set, and can be used by anyone to help defend their devices. In addition, we have done the usual blocklisting of domains/IPs as appropriate and convicting of the hashes associated with this threat to cover those who are protected by the Cisco Security ecosystem. We have reached out to Linksys, Mikrotik, Netgear, TP-Link and QNAP regarding this issue. (Note: QNAP has been aware of certain aspects of VPNFilter and previously done work to counter the threat.)

Finally, we have also shared these indicators and our research with international law enforcement and our fellow members of the Cyber Threat Alliance in advance of this publication so they could move quickly to help counter this threat more broadly.

Recommendations

We recommend that:

- Users of SOHO routers and/or NAS devices reset them to factory defaults and reboot them in order to remove the potentially destructive, non-persistent stage 2 and stage 3 malware.
- Internet service providers that provide SOHO routers to their users reboot the routers on their customers' behalf.
- If you have any of the devices known or suspected to be affected by this threat, it is extremely important that you work with the manufacturer to ensure that your device is up to date with the latest patch versions. If not, you should apply the updated patches immediately.
- ISPs work aggressively with their customers to ensure their devices are patched to the most recent firmware/software versions.

Due to the potential for destructive action by the threat actor, we recommend out of an abundance of caution that these actions be taken for all SOHO or NAS devices, whether or not they are known to be affected by this threat.
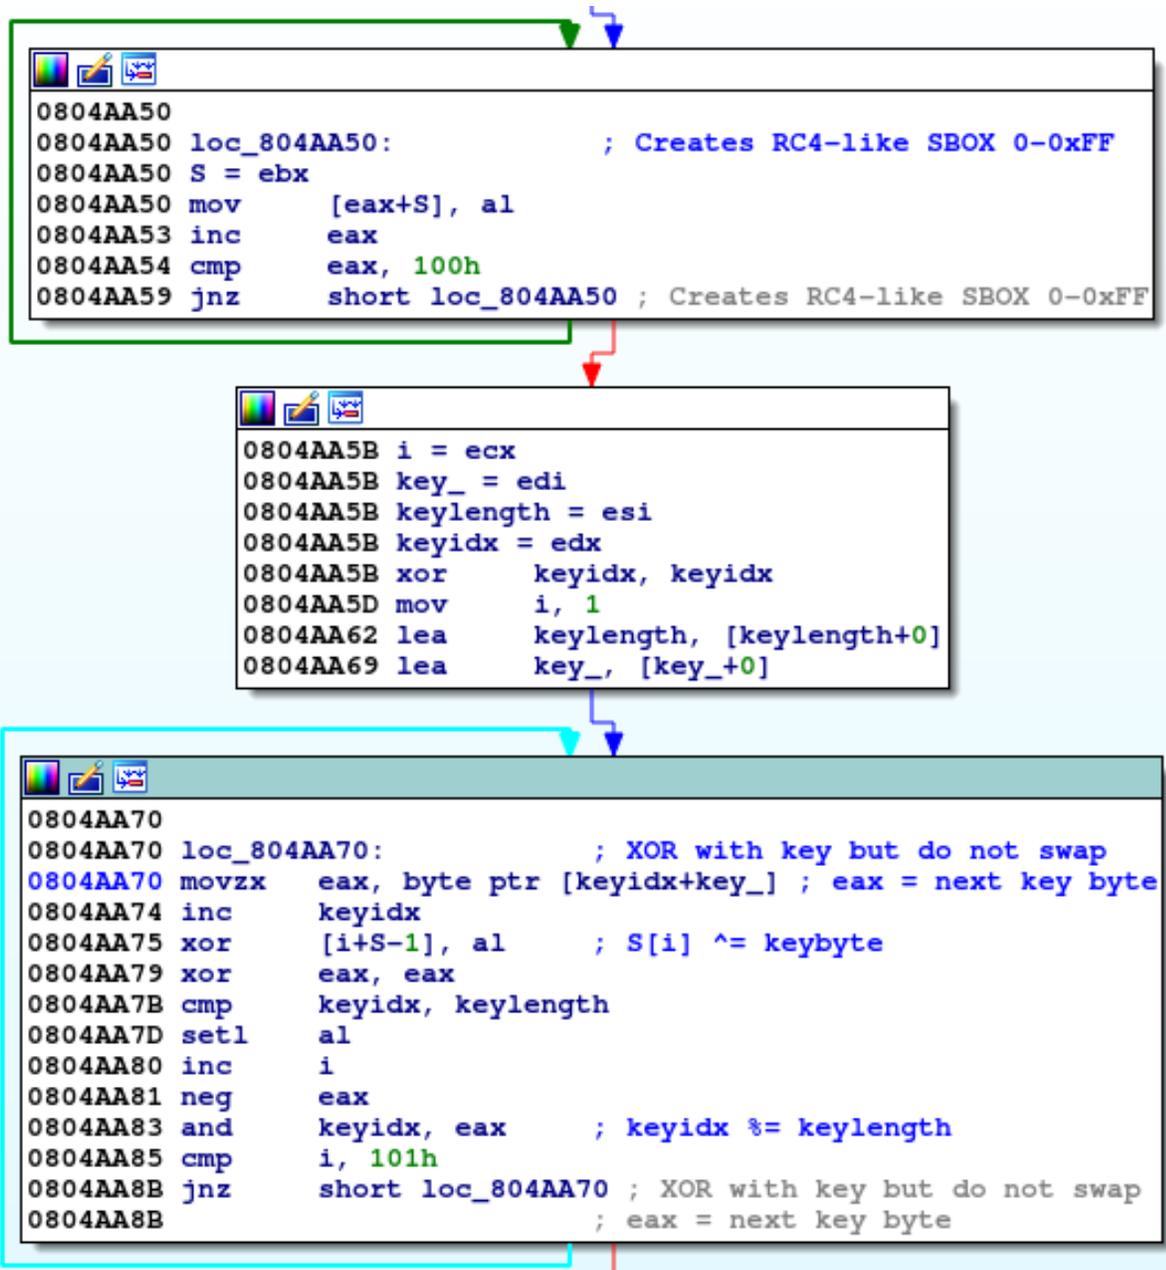
Multi-Stage Technical Details

Exploitation

At the time of this publication, we do not have definitive proof on how the threat actor is exploiting the affected devices. However, all of the affected makes/models that we have uncovered had well-known, public vulnerabilities. Since advanced threat actors tend to only use the minimum resources necessary to accomplish their goals, we assess with high confidence that VPNFilter required no zero-day exploitation techniques.

Stage 1 (persistent loader)

VPNFilter's stage 1 malware infects devices running firmware based on Busybox and Linux, and is compiled for several CPU architectures. The main purpose of these first-stage binaries is to locate a server providing a more fully featured second stage, and to download and maintain persistence for this next stage on infected devices. It is capable of modifying non-volatile configuration memory (NVRAM) values and adds itself to crontab, the Linux job scheduler, to achieve persistence. This is a departure from previous IoT malware, like Mirai, which is ephemeral and disappears with a simple device reboot.

Talos analyzed samples for MIPS and x86 processors. The C2 communication and additional malware downloads occur over Tor or SSL-encrypted connections. While the binaries themselves are not obfuscated beyond being stripped, some strings are stored in an encrypted form, and are only decrypted at runtime. The decryption routine looked suspiciously similar to RC4 in the static analysis, but it looks like the malware authors got the initialization of the S-boxes wrong. During the permutation step, values are XOR'd, but not swapped. Analysis of this RC4 implementation shows that it is identical to the implementation used in BlackEnergy, which is believed by law enforcement agencies to originate with a state actor.

```
0804AA50
0804AA50 loc_804AA50:                  ; Creates RC4-like SBOX 0-0xFF
0804AA50 S = ebx
0804AA50 mov      [eax+S], al
0804AA53 inc      eax
0804AA54 cmp      eax, 100h
0804AA59 jnz      short loc_804AA50 ; Creates RC4-like SBOX 0-0xFF
```

```
0804AA5B i = ecx
0804AA5B key_ = edi
0804AA5B keylength = esi
0804AA5B keyidx = edx
0804AA5B xor      keyidx, keyidx
0804AA5D mov      i, 1
0804AA62 lea      keylength, [keylength+0]
0804AA69 lea      key_, [key_+0]
```

```
0804AA70
0804AA70 loc_804AA70:                  ; XOR with key but do not swap
0804AA70 movzx    eax, byte ptr [keyidx+key_] ; eax = next key byte
0804AA74 inc      keyidx
0804AA75 xor      [i+S-1], al      ; S[i] ^= keybyte
0804AA79 xor      eax, eax
0804AA7B cmp      keyidx, keylength
0804AA7D setl     al
0804AA80 inc      i
0804AA81 neg      eax
0804AA83 and      keyidx, eax      ; keyidx %= keylength
0804AA85 cmp      i, 101h
0804AA8B jnz      short loc_804AA70 ; XOR with key but do not swap
0804AA8B                           ; eax = next key byte
```

The RC4 initialization XORs the values in the permutation phase of the internal state initialization. As you can see in the last basic block, the code doesn't swap the values of S[i] and S[j] (compared to the RC4 pseudo code below).

```
for i from 0 to 255
    S[i] := i
endfor
j := 0
for i from 0 to 255
    j := (j + S[i] + key[i mod keylength]) mod 256
    swap values of S[i] and S[j]
endfor
```

Once the malware has completed initialization, it starts to download pages from the seed URLs. In the MIPS sample cache and all but one URL of the x86 sample, the URLs pointed to Photobucket.com, an image-sharing host. The malware downloads the first image from the gallery the URL is referencing, and then proceeds to extract the download server's IP address. The IP address is extracted from six integer values for GPS latitude and longitude in the EXIF information.

If stage 1 fails to connect to, download an image from, or successfully acquire an IP address via an image from Photobucket, the malware reaches out to a backup domain, toknowall[.]com, to download an image and attempt the same process.

If the attempt to the backup domain fails, stage 1 opens a listener that waits for a specific trigger packet to open a connection for the actor to connect interactively to the device. When the listener opens, it checks its public IP from api.ipify[.]org and stores it for later comparison. Then, when any packet arrives on any port, the listener performs a series of checks to identify a trigger packet. If the packet meets a predefined set of criteria, it will extract an IP address from the packet and attempt a stage 2 download.

Listener actions:

1. Inspects all TCP/IPv4 packets with a SYN flag set
2. Checks that the destination IP matches what it found when the listener opened (Note: if the listener failed to get an IP from api.ipify[.]org it will skip this check)
3. Makes sure the packet has eight or more bytes
4. Scans the data for the bytes \x0c\x15\x22\x2b
5. The bytes directly after that 4-byte marker are interpreted as an IP so \x01\x02\x03\x04 becomes -> 1.2.3[.]4
6. Calls out to the newly received IP as usual for stage 2
7. Confirms that stage 2 is at least 1,001 bytes (Note: this is much smaller than the other callout methods which require the stage 2 to be 100,000 or more)

Stage 2 (non-persistent)

The stage 2 malware first sets up the working environment by creating a modules folder (/var/run/vpnfilterm) and a working directory (/var/run/vpnfilterw). Afterward, it will run in a loop, where it first reaches out to a C2 server, and then executes commands retrieved from

the C2. The command names are encrypted with the same broken RC4 function as in stage 1. Fortunately, older versions of x86 stage 2 sample were very verbose, and debug printed all the steps it performed. Newer versions of the x86 stage 2 did not contain the debug prints, nor did the MIPS sample.

The x86 sample can perform the following operations:

- kill: Overwrites the first 5,000 bytes of /dev/mtdblock0 with zeros, and reboots the device (effectively bricking it).
- exec: Executes a shell command or plugin.
- tor: Sets the Tor configuration flag (0 or 1).
- copy: Copies a file from the client to the server.
- seturl: Sets the URL of the current configuration panel.
- proxy: Sets the current proxy URL.
- port: Sets the current proxy port.
- delay: Sets the delay between main loop executions.
- reboot: Reboots the device if it has been up for more than 256 seconds, and the build name is specified in the parameter.
- download: Downloads a URL to a file. This can be applied to all devices or just a certain build name.

The MIPS sample has the following additional operations:

- stop: Terminate the malware process.
- relay: A misspelled version of the `delay` command from the x86 version.

Until the Tor module is installed, stage 2 will use one or more IPs stored in its configuration as SOCKS5 proxies to Tor and attempt to communicate with a control panel also found in its configuration. Like in stage 1, the communication between the malware and the proxy will connect over a verified SSL connection. When the Tor module is installed, it will connect to .onion domains through the local SOCKS5 proxy provided by the module over plain HTTP instead. We used a fake SOCKS5 proxy, which redirects all traffic to INetSim for analysis.

An example request from the malware to the server:

{

"uq":"px(01:02:03:04:05:06)",

"pv":"pPRXi686QNAPX86",

"ad":"10.0.0.1",

"bv":"0.11.1a/0.3.9qa",

"nn":"YnVpbGRyb290",

"tn":"",

"on":"1"

}

The malware encodes this request into a JSON object, which is then base64-encoded and sent to the path /bin32/update.php in the HTTP POST parameter "me". The user agent used in the request is peculiar (Mozilla/6.1 (compatible; MSIE 9.0; Windows NT 5.3; Trident/5.0)), as a version "Windows NT 5.3" doesn't exist.

- uq: A unique ID for the infected device (the MAC address of the malware's network interface).
- pv: The platform version the malware is running on
- ad: The public IP address of the malware's device
- bv: Version of the stage 1 loader (0.3.9qa) and the stage 2 binary (0.11.1a)
- nn: The node name
- tn: The Tor flag
- on: The onion flag

The server's response to the message:

{

"tr":3060,

"pxs":["217.12.202.40","94.242.222.68","91.121.109.209"],

"tor":"tor 1",

"mds":[]

}

- tr: Sets the delay for the main loop.
- pxs: List of panels to connect to. These are the C2 servers.
- tor: Sets the name and version of the Tor module.
- mds: A list of modules to fetch. Each entry is in the format "<command_id> <module_id> <module_name> <module_args (base64-encoded)>". The malware will download the module from /bin32/update.php by setting the POST form parameter me to the module name with the architecture appended, e.g., tor_i686 for the Tor module, and execute it in each iteration. A blank list of commands (as in the example response above) will clear any existing commands by deactivating them and killing any running processes associated with them.

Stage 3 (non-persistent)

We have analyzed two plugin modules for the malware, a packet sniffer and a communication plugin that allows the malware to communicate over Tor. We assess with high confidence that there are likely several more that we have not yet discovered. Among the initial samples Talos acquired, there was a plugin for the MIPS stage 2, which is a packet sniffer. It intercepts all network traffic through a raw socket and looks for strings used in HTTP basic authentications. Further, it specifically tracks Modbus TCP/IP packets. The resulting log file is placed in the stage 2 working directory, /var/run/vpnfilterw. This allows the attackers to understand, capture, and track the traffic flowing through the device.

The Tor plugin module is partially linked into stage 2, but has a separate Tor executable, which is downloaded to /var/run/tor and run in a process separate from stage 2. The Tor binary looks like the standard Tor client, in the form of a statically linked and stripped binary. It creates a configuration file in /var/run/torrc and a working directory in /var/run/tord.

Conclusion

VPNFilter is an expansive, robust, highly capable, and dangerous threat that targets devices that are challenging to defend. Its highly modular framework allows for rapid changes to the actor's operational infrastructure, serving their goals of misattribution, intelligence collection, and finding a platform to conduct attacks.

The destructive capability particularly concerns us. This shows that the actor is willing to burn users' devices to cover up their tracks, going much further than simply removing traces of the malware. If it suited their goals, this command could be executed on a broad scale, potentially rendering hundreds of thousands of devices unusable, disabling internet access for hundreds of thousands of victims worldwide or in a focused region where it suited the actor's purposes.

While the threat to IoT devices is nothing new, the fact that these devices are being used by advanced nation-state actors to conduct cyber operations, which could potentially result in the destruction of the device, has greatly increased the urgency of dealing with this issue. We call on the entire security community to join us in aggressively countering this threat.

We will continue to monitor VPNFilter and work with our partners to understand the threat as it continues to evolve in order to ensure that our customers remain protected and the public is informed.

IOCs

As stated previously, we highly suspect that there are additional IOCs and versions of this malware that we are not currently aware of. The following list of IOCs comprises what we know as of this date.

Known C2 Domains and IPs

Associated with the 1st Stage

photobucket[.]com/user/nikkireed11/library
photobucket[.]com/user/kmila302/library
photobucket[.]com/user/lisabraun87/library
photobucket[.]com/user/eva_green1/library
photobucket[.]com/user/monicabelci4/library
photobucket[.]com/user/katyperry45/library
photobucket[.]com/user/saragray1/library
photobucket[.]com/user/millerfred/library
photobucket[.]com/user/jeniferaniston1/library
photobucket[.]com/user/amandaseyfried1/library
photobucket[.]com/user/suwe8/library
photobucket[.]com/user/bob7301/library
toknowall[.]com

Associated with the 2nd Stage

91.121.109[.]209
217.12.202[.]40
94.242.222[.]68
82.118.242[.]124
46.151.209[.]33
217.79.179[.]14
91.214.203[.]144
95.211.198[.]231
195.154.180[.]60
5.149.250[.]54
91.200.13[.]76
94.185.80[.]82
62.210.180[.]229
zuh3vcyskd4gipkm[.]onion/bin32/update.php

Known File Hashes

1st Stage Malware

50ac4fcd3fbc8abcaa766449841b3a0a684b3e217fc40935f1ac22c34c58a9ec
0e0094d9bd396a6594da8e21911a3982cd737b445f591581560d766755097d92

2nd Stage Malware

9683b04123d7e9fe4c8c26c69b09c2233f7e1440f828837422ce330040782d17
d6097e942dd0fdc1fb28ec1814780e6ecc169ec6d24f9954e71954eedbc4c70e
4b03288e9e44d214426a02327223b5e516b1ea29ce72fa25a2fcef9aa65c4b0b
9eb6c779dbad1b717caa462d8e040852759436ed79cc2172692339bc62432387
37e29b0ea7a9b97597385a12f525e13c3a7d02ba4161a6946f2a7d978cc045b4
776cb9a7a9f5afbaffdd4dbd052c6420030b2c7c3058c1455e0a79df0e6f7a1d
8a20dc9538d639623878a3d3d18d88da8b635ea52e5e2d0c2cce4a8c5a703db1
0649fda8888d701eb2f91e6e0a05a2e2be714f564497c44a3813082ef8ff250b

3rd Stage Plugins

f8286e29faa67ec765ae0244862f6b7914fcdde10423f96595cb84ad5cc6b344
afd281639e26a717aead65b1886f98d6d6c258736016023b4e59de30b7348719

Self-Signed Certificate Fingerprints

d113ce61ab1e4bfcb32fb3c53bd3cdeee81108d02d3886f6e2286e0b6a006747
c52b3901a26df1680acbfb9e6184b321f0b22dd6c4bb107e5e071553d375c851
f372ebe8277b78d50c5600d0e2af3fe29b1e04b5435a7149f04edd165743c16d
be4715b029cbd3f8e2f37bc525005b2cb9cad977117a26fac94339a721e3f2a5
27af4b890db1a611d0054d5d4a7d9a36c9f52dffeb67a053be9ea03a495a9302
110da84f31e7868ad741bcb0d9f7771a0bb39c44785055e6da0ecc393598adc8
fb47ba27dceea486aab7a0f8ec5674332ca1f6af962a1724df89d658d470348f
b25336c2dd388459dec37fa8d0467cf2ac3c81a272176128338a2c1d7c083c78
cd75d3a70e3218688bdd23a0f618add964603736f7c899265b1d8386b9902526
110da84f31e7868ad741bcb0d9f7771a0bb39c44785055e6da0ecc393598adc8
909cf80d3ef4c52abc95d286df8d218462739889b6be4762a1d2fac1adb2ec2b
044bfa11ea91b5559f7502c3a504b19ee3c555e95907a98508825b4aa56294e4
c0f8bde03df3dec6e43b327378777ebc35d9ea8cfe39628f79f20b1c40c1b412
8f1d0cd5dd6585c3d5d478e18a85e7109c8a88489c46987621e01d21fab5095d
d5dec646c957305d91303a1d7931b30e7fb2f38d54a1102e14fd7a4b9f6e0806
c0f8bde03df3dec6e43b327378777ebc35d9ea8cfe39628f79f20b1c40c1b412

Known Affected Devices

The following devices are known to be affected by this threat. Based on the scale of this research, much of our observations are remote and not on the device, so it is difficult to determine specific version numbers and models in many cases. It should be noted that all of these devices have publicly known vulnerabilities associated with them.

Given our observations with this threat, we assess with high confidence that this list is incomplete and other devices could be affected.

Linksys Devices:

E1200
E2500
WRVS4400N

Mikrotik RouterOS Versions for Cloud Core Routers:

1016
1036
1072

Netgear Devices:

DGN2200
R6400
R7000
R8000
WNR1000
WNR2000

QNAP Devices:

TS251
TS439 Pro
Other QNAP NAS devices running QTS software

TP-Link Devices:

R600VPN

Coverage

Cisco customers are protected by this threat by Cisco Advanced Malware Protection (AMP), Cloud Web Security (CWS), Network Security, ThreatGrid, Umbrella, and Web Security Appliance (WSA). Additionally, StealthWatch and StealthWatch Cloud can be utilized to find devices communicating with the known C2 IP addresses and domains.

In StealthWatch, two items need to be configured to send an alert that there are communications to nefarious IP addresses.

- The first step is to create a new Host Group named "VPNFilter C2" under Outside Hosts using the Java user interface.
- Once this is created, you will likely want to validate that there are no active communications presently occurring.
- This validation can be achieved by right-clicking on the recently created "VPNFilter C2" Host Group and navigating to Top -> Conversations -> Total.
- Once you are viewing these top conversations, you will easily be able to see if there is active traffic.
- In the event that there is no active traffic, an alarm can be created to generate alerts in the event that traffic to or from any of the "VPNFilter C2" hosts is observed.
- This alarm can be configured by creating a custom event and selecting the appropriate hosts or objects in the web user interface.

PRODUCT	PROTECTION
AMP	✔
CWS	✔
Email Security	N/A
Network Security	✔
Threat Grid	✔
Umbrella	✔
WSA	✔

VPNFilter specific Snort detection:

45563 45564 46782 46783

Snort rules that protect against known vulnerabilities in affected devices:

25589 26276 26277 26278 26279 29830 29831 44743 46080 46081 46082 46083 46084 46085 46086 46287 46121 46122 46123 46124 41445 44971 46297 46298 46299 46300 46301 46305 46306 46307 46308 46309 46310 46315 46335 46340 46341 46342 46376 46377 37963 45555 46076 40063 44643 44790 26275 35734 41095 41096 41504 41698 41699 41700 41748 41749 41750 41751 44687 44688 44698 44699 45001 46312 46313 46314 46317 46318 46322 46323 40866 40907 45157

ClamAV Signatures:

Unix.Trojan.Vpnfilter-6425811-0
Unix.Trojan.Vpnfilter-6425812-0
Unix.Trojan.Vpnfilter-6550590-0
Unix.Trojan.Vpnfilter-6550591-0
Unix.Trojan.Vpnfilter-6550592-0