

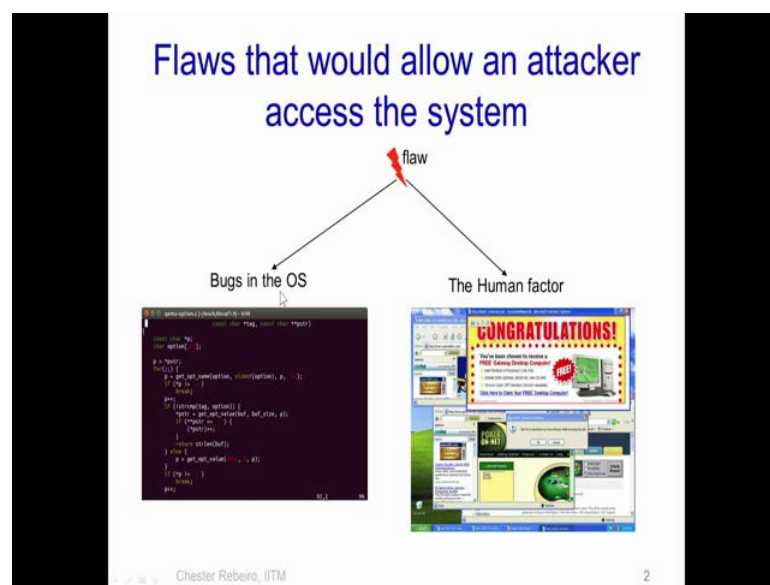
Introduction to Operating Systems
Prof. Chester Rebeiro
Department of Computer Science and Engineering
Indian Institute of Technology, Madras

Week – 08
Lecture – 37
Operating System Security (Buffer Overflows)

Hello. In this video, we will talk about Buffer Overflows. Essentially, Buffer Overflows is the vulnerability in the system and it is not just restricted to the operating system, but it could be pertaining to any application that runs in the system. Now buffer overflows is vulnerability and that allows malicious applications to enter into the systems, even though they do not have a valid access essentially it would allow unauthorized access into the system.

Let us look at Buffer Overflows in this particular lecture.

(Refer Slide Time: 01:02)



So, when we look at how an unauthorized user or an unauthorized attacker could gain access into the system. So, we see that it is just by flaws present. There are two types of flaws that a system can have. One it could have bugs present in the application or the operating system in this particular case. Or it could have flaws due to the human factor

when we for instance browse the internet where we see many such web pages opening and prompting us to click on particular things which would take us to may be a malicious website and a result of that would cause malicious applications to be downloaded into the system.

Another one, which is more pertaining to the operating system, is when there are bugs in the operating system code. Now, modern day operating systems especially the ones that we typically use on a desktop and servers are extremely large pieces of code. For instance the current Linux kernel has over ten million lines of code and all these codes are obviously, written by programmers and will have numerous bugs. So, these bugs are not very easy to detect; however, if an attacker decides to look he could find such a bug and he could then exploit this bug in the operating system to gain access into the OS.

And as you know that once the attacker gains access into the OS, he will be able to do various things like he will be able to execute various components of the operating system code, he could control all the resources present in the system, he could also control which users execute in the system and so on. Thus, an unauthorized access through a bug in the operating system is a very critical aspect.

(Refer Slide Time: 03:20)

Program Bugs that can be exploited

- Buffer overflows
 - In the stack
 - In the heap
 - Return-to-libc attacks
- Double frees
- Integer overflows
- Format string bugs

Chester Rebeiro, IITM 3

There are a number of bugs that an attacker can exploit in order to gain unauthorized access into the operating system. So, here is the list of some of them. So, they could be buffer overflows in the stack of the program or in the OS in the heap, they may be something known as Return-to-libc attacks. There are double frees, essentially this occurs when a single memory location which is dynamically allocated through something like a malloc gets freed more than once, there are integer overflow bugs, and there are format string bugs.

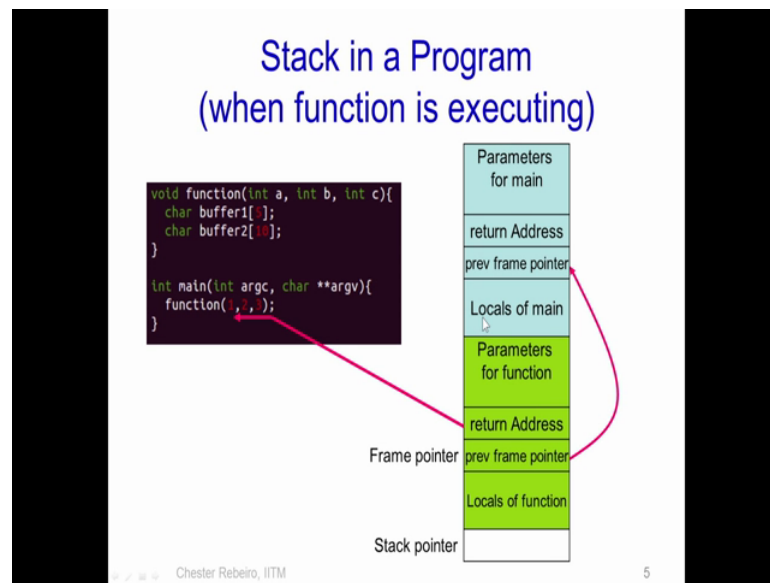
There are essentially numerous different ways that bugs can be exploited by an attacker to enter into the system. So, what we will be seeing today are the bugs in the stack and something known as a Return-to-libc attack which essentially is a variant of the buffer overflow attack in the stack.

(Refer Slide Time: 04:22)



So, in order to understand how the buffer overflows work in the stack, we first need to know how a stack is managed. Let us see how the user stack of a process is managed.

(Refer Slide Time: 04:38)



Let us say we take this very simple example which has two functions the main; and in this main function, we invoke another function with parameters 1, 2 and 3. And this function just allocates two buffers - buffer 1 of 5 bytes and buffer 2 of 10 bytes. So, as we know when we execute this program in the system, the operating system creates a process comprising of various things like the instruction area containing the text or the various instructions of this particular program the data section the heap as well as the stack.

The stack in particular is used for passing parameters from one function to another and it is also used to store local variables. Let us see how the stack is used in this particular example. Let us say that this is the stack and this stack corresponds to when the main function is executing. Now when main wants to invoke this function over here that is function 1, 2 and 3 it begins to push something on to the stack. So, what is pushed onto the stack you will see now?

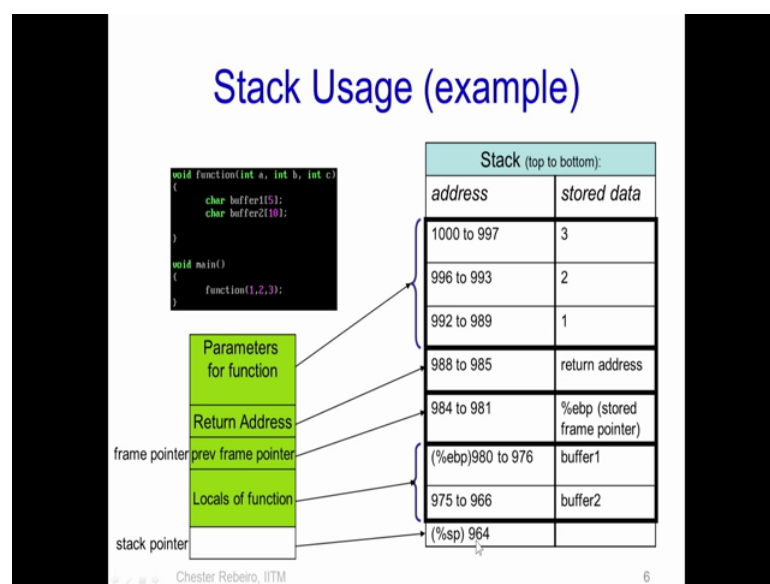
So, first the 3 parameters 1, 2 and 3 which are passed from main two function are pushed on to the stack that is parameters for function 1, 2 and 3 would be pushed onto the stack, then the return address is pushed on to the stack. So, this return address will point to the instruction that follows this function invocation. As we know in order to invoke function

in an x86 processor, the instruction that is used is the call. The return address will point to the next instruction following the call.

So, after the return address something known as the previous frame pointer is pushed onto the stack, this frame pointer points to the frame corresponding to the main function. So, this is the frame which is used when the function is executed, while this frame is used when main is executed. Now after the previous frame pointer is used the local variables, which are defined in function, are then allocated. In this case, we have two character arrays which are allocated 1 is of size 5 and the other is of size 10 bytes.

Besides all of this, we have 2 CPU registers which are used to manage this stack pointer. 1 is the frame pointer which is typically the register bp in Intel x86, and the other one is the stack pointer or sp in the x86 nomenclature. Now the frame pointer points to the current functions frame. So, it actually points to this particular thing corresponding to the frame for function. Now after this function completes its execution and returns, this previous frame pointer is loaded into the register bp, therefore the frame pointer will then point over here that is the frame corresponding to the function main. Now the stack pointer on the other hand points to the bottom of the stack.

(Refer Slide Time: 08:37)



Now, let us look at this in more detailed. Let us say that this is the stack and this is the address for the various stack locations and this is the data stored in that particular address. Let us assume that the top of the stack is 1000, and it decrements downwards. So, this was the stack corresponding to the function when it is invoked. So, we first see that there are the parameters that are passed to the function are pushed onto the stack this is the parameters 3, 2 and 1 which are pushed onto the stack. So, we note that each of these parameters since they have defined as integer in this function is given 4 bytes. The integer a, which is passed to function would start at the address location 997; and from there be 4 bytes 997, 998, 999 and 1000. Similarly, the second and third parameters also take 4 bytes.

The return address for this function at essentially the point at which the function has to return is also given 4 bytes. While the base pointer since it is a 32 system is also given 4 bytes, then we have the buffer 1 which is allocated as a local of the function, which is given 5 bytes 976 to 980; and then buffer 2 is allocated 10 bytes. So, these two arrays are the locals of the function the base pointer points to this particular location and the stack pointer points over here to the address number 964.

(Refer Slide Time: 10:42)

Stack Usage Contd.

```

void function(int a, int b, int c)
{
    char buffer1[5];
    char buffer2[10];
}

void main()
{
    function(1,2,3);
}

```

What is the output of the following?

- printf("%x", buffer2) : 966
- printf("%x", &buffer2[10])
976 → buffer1

Therefore buffer2[10] = buffer1[0]
A BUFFER OVERFLOW

Stack (top to bottom):	
address	stored data
1000 to 997	3
996 to 993	2
992 to 989	1
988 to 985	return address
984 to 981	%ebp (stored frame pointer)
(%ebp)980 to 976	buffer1
975 to 966	buffer2
(%sp) 964	

Chester Rebeiro, IITM 7

Now, let us look at some very simple aspects. Let us say what would happen if we print

this particular line, so printf percentage x buffer 2. So, as we know buffer 2 corresponds to the address of this particular array, this particular printf statement would print the address of buffer 2. So, if we look up the stack we see that the start address of buffer 2 is 966; therefore, this printf will print 966.

Now what happens if we do something like this printf ampersand buffer 2 of 10? So, we know that buffer 2 is of 10 bytes and will have indexes from 0 to 9; now buffer 2 of 10 is 966 plus 10 which is 976. So, what is going to be printed over here is - 976. Now it so happens that 976, is outside the region of buffer 2, in fact, 976 is in buffer 1. Therefore, what we are getting now is that we are printing an address which is outside buffer 2, and this is what is known as a buffer overflow.

Essentially, we have defined a buffer of 10 bytes, but we are accessing data which is outside the buffer 2 area. So, we are accessing the 10th, 11th, 12th and so on byte. So, this is known as a buffer overflow. Now, what we will see next is how this buffer overflow can be exploited by an attacker, and how an attacker could then force a system to execute his own code.

(Refer Slide Time: 12:39)

Modifying the Return Address

buffer2[19] =
&arbitrary memory location

Stack (top to bottom):	
address	stored data
1000 to 997	3
996 to 993	2
992 to 989	1
988 to 985	Arbitrary Location
984 to 981	%ebp (stored frame pointer)
(%ebp)980 to 976	buffer1
976 to 966	buffer2
(%sp) 964	

Chester Rebeiro, IITM 8

Now one important thing from the attackers' perspective is the return address. If the


attacker could somehow fill this buffer 2 in such a way that he would cause a buffer overflow, and modify this particular return address then let us see what would happen. Let us say buffer he makes this particular statement. So, buffer 2 of 19 is some arbitrary memory location. So, what the attacker is doing he is that he is forcing this buffer 2 to overflow and he is overflowing it in such a way that the return address which was stored onto the stack is replaced with his own filled location.

After the function completes executing, it would look into this location and instead of getting the valid return address it would get this arbitrary location, and then it would go to this arbitrary location and start to execute code. So, what we would see is that instead of returning to the main function as would be expected in the normal program, since the attacker has changed this return address to some arbitrary location the processor would then cause this instructions corresponding to this arbitrary location to be fetched and execute it. So, now, it looks quite obvious what the attacker could do in order to create an attack.

(Refer Slide Time: 14:19)

Stack (top to bottom):	
address	stored data
1000 to 997	3
996 to 993	2
992 to 989	1
988 to 985	ATTACKER'S code pointer
984 to 981	%ebp (stored frame pointer)
(%ebp)980 to 976	buffer1
976 to 966	buffer2
(%sp) 964	

Now that we seen how buffer overflows can skip an instruction, We will see how an attacker can use it to execute his own code (exploit code)



Chester Rebeiro, IITM 9

Essentially what the attacker is going to do is that he is going to change this return address the valid return address present in the stack with a pointer to an attack code. Therefore, when the function returns instead of taking the standard return address, it


overflow the buffer with BA. So, he keeps overflowing the buffer with BA, and when this happens at one particular case the returned address present in the stack is changed from the valid return address to B A.

Thus, when the function returns what the CPU is going to see is the address BA, in the return address location thus it is going to take this address BA and start executing code from that. So, since BA corresponds to the address where the exploit code is present, the CPU would then begin to execute this exploit code. And in this way, the attacker could force the CPU or the processor to execute the exploit code.

(Refer Slide Time: 17:33)

Exploit Code

- Lets say the attacker wants to spawn a shell
- ie. do as follows:



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

void main()
char *name[2];

name[0] = "/bin/sh"; /* exe filename */
name[1] = NULL; /* exe arguments */
execve(name[0], name, NULL);
exit(0);
```

- How does he put this code onto the stack?

Chester Rebeiro, IITM 11

So, now we will see how a one particular attack code is created, and how an attacker can force an application or an operating system to execute that exploit code. Let us take a very simple example of the exploit code, which is shown over here. Essentially this exploit code which we call as shell code does nothing but only executes a particular shell, the shell is specified by slash bin slash sh.

And this particular function, `execve` is invoked in order to execute this shell the parameters are `name 0`, which comprises of the executable name, and there is `name 1` which is essentially a null terminated string. So, we will see how this particular code can be

forced to be executed by an unauthorized attacker. The first question that needs to be asked is how does this attacker manage to put this code onto the stack.

(Refer Slide Time: 18:48)

Step 1 : Get machine codes

```
include cstdio.h
include cstdlib.h

void main()
char name[2];

name[0] = "A/a/a/h"; /* exe filename */
name[1] = NULL; /* exe arguments */
execl(name[0], name, NULL);
exit(0);
}
```

```
00000000 main():
0: 55                push  %ebp
1: 89 e5             mov   %esp,%ebp
3: eb 1e             jmp   23 <main+0x23>
5: 5c                pop   %esi
6: 89 76 08          mov   %esi,0x8(%esi)
7: c6 4b 07 00      movb  $0x0,0x7(%esi)
4: c7 4b 0c 00 00 00 movl  $0x0,0xc(%esi)
14: b8 0b 00 00 00   mov   $0xb,%eax
13: 89 f3             mov   %esi,%ebx
15: 84 4e 08          lea  0x8(%esi),%ecx
1e: 84 56 0c          lea  0xc(%esi),%ebx
21: c4 80             int   $0x80
23: e8 44 ff ff ff   call  5 <main+0x5>
```

```
void main(void){
asm
"movl $1f,%esi;"
"movl %esi,0x8(%esi);"
"movb $0x0,0x7(%esi);"
"movl $0x0,%eax;"
"movl %esi,%ebx;"
"leal 0x8(%esi),%ecx;"
"leal 0xc(%esi),%ebx;"
"int $0x80;"
"section .data;"
".string \"A/a/a/h\" \",\"
"section .text;"
};
}
```

- objdump -disassemble-all shellcode.o
- Get machine code: "eb 1e 5e 89 76 08 c6 46 07 00 c7 46 0c 00 00 00 00 b8 0b 00 00 00 89 f3 8d 4e 08 8d 56 0c cd 80 cd 80"
- If there are 00s replace it with other instructions

Chester Rebeiro, IITM 12

The first step in doing so is that the attacker needs to obtain the binary data corresponding to this particular program. In order to do this, what the attacker does is that he will re write this program in assembly code. So, this assembly code as we see here does exactly the same thing as done by this program.

The next thing what the attacker would do is to compile this particular assembly code and get what is known as the object dump. The object dump is obtained by running this particular command thing. So, he first compiles this particular code to get what is known as the shell code dot o which is the object file, and then he will run this particular command which is of objdump disassemble all shellcode dot o to get this particular file.

Now, what is important for us over here is this particular column or the second column. The numbers what you see over here the hexadecimal numbers are in fact, the machine code for this program. The numbers like eb, 1e, 5e, 89, 26, 08, and so on correspond to the machine code of this particular program. In other words, if the attacker manages to put this machine code onto the stack and is able to force execution to this particular

machine code, then the attacker would be able to execute the shell as required.

The machine code is shown over here and one thing which is required for this particular attack is to replace all the zeros present in this machine code with some other instructions, so that you do not have any zeros present over here.

(Refer Slide Time: 20:48)

The slide displays the following C code snippet:

```
char large_string[128];  
char buffer[48]; ← Defined on stack  
  
0  
0  
0  
0  
0  
  
strcpy(buffer, large_string);
```

The slide also includes a footer with the text "Chester Rebeiro, IITM" and the number "13".

The next thing is to scan the entire application in order to find one location which can be exploited for a buffer overflow. So, essentially the requirements for a buffer overflow is that the attacker finds in the application code a command such as this a string copy buffer comma large string where buffer is a small array and it is defined locally in the stack while the large string is a much larger array. So, as we know the way this particular function strcpy works is that the large string gets copied to buffer and this copying will continue byte by byte until there is a slash 0, which is found in large string; in which case the strcpy will complete executing and will return.

Let us assume that the attacker has found such a case where we have the buffer a large string and a string copy, and the buffer is a small array defined onto the stack, and how does the attacker make use of this.

(Refer Slide Time: 22:46)

Step 3 (contd) : Fill up Large String with BA

```
char large_string[128];  
char buffer[48]; ← Address of buffer is BA
```

large_string

shellcode	BA	BA	BA	BA	BA	BA	BA	BA	BA
-----------	----	----	----	----	----	----	----	----	----

Chester Rebeiro, IITM 15

Then he computes what the address of the buffer should be, and fills the remaining part of the large string with the buffer address.

(Refer Slide Time: 23:00)

Final state of Stack

- Copy large string into buffer

```
strcpy(buffer, large_string);
```
- When strcpy returns the exploit code would be executed

buffer

BA
BA
BA
BA
BA
BA
BA
BA
shellcode

large_string

shellcode	BA	BA	BA	BA	BA	BA	BA	BA	BA
-----------	----	----	----	----	----	----	----	----	----

BA buffer Address

Chester Rebeiro, IITM 16

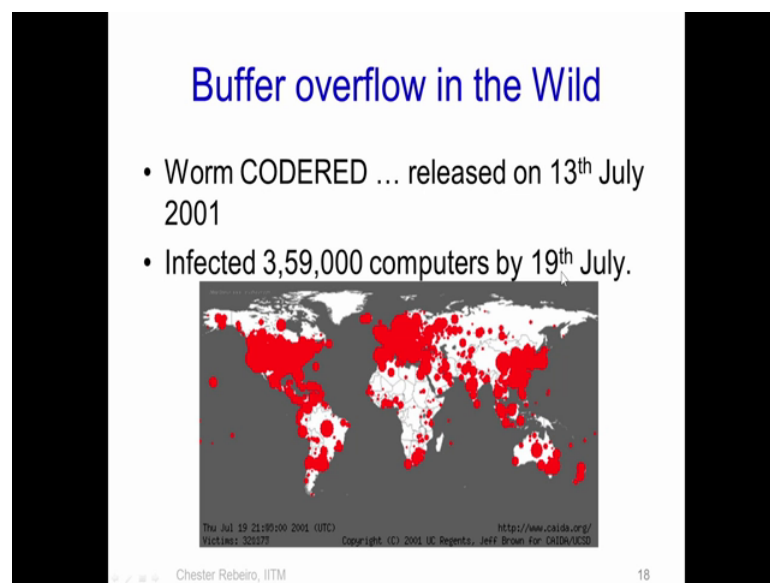
Now he needs to force this string copy to execute with this buffer and the large string which he has just created as shown over here. Now as result of this string copy being

Next, let us assume that somewhere in the application there is this particular code, we have the large string and we have short string, which is buffer, which is a local array a locally defined array and therefore, gets created onto the stack.

So, what we first do is somehow manage to fill the long string or the large string with the address of buffer. So, if you recollect, this is the BA paths which are present, then we will copy the shell code on to the large string. So, we have created this large string in the format that we require. In the first part of this large string is that the shell code, and then it is followed by the buffer return address. And then if there is a function like string copy which copies large string into buffer it will result in a buffer overflow to occur and instead of the function returning to this particular point soon after string copy.

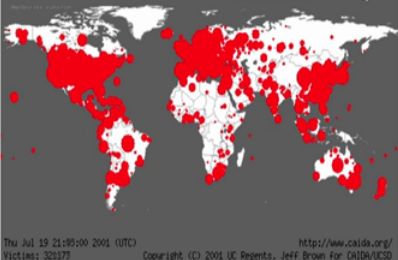
On the other hand, execute this particular shell code and cause this shell specified by this command to be executed. So, if we actually see this if we execute gcc overflow dot c or this is called overflow 1 dot c and run dot slash a dot out. Instead of just doing this string copy and exiting, this particular program created a new shell due to the exploit code that is executing.

(Refer Slide Time: 26:44)



Buffer overflow in the Wild

- Worm CODERED ... released on 13th July 2001
- Infected 3,59,000 computers by 19th July.



Thu Jul 19 21:00:00 2001 (UTC) <http://www.caida.org/>
Riches: 32473 Copyright (C) 2001 UC Regents, Jeff Broin for CAIDA/UCSD

Chester Rebeiro, IITM 18

So, buffer overflows are an extremely well known bug and extremely exploited by

