

Linux IR - AI-Assisted Malware Analysis

Taz Wake | Halkyn Consulting Ltd | Originally published 2024-06-06

Introduction

Incident response often has to be fast. We are chasing an active attacker and trying to get control of a situation before anything gets worse. It is not often that we have the luxury of time.

This means that *sometimes* we have to find "quick and dirty" ways to resolve issues. In turn, this often carries some sacrifices - we give up a deep understanding and, instead, rely on a high-level review to make decisions.

Where this becomes most apparent on Linux intrusions, is dealing with malware. Even if we have the right skills and knowledge to take apart complex ELF files, it is rare we can put aside enough time to get value.

In this article, I will look at ways you can get a superficial assessment by running some basic commands and sharing the data with ChatGPT (or any AI/LLM platform).


Note 1: None of this is a replacement for having trained, capable, and experienced reverse engineers available. The "quick and dirty" approach is just to help you deal with the early stage of the incident. You absolutely need a deeper understanding at some point.

Note 2: I am going to use ChatGPT here but you need to consider your own situation. You need to make sure you have appropriate approvals to upload data to public platforms and you need to understand enough that you can identify where the platform is making a mistake (or just plain old lying to you). This is 100% not a replacement for skilled staff, it is just a way to improve your efficiency.



AI platforms have limits. They can be great for quick assessments but rarely stand up to scrutiny...

Overview



High-level workflow

The high-level workflow basically has four steps:

- Discover a suspicious file during an incident.
- Run basic analysis commands, save the output as a text file.
- Share output with your AI/LLM model and provide a prompt to support your analysis.
- Review output and act on any new knowledge.

The good news is that a lot of this is definitely scriptable and you could easily build this into your SOAR application with API calls to the AI/LLM platform. In this article we will focus on the basic commands to run and some suggested prompts.

A lot of this process will rely on you keeping good notes during your IR work. It should go without saying that this is critically important!

Basic Analysis Commands

In this article, I will use `malware.elf` as a placeholder for the file you are analysing.

Basic Information

- Note the filename, where it was discovered and any immediately relevant information or clues about its behaviour.
- Collect file metadata

This is an example of what the output should look like, although in use this would be redirected to a text file.

```
root@siftworkstation:~/malware# file malware.elf
malware.elf: ELF 64-bit LSB executable, x86-64, version 1 (SYSV), dynamically linked, interpreter /lib64/ld-linux-x86-64.so.2,
missing section headers at 105168
```

```
root@siftworkstation:~/malware# sha1sum malware.elf
7f59204b53cd7fab4ce40e73a97115130e7382ee  malware.elf
```

Static Analysis

readelf

The `readelf` command is used to extract information from an ELF file. It provides details on things such as the section headers, program headers, symbols etc. It has a range of command line arguments which allow you specify what specific items you are interested in, but for our purposes we will use the **-a** argument to collect all headers.

The output which is redirected into the text file should look something like this:

```

root@siftworkstation:~/malware# readelf -a malware.elf
readelf: Error: ELF Header:
Reading 1792 bytes extends past end of file for section headers
  Magic:   7f 45 4c 46 02 01 01 00 00 00 00 00 00 00 00 00
  Class:                               ELF64
readelf: Error:   Data:                               2's complement, little endian
Section headers are not available!
  Version:                               1 (current)
  OS/ABI:                               UNIX - System V
  ABI Version:                           0
  Type:                                  EXEC (Executable file)
  Machine:                               Advanced Micro Devices X86-64
  Version:                               0x1
  Entry point address:                   0x401530
  Start of program headers:              64 (bytes into file)
  Start of section headers:              103440 (bytes into file)
  Flags:                                 0x0
  Size of this header:                   64 (bytes)
  Size of program headers:               56 (bytes)
  Number of program headers:              9
  Size of section headers:               64 (bytes)
  Number of section headers:              28
  Section header string table index: 27

Program Headers:
  Type           Offset             VirtAddr           PhysAddr
                 FileSiz              MemSiz              Flags  Align
PHDR             0x0000000000000040 0x0000000000004000 0x000000000000400040
                 0x00000000000001f8 0x00000000000001f8  R E    0x8
INTERP           0x0000000000000238 0x000000000000400238 0x000000000000400238
                 0x000000000000001c 0x000000000000001c  R     0x1
[Requesting program interpreter: /lib64/ld-linux-x86-64.so.2]

```

Example readelf output

objdump

Objdump is similar but different. Again, there are a range of options we can use to specify which objects we want, but for speed and simplicity, we will use the -D argument to disassemble all the sections, not just the ones expected to contain instructions. Please be aware that this is very noisy and might result in several hundred thousand lines of output. It might also result in nothing depending on how the malware is compiled.

Disassembly of section .interp:

0000000000400238 <.interp>:

```
400238: 2f          (bad)
400239: 6c          insb    (%dx),%es:(%rdi)
40023a: 69 62 36 34 2f 6c 64 imul    $0x646c2f34,0x36(%rdx),%esp
400241: 2d 6c 69 6e 75    sub    $0x756e696c,%eax
400246: 78 2d        js     400275 <_ZSt15__once_callable@@Base+0x40025d>
400248: 78 38        js     400282 <_ZSt15__once_callable@@Base+0x40026a>
40024a: 36 2d 36 34 2e 73    ss sub $0x732e3436,%eax
400250: 6f          outsl   %ds:(%rsi),(%dx)
400251: 2e 32 00     cs xor (%rax),%al
```

Disassembly of section .note.ABI-tag:

0000000000400254 <.note.ABI-tag>:

```
400254: 04 00          add    $0x0,%al
400256: 00 00          add    %al,(%rax)
400258: 10 00          adc    %al,(%rax)
40025a: 00 00          add    %al,(%rax)
40025c: 01 00          add    %eax,(%rax)
40025e: 00 00          add    %al,(%rax)
400260: 47 00          add    %eax,(%rax)
```

Example Objdump output

strings

Often overlooked, strings is an excellent way of identifying human-readable content in a file. You can specify the minimum length string and, as a general guide, I find starting with -n8 is a good approach.

```
uname -a && echo " | " && hostname
fatal error, no cfg!
esxcli --formatter=csv --format-param=fields=="WorldID,DisplayName" vm process list | awk -F "\"*,\"" '{system("esxcli vm proc
ess kill --type=force --world-id=" $1)}'
Error create note in dir %s
Error parse cfg
fatal error malloc enc
File is encrypted by data %s
[%s] already encrypted
[%s] is protected by os
Encrypting [%s]
File [%s] was NOT encrypted
File [%s] was encrypted
iji iji iji iji iji iji tYiji iji iji ifi iji iji iji
iji iji iji iji iji iji -----ji iji ifi iji iji iji
iji iji iji iji iji ENCRYPTED |ji iji ifi iji iji iji
iji iji iji iji iji| - - - - -ji iji ifi iji iji iji
iji iji iji iji iji| %08d |ji iji ifi iji iji iji
iji iji iji iji iji| FILES |ji iji ifi iji iji iji
iji iji iji iji iji| 'ji iji ifi iji iji iji
Usage example: elf.exe --path /vmfs/ --threads 5
without --path encrypts current dir
--silent (-s) use for not stoping VMs mode
!!!BY DEFAULT THIS SOFTWARE USES 50 THREADS!!!
Unable to find path argument
Key initialization error ... something wrong with config
Using silent mode, if you on esxi - stop VMs manualy
ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789+/
/dev/urandom
Error open urandm line %d!
Error read urandm line %d!
```

Example output from running strings on a malware sample

ldd

Finally a quick check of any libraries the binary calls with the ldd command.

This varies in effectiveness, with lots of malware families (especially ransomware) coming up empty here.

```
root@siftworkstation:~/malware# ldd suspicious
linux-vdso.so.1 (0x00007ffe8e8c3000)
libpthread.so.0 => /lib/x86_64-linux-gnu/libpthread.so.0 (0x00007f1b0a889000)
libm.so.6 => /lib/x86_64-linux-gnu/libm.so.6 (0x00007f1b0a7a2000)
libc.so.6 => /lib/x86_64-linux-gnu/libc.so.6 (0x00007f1b0a57a000)
/lib64/ld-linux-x86-64.so.2 (0x00007f1b0a8a8000)
```

ldd output from a SIDEWALK malware sample

Dynamic Analysis

Next is checking how the application works. This does present an element of risk as for strace and ltrace, the binary has to be made executable. **Always approach any analysis with caution.**

gdb

You can use gdb (the "Gnu Debugger") to get a rapid assessment of some key components of the file. You can do this manually, or for speed, we can use some command line arguments to extract commonly analysed elements.

For our purposes this command will be useful, *most* of the time.

This is an example of how the output might look.

```
(No debugging symbols found in suspicious)
Symbols from "/root/malware/suspicious".
Local exec file:
  `/root/malware/suspicious', file type elf64-x86-64.
Entry point: 0x45acb4
0x000000000400238 - 0x000000000400254 is .interp
0x000000000400254 - 0x000000000400274 is .note.ABI-tag
0x000000000400278 - 0x000000000408bf4 is .hash
0x000000000408bf8 - 0x000000000424468 is .dynsym
0x000000000424468 - 0x0000000004546a8 is .dynstr
0x0000000004546a8 - 0x000000000456b5c is .gnu.version
0x000000000456b60 - 0x000000000456c30 is .gnu.version_r
0x000000000456c30 - 0x0000000004587d8 is .rela.dyn
0x0000000004587d8 - 0x000000000459c00 is .rela.plt
0x000000000459c00 - 0x000000000459c18 is .init
```

Example output from a SIDEWALK malware sample

For some malware samples this won't work - especially if they have a lot of protection mechanisms built in, so don't be too surprised if the output is effectively blank. A good reverse engineer will work around this but that is outside the scope of our rapid assessment.

```
No symbol table is loaded. Use the "file" command.  
All defined functions:  
All defined variables:  
No stack.
```

Example from a Sodinokibi ransomware sample



While gdb is generally safe to run, the next two commands do significantly increase the risks. **The file must be executable and the binary is running on the system.** If this is a malware sample, it is likely to be able to run its payloads.

Only carry out these steps if you have an isolated analysis environment or sandbox where the impact from (for example) deploying ransomware would be limited.

strace

This command attaches a debugger to the running process and records system calls. At the most basic use, it will monitor the file until it exits which can be a considerable time. It is often a good idea to use the timeout command to avoid any system hangs.

In this example, strace will write to an output file (-o), follow child processes (-f) and add a timestamp with millisecond precision to each entry. It will trace all system calls (-e trace=all) and capture up to 256 characters of strings.

In use, you have to specify the path correctly so either use full paths or strings to point to objects in the same folder.

Alternatively, you can run the command with a timeout. This makes it much faster to run but can miss malicious activity if there are delays built in (a common anti-analysis technique).

This example will kill the process after 10 seconds (timeout 10) and will summarise the trace (-c). The summary causes the output to lose millisecond precision timestamps.

[illegible]

Example of strace output against a SIDEWALK malware sample with millisecond timestamps

ltrace

ltrace is similar to strace in that it runs the file, however, this time it intercepts and records dynamic library calls.

In this example, the commands are:

- o output file.
- f follow child processes
- S show system calls
- C translate C++ symbols to make the output more readable
- tt microsecond timestamps

This is also a good candidate for using the `timeout` command to prevent it hanging for long periods of time.

```

6496 18:20:55.173653 SYS_brk(0) = 0x14cb000
6496 18:20:55.174861 SYS_arch_prctl(0x3001, 0x7ffe12c513e0, 0x7fddce6b23c0, 1) = -22
6496 18:20:55.175140 SYS_mmap(0, 8192, 3, 34) = 0x7fddce69100
6496 18:20:55.175334 SYS_access("/etc/ld.so.preload", 04) = -2
6496 18:20:55.175565 SYS_openat(0xffffffff9c, 0x7fddce6c121b, 0x80000, 0) = 3
6496 18:20:55.176120 SYS_newfstatat(3, 0x7fddce6c1ee9, 0x7ffe12c50530, 4096) = 0
6496 18:20:55.176603 SYS_mmap(0, 0x17037, 1, 2) = 0x7fddce67900
6496 18:20:55.176819 SYS_close(3) = 0
6496 18:20:55.177006 SYS_openat(0xffffffff9c, 0x7fddce6914b0, 0x80000, 0) = 3
6496 18:20:55.177460 SYS_read(3, "\177ELF\002\001\001", 832) = 832
6496 18:20:55.177810 SYS_newfstatat(3, 0x7fddce6c1ee9, 0x7ffe12c50600, 4096) = 0
6496 18:20:55.178183 SYS_mmap(0, 0x4028, 1, 2050) = 0x7fddce67400
6496 18:20:55.178615 SYS_mmap(0x7fddce675000, 4096, 5, 2066) = 0x7fddce675000
6496 18:20:55.178990 SYS_mmap(0x7fddce676000, 4096, 1, 2066) = 0x7fddce676000
6496 18:20:55.179344 SYS_mmap(0x7fddce677000, 8192, 3, 2066) = 0x7fddce677000
6496 18:20:55.179635 SYS_close(3) = 0
6496 18:20:55.179830 SYS_openat(0xffffffff9c, 0x7fddce6919d0, 0x80000, 0) = 3
6496 18:20:55.180022 SYS_read(3, "\177ELF\002\001\001\003", 832) = 832
6496 18:20:55.181249 SYS_newfstatat(3, 0x7fddce6c1ee9, 0x7ffe12c505e0, 4096) = 0
6496 18:20:55.181735 SYS_mmap(0, 0xe6108, 1, 2050) = 0x7fddce58d000
6496 18:20:55.182056 SYS_mmap(0x7fddce59b000, 0x7c000, 5, 2066) = 0x7fddce59b000
6496 18:20:55.182385 SYS_mmap(0x7fddce617000, 0x5b000, 1, 2066) = 0x7fddce617000
6496 18:20:55.182615 SYS_mmap(0x7fddce672000, 8192, 3, 2066) = 0x7fddce672000
6496 18:20:55.182869 SYS_close(3) = 0
6496 18:20:55.183033 SYS_openat(0xffffffff9c, 0x7fddce691ed0, 0x80000, 0) = 3
6496 18:20:55.183256 SYS_read(3, "\177ELF\002\001\001\003", 832) = 832

```

Example ltrace output from a SIDEWALK malware sample

Scripting it

Now, when we have a situation where we run the same set of commands each time, then it is crying out to be scripted. This is no exception.

A starter example using a bash script to run the commands above is available at <https://for577.com/analysiscript>

This is not likely to be perfect in every environment, instead, it should be seen as a starting point to build your own. If you have access to an AI/LLM with an API gateway you could even look to automate the submission and response parts.

Working with "AI"

Once you have carried out the basic analysis, it is time to ask your AI platform for help. The exact syntax and process here will depend on your investigation (and the platform) but, in general, it works best to provide the evidence and ask for assistance in analysing the files.

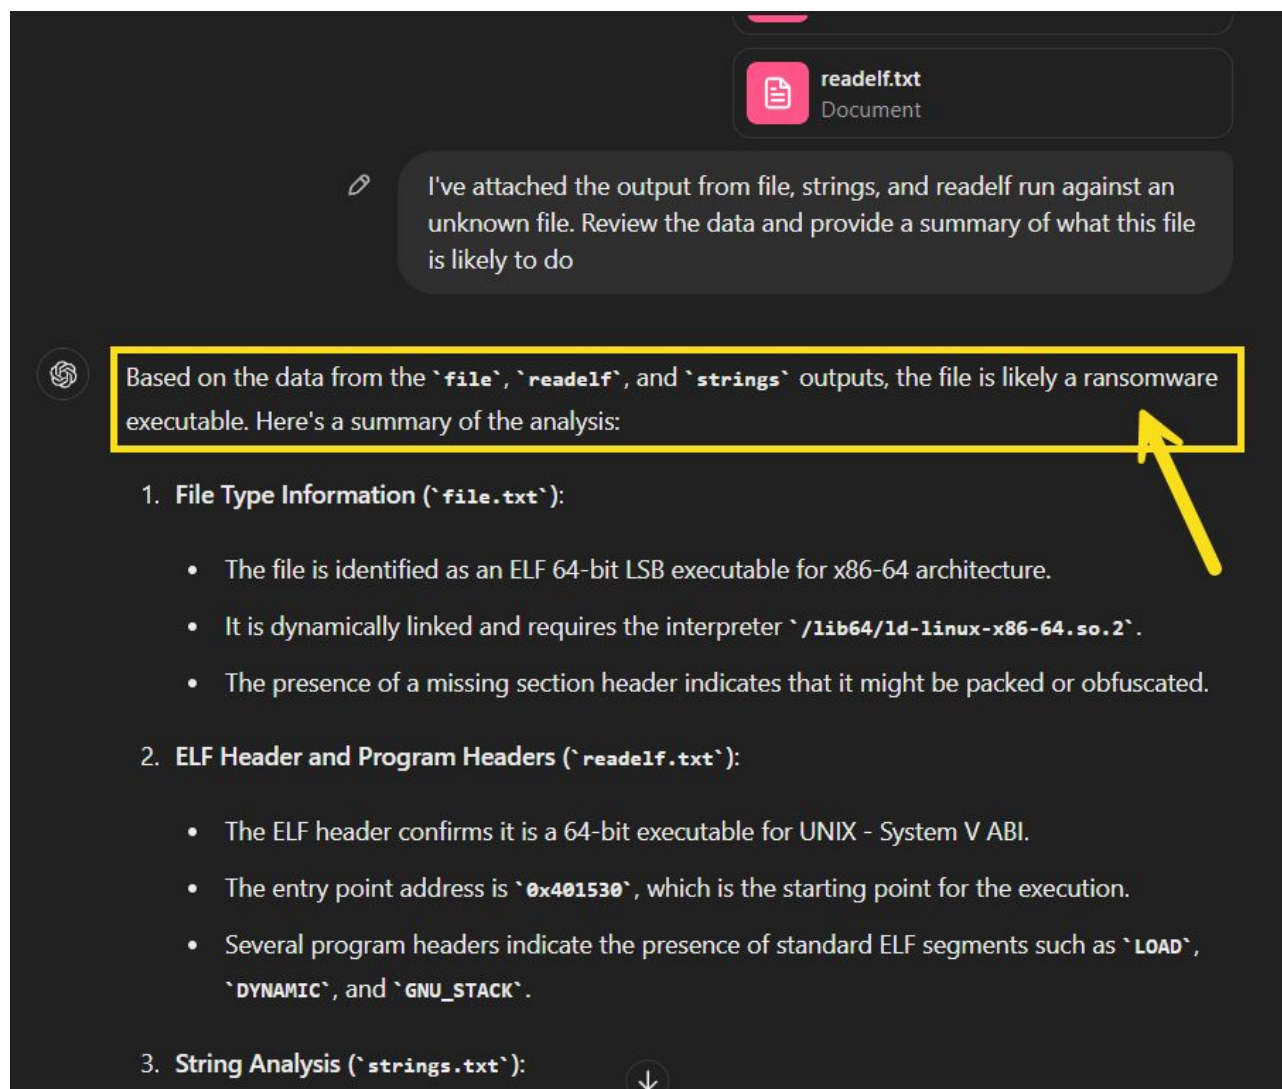
Most LLM tools will struggle to make a definite "malware/not malware" assessment and you need to be wary of their answers. Instead, it is better to have the tool explain what the software does and then make the determination yourself.

Some example prompts which have worked well for me are:

I've attached the output from file, strings, and readelf run against an unknown file. Review the data and provide a summary of what this file is likely to do

(This was for a suspected ransomware sample and no dynamic analysis took place)

With this information, ChatGPT was able to quickly identify the file as probably ransomware. In this case, it was correct and the sample was Sodinokibi (<https://www.virustotal.com/gui/file/a322b230a3451fd11dcfe72af4da1df07183d6aaf1ab9e062f0e6b14cf6d23cd>)



Correctly identified Sodinokibi

It is worth noting that this detection was almost entirely down to the strings data.

Another example is:

Attached is the output from gdb, strings, file, ldd, readelf, strace and ltrace for a suspicious file. Read them and provide a summary of what this application is likely to do. Also provide a yara rule to hunt for other samples.

With this sample, objdump was 15mb in size and rejected by ChatGPT. However, the other files were sufficient for a good determination. Although the response was verbose, it did correctly identify that this was a C2 implant.



Summary of the Suspicious Application

Based on the analysis of the files provided, here is a summary of what the suspicious application is likely to do:

- General Functionality:**
 - The application appears to handle various HTTP requests and responses, indicating some networking capabilities.
 - It has functionalities related to file management, such as uploading and downloading files, suggesting potential use in data exfiltration or malware propagation.
- Networking and Communication:**
 - There are references to handling HTTP requests, setting proxies, and dealing with HTTP responses, indicating it can communicate over the internet using HTTP protocols.
 - It also includes functions for SOCKS proxy management, which can be used to anonymize or hide network traffic.
- Task Scheduling and Execution:**
 - The application includes a module for scheduling tasks (Cron jobs), which can be used to execute specific commands or payloads at predefined times.
 - Functions related to handling scheduled jobs, such as adding, updating, and deleting tasks, are present.
- Data Handling and Transformation:**
 - There are functions for encoding and decoding data (e.g., Base64 encoding), which can be used to obfuscate or transform data.
 - It also includes file handling operations like reading and writing files, which are crucial for manipulating data on the disk.
- Potential Malicious Activities:**
 - Given the presence of HTTP request handling, file management, and proxy support, the application may be used for:
 - Data exfiltration: Uploading sensitive information to a remote server.
 - Command and Control (C2) communication: Receiving instructions from a remote server.
 - Propagation: Downloading and executing additional malicious payloads.

YARA Rule for Detection

Here is a YARA rule that can be used to detect other samples of this application:

yaraCopy code

```
rule Suspicious_Application {
  meta:
    description = "Detects suspicious application with HTTP, file management, and proxy support"
    author = "Your Name"
    date = "2024-06-06"
    version = "1.0"

  strings:
    $http_request = "HttpRequest::Send" fullword ascii
    $proxy_set = "HttpRequest::SetProxy" fullword ascii
    $socks_mgr = "SocksProxyMgr::" fullword ascii
    $upload_init = "FileMgr::UploadFileInit" fullword ascii
    $cron_parser = "CronParser::Parse" fullword ascii
    $task_scheduler = "TaskSchedulerMod::HandleAddTask" fullword ascii

  condition:
    all of them
}
```

This rule matches on specific function names found within the suspicious application, focusing on HTTP communication, proxy management, file handling, and task scheduling capabilities. Adjustments to the strings and condition can be made based on more detailed analysis or specific needs.

ChatGPT output based on analysis of a SIDEWALK implant.

The Yara rules might need some tweaking but definitely provide a good starting point for incident responders.

Conclusion

Using an AI/LLM can, in the right hands, speed up the incident response cycle and help free up DFIR staff for other tasks. It is important to remember that it does need skilled, knowledgeable staff to get good results.

In the examples here, it took under four minutes to collect the data on the Sodinokbi sample and have a determination, without using API calls. The SIDEWALK took slightly longer (six minutes) but did generate more useful data with dynamic analysis.

It is definitely worth adding this to your DFIR tool box if your organisation allows you to submit data to AI platforms, or if you have an internal AI tool.

About Halkyn Consulting

Halkyn Consulting Ltd is a UK-based security and risk management consultancy providing 24/7 incident response, digital forensics, threat hunting, virtual security leadership (vCISO) and security compliance support to organisations worldwide. We deliver clear, practical, vendor-independent advice based on real-world threats — no hype, no upsell.

More free security resources are available at www.halkynconsulting.co.uk/security-resources.

Web: www.halkynconsulting.co.uk | **Email:** info@halkynconsulting.co.uk | **Tel:** +44 1244 940858