

Linux Rekoobe Operating with New, Undetected Malware Samples

 intezer.com/blog-linux-rekoobe-operating-with-new-undetected-malware-samples/

January 20, 2020



Written by Ignacio Sanmillan - 20 January 2020



Get Free Account

[Join Now](#)

Introduction

Our research team has identified new versions of an old Linux malware known as Rekoobe, a minimalistic trojan with a complex CNC authentication protocol originally targeting SPARC and Intel x86, x86-64 systems back in 2015.

The new malware samples have lower detection rates than their predecessors. We believe this malware ceased its operation in 2016 after it was reported, however, based on our findings we can estimate the operators behind Rekoobe have resumed their operations utilizing a newer version of the malware.

Based on our research, we can estimate the new samples have been in the wild since 2018.

The screenshot shows the VirusTotal interface for a file. On the left, a circular progress indicator shows '1 / 58' detections. A red warning icon and text state 'One engine detected this file'. The file's SHA-256 hash is '2e2dc0328f6c19b033bb19c24e59e354e519606958afb93fd8049d162a8e3edd'. The file size is '756.38 KB' and it was detected on '2019-07-23 03:43:21 UTC' (5 months ago). The file type is identified as 'elf'. A 'Community Score' section is partially visible with a question mark icon.

We will present a brief technical analysis of the new Rekoobe samples and explain why we believe the new samples have lower detection rates despite previous versions of Rekoobe being well-detected by different security vendors.

Technical Analysis

Linux Rekoobe was first reported by DrWeb in 2015. We believe this malware resumed its operation some time between 2018 and 2019, based on CNC reverse DNS intelligence from RiskIQ, along with sample information from VirusTotal concerning the new variants.

RESOLUTIONS ⓘ

1 - 2 of 2 ▾ Sort : Last Seen Descending ▾ 25 / Page ▾ [Download](#) [Copy](#)

Resolve	First	Last	Source	Tags
<input type="checkbox"/> huawel.site	2019-08-22	2020-01-12	riskiq	
<input type="checkbox"/> 96.45.187.113.16clouds.com	2019-10-03	2019-12-03	riskiq	

1 - 2 of 2 ▾



It's important to mention that the new variants are statically compiled ELF binaries while older variants were dynamically compiled.

This undeniably implies that on a code level these two generations of the same malware are different. Since the compilation flags used by the GCC compiler differ between dynamically and statically linked code, the compiler will generate different code accordingly.

We have also taken into consideration that if indeed these new implants were generated after a three to four year gap, the GCC compiler version would also differ. Consequently, the compiler will generate different code on an assembly level even though additional compilation flags may not have been explicitly used.

It's important to mention that the authors of Linux Rekoobe have removed every attributive string from older variants in their new samples. This explains why string-based signatures have not been able to detect the new version of the malware.

In previous Rekoobe variants, the malware would initially collect some preliminary configuration saved in disk, masquerading this configuration to be a shared object, as previously reported by DrWeb researchers. However, in the new variants, the need for this configuration file has been completely removed, having hard-coded the subject artifacts previously dependent on the configuration file.

Old Rekoobe	New Rekoobe
<pre> mov eax, eax mov dword ptr [esp], offset aSecret ; "SECRET" call read_config_var mov ds:q_Secret, eax mov dword ptr [esp], offset aMagic ; "MAGIC" call read_config_var mov ds:dword_8053A44, eax mov dword ptr [esp], offset aProxyhost ; "PROXYHOST" call read_config_var mov ds:dword_8053A40, eax mov dword ptr [esp], offset aProxyport ; "PROXYPORT" call read_config_var mov dword ptr [esp+8], 0Ah ; base mov dword ptr [esp+4], 0 ; endptr mov [esp], eax ; nptr call _strtol mov ds:dword_8053A48, eax mov dword ptr [esp], offset aUsername ; "USERNAME" call read_config_var mov ds:dword_8053A44, eax mov dword ptr [esp], offset aPassword ; "PASSWORD" call read_config_var mov ds:dword_8054AD0, eax mov dword ptr [esp], offset aEndpoint ; "ENDPOINT" call read_config_var mov ds:dword_8054ADC, eax mov dword ptr [esp], offset aServerPort ; "SERVER_PORT" call read_config_var mov dword ptr [esp+8], 0Ah ; base mov dword ptr [esp+4], 0 ; endptr mov [esp], eax ; nptr call _strtol mov ds:dword_8054AD8, eax mov dword ptr [esp], offset aConnectBackDel ; "CONNECT_BACK_DELAY" call read_config_var mov dword ptr [esp+8], 0Ah ; base mov dword ptr [esp+4], 0 ; endptr mov [esp], eax ; nptr call _strtol mov ds:seconds, eax call _fork </pre>	<pre> envp= dword ptr 10h ; __unwind { lea ecx, [esp+4] and esp, 0FFFFFFF0h push dword ptr [ecx-4] push ebp mov ebp, esp push edi push esi push ebx push ecx sub esp, 2Ch mov ebx, [ecx+4] mov eax, large gs:14h mov [ebp+var_1C], eax xor eax, eax mov edx, [ebx] mov ecx, 0FFFFFFFh mov edi, edx repne scasb mov eax, ecx not eax sub eax, 1 push eax ; int push 0 ; int push edx ; int call memset mov eax, [ebx] mov dword ptr [eax], 'bil/' mov dword ptr [eax+4], 'sys/' mov dword ptr [eax+8], 'dmet' mov dword ptr [eax+0Ch], 'sys/' mov dword ptr [eax+10h], 'dmet' mov dword ptr [eax+14h], 'edu-' mov word ptr [eax+18h], 'dv' mov byte ptr [eax+1Ah], 0 call fork add esp, 10h test eax, eax jns loc_804D3FA </pre>

Rewriting argv[0]

Main function comparison between Rekoobe variants

Among the hardcoded artifacts we can find the CNC IP and port, in addition to the shared 'secret' used to authenticate the CNC and client.

A new feature was added to the new Rekoobe variants to rewrite argv[0], as an attempt to rename the process name (as shown in the picture above) since some forensics tools do retrieve the process name from this location.

The names to rewrite argv[0] chosen by the authors were:

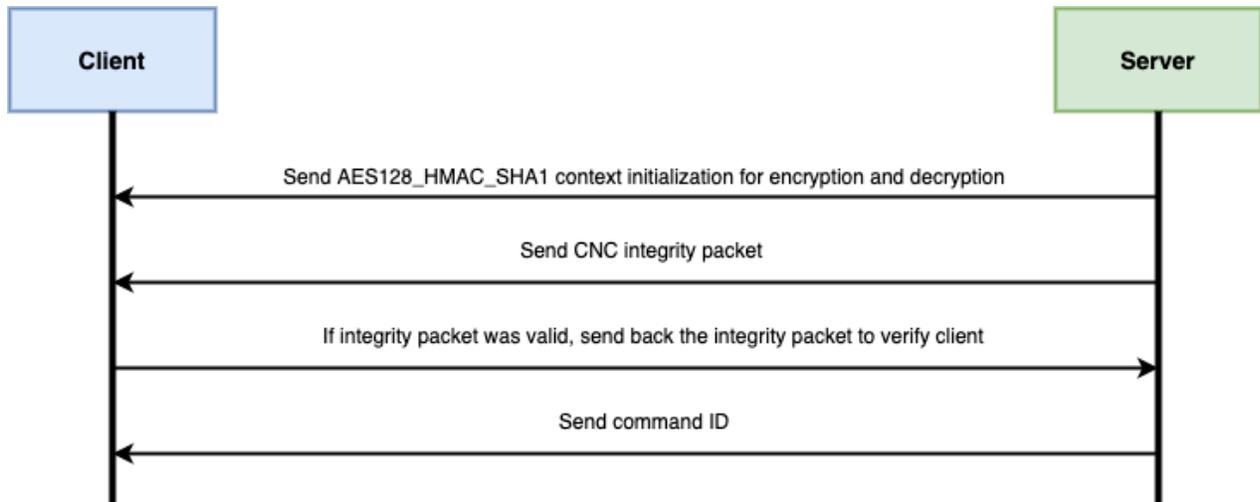
- /lib/sys/temd/sys/temd-udev
- /bin/bash

Old Rekoobe	New Rekoobe
<pre> mov eax, large var_14 mov [ebp+var_14], eax mov eax, eax mov dword ptr [esp], offset aSecret ; "SECRET" call read_config_var mov ds:q_Secret, eax loc_804E3F5: mov eax, ds:dword_8054AD4 mov [esp+4], eax ; buf mov eax, ds:fd mov [esp], eax ; fd call write mov dword ptr [esp], 3 ; seconds call alarm mov eax, ds:q_Secret mov [esp+4], eax ; int mov eax, ds:fd mov [esp], eax ; fd call cnc_handshake cmp eax, 1 jns short loc_804E459 </pre>	<pre> test edx, edx jns loc_804D3FF sub esp, 0Ch push 3 call alarm add esp, 8 push a_idontknow ; "idontknow" push esi ; fd call cnc_handshake add esp, 10h cmp eax, 1 jns short loc_804D361 </pre>

Shared secret retrieval comparison between different Rekoobe versions

Another noticeable difference is that in previous variants a magic value was retrieved from the configuration file and sent to the CNC prior to the handshake mechanism. It seems that the authors removed this preliminary packet on the new Rekoobe variants.

The remaining code resembles previous variants of Rekoobe. The following is a simplified overview of the network protocol used by the new Rekoobe samples, mostly shared with older Rekoobe variants:



The following is a brief description of the most relevant steps in the authentication mechanism:

1. The client will read a stream of 40-bytes from the server. This packet will be divided into two blocks of 20-bytes that will be utilized to initialize two AES128_HMAC_SHA1 contexts. The HMAC SHA1 pair will be generated against each of the 20-byte streams using a given shared secret ("idontknow" hardcoded string in newer variants) and they will be used as AES128 keys for encryption and decryption of future packets.

These computed keys will be used throughout the communication process to decrypt the received packets and to encrypt the packets to be sent by the client with different keys. This implies that if this preliminary packet is not retrieved and the shared secret is not known, then the remaining traffic will unlikely be decryptable.

```

mov     dword ptr [esp+0h+salt1], eax
sub     esp, 4
lea     eax, [esp+40h+salt1]
push   eax                ; salt
push   edi                ; secret
push   offset g_Hmac_1 ; buffer
call   aes128_hmac_sha1
add     esp, 0Ch
lea     eax, [esp+40h+salt2]
push   eax                ; salt
push   edi                ; secret
push   offset g_Hmac_2 ; buffer
call   aes128_hmac_sha1
add     esp, 0Ch
lea     eax, [esp+40h+recv_size]
push   eax                ; arg_8
push   offset g_Rcv_buffer ; recv_buff
push   esi                ; fd
call   cnc_recv_data
add     esp, 10h

```

```

lea     eax, [ebx+214h]
mov     dword ptr [ebx+214h], 36363636h
mov     dword ptr [eax+4], 36363636h
mov     dword ptr [eax+8], 36363636h
mov     dword ptr [eax+0Ch], 36363636h
mov     dword ptr [eax+10h], 36363636h
mov     dword ptr [eax+14h], 36363636h
mov     dword ptr [eax+18h], 36363636h
mov     dword ptr [eax+1Ch], 36363636h
mov     dword ptr [eax+20h], 36363636h
mov     dword ptr [eax+24h], 36363636h
mov     dword ptr [eax+28h], 36363636h
mov     dword ptr [eax+2Ch], 36363636h
mov     dword ptr [eax+30h], 36363636h
mov     dword ptr [eax+34h], 36363636h
mov     dword ptr [eax+38h], 36363636h
mov     dword ptr [eax+3Ch], 36363636h
mov     dword ptr [ebx+254h], 5C5C5C5Ch
mov     dword ptr [ebx+25Ch], 5C5C5C5Ch
mov     dword ptr [ebx+260h], 5C5C5C5Ch
mov     dword ptr [ebx+264h], 5C5C5C5Ch
mov     dword ptr [ebx+268h], 5C5C5C5Ch
mov     dword ptr [ebx+26Ch], 5C5C5C5Ch
mov     dword ptr [ebx+270h], 5C5C5C5Ch
mov     dword ptr [ebx+274h], 5C5C5C5Ch
mov     dword ptr [ebx+278h], 5C5C5C5Ch
mov     dword ptr [ebx+27Ch], 5C5C5C5Ch
mov     dword ptr [ebx+280h], 5C5C5C5Ch
mov     dword ptr [ebx+284h], 5C5C5C5Ch
mov     dword ptr [ebx+288h], 5C5C5C5Ch
mov     dword ptr [ebx+28Ch], 5C5C5C5Ch
mov     dword ptr [ebx+290h], 5C5C5C5Ch

```

opad ipad hmac padding

AES128_HMAC_SHA1 context generation

2. After AES128_HMAC_SHA1 context initialization is achieved there will be a CNC authentication procedure similar to a Challenge-Handshake Authentication Protocol (CHAP). The client will read an additional stream of 16-bytes that will then decrypt using AES128 with the corresponding and previously generated SHA1 hash as key. It will xor decode the stream to retrieve information of the next packet to be read. This process will be repeated every time the client and server sends an additional packet.

3. Every other packet will be subject to HMAC integrity checks to verify the integrity of the packets. Two layers of SHA1 will be computed against the subject packet with two additional salts. This computed SHA1 hash will be then compared with the first 20-bytes of the received packet which contains the pre-computed SHA1 of the AES encrypted packet's payload.

If the computed SHA1 hash does not match with the hardcoded hash delivered in the packet, the client will cease execution since it would imply that the integrity of that packet has been compromised; otherwise the client will proceed to AES128 decrypt the packet's payload and then apply an xor layout.

```

call    shal_init
add     esp, 0Ch
push    40h
push    offset g_Shal_val1
push    esi
call    shal_update
add     esp, 0Ch
push    ebx                ; int
push    offset g_Rcv_buffer
push    esi
call    shal_update
add     esp, 8
lea     ebx, [esp+0D4h+computed_shal_hash]
push    ebx                ; a2
push    esi                ; a1
call    shal_final
mov     [esp+0DC+var_DC], esi
call    shal_init
add     esp, 0Ch
push    40h
push    offset g_Shal_val2
push    esi
call    shal_update
add     esp, 0Ch
push    14h                ; int
push    ebx
push    esi
call    shal_update
add     esp, 8
push    ebx                ; a2
push    esi                ; a1
call    shal_final
add     esp, 0Ch
push    14h                ; a3
push    ebx                ; a2
lea     eax, [esp+0D8h+packet_shal]
push    eax                ; a1
call    _memcmp
add     esp, 10h
test    eax, eax
jz      short loc_8048FAF

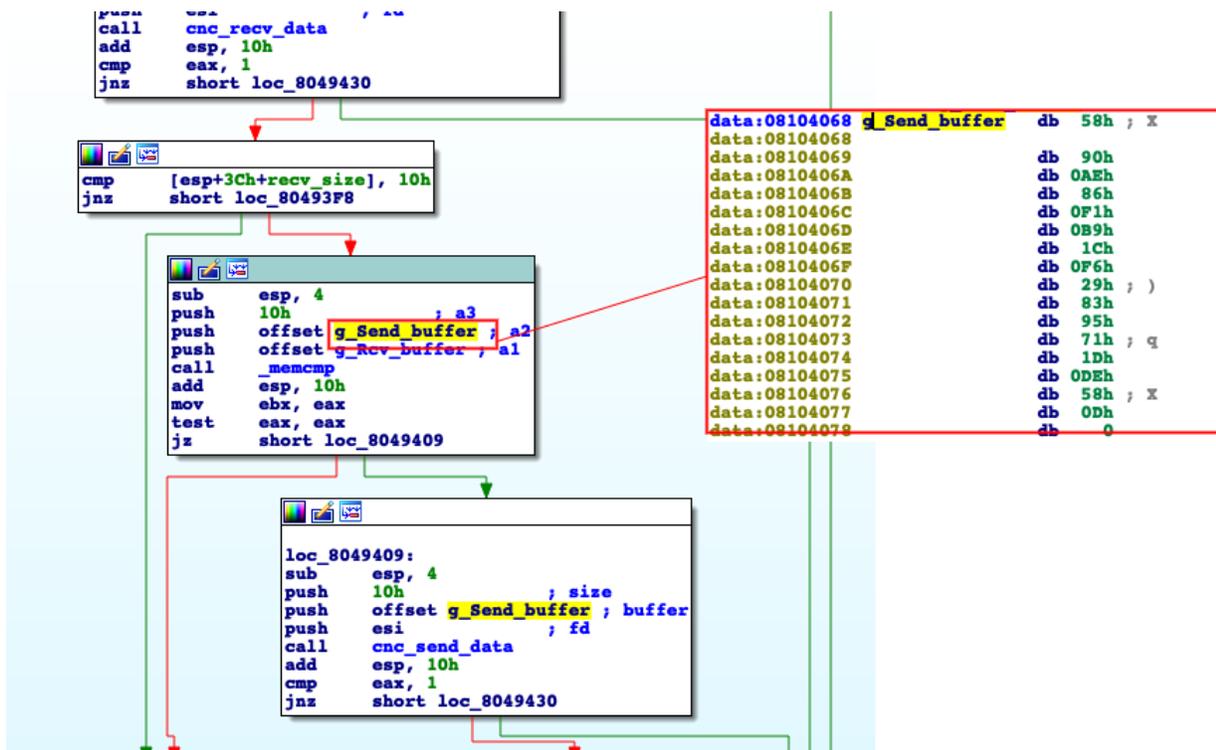
```

SHA1 scheme applied to encrypted packet and comparison to hardcoded SHA1

4. After this packet has been decrypted, the contents will be compared against a hardcoded byte sequence of 16-bytes in order to verify the integrity of the server.

Once again, if this sequence does not match, the handshake will fail.

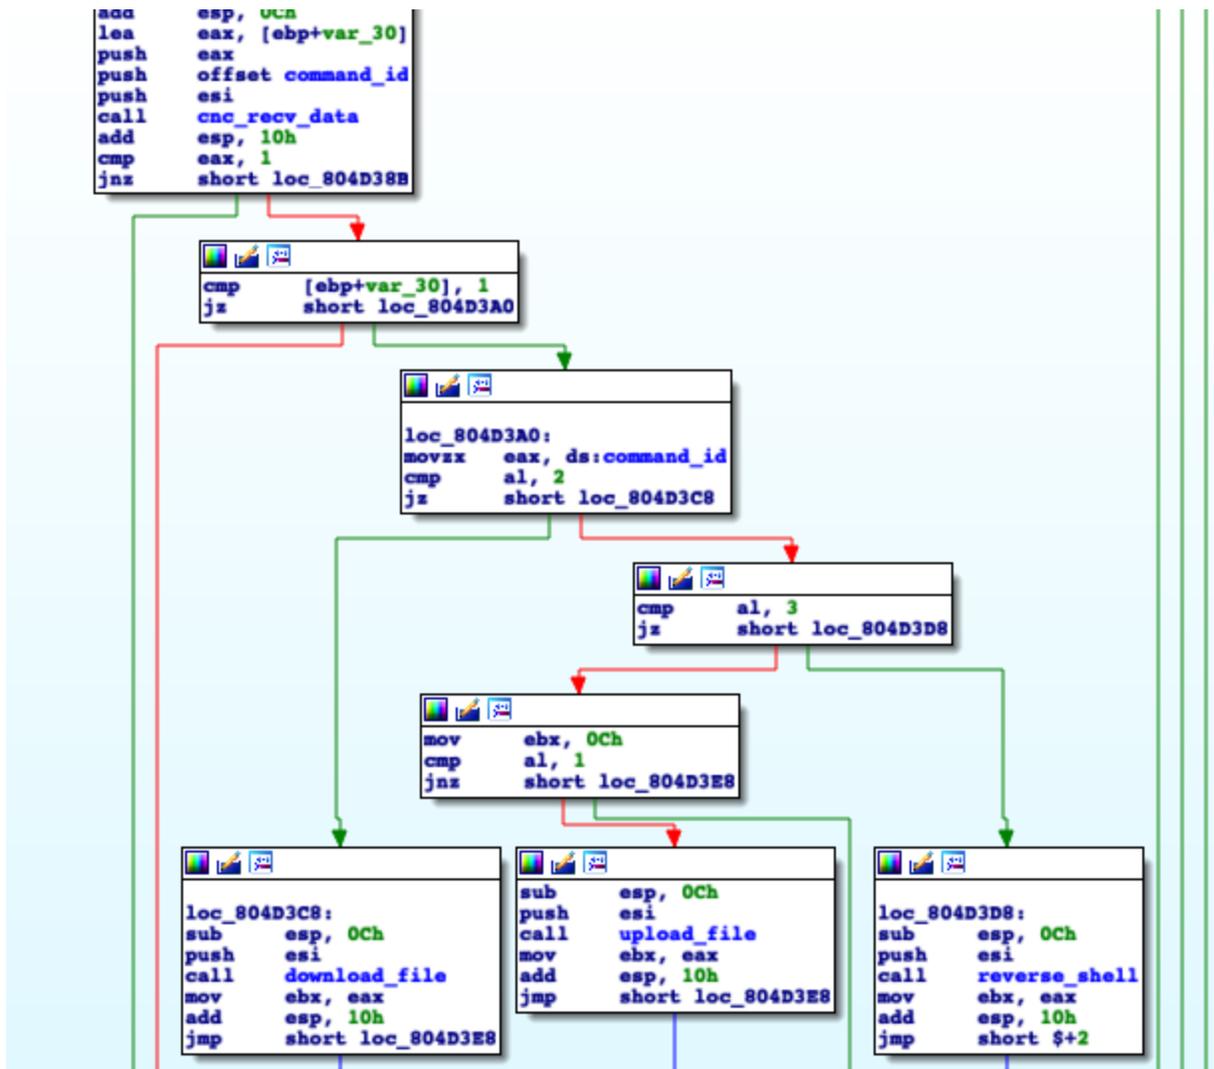
5. If the sequence is correct, then this byte sequence will be AES128 encrypted with the corresponding HMAC SHA1 as key, two layers of SHA1 will be computed with its correspondent salts against the encrypted payload, and this SHA1 hash will be written on the first 20-bytes of the packet itself to then be sent back to the CNC as the last verification step, in this case to verify the integrity of the client.



Hardcoded stream of 16-bytes the CNC and client authentication is based on

We assume this field would be an ideal choice for also identifying the campaign ID by the operators, although it seems the value of this field was shared with previous variants.

6. If the handshake was successful, the client will proceed to listen for a command from the CNC. The malware supports the same three different commands as it did in previous variants, those being file upload, file download, and a reverse shell. The following picture shows the command management implemented by the client after command packet has been successfully decrypted:



Conclusion

We have provided a brief technical analysis of the new Linux Rekoobe samples, highlighting some of the differences between these variants and previous samples. We have also provided an overview of the network protocol.

We have provided several reasons for why the malware has gone undetected, even though the code base of these new variants doesn't appear to have been heavily modified on a source code level from the original Rekoobe samples.

We do not believe this malware has been consistently operational since late 2015. In contrast, we believe the malware may have operated intermittently over small periods of time, since the newly discovered samples appear to be created in recent years, and there appears to be a gap in 2017 where additional Rekoobe samples have not been found.

We have indexed the code from the new Linux Rekoobe variants in our Genetic [Malware Analysis](#) platform [Intezer Analyze](#), and we have published a new [YARA](#) rule in order to help the community to detect this threat.

We expect Linux threats to pose a significant challenge to enterprise cloud security in the near future. We have just released our new cloud security product, Intezer Protect, which is based on our Genetic Malware Analysis technology and provides native cloud protection. For more information, visit <https://www.intezer.com//intezer-protect/>.

IOCs

80e5fec19843c32c6c3fc38aabdeb428c339b0dfce28023529144405b9c72b33
C9eb46d00e11acb354b518f725412b88c69cc511ec8d5bd3cb03c1740f8a2936
2e2dc0328f6c19b033bb19c24e59e354e519606958afb93fd8049d162a8e3edd
E63c2e35a41c51e33b246f5b60c5d1b8da0d8c50bf7ec592383b61818217e8d7
7148ae1ab45e17889915100fdc203fe7941d8e9b946d44a3989ab8baeb6066e1
1d0591049a65db6508a9517f72954541ef6b5a7fe9153c5edcb1bac1b70b991c
4B45E601D480124C38BE06A706F7D8F4
F34119A442651945D5EFB33DB8901D9B
7xin.bitscan[.]win
huawei[.]site
96.45.187[.]113
119.3.22[.]174



Ignacio Sanmillan

Nacho is a security researcher specializing in reverse engineering and malware analysis. Nacho plays a key role in Intezer's malware hunting and investigation operations, analyzing and documenting new undetected threats. Some of his latest research involves detecting new Linux malware and finding links between different threat actors. Nacho is an adept ELF researcher, having written numerous papers and conducting projects implementing state-of-the-art obfuscation and anti-analysis techniques in the ELF file format.