# The curious case of the 7777-Botnet

By Gi7w0rm

Published: 2023-10-20 · Archived: 2026-04-05 18:21:21 UTC

Hello there and welcome back again to yet another blog post. Today, I am reporting on something I have been investigating for a while now. An alledged botnet which at its peak included more than 16.000 infected devices has been observed in targeted attacks in the US, UK, and France. There is even a link to organized cybercrime. But in the end, it still remains a mystery…

**Summary:**

In recent months an unknown botnet has been observed brute forcing Microsoft Azure instances via Microsoft Azure PowerShell bruteforcing. The botnet has a unique pattern of opening port 7777 on infected devices, returning an “xlogin:” message. The botnet has been used for low-volume attacks against targets of all industry sectors at a global scale, almost exclusively targeting C-Level employee logins. Due to the very low volume of around 2–3 login requests per week, the botnet is able to evade most security solutions. An attribution is not possible with the current insights. — h-o-w-e-v-e-r-,--a-l-o-o-s-e-c-o-n-n-e-c-t-i-o-n- -w-i-t-h- -e-i-t-h-e-r- -U-N-C-3-9-4-4- -/- -S-t-o-r-m- — 0-8-7-5- -/- -S-c-a-t-t-e-r-e-d- -S-p-i-d-e-r- -o-r- -t-h-e- -L-a-z-a-r-u-s- -g-r-o-u-p- -c-a-n- -b-e- -m-a-d-e-. — *

**Chapter 1: A New Contact**

As with most of my current Cyber Threat Intelligence endeavors, this case starts with someone reaching out to me for help. On the 19th of July of this year, I was contacted by an employee of the British cybersecurity company Goldilock, who sent me the following message:

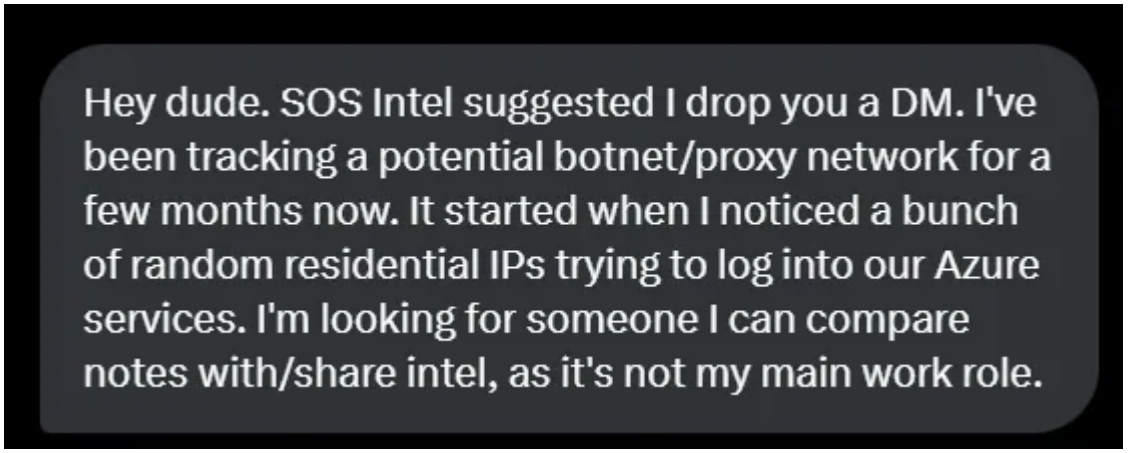


Figure 1: The beginning

What the researcher at Goldilock had found was a strange set of IP addresses that time and time again attempted (and failed) to log into their Microsoft Azure/Exchange environment, in what appeared to be a very very slow and continuous brute-force attempt. This attack had the following characteristics:

- An average of 2–3 attempts per week on a single account was observed, interspersed with occasional 7 to 10-day hiatuses
- Each IP address is used uniquely for each user and is not recycled for different accounts
- The login attempts all attempted to abuse “Microsoft Azure PowerShell” according to Azure Portal Sign-In logs, failing at the credential stage.

What is interesting here is that this particular attack had not been identified by any security measure in place, neither the company's SIEM nor its Microsoft Security implementations. The attack was simply so low in volume that the security appliances did not see it as unusual. Only when the employee noticed a strange login from a country they had no business in, this attack was discovered.

## Chapter 2: A botnet?!

After some digging, the employee was able to uncover around 50 different IPs that had taken part in this attack. In a first effort to uncover the source of this attack, the employee put them into the Shodan Search engine. And indeed they all shared one similarity: An open port 7777!

Here is a random example of how one of the identified devices looks in Shodan:

Press enter or click to view image in full size

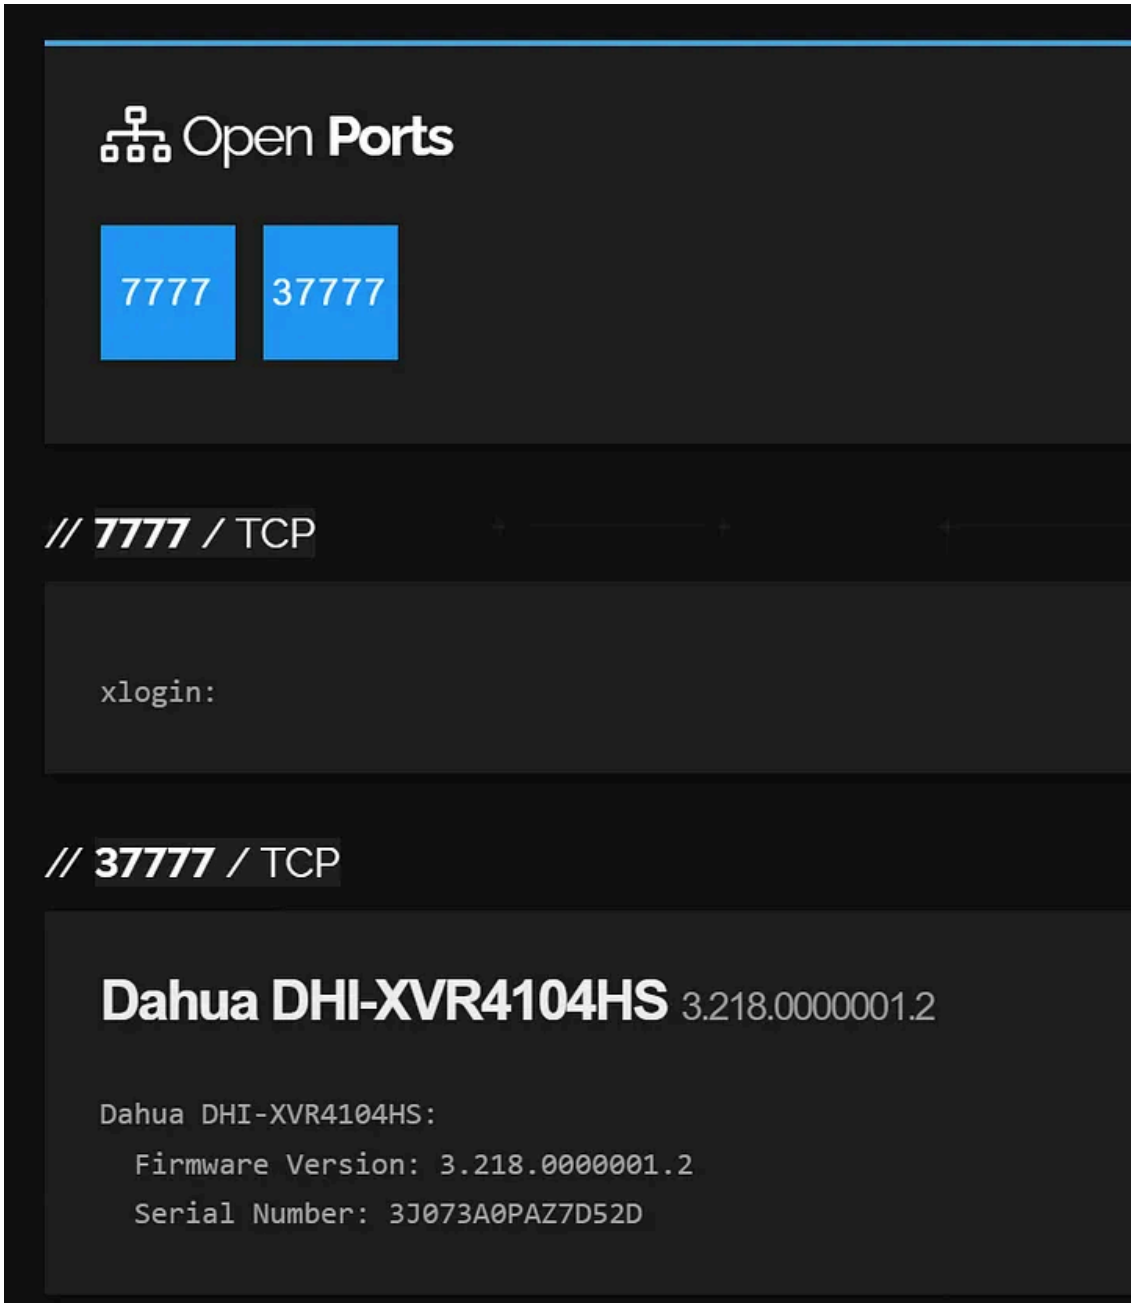


Figure 2: An example device forming part of this Botnet

As you can see, port 7777 is open and showing a mysterious “xlogin:” message. Luckily enough, this particular pattern is pretty unique. So we can use Shodan to pivot on it via the following query:

As of today, 07.10.2023, the following info can be obtained using this search:

Press enter or click to view image in full size


Figure 3: Shodan Report for our alleged Botnet

As you can see, a total amount of 3996 devices return the identified pattern. However, this is far from the original size of this pattern. When we first looked at the data presented by Shodan, we were looking at 9345 IPs. Indeed, Shodan's history feature gives us a pretty telling history graph on the matter:

Press enter or click to view image in full size


Figure 4: Historical Pattern Analysis

As can be seen, somewhere between June and July 2022, the pattern of this botnet started appearing. It quickly amounted to 16.108 devices in a 2-month timespan reaching its peak distribution around August 2022. From there, the number of devices reflecting this pattern continuously decreased. However, there is a sudden strong decrease

between August and September 2023. It is yet unknown what is the cause of this. But there is an explanation for the constant decrease in numbers. We will get there later.

To no one's real surprise, the devices showing the above-mentioned behavior all seem to be IoT devices many of them connected via a residential IP. Among others, we identified HKVision software, Software related to TP-Link Routers, and Dahua digital video recorders. One of the most affected devices seems to be the TL-WR841N router which accounts for most of the infected devices in Bulgaria. We were able to identify that one of Bulgaria's largest ISPs issues this device to most of their clients. We will have a closer look at this device later on.

Observational sidenote: Curiously, the strong decrease in numbers coincides with the Qakbot Takedown at the end of August 2023.

**Chapter 3: Expanding Comms**

Well, after we jointly had mapped out this information about the supposed botnet, we attempted several things to find out more.

First of all, I obviously send out a tweet. I knew something was odd and I knew where to look, but what to look for? Without direct access to an infected device, it is hard to get ahold of any malware or identify any script used to spread malicious code. So to gather some ideas I sent out a tweet.

And while no one really had a good answer to my question, there was another interesting side effect. Shortly after making this tweet, a researcher going by the handle B1RD_D06 from the United States reached out to me because based on the information he was able to identify a threat that he himself had identified attacking their network during recent weeks. After some conversation, he disclosed that he was at that time working at a big company in the US Energy sector. He had shared his observations with a community of cybersecurity-responsible individuals in the energy sector and in fact, several other US companies in this sector were affected as well. The attacks were all targeted at members of the individual companies' C-Level members, which hints that this attack is somewhat more sophisticated than your average spray-and-pray brute-force attack.

I was then also contacted by a researcher from the French security company Intrinsec. Among other services, Intrinsec provides 24/7 SoC Monitoring and the particular researcher was interested in using the information about the botnet to increase the security of their clients and in turn to provide me with some helpful information on their insights into this particular threat. Besides that, we discussed the idea of reaching out to victims of the IoT attack chain, as to see if we could extract valuable information from the infected devices. In fact, the researcher from Goldilock had the exact same idea and in a later part of this blog post, we will have a look at our attempts at physically reverse-engineering an infected router. However, I first want to point out another observation that Intrinsec made after I had provided a list of IPs associated with this particular threat:

Press enter or click to view image in full size

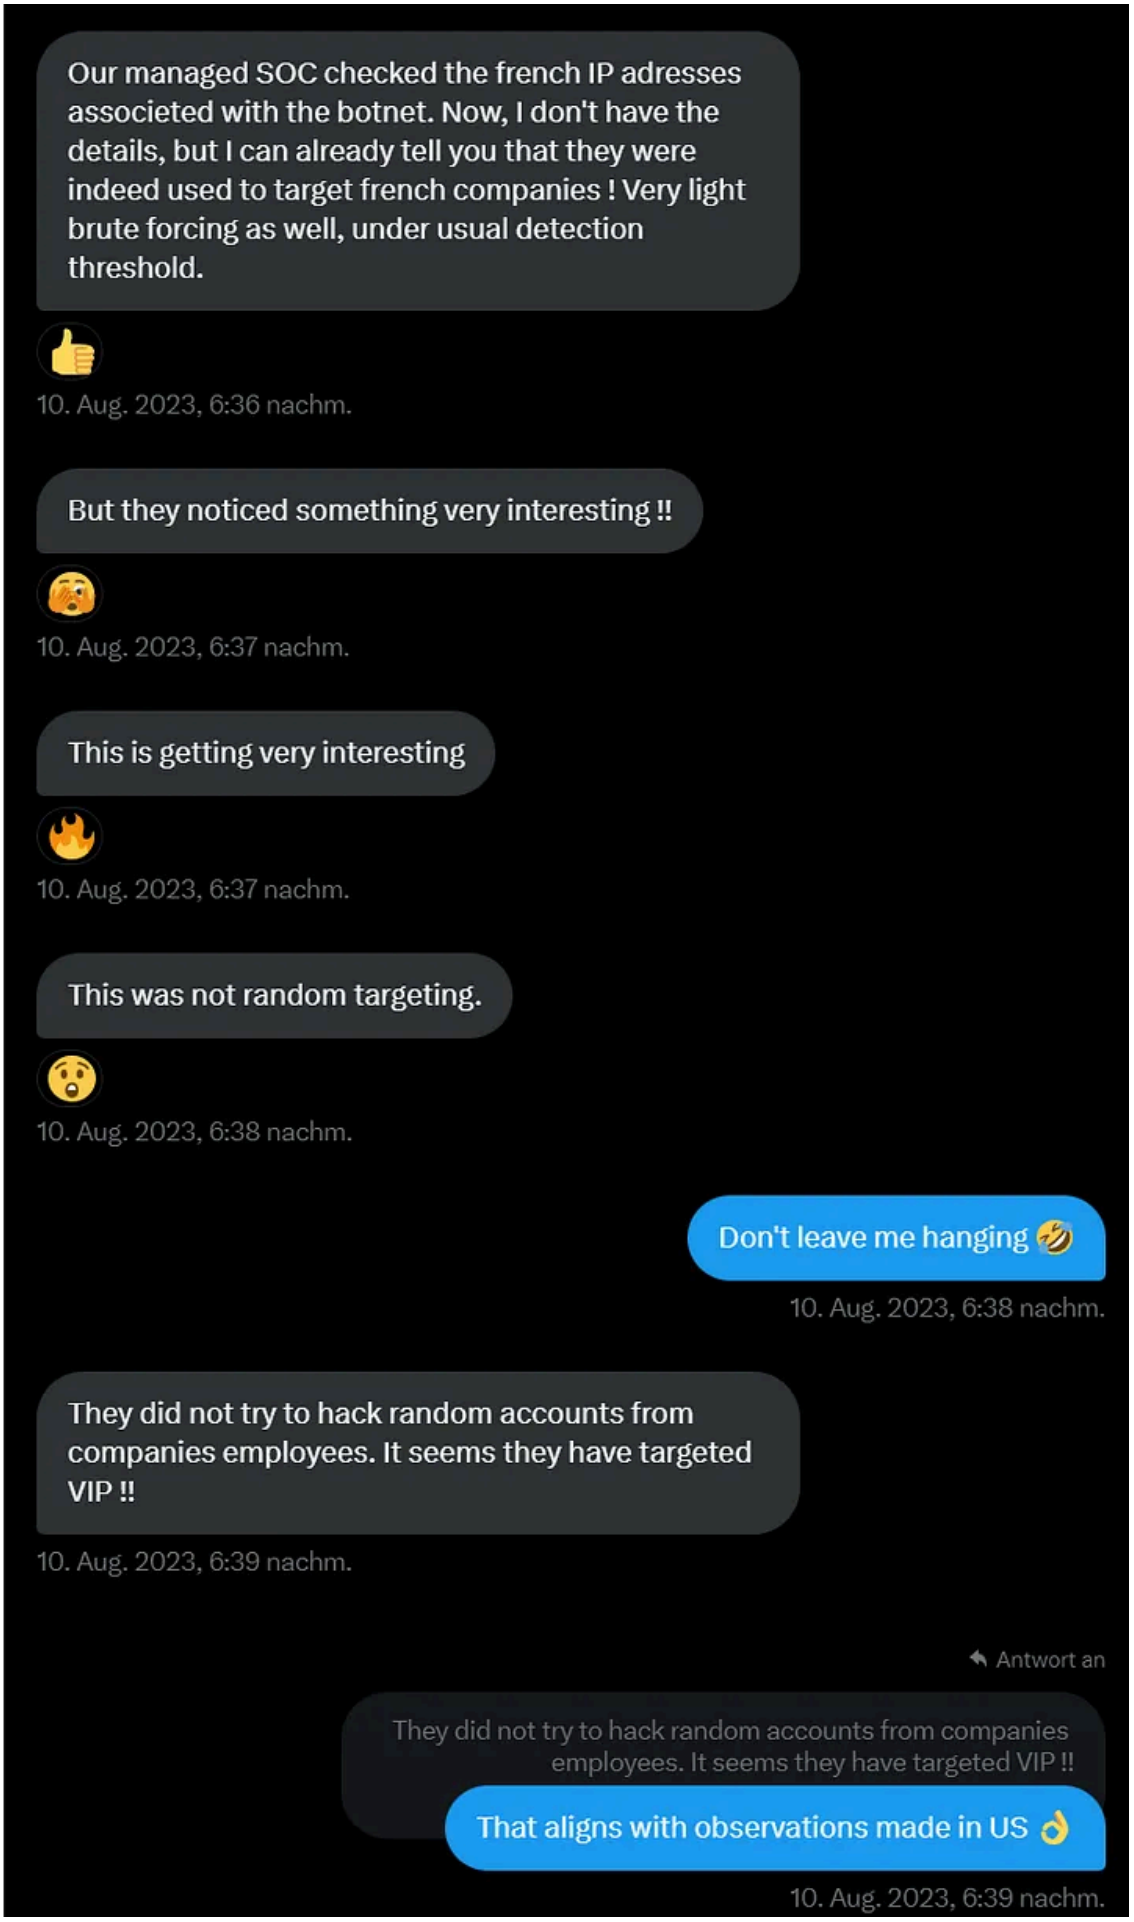
Our managed SOC checked the french IP adresses associeted with the botnet. Now, I don't have the details, but I can already tell you that they were indeed used to target french companies ! Very light brute forcing as well, under usual detection threshold.
10. Aug. 2023, 6:36 nachm.
But they noticed something very interesting !!
10. Aug. 2023, 6:37 nachm.
This is getting very interesting
10. Aug. 2023, 6:37 nachm.
This was not random targeting.
10. Aug. 2023, 6:38 nachm.
Don't leave me hanging
10. Aug. 2023, 6:38 nachm.
They did not try to hack random accounts from companies employees. It seems they have targeted VIP !!
10. Aug. 2023, 6:39 nachm.
Antwort an
They did not try to hack random accounts from companies employees. It seems they have targeted VIP !!
That aligns with observations made in US
10. Aug. 2023, 6:39 nachm.

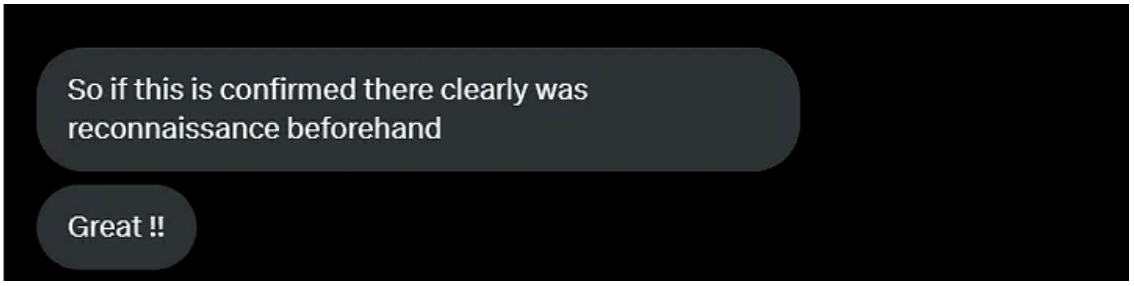


Figure 5: Intrinsec observed attacks (sorry for the emoji spam, I was excited ^^ )

So indeed after some further discussion with the SoC employees of Intrinsec, we were able to confirm that this botnet had also targeted several companies in France, using the attack method of low volume “Microsoft Azure PowerShell” brute-forcing targeting exclusively VIPs. However, Intrinsec was further able to confirm that there was no specific business sector targeted. Based on their findings and the rather limited insight I have into other business sectors in the US, I do believe the botnet does not target the Energy Sector exclusively but that other sectors are targeted as well. The goal does seem to be to compromise high-value targets, which could be an indication of an actor with financial motives.

A cross-check with the Greynoise platform, which collects and marks known botnet attacks and activity revealed a complete absence of these IPs from their tracker. This is a further indication that the observed botnet is not broadly targeting in a spray-and-pray-like pattern but actually working in a more targeted manner.

**Chapter 4: Physical Reversing Attempt**

Early on in this research, the researcher from Goldilock and I did agree that one of the best things to do to identify the malware used in this botnet attack was to get ahold of an infected device. With a bit of digging into the retrieved list of IP addresses from Shodan, we were able to identify several affected entities both in France and the UK which were in the area of the Goldilock and Intrinsec locations. And indeed after some digging and sending out several emails, the researcher at Goldilock was able to retrieve a router with the above-mentioned characteristics. As Goldilock itself is a producer of security hardware, my contact had sufficient expertise in trying to physically reverse engineer the device.

Press enter or click to view image in full size

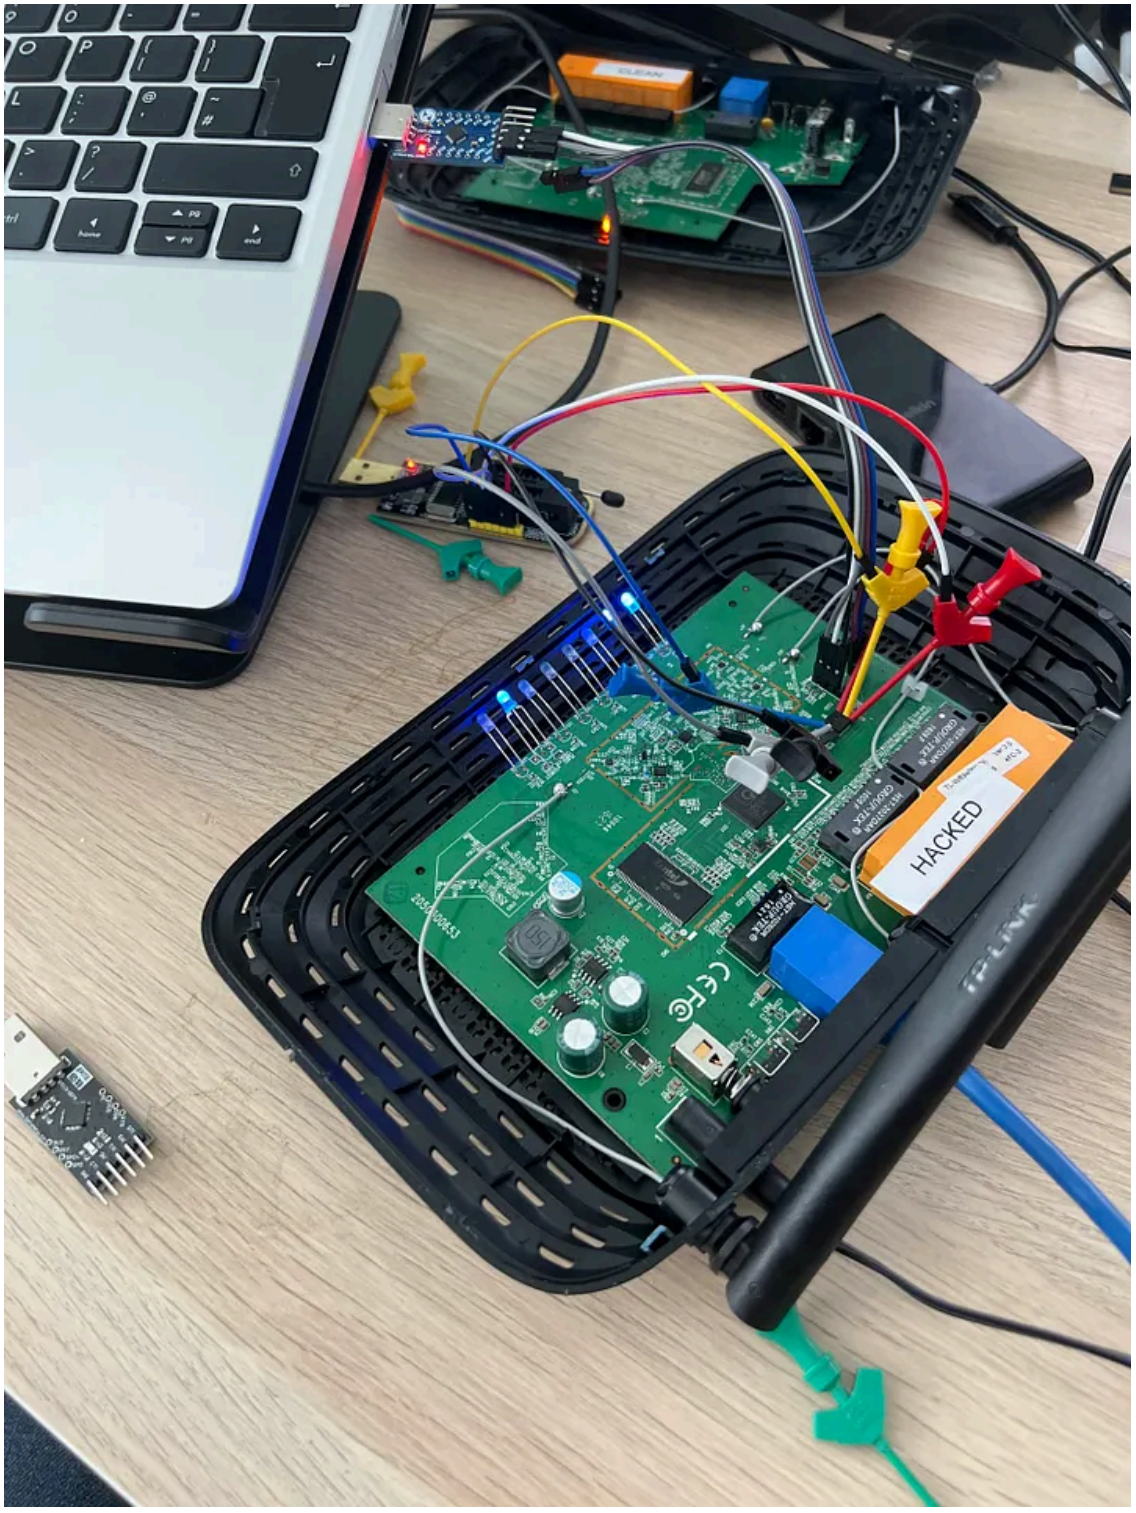


Figure 6: Physical reverse attempt of an infected TP-Link router

While this makes for some pretty neat images, we soon arrived at a sad conclusion: There was no evidence we could recover, which could be associated with an attack. In fact, after setting up the router at Goldilock's lab, there was no port 7777 open on the infected device. It was all gone.

## Get Gi7w0rm's stories in your inbox

Join Medium for free to get updates from this writer.

Remember me for faster sign in

To explain this, we need to look at the hardware level of IoT devices. The TP-Link TL-WR940N, like many of the basic models we noticed were infected, operates on an entry-level ARM/RISC-based System on Chip, in this case, the Qualcomm Atheros TP9343 processor. This system has a mere 4MB of persistent flash storage, most of it being occupied by the firmware and operating system.
The operating system used is a customized Linux Build Root image which is loaded into the router's RAM through SquashFS. As a result, if a Remote Code Execution (RCE) attack occurs, any malicious software or commands will only last until the router is turned off and on again. Once restarted, there's no evidence of the breach.
This could shed light on why we've seen a drop in the number of these routers over the past 6–8 months. Typically, users might restart their routers every few months, especially if they're experiencing connectivity or performance problems.

We did not get any further intel, but I share this information to make sure that other researchers who might look into this issue take the appropriate care when trying to analyze an infected device. Mistakes were made.

## Chapter 5: An attempt at attribution

So as you can see, after several months of research, our findings are still based mostly on IP addresses in Logs attempting brute force. Neither were we able to recover the associated malware nor were we able to recover enough evidence to identify the people behind it. However, there is one interesting observation that I have kept for the end of this post.
You see, while investigating the botnet, there was one IP which stood out:
**45.61.136[.]133**

When looking at this IP using the Shodan history tab, we can see that in January this year, the IP had a very similar pattern sitting on Port 7777:

Press enter or click to view image in full size

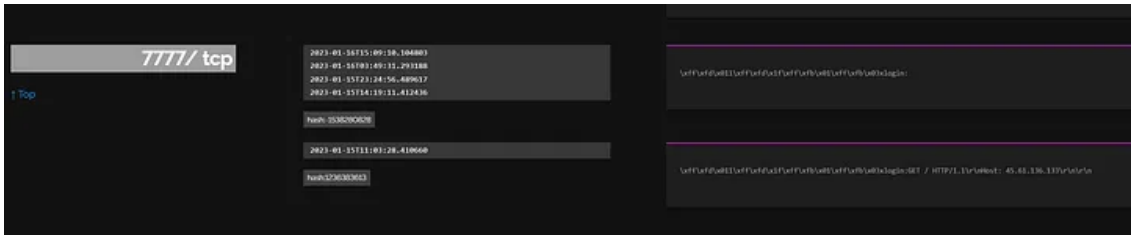


Figure 7: Port 7777 on a WebServer?

This particular IP stands out because it is the only IP we could identify, which is not an IoT device but an actual web server at a hosting company. Also the "xlogin:" pattern is slightly off according to Intrinsec CTI team, who determined a 1 Byte difference in the communication pattern. The fact that this is the only non-IoT device we have been observing makes me believe that it might have been a C2 Server for the botnet in or around January of this year.

***outdated paragraph***

*However, this leads to an interesting correlation. The IP **45.61.136[.]133** actually appears in a Threat Intelligence report by ReversingLabs.*

*The report, which was released in August 2023, discusses a Supply Chain attack on the Python Package Index (PyPI) repository. The attackers uploaded several different malicious Python packages which were designed to download additional malicious instructions after being imported into other projects.*

*In their article, ReversingLabs states that they attribute their observed campaign to the Lazarus actor based on the following observations:*

Press enter or click to view image in full size

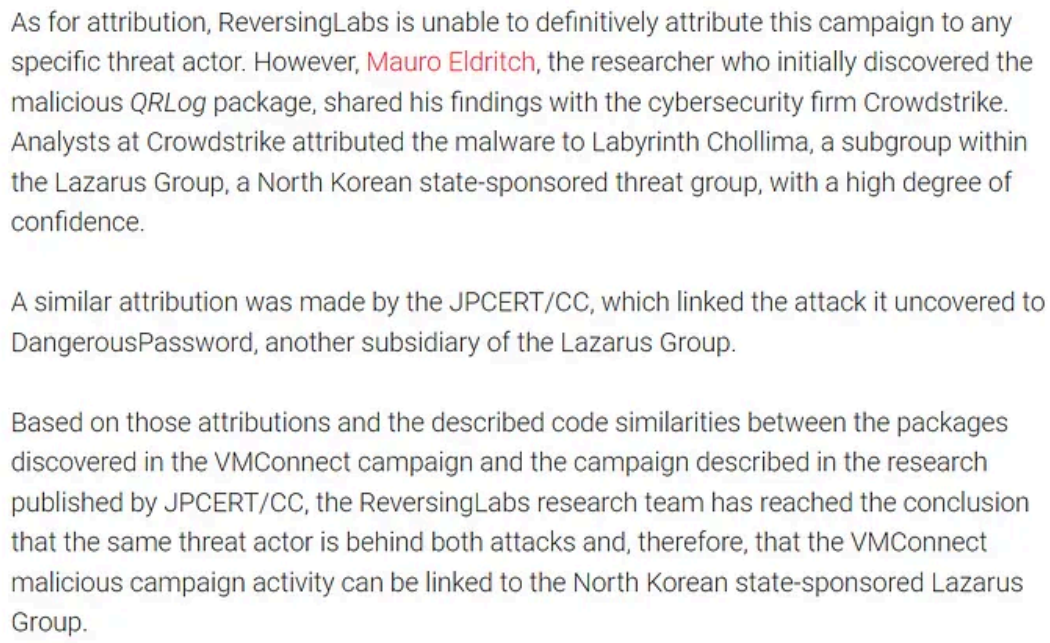
As for attribution, ReversingLabs is unable to definitively attribute this campaign to any specific threat actor. However, Mauro Eldritch, the researcher who initially discovered the malicious *QRLog* package, shared his findings with the cybersecurity firm Crowdstrike. Analysts at Crowdstrike attributed the malware to Labyrinth Chollima, a subgroup within the Lazarus Group, a North Korean state-sponsored threat group, with a high degree of confidence.

A similar attribution was made by the JPCERT/CC, which linked the attack it uncovered to DangerousPassword, another subsidiary of the Lazarus Group.

Based on those attributions and the described code similarities between the packages discovered in the VMConnect campaign and the campaign described in the research published by JPCERT/CC, the ReversingLabs research team has reached the conclusion that the same threat actor is behind both attacks and, therefore, that the VMConnect malicious campaign activity can be linked to the North Korean state-sponsored Lazarus Group.

*Figure 8: Lazarus?*

*However, here is where it gets a bit tricky. After reading this report I actually reached out to some contact at CrowdStrike about the attribution of this particular IP. And what I was told is that the CrowdStrike CTI product actually has this IP attributed to a completely different group:*
*UNC3944 / Storm-0875 / Scattered Spider. This didn't really make things easier. So I decided to double-check with a researcher from Microsoft MSTIC. They confirmed that the IP Address in question is attributed to UNC3944 / Storm-0875 / Scattered Spider. And as if this was not enough, RecordedFuture, another big Threat Intelligence vendor has it under Lazarus (Labyrinth Chollima). So that's a solid 2 vs 2 on the attribution site. And even if we would have a clear picture of whom to attribute the IP to, it wouldn't be safe to say the potential Botnet control server was run by this group or if the server was rented by yet another entity at the time of its malicious involvement in the Botnet activity.***outdated paragraph end***

**Update: After releasing this article, I was contacted by a fellow IT Security researcher. The researcher pointed out that at the time of the observed pattern, the IP 45.61.136[.]133 was owned by them. The researcher had indeed observed the same botnet pattern and activity and had**

**therefore made an attempt to copy the pattern. Their goal was to collect additional intelligence on the botnet's activity by trying to mimic a compromised device. Sadly their attempt was unsuccessful, likely because the botnet has some sort of C2 Server that actively registers new devices which is why his fake device did not receive any commands. The attribution made above is therefore obsolete.**

So yet again, our investigation had run into a dead-end. Even after more weeks of searching and some people at RecordedFuture and Microsoft running additional queries, no further intelligence was uncovered.

Thank you for reading this blog post. I really appreciate your time and hope you have learned something of value. Also thank you to:

- Goldilock
- the researchers of the Intrinsec SoC and their CTI team
- Dunstable Toblerone
- Fr0gger
- B1RD_D06
- SOSIntel
- aejleslie

who all to some extent took a role in investigating the 7777Botnet.

If any additional findings are based on this blog post I would love to be tagged on it as I am still very curious about this threat. Consider following me here (to get notified on upcoming posts) or on Twitter.

Cheers ❤

*Update 20.10.2023: Added a paragraph in Chapter 5 to reflect further intelligence gathered in regard to attribution. Please read the highlighted section!

---