Binary Reverse Engineering And Analysis

Course 4: Dynamic Analysis

Caragea Radu

March 9, 2021

Recap

- Last time: dissecting executables
- Today we study "moving targets"
  - From executable to process
  - Tracing unknown binaries
  - Modifying control flow

But... why?

- Can drastically reduce static analysis time
- Can uncover subtle vulnerabilities inside/outside the code
- Can uncover vulnerabilities unrelated to the actual code (!)

Example 1: side-channels

University Science

assisted by Lev Pachmanov and numerous others

Summary

Many computers emit a high-pitched noise during operation, due to vibration in some of their electronic components. These acoustic emanations are more than a nuisance: they can convey information about the software running on the computer and, in particular, leak sensitive information about security-related computations. In a preliminary presentation, we have shown that different RSA keys induce different sound patterns.

Here, we describe a new *acoustic cryptanalysis* key extraction attack, applicable to GnuPG's current implementation of RSA. The attack can extract full 4096-bit RSA decryption keys from laptop computers (of various models), within an hour, using the sound generated by the computer during the decryption of some chosen ciphertexts. We experimentally demonstrate that such attacks can be carried out, using either a plain mobile phone placed next to the computer, or a more sensitive microphone placed 4 meters away.

More info: `https://m.tau.ac.il/~tromer/acoustic`

Example 2: disappearing security measures (before)

https://godbolt.org/z/QMZxYe

```c
#include <stdio.h>
#include <string.h>

void secure_use_password(char* buf){
    printf("TODO, use password %s\n",buf );
}

void secure_get_user_password(void)
{
    char pwd[64];
    fgets(pwd, 64, stdin);
    secure_use_password(pwd);

    //wipe the password from the memory
    memset(pwd, 0, sizeof(pwd));
}
```

```
.LC0:
        .string "TODO, use password %s\n"
secure_use_password(char*):
        push    rbp
        mov     rbp, rsp
        sub     rsp, 16
        mov     QWORD PTR [rbp-8], rdi
        mov     rax, QWORD PTR [rbp-8]
        mov     rsi, rax
        mov     edi, OFFSET FLAT:.LC0
        mov     eax, 0
        call    printf
        nop
        leave
        ret
secure_get_user_password():
        push    rbp
        mov     rbp, rsp
        sub     rsp, 64
        mov     rdx, QWORD PTR stdin[rip]
        lea     rax, [rbp-64]
        mov     esi, 64
        mov     rdi, rax
        call    fgets
        lea     rax, [rbp-64]
        mov     rdi, rax
        call    secure_use_password(char*)
        lea     rax, [rbp-64]
        mov     edx, 64
        mov     esi, 0
        mov     rdi, rax
        call    memset
        nop
```

Example 2: disappearing security measures (after)

https://godbolt.org/z/3EyZXQ

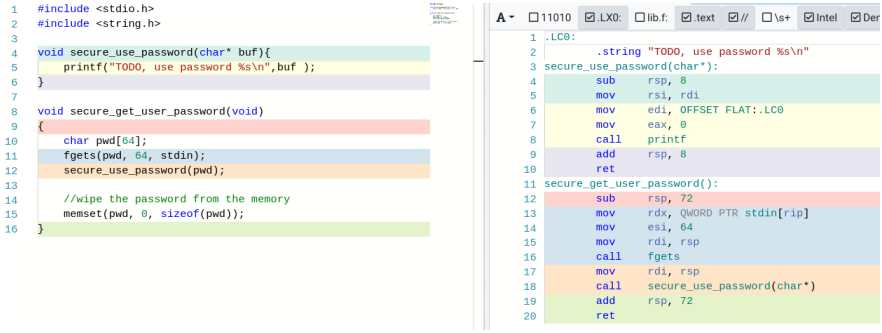

```c
#include <stdio.h>
#include <string.h>

void secure_use_password(char* buf){
    printf("TODO, use password %s\n",buf );
}

void secure_get_user_password(void)
{
    char pwd[64];
    fgets(pwd, 64, stdin);
    secure_use_password(pwd);

    //wipe the password from the memory
    memset(pwd, 0, sizeof(pwd));
}
```

```
.LC0:
        .string "TODO, use password %s\n"
secure_use_password(char*):
        sub     rsp, 8
        mov     rsi, rdi
        mov     edi, OFFSET FLAT:.LC0
        mov     eax, 0
        call    printf
        add     rsp, 8
        ret
secure_get_user_password():
        sub     rsp, 72
        mov     rdx, QWORD PTR stdin[rip]
        mov     esi, 64
        mov     rdi, rsp
        call    fgets
        mov     rdi, rsp
        call    secure_use_password(char*)
        add     rsp, 72
        ret
```

Executables

- Start as files on the filesystem
- As seen last time, executables carry loading information
- But what happens when we `run` the executable?

OS Kernel

- Provides a separate address space from other processes
- Provides randomization where compatible (TBD)
- Provides expandable stack space, heap space
- Passes control to a suitable loader (interpreter)

Loaders

- Parse the file structure
- Copy segment contents into memory
- Expand sparse segments
- Set adequate permissions to each segment
- Do the same for any linked libraries needed
- Pass control to the address specified in the header

Linux Address Space Layout (1/2)

- Static Executable

```
Temporary breakpoint 1, 0x0000000000401c3a in main ()
gdb-peda$ vmmap
Start              End                Perm      Name
0x00400000         0x00401000         r--p      /ctf/unibuc/curs/curs_04/demo_01_linux_memory/hello_static
0x00401000         0x00495000         r-xp      /ctf/unibuc/curs/curs_04/demo_01_linux_memory/hello_static
0x00495000         0x004ba000         r--p      /ctf/unibuc/curs/curs_04/demo_01_linux_memory/hello_static
0x004bb000         0x004c1000         rw-p      /ctf/unibuc/curs/curs_04/demo_01_linux_memory/hello_static
0x004c1000         0x004e5000         rw-p      [heap]
0x00007ffff7ffa000 0x00007ffff7ffd000 r--p      [vvar]
0x00007ffff7ffd000 0x00007ffff7fff000 r-xp      [vdso]
0x00007fffffde000 0x00007ffffffff000 rw-p      [stack]
gdb-peda$
```

Linux Address Space Layout (2/2)

- Dynamic Executable

```
Temporary breakpoint 1, 0x00000000004011e2 in main ()
gdb-peda$ vmmap
Start              End                Perm      Name
0x00400000         0x00401000         r--p      /ctf/unibuc/curs/curs_04/demo_01_linux_memory/hello_dynamic
0x00401000         0x00402000         r-xp      /ctf/unibuc/curs/curs_04/demo_01_linux_memory/hello_dynamic
0x00402000         0x00403000         r--p      /ctf/unibuc/curs/curs_04/demo_01_linux_memory/hello_dynamic
0x00403000         0x00404000         r--p      /ctf/unibuc/curs/curs_04/demo_01_linux_memory/hello_dynamic
0x00404000         0x00405000         rw-p      /ctf/unibuc/curs/curs_04/demo_01_linux_memory/hello_dynamic
0x00007ffff7dc6000 0x00007ffff7de8000 r--p      /lib/x86_64-linux-gnu/libc-2.28.so
0x00007ffff7de8000 0x00007ffff7f30000 r-xp      /lib/x86_64-linux-gnu/libc-2.28.so
0x00007ffff7f30000 0x00007ffff7f7c000 r--p      /lib/x86_64-linux-gnu/libc-2.28.so
0x00007ffff7f7c000 0x00007ffff7f7d000 ---p      /lib/x86_64-linux-gnu/libc-2.28.so
0x00007ffff7f7d000 0x00007ffff7f81000 r--p      /lib/x86_64-linux-gnu/libc-2.28.so
0x00007ffff7f81000 0x00007ffff7f83000 rw-p      /lib/x86_64-linux-gnu/libc-2.28.so
0x00007ffff7f83000 0x00007ffff7f87000 rw-p      mapped
0x00007ffff7f87000 0x00007ffff7f89000 rw-p      mapped
0x00007ffff7fd0000 0x00007ffff7fd3000 r--p      [vvar]
0x00007ffff7fd3000 0x00007ffff7fd5000 r-xp      [vdso]
0x00007ffff7fd5000 0x00007ffff7fd6000 r--p      /lib/x86_64-linux-gnu/ld-2.28.so
0x00007ffff7fd6000 0x00007ffff7ff4000 r-xp      /lib/x86_64-linux-gnu/ld-2.28.so
0x00007ffff7ff4000 0x00007ffff7ffc000 r--p      /lib/x86_64-linux-gnu/ld-2.28.so
0x00007ffff7ffc000 0x00007ffff7ffd000 r--p      /lib/x86_64-linux-gnu/ld-2.28.so
0x00007ffff7ffd000 0x00007ffff7ffe000 rw-p      /lib/x86_64-linux-gnu/ld-2.28.so
0x00007ffff7ffe000 0x00007ffff7fff000 rw-p      mapped
0x00007ffffffde000 0x00007ffffffff000 rw-p      [stack]
gdb-peda$
```

Windows Address Space Layout

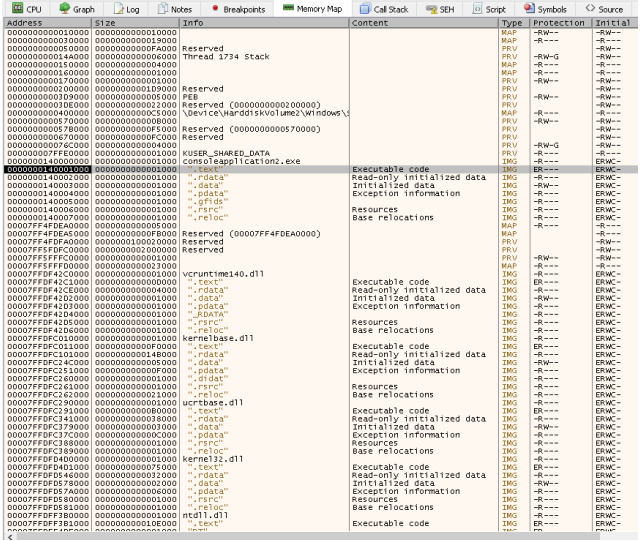

CPU | Graph | Log | Notes | Breakpoints | Memory Map | Call Stack | SEH | Script | Symbols | Source

Address	Size	Info	Content	Type	Protection	Initial
0000000000010000	0000000000010000			MAP	-RW--	-RW--
0000000000030000	0000000000019000			MAP	-R---	-R---
0000000000050000	0000000000FA000	Reserved		PRV		-RW--
000000000014A000	0000000000006000	Thread 1734 Stack		PRV	-RW-G	-RW--
0000000000150000	0000000000004000			MAP	-R---	-R---
0000000000160000	0000000000001000			MAP	-R---	-R---
0000000000170000	0000000000001000			PRV	-RW--	-RW--
0000000000200000	000000001D9000	Reserved		PRV		-RW--
000000001D9000	0000000000005000	PEB		PRV	-RW--	-RW--
000000003DE000	0000000000022000	Reserved (0000000000200000)		PRV		-RW--
0000000000400000	000000000C5000	\Device\HarddiskVolume2\Windows\S		MAP	-R---	-R---
0000000000570000	0000000000B000			PRV	-RW--	-RW--
000000000057B000	0000000000F5000	Reserved (0000000000570000)		PRV		-RW--
0000000000670000	000000000FC000	Reserved		PRV		-RW--
000000000076C000	0000000000004000			PRV	-RW-G	-RW--
00000007FFE0000	0000000000001000	KUSER_SHARED_DATA		PRV	-R---	-R---
0000000140000000	0000000000001000	consoleapplication2.exe		IMG	-R---	ERWC-
0000000140001000	0000000000001000	".text"	Executable code	IMG	ER---	ERWC-
0000000140002000	0000000000001000	".rdata"	Read-only initialized data	IMG	-R---	ERWC-
0000000140003000	0000000000001000	".data"	Initialized data	IMG	-RW--	ERWC-
0000000140004000	0000000000001000	".pdata"	Exception information	IMG	-R---	ERWC-
0000000140005000	0000000000001000	".gfids"		IMG	-R---	ERWC-
0000000140006000	0000000000001000	".rsrc"	Resources	IMG	-R---	ERWC-
0000000140007000	0000000000001000	".reloc"	Base relocations	IMG	-R---	ERWC-
00007FF4FDEA0000	0000000000005000			MAP	-R---	-R---
00007FF4FDEA5000	0000000000FB000	Reserved (00007FF4FDEA0000)		MAP		-R---
00007FF4FDFA0000	0000000100020000	Reserved		PRV		-RW--
00007FF5FDFC0000	0000000002000000	Reserved		PRV		-RW--
00007FF5FFFC0000	0000000000001000			PRV	-RW--	-RW--
00007FF5FFFD0000	0000000000023000			MAP	-R---	-R---
00007FFDF42C0000	0000000000001000	vcruntime140.dll		IMG	-R---	ERWC-
00007FFDF42C1000	0000000000000000	".text"	Executable code	IMG	ER---	ERWC-
00007FFDF42CE000	0000000000004000	".rdata"	Read-only initialized data	IMG	-R---	ERWC-
00007FFDF42D2000	0000000000001000	".data"	Initialized data	IMG	-RW--	ERWC-
00007FFDF42D3000	0000000000001000	".pdata"	Exception information	IMG	-R---	ERWC-
00007FFDF42D4000	0000000000001000	"_RDATA"		IMG	-R---	ERWC-
00007FFDF42D5000	0000000000001000	".rsrc"	Resources	IMG	-R---	ERWC-
00007FFDF42D6000	0000000000001000	".reloc"	Base relocations	IMG	-R---	ERWC-
00007FFDFC010000	0000000000001000	kernelbase.dll		IMG	-R---	ERWC-
00007FFDFC011000	000000000F0000	".text"	Executable code	IMG	ER---	ERWC-
00007FFDFC101000	000000000014B000	".rdata"	Read-only initialized data	IMG	-R---	ERWC-
00007FFDFC24C000	0000000000005000	".data"	Initialized data	IMG	-RW--	ERWC-
00007FFDFC251000	0000000000000F000	".pdata"	Exception information	IMG	-R---	ERWC-
00007FFDFC260000	0000000000001000	".didat"		IMG	-R---	ERWC-
00007FFDFC261000	0000000000001000	".rsrc"	Resources	IMG	-R---	ERWC-
00007FFDFC262000	0000000000021000	".reloc"	Base relocations	IMG	-R---	ERWC-
00007FFDFC290000	0000000000001000	ucrtbase.dll		IMG	-R---	ERWC-
00007FFDFC291000	00000000000B0000	".text"	Executable code	IMG	ER---	ERWC-
00007FFDFC341000	0000000000038000	".rdata"	Read-only initialized data	IMG	-R---	ERWC-
00007FFDFC379000	0000000000003000	".data"	Initialized data	IMG	-RW--	ERWC-
00007FFDFC37C000	000000000000C000	".pdata"	Exception information	IMG	-R---	ERWC-
00007FFDFC388000	0000000000001000	".rsrc"	Resources	IMG	-R---	ERWC-
00007FFDFC389000	0000000000001000	".reloc"	Base relocations	IMG	-R---	ERWC-
00007FFDFD400000	0000000000001000	kernel32.dll		IMG	-R---	ERWC-
00007FFDFD401000	0000000000075000	".text"	Executable code	IMG	ER---	ERWC-
00007FFDFD476000	0000000000021000	".rdata"	Read-only initialized data	IMG	-R---	ERWC-
00007FFDFD578000	0000000000002000	".data"	Initialized data	IMG	-RW--	ERWC-
00007FFDFD57A000	0000000000006000	".pdata"	Exception information	IMG	-R---	ERWC-
00007FFDFD580000	0000000000001000	".rsrc"	Resources	IMG	-R---	ERWC-
00007FFDFD581000	0000000000001000	".reloc"	Base relocations	IMG	-R---	ERWC-
00007FFDFF3B0000	0000000000001000	ntdll.dll		IMG	-R---	ERWC-
00007FFDFF3B1000	0000000000010E000	".text"	Executable code	IMG	ER---	ERWC-

How do processes inter-communicate?

How do processes inter-communicate?

- Shared memory
- Message queues
- Pipes
- Sockets
- Synchronization

How do processes inter-communicate?

- Shared memory
- Message queues
- Pipes
- Sockets
- Synchronization
- Direct access (used by debugging processes)

Linux debug methods (ptrace syscall)

- Attach to a process (called tracee)
- Read/write memory from tracee
- Read/write CPU registers from tracee
- Single step (one CPU instruction at a time)
- Start/stop/continue execution
- Handle breakpoints

Linux low-level debugging

- Debuggers mainly use ptrace
- We study GDB plus a plugin (PEDA)

Windows debug methods (separate syscalls)

- Attach to a process (OpenProcess)
- Read/write memory from tracee (ReadProcessMemory/WriteProcessMemory)
- Read/write CPU registers from tracee (GetThreadContext)
- Start/stop/continue execution (DebugBreakProcess)
- Handle breakpoints (WaitForDebugEvent/ContinueDebugEvent)

Windows low-level debugging

- Windbg is the most powerful but hard to learn
- X64dbg is a decent debugger handling 32/64

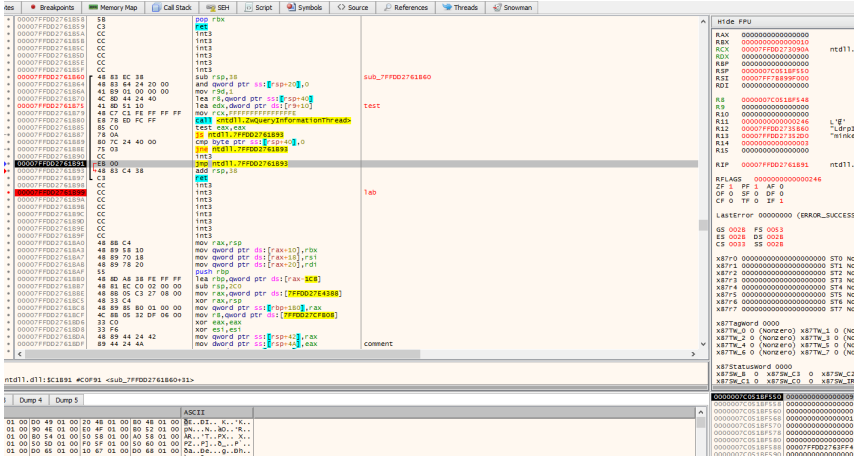

Fundamental tasks in a debugger wrt RE

- Interrupt (break) execution at a certain point in the code
- Inspect/modify virtual memory state/contents
- Inspect/modify CPU registers
- Analyze the call stack

Alternatives

- Processes can also be instrumented
- Intel PIN (Linux/Windows)
- Add extra code in the same address space
- More power, harder to detect, more complexity

Practice

- Any Questions?
- `http://pwnthybytes.ro/unibuc_re/04-lab.html`