

Analyzing ELF/Sshinjector.A!tr with a Human and Artificial Analyst | FortiGuard Labs

By Axelle Apvrille

Published: 2025-02-04 · Archived: 2026-04-05 20:01:47 UTC

Affected Platform: Linux

Impacted Users: Linux-based network appliances or IoT

Impact: Data exfiltration

Severity Level: Medium

ELF/Sshinjector.A!tr is a collection of malware that can be injected into the SSH daemon. Samples of this malware collection surfaced around mid-November 2024. While we have a good amount of threat intelligence on them (e.g., they are [attributed to the DaggerFly espionage group](#)), nobody seems to have looked into what they actually do. In this blog post, we will focus on the reverse engineering of the attack's binaries and how this reverse engineering was achieved.

Reverse Engineering of ELF/Sshinjector.A!tr

The attack uses several binaries:

- A dropper checks if the host is infected. If not, it drops all malicious binaries (see Figure 1) at the right places.
- A malicious SSH library named *libsshd.so* communicates with a remote bot master and will typically exfiltrate information.
- Several other infected binaries (*mainpasteheader*, *selfrecoverheader*,...) ensure the host remains infected (malware persistence).

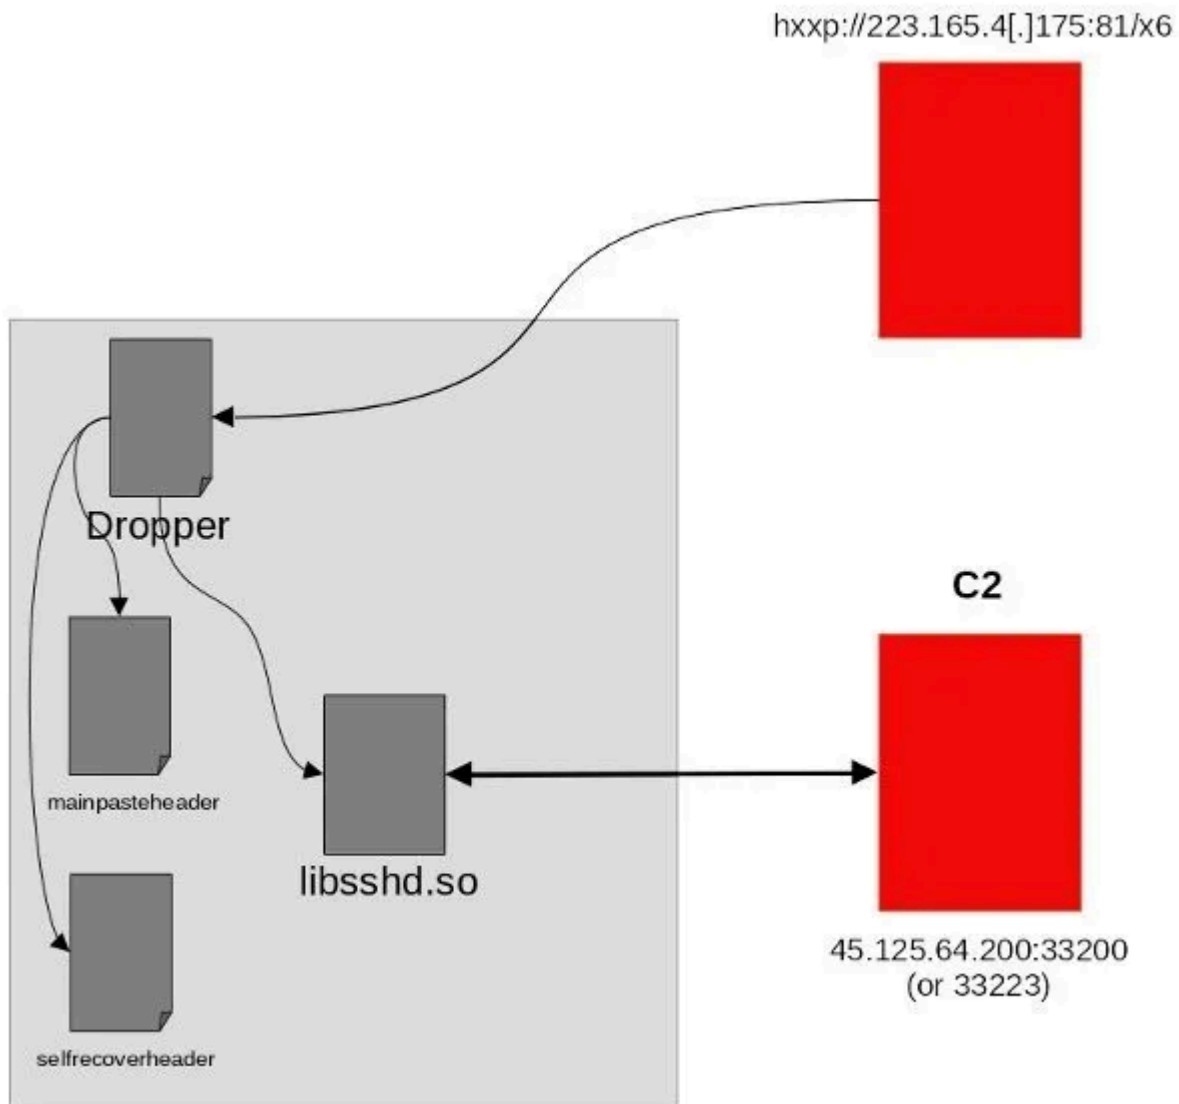


Figure 1: Overview of ELF/Sshdinjector

More precisely, the dropper checks if it is being run under root privileges and, if not, exits. It then checks whether the host is infected by searching for a file named */bin/lsxxxxsswwdd11vv* containing the word *WATERDROP*. If the host is not yet infected, it attempts to overwrite the legitimate binaries *ls*, *netstat*, and *crond* with infected binaries (*/bin/lsxxxxsswwdd11vv*, *selfrecoverheader*, *mainpasteheader*,...). Finally, it searches the SSH daemon and, if necessary, infects it with *libsshd.so*.

The malicious payload is contained in *libsshd.so*. The main element is found in a function named “*haha*.” It also creates two other threads from the functions “*heihei*” and “*xixi*”. Those three names refer to laughing in Chinese. Function “*xixi*” checks whether it has access to */root/intensify-mm-inject/xxx*, in which case it will kill and restart both the SSH and Cron daemons. Function “*heihei*” connects to the remote C2 (hard-coded IP address 45.125.64[.]200, port 33200 or 33223) and listens for incoming commands.

Command Id	Description
1	SERVER_REQ_BASE_INFO. Exfiltrates uname, MAC address etc to C2
2	List running services, by listing files in <i>/etc/init.d</i>
3	Reads users from <i>/etc/shadow</i>
4	Lists running process
5	Tests access to <i>/var/log/dmesg</i>
6	Tests access to <i>/tmp/fcontr.xml</i>
7	Lists a given directory
8	File transfer
9	Opens a shell terminal
10	Executes a command in the terminal
11	Unloads and exits the malicious process
12	Removes a file
13	Renames a file
1000	SERVER_RET_ONLINE_ACK

0x80000001	Client status change notification. It sends base info, service list, read <i>/etc/shadow</i> .
------------	--

Communication with the C2 uses its own protocol. All packets include a hard-coded UUID (*a273079c-3e0f-4847-a075-b4e1f9549e88*), an identifier (*afa8dcd81a854144*), and the response to the command.

AI-Assisted Malware Analysis

Reverse engineering was performed using Radare2, assisted by Generative AI through the [Radare2 extension](#) “r2ai.”

This study shows that AI provides excellent insights into the malware, delivering high-quality source code that complements the output obtained from a standard decompiler.

For example, I used **r2ai in “auto” mode**. In this mode, the user asks the AI a question, and the AI automatically performs the necessary steps with the radare2 disassembler to answer. This is particularly helpful for users who don’t know Radare2 well.

```
[r2ai:0x00405e32]> ' this binary checks that various other binaries are infected (e.g. sshd, ls, netstat...) and if not it copies infected version to replace them. This binary does not contain the infected versions, i.e we do not know what the infected versions do. Is this correct?
Let me verify this by:
1. Looking at the binary copying patterns
2. Checking if the infected versions are stored within the binary
> iz
nth  paddr      vaddr      len  size  section  type  string
-----
0  0x00024090  0x00424090  17   18    .rodata  ascii  open %s fail %d \n
1  0x000240a2  0x004240a2  14   15    .rodata  ascii  stat fail %d \n
2  0x000240b1  0x004240b1  19   20    .rodata  ascii  ./selfrecoverheader
3  0x000240c5  0x004240c5  17   18    .rodata  ascii  ./mainpasteheader
4  0x000240d7  0x004240d7  9    10    .rodata  ascii  ./nrplant
5  0x000240e1  0x004240e1  12   13    .rodata  ascii  ./libsshd.so
```

Figure 2: r2ai runs in auto mode and automatically issues r2 command “iz” to start working on the question.

In this screenshot, we see the AI automatically searches for strings in the binary, via r2’s command “iz”.

The overview of the dropper is excellent. **AI excels in reading large quantities of information and summarizing them.**

```
Yes, your analysis is correct! Here's what we can confirm:

1. The binary acts as an infector/installer that checks and replaces various system binaries:
  - ls ("/bin/lsxxxxsswwdd11vv")
  - netstat ("/bin/netstat")
  - sshd ("/usr/sbin/sshd")
  - crond ("/usr/sbin/crond" or "/usr/sbin/cron")

2. The binary does NOT contain the infected versions within itself:
  - Looking at the sections and strings, we don't see any large embedded binaries
  - The strings show paths to the infected versions (e.g., "/bin/lsxxxxsswwdd11vv") but not the actual malicious code
  - The error messages indicate it expects to find these infected versions elsewhere (e.g., "fail:netstat has already plant or not exist")

3. The program's behavior:
  - Checks if binaries are "infected" (using integrity checks)
  - If a binary is not infected, it tries to replace it with an infected version
  - Shows messages like "already plant" when it finds already infected binaries
  - Has error handling for cases where infected versions are not found

4. Additional interesting details:
  - Uses a shared library "libsshd.so" for some functionality
  - Has anti-tampering mechanisms
  - Includes root checking ("no root!" message)
  - Attempts to disable SELinux ("setenforce 0")

So this appears to be an installer/infector component of a larger malware system, where the actual malicious binaries are distributed separately. The program acts as the mechanism to check for and deploy these infected versions, but doesn't contain them itself.
```

Figure 3: The AI summarizes quite well the behavior of the malware.

We can then ask the AI to decompile the main. **While the AI-generated source code is easy to understand, its details are not always correct.** By comparison, source code produced by decompilers is often difficult to read but is accurate. Because of this, it is important to remember that these approaches complement each other and are ideally viewed side by side.

```
int main(int argc, char **argv) {
    // Set global flag
    global_control_flag = 1;

    // Process initial file setup
    process_file_setup(argv[0], global_data, 4, filename); // AXELLE REMARK: not detailed enough

    // Check for existing installation
    snprintf(command_buffer, 100, "%s", NRPLANT_PATH);
    system(command_buffer); // AXELLE REMARK: wrong. Does not use "system" but searches for WATERDROP to see if infected
    remove(filename);
    // AXELLE REMARK: g_use_malicious_ls = 1 if the program was found to be infected

    // Initialize injection system
    if (initialize_injection()) {
        // Check for previous installation
        if (check_watermark("WATERDROP") != 0) {
            printf("fail!!! netstat have already plant?\n");
            return -1;
        }

        // Handle mainpaste header
        if (!setup_header("./mainpasteheader")) {
            printf("setuid fail!\n");
            return -1;
        }

        // Copy and setup recovery files
        copy_file("./selfrecoverheader", "./selfrecoverheader1");
        copy_file("./mainpasteheader", "./mainpasteheader1");

        // Handle LS binary replacement
        copy_file("/bin/lsxxxxsswwdd11vv", "./ls");
        inject_binary("./selfrecoverheader1", "./ls", 1); // AXELLE REMARK: wrong, this copy file + name replacement, not injection.
        link_binary("./mainpasteheader1", "./ls", 0); // AXELLE REMARK: wrong, this is file copy...

        // Check if we should use malicious ls
        if (g_use_malicious_ls) {
            copy_file("/bin/lsxxxxsswwdd11vv", "./ls");
        } else {
            printf("ls has been ignor in this version!\n");
        }
        remove("./ls");

        // Handle netstat replacement
        copy_file("./mainpasteheader", "./mainpasteheader2");
    }
}
```

Figure 4: This source code was generated by the AI, via r2ai. It is globally correct, readable and useful. Only comments marked “AXELLE REMARK” are my own, and highlight a few errors of the AI.

While AI performs very well, there are many cases in which it does not produce a satisfactory answer, at least not at first.

The most common issue is **hallucination**, wherein the AI *invents* something that isn’t true. Worse, it’s not always easy for a human analyst to spot hallucinations because the AI can sound very convincing. For example, in the AI-generated code below, the AI completely created an upload and a download command that is pure invention.

```
case SERVER_REQ_FILE_DOWNLOAD:
    std::string file_path = getPacketString(packet_data, &index);
    handleFileDownload(pid, client_id, proc_id, taskid, file_path);
    break;

case SERVER_REQ_FILE_UPLOAD:
    std::string src = getPacketString(packet_data, &index);
    std::string dst = getPacketString(packet_data, &index);
    handleFileUpload(pid, client_id, proc_id, taskid, src, dst);
    break;
```

Figure 5: An example of AI hallucination: the botnet does not have any FILE_DOWNLOAD nor FILE_UPLOAD command. This is an invention, misunderstanding an existing “file copy” feature.

Another frequent issue is **extrapolation**, where the AI does not totally invent something but *extrapolates* it. For example, the AI says the malware “manipulates” the MAC address. This is far-fetched. While it *creates* a string containing the host's MAC address and exfiltrates it, there is no modification of the MAC address. In a related example, the AI claims the malware hides its network communications. It does not. It hides on the OS by infecting common binaries such as netstat, but it does not attempt to hide the communication itself.

```
2. **Communication Patterns**:  
- The malware appears to use socket operations for both listening and connecting  
- It includes functionality to manipulate network interfaces (using `ioctl` calls)  
- Contains MAC address manipulation code (seen in the formatting of MAC addresses with `%2.2X:%2.2X:%2.2X:%2.2X:%2.2X:%2.2X`)  
3. **Stealth Mechanisms**:  
- The malware modifies `netstat` to hide its network connections  
- It replaces system binaries that could detect its network activity  
- Uses local socket manipulation to potentially hide its network presence
```

WRONG

Figure 6: AI extrapolation. The sentences outlined in red have been largely exaggerated by the AI.

Yet another issue is **omissions**. Omissions are the downside of AI’s power to summarize situations. Its summaries often lack the details a human would find important. For example, at some point, the malware tests access to a file named `/tmp/fcontr.xml`. Despite this being absolutely clear in the assembly, **the first version of AI-generated code completely eluded this part**. The solution to this issue is to **ask again by modifying the question/prompt** for the AI. In this case, I simply added to the end of the prompt: “Please pay attention to what is around `fcontr.xml`” and it solved the problem. Of course, this requires knowing that something had been omitted in the first place.

In fact, **interactions with AI are seldom perfect in a single shot**. Rather, they could be compared to a discussion with a capable colleague with impressive knowledge and intelligence but less intuition and experience. For this research, I kept my disassembler open. I used it several times to check for hallucinations, assist the AI (!) when it failed to find correct addresses or cross-references, or guide it to look into interesting parts.

Conclusion

While disassemblers and decompilers have improved over the last decade, this cannot be compared to the level of innovation we are seeing with AI. This is **outstanding!**

AI is particularly good at providing overviews of samples and generating easy-to-understand source code.

Fortunately—or not?—**AI cannot work alone and must be piloted and complemented by competent human analysis to spot hallucinations** (the most dangerous issues), **refine questions, identify omissions, or guide** the AI in the most interesting direction.

I haven’t discussed *language models yet*. Obviously, r2ai’s results depend on the language model used. Language models are configurable, and we can conveniently switch from one to another, whether a local model or a remote one, free access or paid. The results from this article were mainly obtained using Claude 3.5 Sonnet 2024-10-22.

Last but not least, this blog post was written without AI assistance ;-)

Fortinet Protections

Fortinet customers are already protected from this malware variant through our AntiVirus as follows: FortiGuard Labs detects the sample with the following AV signatures:

ELF/Sshinjector.A !tr and Linux/Agent.ACQ!tr

The [FortiGuard AntiVirus service](#) is supported by FortiGate, FortiMail, FortiClient, and FortiEDR. Fortinet EPP customers running current AntiVirus updates are also protected.

IOCs

94e8540ea39893b6be910cfee0331766e4a199684b0360e367741facca74191f

0e2ed47c0a1ba3e1f07711fb90ac8d79cb3af43e82aa4151e5c7d210c96baebb

6d08ba82bb61b0910a06a71a61b38e720d88f556c527b8463a11c1b68287ce84

Source: <https://www.fortinet.com/blog/threat-research/analyzing-elf-sshinjector-with-a-human-and-artificial-analyst>