CS 213, Fall 2001
Lab Assignment L3: The Buffer Bomb
Assigned: Sep. 27, Due: Thurs., Oct. 4, 11:59PM

Randy Bryant (`Randy.Bryant@cs.cmu.edu`) is the lead person for this assignment.

Introduction

This assignment helps you develop a detailed understanding of the calling stack organization on an IA32 processor. It involves applying a series of *buffer overflow attacks* on an executable file `bufbomb` in the lab directory.

Note: In this lab, you will gain firsthand experience with one of the methods commonly used to exploit security weaknesses in operating systems and network servers. Our purpose is to help you learn about the runtime operation of programs and to understand the nature of this form of security weakness so that you can avoid it when you write system code. We do not condone the use of these or any other form of attack to gain unauthorized access to any system resources. There are criminal statutes governing such activities.

Logistics

You may work in a group of up to two people in solving the problems for this assignment. The only "hand-in" will be an automated logging of your successful attacks. Any clarifications and revisions to the assignment will be posted on the course Web page.

Hand Out Instructions

In the directory

`/afs/cs.cmu.edu/academic/class/15213-f01/L3`

you will see the files for three programs:

MAKECOOKIE**:** Generates a "cookie" based on your team name.

BUFBOMB**:** The code you will attack.

SENDSTRING**:** A utility to help convert between string formats.

All of these programs are compiled to run on the Fish machines.

In the following, we will assume that you have defined the lab directory to be on your execution path. You can do this by executing the following command:

```
unix> setenv PATH /afs/cs.cmu.edu/academic/class/15213-f01/L3:$PATH
```

Team Name and Cookie

You should create a team name for the one or two people in your group of the following form:

- "ID" where ID is your Andrew ID, if you are working alone, or
- "ID_1+ID_2" where ID_1 is the Andrew ID of the first team member and ID_2 is the Andrew ID of the second team member.

You should choose a consistent ordering of the IDs in the second form of team name. Teams "`ac00+bovik`" and "`bovik+ac00`" are considered distinct. **You must follow this scheme for generating your team name. Our grading program will only give credit to those people whose Andrew IDs can be extracted from the team names.**

A *cookie* is a string of eight hexadecimal digits that is (with high probability) unique to your team. You can generate your cookie with the `makecookie` program giving your team name as the argument. For example:

```
unix> makecookie ac00+bovik
0x78327b66
```

In three of your four buffer attacks, your objective will be to make your cookie show up in places where it ordinarily would not.

The BUFBOMB Program

The BUFBOMB program reads a string from standard input with a function `getbuf` having the following C code:

```
1 int getbuf()
2 {
3     char buf[12];
4     Gets(buf);
5     return 1;
6 }
```

The function `Gets` is similar to the standard library function `gets`—it reads a string from standard input (terminated by '`\n`' or end-of-file) and stores it (along with a null terminator) at the specified destination. In this code, the destination is an array `buf` having sufficient space for 12 characters.

Neither `Gets` nor `gets` have any way to determine whether there is enough space at the destination to store the entire string. Instead, they simply copy the entire string, possibly overrunning the bounds of the storage allocated at the destination.

If the string typed by the user to `getbuf` is no more than 11 characters long, it is clear that `getbuf` will return 1, as shown by the following execution example:

```
unix> bufbomb
Type string: howdy doody
Dud: getbuf returned 0x1
```

If we type a longer string, typically an error occurs:

```
unix> bufbomb
Type string: This string is too long
Ouch!: You caused a segmentation fault!
```

As the error message indicates, overrunning the buffer typically causes the program state to be corrupted, leading to a memory access error. Your task is to be more clever with the strings you feed BUFBOMB so that it does more interesting things. These are called *exploit* strings.

BUFBOMB takes several different command line arguments:

`-t` ***TEAM*:** Operate the bomb for the indicated team. You should always provide this argument for several reasons:

- It is required to log your successful attacks.
- BUFBOMB determines the cookie you will be using based on your team name, just as does the program MAKECOOKIE.
- We have built features into BUFBOMB so that some of the key stack addresses you will need to use depend on your team's cookie.

`-h`**:** Print list of possible command line arguments

`-n`**:** Operate in "Nitro" mode, as is used in Level 3 below.

Your exploit strings will typically contain byte values that do not correspond to the ASCII values for printing characters. The program SENDSTRING can help you generate these *raw* strings. It takes as input a *hex-formatted* string. In this format, each byte value is represented by two hex digits. For example, the string "`012345`" could be entered in hex format as "`30 31 32 33 34 35`," since the ASCII code for decimal digit x is `0x3`x. Non-hex digit characters are ignored, including the blanks in the example shown.

If you generate a hex-formatted exploit string in the file `exploit.txt`, you can apply the raw string to BUFBOMB in several different ways:

1. You can set up a series of pipes to pass the string through SENDSTRING.

   ```
   unix> cat exploit.txt | sendstring | bufbomb -t bovik
   ```

2. You can store the raw string in a file and use I/O redirection to supply it to BUFBOMB:

   ```
   unix> sendstring < exploit.txt > exploit-raw.txt
   unix> bufbomb -t bovik < exploit-raw.txt
   ```

   This approach can also be used when running BUFBOMB from within GDB:

   ```
   unix> gdb /afs/cs.cmu.edu/academic/class/15213-f01/L3/bufbomb
   (gdb) run -t bovik < exploit-raw.txt
   ```

One important point: your exploit string must not contain byte value `0x0A` at any intermediate position, since this is the ASCII code for newline ('`\n`'). When `Gets` encounters this byte, it will assume you intended to terminate the string. SENDSTRING will warn you if it encounters this byte value.

When you correctly solve one of the levels, BUFBOMB will automatically send an email notification to our grading server. The server will test your exploit string to make sure it really works, and it will update the lab web page indicating that your team (listed by cookie) has completed this level.

Unlike the bomb lab, there is no penalty for making mistakes in this lab. Feel free to fire away at BUFBOMB with any string you like.

Level 0: Sparkler (15 pts)

The function `getbuf` is called within BUFBOMB by a function `test` having the following C code:

```
1  void test()
2  {
3      int val;
4      val = getbuf();
5      if (val == cookie) {
6          printf("Boom!: getbuf returned 0x%x\n", val);
7          validate(2);
8      }
9      else {
10         printf("Dud: getbuf returned 0x%x\n", val);
11     }
12 }
```

When `getbuf` executes its return statement (line 5 of `getbuf`), the program ordinarily resumes execution within function `test` (at line 5 of this function).

Within the file `bufbomb`, there is a function `fizz` having the following C code:

```
void fizz()
{
    printf("Fizz!: You called fizz()\n");
    validate(0);
    exit(0);
}
```

Your task is to get BUFBOMB to execute the code for `fizz` when `getbuf` executes its return statement, rather than returning to `test`. You can do this by supplying an exploit string that overwrites the stored return pointer in the stack frame for `getbuf` with the address of the first instruction in `fizz`. Note that your exploit string may also corrupt other parts of the stack state, but this will not cause a problem, since `fizz` causes the program to exit directly.

Some Advice:

- All the information you need to devise your exploit string for this level can be determined by examining a diassembled version of BUFBOMB.
- Be careful about byte ordering.
- You might want to use GDB to step the program through the last few instructions of `getbuf` to make sure it is doing the right thing.

Level 1: Firecracker (25 pts)

Within the file `bufbomb` there is also a function `bang` having the following C code:

```
void bang(int val)
{
    if (val == cookie) {
        printf("Bang!: You called bang(0x%x)\n", val);
        validate(1);
    } else
        printf("Misfire: You called bang(0x%x)\n", val);
    exit(0);
}
```

Similar to Level 0, your task is to get BUFBOMB to execute the code for `bang` rather than returning to `test`. In this case, however, you must make it appear to `bang` as if you have passed your cookie as its argument. You can do this by encoding your cookie in the appropropriate place within your exploit string.

Some Advice: Note that the program won't really call `bang`—it will simply execute its code. This has important implications for where on the stack you want to place your cookie.

Level 2: Dynamite (60 pts)

A much more sophisticated form of buffer attack involves supplying a string that encodes actual machine instructions. The exploit string then overwrites the return pointer with the address of these instructions. When the calling function (in this case `getbuf`) executes its `ret` instruction, the program will start executing the instructions on the stack rather than returning. With this form of attack, you can get the program to do almost anything. The code you place on the stack is called the *exploit* code. This style of attack is tricky, though, since you must: 1) get machine code onto the stack, 2) set the return pointer to the start of this code, and 3) avoid corrupting other parts of the stack state.

Your job for this level is to supply an exploit string that will cause `getbuf` to return your cookie, rather than the value 1. You can see in the code for `test` that this will cause the program to go "`Boom!`." Your exploit code should set your cookie as the return value, restore any corrupted state, push the correct return location on the stack, and execute a `ret` instruction to really return to `test`.

Some Advice:

- In order to overwrite the return pointer, you must also overwrite the saved value of `%ebp`. However, it is important that this value is correctly restored before you return to `test`. You can do this by either 1) making sure that your exploit string contains the correct value of the saved `%ebp` in the correct position, so that it never gets corrupted, or 2) restore the correct value as part of your exploit code.
- You can use GDB to get the information you need to construct your exploit string. Set a breakpoint within `getbuf` and run to this breakpoint. Determine parameters such as the saved return address and the saved value of `%ebp`.
- Determining the byte encoding of instruction sequences by hand is tedious and prone to errors. You can let tools do all of the work by writing an assembly code file containing the instructions and data you want to put on the stack. Assemble this file with GCC and disassemble it with OBJDUMP. You should be able to get the exact byte sequence that you will type at the prompt. (A brief example of how to do this is included at the end of this writeup.)
- Keep in mind that your exploit string depends on your machine, your compiler, and even your team's cookie. Do all of your work on a Fish machine, and make sure you include the proper team name on the command line to BUFBOMB.

Once you complete this level, pause to reflect on what you have accomplished. You caused a program to execute machine code of your own design. You have done so in a sufficiently stealthy way that the program did not realize that anything was amiss.

Level 3: Nitroglycerin (10 pts)

If you have completed the first three levels, you have earned 100 points. You have mastered the principles of the runtime stack operation, and you have gained firsthand experience with buffer overflow attacks. We consider this a satisfactory mastery of the material. You are welcome to stop right now.

The next level is for those who want to push themselves beyond our baseline expectations for the course, and who want to face a challenge in designing buffer overflow attacks that arises in real life. This part of the assignment only counts 10 points, even though it requires a fair amount of work to do, so don't do it just for the points.

From one run to another, especially by different users, the exact stack positions used by a given procedure will vary. One reason for this variation is that the values of all environment variables are placed near the base of the stack when a program starts executing. Environment variables are stored as strings, requiring different amounts of storage depending on their values. Thus, the stack space allocated for a given user depends on the settings of his or her environment variables. Stack positions also differ when running a program under GDB, since GDB uses stack space for some of its own state.

In the code that calls `getbuf`, we have incorporated features that stabilize the stack, so that the position of `getbuf`'s stack frame will be consistent between runs. This made it possible for you to write an exploit string knowing the exact starting address of `buf` and the exact saved value of `%ebp`. If you tried to use such an exploit on a normal program, you would find that it works some times, but it causes segmentation faults at other times. Hence the name "dynamite"—an explosive developed by Alfred Nobel that contains stabilizing elements to make it less prone to unexpected explosions.

For this level, we have gone the opposite direction, making the stack positions even less stable than they normally are. Hence the name "nitroglycerin"—an explosive that is notoriously unstable.

When you run BUFBOMB with the command line flag "`-n`," it will run in "Nitro" mode. Rather than calling the function `getbuf`, the program calls a slightly different function `getbufn`:

```
int getbufn()
{
    char buf[512];
    Gets(buf);
    return 1;
}
```

This function is similar to `getbuf`, except that it has a buffer of 512 characters. You will need this additional space to create a reliable exploit. The code that calls `getbufn` first allocates a random amount of storage on the stack (using library function `alloca`) that ranges between 0 and 127 bytes. Thus, if you were to sample the value of `%ebp` during two successive executions of `getbufn`, you would find they differ by as much as ± 127.

In addition, when run in Nitro mode, BUFBOMB requires you to supply your string 5 times, and it will execute `getbufn` 5 times, each with a different stack offset. Your exploit string must make it return your cookie each of these times.

Some Advice:

- You can use the program SENDSTRING to send multiple copies of your exploit string. If you have a single copy in the file `exploit.txt`, then you can use the following command:

  ```
  unix> cat exploit.txt | sendstring -n 5 | bufbomb -n -t bovik
  ```

You must use the same string for all 5 executions of `getbufn`. Otherwise it will fail the testing code used by our grading server.

- The trick is to make use of the `nop` instruction. It is encoded with a single byte (code `0x90`). You can place a long sequence of these at the beginning of your exploit code so that your code will work correctly if the initial jump lands anywhere within the sequence.
- You will need to restore the saved value of `%ebp` in a way that is insensitive to variations in stack positions.

Logistical Notes

Hand in occurs automatically whenever you correctly solve a level. The program sends email to our grading server containing your team name (be sure to set the "`-t`" command line flag properly), and your exploit string to the grading server. You will be informed of this by BUFBOMB. Upon receiving the email, the server will validate your string and update the lab web page. You should check this page a few minutes after your submission to make sure your string has been validated. [If you really solved the level, your string *should* be valid.]

Note that each level is graded individually. You do not need to do them in the specified order, but you will get credit only for the levels for which the server receives a valid message.

Have fun!

Generating Byte Codes

Using GCC as an assembler and OBJDUMP as a disassembler makes it convenient to generate the byte codes for instruction sequences. For example, suppose we write a file `example.s` containing the following assembly code:

```
# Example of hand-generated assembly code
        pushl $0x89ABCDEF       # Push value onto stack
        addl $17,%eax           # Add 17 to %eax
        .align 4                # Following will be aligned on multiple of 4
        .long   0xFEDCBA98      # A 4-byte constant
        .long   0x00000000      # Padding
```

The code can contain a mixture of instructions and data. Anything to the right of a '`#`' character is a comment. We have added an extra word of all 0s to work around a shortcoming in OBJDUMP to be described shortly.

We can now assemble and disassemble this file:

```
unix> gcc -c example.s
unix> objdump -d example.o > example.d
```

The generated file `example.d` contains the following lines

```
0:   68 ef cd ab 89          push   $0x89abcdef
5:   83 c0 11                add    $0x11,%eax
8:   98                      cwtl                    Objdump tries to interpret
9:   ba dc fe 00 00          mov    $0xfedc,%edx     these as instructions
```

Each line shows a single instruction. The number on the left indicates the starting address (starting with 0), while the hex digits after the '`:`' character indicate the byte codes for the instruction. Thus, we can see that the instruction `pushl $0x89ABCDEF` has hex-formatted byte code `68 ef cd ab 89`.

Starting at address 8, the disassembler gets confused. It tries to interpret the bytes in the file `example.o` as instructions, but these bytes actually correspond to data. Note, however, that if we read off the 4 bytes starting at address 8 we get: `98 ba dc fe`. This is a byte-reversed version of the data word `0xFEDCBA98`. This byte reversal represents the proper way to supply the bytes as a string, since a little endian machine lists the least significant byte first. Note also that it only generated two of the four bytes at the end with value `00`. Had we not added this padding, OBJDUMP gets even more confused and does not emit all of the bytes we want.

Finally, we can read off the byte sequence for our code (omitting the final 0's) as:

```
68 ef cd ab 89 83 c0 11 98 ba dc fe
```