Routers from the Underground: Exposing AVrecon

blog.lumen.com/routers-from-the-underground-exposing-avrecon/

July 12, 2023

Black Lotus Labs Posted On July 12, 2023

0

10.9K Views

Executive Summary

Lumen Black Lotus Labs® identified another multi-year campaign involving compromised routers across the globe. This is a complex operation that infects small-office/home-office (SOHO) routers, deploying a Linux-based Remote Access Trojan (RAT) we've dubbed "AVrecon." Apart from a single reference to AVrecon in May 2021, the malware has been operating undetected for more than two years. Black Lotus labs performed an extensive analysis documenting the malware functionality, its size, and how it fits into the cybercrime ecosystem.

We assess the purpose of the campaign appears to be the creation of a covert network to quietly enable a range of criminal activities from password spraying to digital advertising fraud. Due to the surreptitious nature of the malware, owners of infected machines rarely notice any service disruption or loss of bandwidth. This assessment is based on observed telemetry and the analysis of functionality in the binary that allows the actor to interact with a remote shell and deploy subsequent binaries. Using Lumen's global network visibility, Black Lotus Labs has determined the composition of a network that has infiltrated more than 70,000 machines, gaining a persistent hold in more than 40,000 IPs in more than 20 countries. The use of encryption prevents us from commenting on the results of successful password spraying attempts; however, we have null-routed the command and control (C2) nodes and impeded traffic through the proxy servers, which rendered the botnet inert across the Lumen backbone.

Introduction

Black Lotus Labs has tracked the abuse of networking equipment in various campaigns from nation-state to cybercrime and even hacktivism. On June 13, 2023, the U.S. Cybersecurity and Infrastructure Security Agency (CISA) warned that exploitation of networking devices such as SOHO routers can enable adversaries to stand up actor-controlled infrastructure globally or gain unrestricted access to organizational networks. This is a serious threat because these devices typically live outside the traditional security perimeter, which significantly decreases the ability to detect malicious activity.

In our latest investigation, Black Lotus Labs discovered one of the largest botnets targeting small-office/home-office (SOHO) routers seen in recent history. It has been used to create residential proxy services to shroud malicious activity such as password spraying, web-traffic proxying and ad fraud. Surpassing Qakbot in scale, AVrecon operated undetected for more than two years. While there were references on Twitter in May 2021 by @SethKingHi, no other research showed the breadth of this operation or thoroughly documented its functionality.

Upon infection, the threat actor enumerates the victim's SOHO router, then sends that information back to the embedded C2 domain. From there, the infected system is ordered to begin interacting with a separate set of servers, which we refer to as second stage C2 servers. Using both Lumen's proprietary and commercial telemetry, we discovered 15 unique second stage C2s. Based on information associated with their x.509 certificates, we assess that some of these second stage C2s have been active since at least October 2021. We took a 28-day snapshot of the second stage servers and found more than 70,000 distinct IP addresses communicating with them. We then investigated how many machines were persistently infected – meaning they communicated with one of the second stage servers for two or more days within the 28-day window – and we identified 41,000 nodes.

We have seen a number of SOHO-based botnets like Chaos leverage their access for more aggressive activity like deploying crypto-miners or launching DDoS attacks. These activities tend to draw attention, as mining creates performance issues from overworked CPUs, and DDoS often results in abuse complaints. The AVrecon campaign found success because:

1. The target machines are primarily networking equipment that do not offer standard endpoint detection and response (EDR) solutions.
2. We suspect the threat actor focused on the type of SOHO devices users would be less likely to patch against common vulnerabilities and exposures (CVEs). Instead of using this botnet for a quick payout, the operators maintained a more temperate approach and were able to operate undetected for more than two years.

While there were two brief references to this malware family in 2021, there has not been a comprehensive write-up covering the extent of its operations. This report describes the malware functionality and embedded capabilities, examining how the remote access trojan interacts with the various C2 servers. Finally, we explore the scope of the botnet's spread and the activity stemming from it.

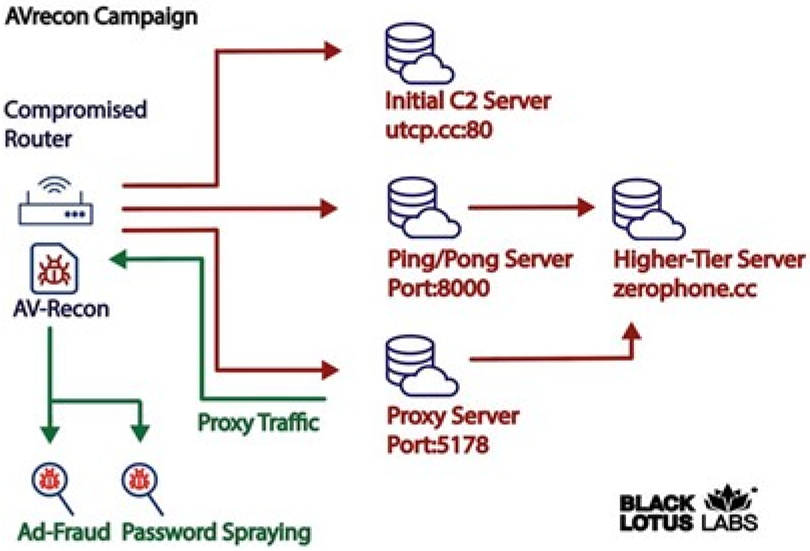

Figure 1: AVrecon campaign

Malware Analysis

The Black Lotus Labs team's analysis revealed that AVrecon was written in C and targeted ARM-embedded devices. The choice of that programming language allows for portability, and, later in our investigation, we discovered that the malware had been compiled for different architectures.

Once the AVrecon RAT is deployed on the infected system, its actions can generally be placed into three categories:

- It checks to see if other instances of the malware are already running on the system.
- It gathers host-based information.
- It builds the parameters of the C2 channel.

Checking for other instances

First, the malware checks for any instances of itself on the host machine by searching for existing processes on port 48102 and opening a listener on that port. If the bind is successful, it writes the process ID to the file *jid.pid* in the *tmp* directory. If there is already a service running on port 48102, the malware kills any existing processes that don't match the current process ID, then moves on to kill any process bound to port 48102. We have not seen the malware utilize any other ports on compromised machines during installation. In an interesting final act, should none of the previous attempts succeed, AVrecon has instructions to delete itself entirely from the host machine.

Gathering host-based information

AVrecon proceeds to collect host-based information about the machine, including the device *uname* (kernel information), CPU, memory usage, *bin* path where it's running, and *hostname*. When it has obtained what it needs, pre-built functions spawn a remote shell to execute commands, download subsequent binaries, and configure a proxy. The full list is shown by command number:

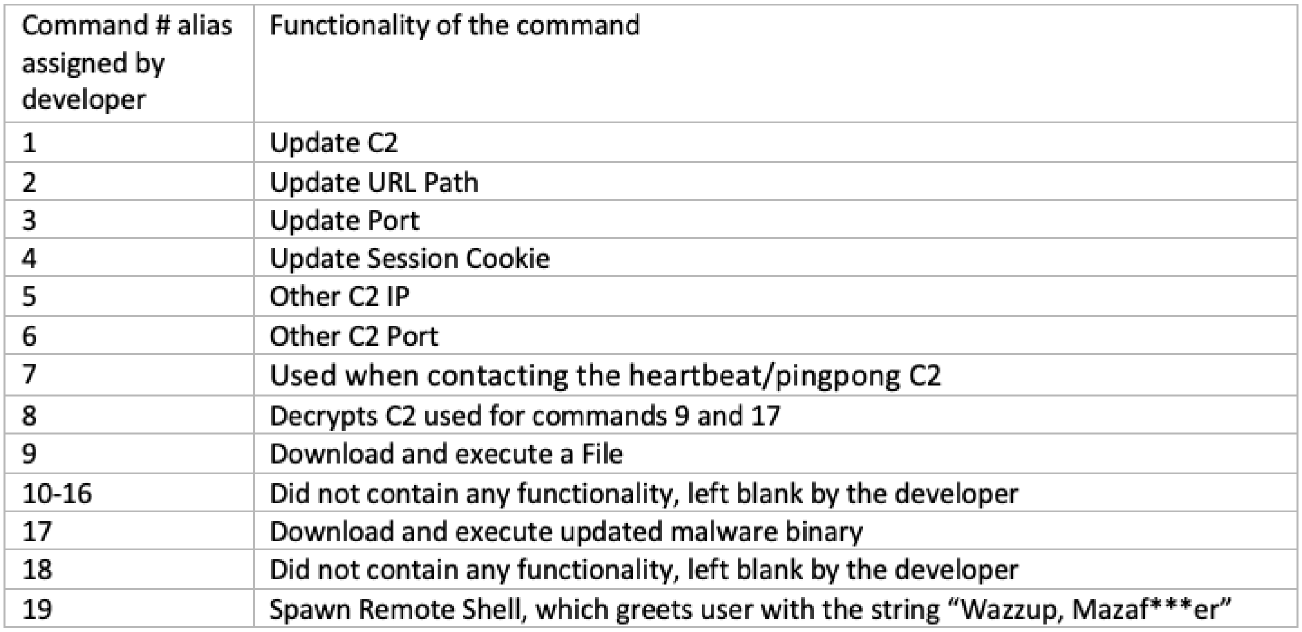

Command # alias assigned by developer	Functionality of the command
1	Update C2
2	Update URL Path
3	Update Port
4	Update Session Cookie
5	Other C2 IP
6	Other C2 Port
7	Used when contacting the heartbeat/pingpong C2
8	Decrypts C2 used for commands 9 and 17
9	Download and execute a File
10-16	Did not contain any functionality, left blank by the developer
17	Download and execute updated malware binary
18	Did not contain any functionality, left blank by the developer
19	Spawn Remote Shell, which greets user with the string "Wazzup, Mazaf***er"

Building C2 parameters

With freedom to move on the infected system, the malware begins constructing the C2 parameters. While the malware is shipped with an embedded configuration for C2 communications, it first checks to see if a configuration already exists on the infected host from prior communications and, if not, uses parameters for the C2 sent in the malware’s configuration. To check if C2 configurations exist, the binary searches for legitimate files named either *nvram* or *xmldbc* in the *bin* and *sbin* directories. It then searches for variables previously configured by the malware in those files — memasik, domik, urlik and portik. If either of these programs is found with the set variables, the binary extracts the encrypted C2 components from *memasik*, *domik*, *urlik*, and *portik* variables, including session cookies from any previous session, domain, URL and port information. If any of the four variables are not found, the malware attempts to read and decrypt configuration from */tmp/dnssmasq.pid*. Finally, if that fails, it decrypts the hardcoded domain, URL and port from the embedded C2 configuration.

Network-based Communications

The binary proceeds with a GET request to the C2 using the configuration extracted in the previous process. The GET request contains various headers with encrypted details about the infected machine. The GET URL and host field are listed below, along with a sample of the decrypted information being sent to the C2.

Sample Initial C2 Beacon

```
GET /lumi/track.php?pet=maral&age=1&up=1112 HTTP/1.1
Host: utcp[.]cc
```

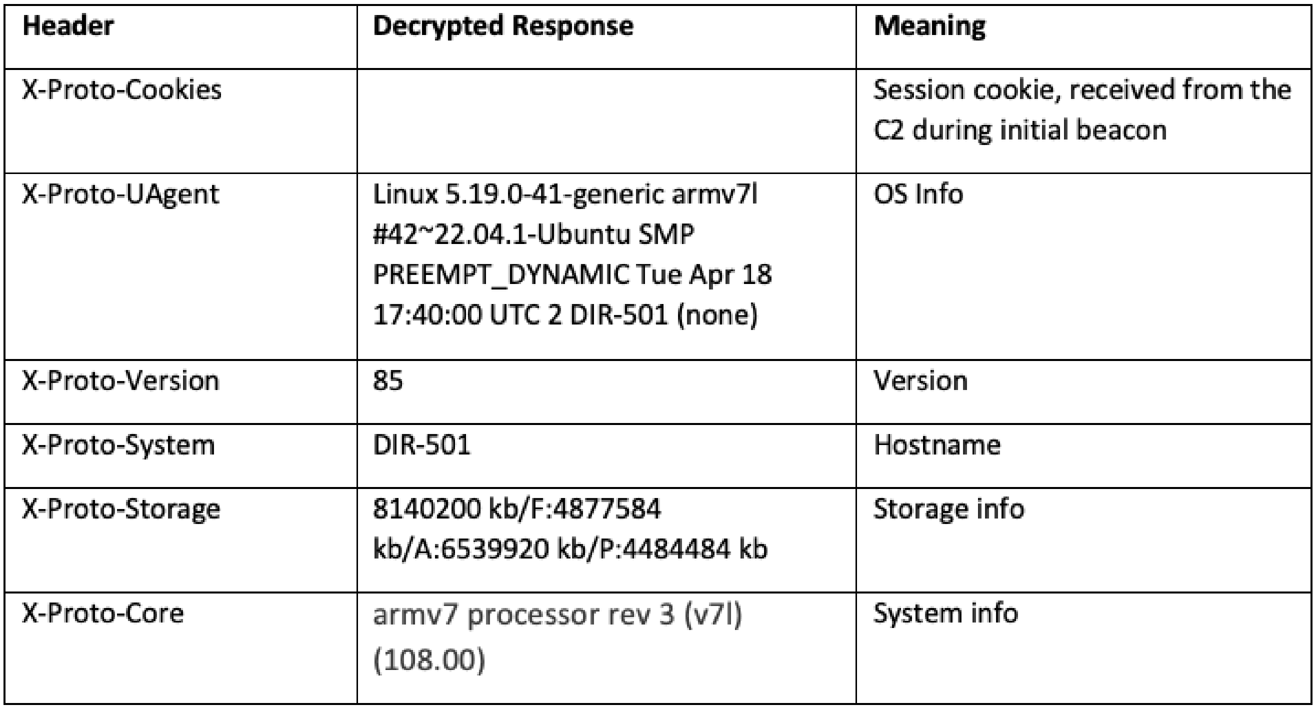

Header	Decrypted Response	Meaning
X-Proto-Cookies		Session cookie, received from the C2 during initial beacon
X-Proto-UAgent	Linux 5.19.0-41-generic armv7l #42~22.04.1-Ubuntu SMP PREEMPT_DYNAMIC Tue Apr 18 17:40:00 UTC 2 DIR-501 (none)	OS Info
X-Proto-Version	85	Version
X-Proto-System	DIR-501	Hostname
X-Proto-Storage	8140200 kb/F:4877584 kb/A:6539920 kb/P:4484484 kb	Storage info
X-Proto-Core	armv7 processor rev 3 (v7l) (108.00)	System info

The response from the C2 is a series of commands from the list above and decrypted parameters. A sample is below. To illustrate the functionality, in the second response parameter, the command is #1 with an encrypted parameter of 757A6D7E336D6D which decrypts to utcp[.]cc. This indicates the malware should update to the C2 to utcp[.]cc.

Initial C2 Response

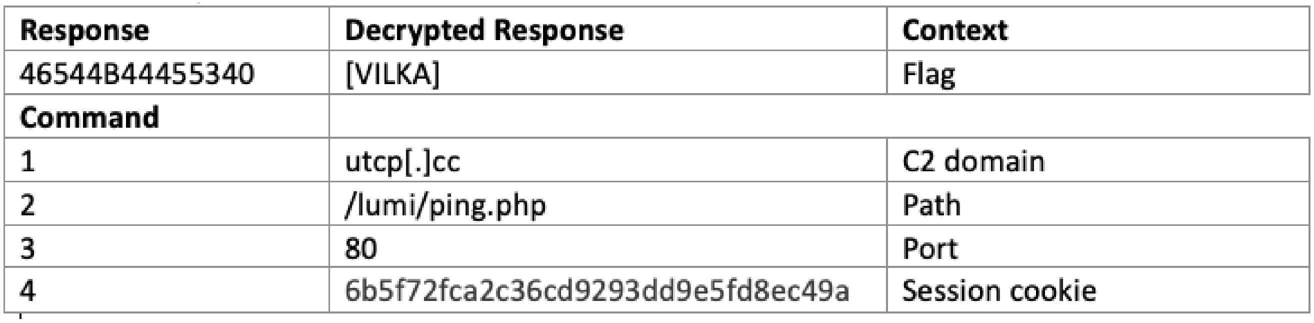

Response	Decrypted Response	Context
46544B44455340	[VILKA]	Flag
Command		
1	utcp[.]cc	C2 domain
2	/lumi/ping.php	Path
3	80	Port
4	6b5f72fca2c36cd9293dd9e5fd8ec49a	Session cookie

Using these new configuration parameters, the malware reaches out to the configured C2. When the malware connects to this URL, a response containing information about an additional C2 is returned, intended to direct future responses.

Response from /lumi/ping.php

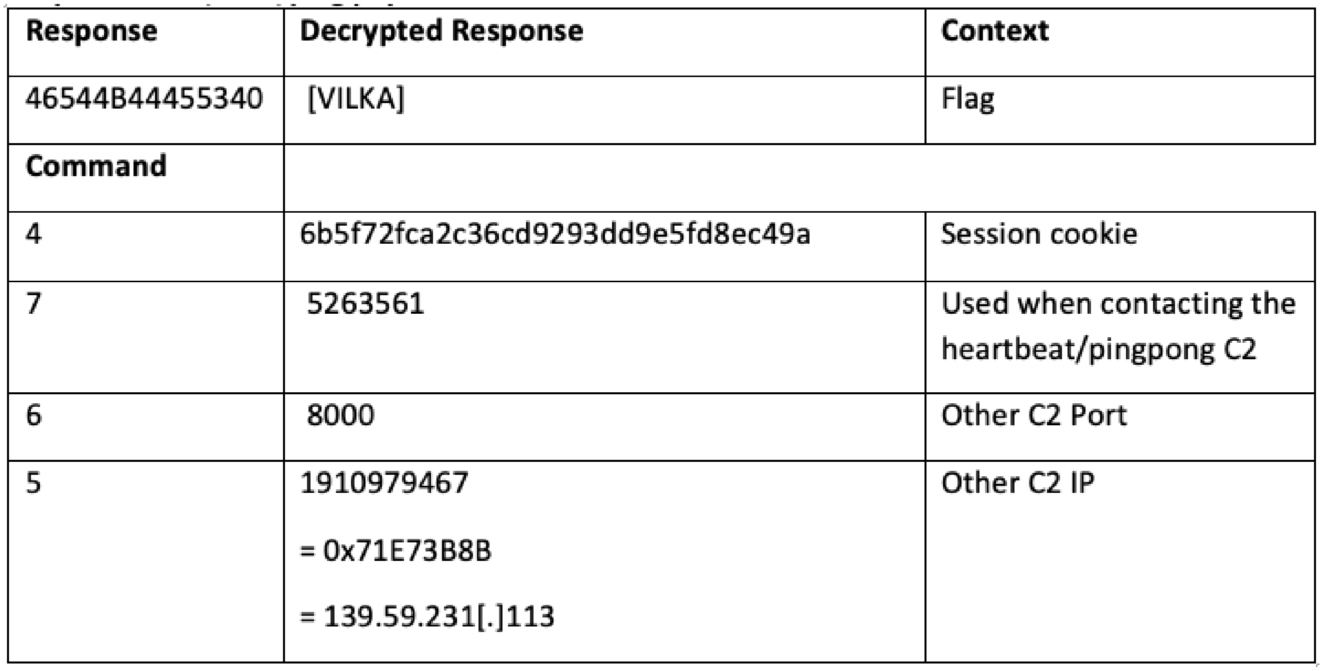

Response	Decrypted Response	Context
46544B44455340	[VILKA]	Flag
Command		
4	6b5f72fca2c36cd9293dd9e5fd8ec49a	Session cookie
7	5263561	Used when contacting the heartbeat/pingpong C2
6	8000	Other C2 Port
5	1910979467 = 0x71E73B8B = 139.59.231[.]113	Other C2 IP

Once the agent receives a command to communicate with the next C2 over port 8000, it immediately checks in with that C2 and enters a loop where it sends the word “ping” and awaits the response “pong.” It repeats this loop until it either expires or receives subsequent tasking from the new second-stage C2.

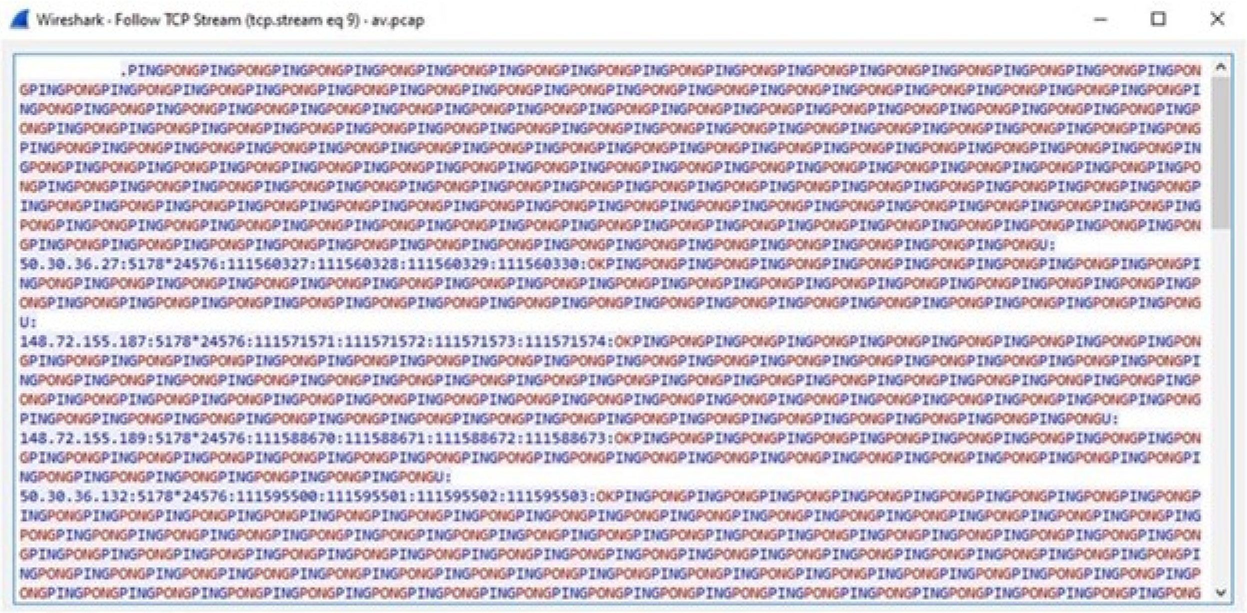

Wireshark - Follow TCP Stream (tcp.stream eq 9) - av.pcap

```
           .PINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPON
GPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPI
NGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGP
ONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONG
PINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPIN
GPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPO
NGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGP
INGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPING
PONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPON
GPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGU:
50.30.36.27:5178*24576:111560327:111560328:111560329:111560330:OKPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPI
NGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGP
ONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONG
U:
148.72.155.187:5178*24576:111571571:111571572:111571573:111571574:OKPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPON
GPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPI
NGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGP
ONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONG
PINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGU:
148.72.155.189:5178*24576:111588670:111588671:111588672:111588673:OKPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPON
GPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPI
NGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGU:
50.30.36.132:5178*24576:111595500:111595501:111595502:111595503:OKPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGP
INGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPING
PONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPON
GPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPI
NGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGP
ONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONGPINGPONG
```

Figure 2: Example of the PING/PONG traffic

Based off one example, we observed the agent get redirected to another server and told to communicate over port 5178. Some of the C2s that were supplied by the second stage C2 server include 148.72.155[.]112, 148.72.155[.]189 and 139.59.231[.]113. Though our global telemetry, we were able to enumerate other nodes that exhibited the same behavior characteristics, resulting in a total of 15 second-stage nodes.

Observed Activity

Our analysis indicates the threat actor used the infected machines to click on various Facebook and Google ads, and to interact with Microsoft Outlook. We suspect the first activity to be part of an advertising fraud effort, and the second activity is likely password spraying and/or data exfiltration. This global network of compromised SOHO routers gives cyber criminals the ability to bypass some standard network-based detection tools, especially those based on geolocation, autonomous system-based blocking, or IP address-based rate limiting.

Global Telemetry Analysis

Bot Analysis

Working from all second-stage IP addresses we could attribute to AVrecon, we searched our telemetry for connections over some of the previously mentioned ports such as 80, 8000 and 5178. When we compiled a list of distinct IP addresses that made at least one connection to one of the 15 servers, the results identified approximately 70,000 unique IP addresses.

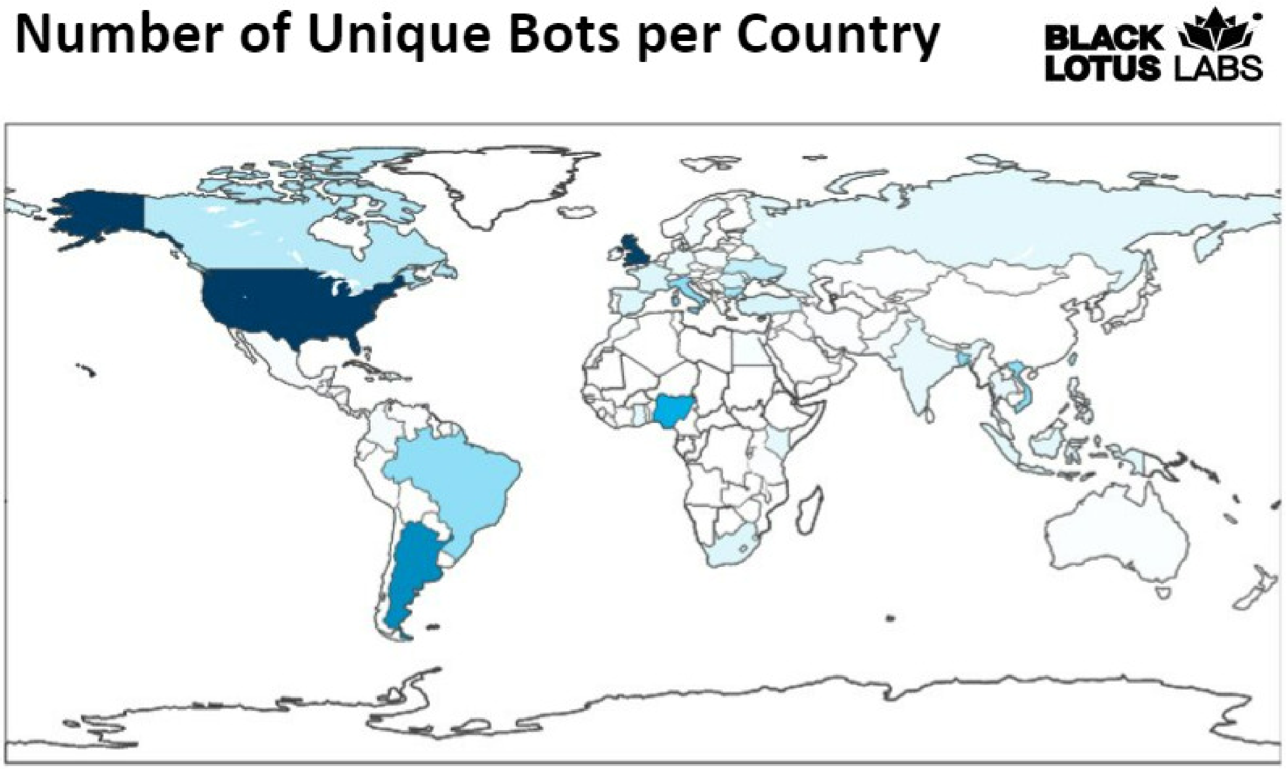

Figure 3: Global distribution of the bots.

Once we identified the global distribution, we sought to determine the average lifespan of the bots by calculating how often they communicate with one of the C2 nodes.

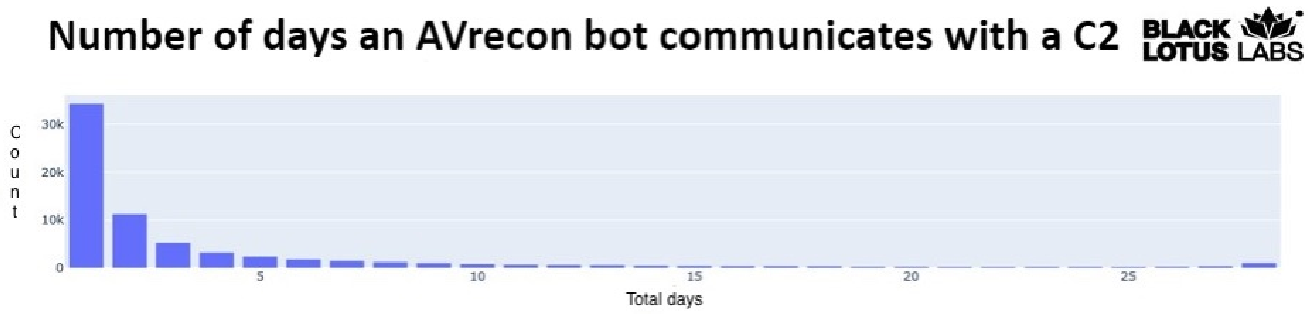

Figure 4: Graph showing the number of bots with total number of days they communicate with a C2 over a 28-day period.

Our analysis revealed that most IPs – more than 45,000 – only communicate with one of the C2s for one to two days. Since we knew there was no persistence mechanism inherent to the malware, we suspect these devices were most likely infected and remediated or otherwise abandoned. We characterized infections lasting more than two days as persistent, which gives us a botnet size of roughly 41,100 devices. We found the mean lifespan of a bot was just over seven days. This means the threat actor had a network of substantial size, in which any device could be used to tunnel malicious traffic at any time.

Command and Control Analysis

After inspecting traffic patterns at the bot-level, we examined the telemetry associated with the core set of 15 second stage servers. We found that, of this group, 12 communicated with one higher tier server located at 51.15.19[.]245 for 30 consecutive days. Using a tiered architecture for the backend is a tactic Black Lotus Labs has observed with other prominent botnets such as Emotet and Qakbot.

One rather interesting data point stuck out: when we scanned the higher-tier IP address on ports 25 and 465, it displayed the domain zerophone[.]cc. This domain was mentioned in the historical write-up on AVrecon from 2021 as part of the original analysis. The domain had a sample associated with it, which was compiled not only for ARM, but for MIPS and MIPSEL as well. This indicates the threat actor was targeting capabilities for the major SOHO router architectures on the market.

Conclusion

Black Lotus Labs remains on the cutting edge of monitoring network equipment and tracking large-scale botnets. This report on AVrecon unifies these two focus areas together to help better secure the internet ecosystem. The manner of attack seems to focus predominantly on stealing bandwidth – without impacting end-users – in order to create a residential proxy service to help launder malicious activity and avoid attracting the same level of attention from Tor-hidden services or commercially available VPN services. This class of cybercrime activity threat may evade detection because it is less likely than a crypto-miner to be noticed by the owner, and it is unlikely to warrant the volume of abuse complaints that internet-wide brute-forcing and DDoS-based botnets typically draw.

Black Lotus Labs has null-routed the AVrecon C2s across the Lumen global backbone and added the indicators of compromise (IoCs) from this campaign into the threat intelligence feed that fuels the Lumen Connected Security portfolio. We will continue to monitor new infrastructure, targeting activity and expanding TTPs, and we will continue to collaborate with the security research community to share findings related to this activity.

We encourage the community to monitor for and alert on these and any similar IoCs. We also advise the following:

Corporate Network Defenders:

- Continue to look for attacks on weak credentials and suspicious login attempts, even when they originate from residential IP addresses which bypass geofencing and ASN-based blocking.

- While we have not observed subsequent exploitation directed towards the adjacent LAN thus far, the threat actors do have the ability to spawn a remote shell and deploy subsequent modules. While this type of activity has historically been observed in more advanced campaigns such as ZuoRat and HiatusRat, it remains an area ripe for exploitation.
- Protect cloud assets from communicating with bots that are attempting to perform password spraying attacks, and begin blocking IoCs with Web Application Firewalls.

Consumers with SOHO routers:

> Follow best practices by regularly rebooting routers and installing security updates and patches. Users should leverage properly configured and updated EDR solutions on hosts and regularly update software consistent with vendor patches where applicable.

For additional IoCs associated with this campaign, please visit our GitHub page.

If you would like to collaborate on similar research, please contact us on Twitter @BlackLotusLabs.

This analysis was performed by Danny Adamitis and Steve Rudd. Technical editing by Ryan English.

This information is provided “as is” without any warranty or condition of any kind, either express or implied. Use of this information is at the end user’s own risk.

Post Views: 10,891