

Information Security - 5 - Secure Systems Engineering
Prof. Chester Rebeiro
Indian Institute of Technology, Madras
Demonstration of a Return Oriented Programming (ROP) Attack

Hello and welcome to this demonstration in the course for secure system engineering, in this demonstration we will looking at ROP attack, now ROP attacks are extremely difficult to do, so we will be just glimpsing about one particular assignment which was submitted by some of the students in the course when I was actually teaching it.

(Refer Slide Time: 0:37)

The image consists of two screenshots of a Linux terminal window. The top screenshot shows the compilation of a ROP attack exploit. The terminal prompt is `root@kali:~/rop_attack`. The user enters `gcc rop_attack.c -o rop_attack`. The terminal output shows the compilation process, including the use of `gcc` and the resulting executable `rop_attack`. The bottom screenshot shows the execution of the exploit. The user enters `./rop_attack`. The terminal output shows the execution of the exploit, including the use of `system()` and the resulting output `root@kali:~/rop_attack`.

```
root@kali:~/rop_attack# gcc rop_attack.c -o rop_attack
root@kali:~/rop_attack# ./rop_attack
root@kali:~/rop_attack#
```

```

ngpt@ubuntu:~/ngpt_codes/module4/NOP
ngpt$ gcc -std=c99 -c *.c
ngpt$ gcc -std=c99 -o ngpt_codes/module4/NOP *.c
ngpt$ ./ngpt_codes/module4/NOP
Input 19 words:
aaaa
bbbb
cccc
ddddd
eeee
ffff
gggg
hhhhh
iiii
jjjjjjj
Merge and the first characters from the 19 words concatenated:
hcdkfgjkl
Null pointers ( 01175022 and 01175022)
value in gbb is 9
ngpt@ubuntu:~/ngpt_codes/module4/NOP vtn tut2$

```

```

ngpt@ubuntu:~/ngpt_codes/module4/NOP
ngpt$ gcc -std=c99 -c *.c
ngpt$ gcc -std=c99 -o ngpt_codes/module4/NOP *.c
ngpt$ ./ngpt_codes/module4/NOP
Input 19 words:
aaaa
bbbb
cccc
ddddd
eeee
ffff
gggg
hhhhh
iiii
jjjjjjj
Merge and the first characters from the 19 words concatenated:
hcdkfgjkl
Null pointers ( 01175022 and 01175022)
value in gbb is 9
ngpt@ubuntu:~/ngpt_codes/module4/NOP vtn tut2$

```

```

ngpt@ubuntu:~/ngpt_codes/module4/NOP
ngpt$ gcc -std=c99 -c *.c
ngpt$ gcc -std=c99 -o ngpt_codes/module4/NOP *.c
ngpt$ ./ngpt_codes/module4/NOP
Input 19 words:
aaaa
bbbb
cccc
ddddd
eeee
ffff
gggg
hhhhh
iiii
jjjjjjj
Merge and the first characters from the 19 words concatenated:
hcdkfgjkl
Null pointers ( 01175022 and 01175022)
value in gbb is 9
ngpt@ubuntu:~/ngpt_codes/module4/NOP vtn tut2$

```

```
mpit_0x@ubuntu:~/Desktop/0x00000000
Terminal
mpit_0x@ubuntu:~/Desktop/0x00000000$ gcc -std=c99 -g -fstack-protector -static test2.c -o test2
mpit_0x@ubuntu:~/Desktop/0x00000000$ ./test2
[
]
let glib;
char rollnumber[10] = "1111111111";
char rollnumber2[10] = "1111111111";
void concatenate_first_chars()
{
    struct {
        char word_buf[10];
        int i;
        char* cat_pointer;
        char cat_buf[10];
    } locals;
    locals.cat_pointer = locals.cat_buf;
    printf("Input 10 words:\n");
    for(locals.i=0; locals.i<10; locals.i++)
    {
        // Read from stdin
        if(fgets(locals.word_buf, 1024, stdin) == 0) { locals.word_buf[0] = '\0';
        printf("Failed to read words");
        return;
        }
        // Copy first char from word to next location to concatenated by
        *locals.cat_pointer = *locals.word_buf;
        locals.cat_pointer++;
    }
    // Even if something goes wrong, there's a null byte here
    // preventing buffer overflow
    locals.cat_buf[10] = '\0';
    printf("Now see the first characters from the 10 words concatenated:\n");
    printf("%s", locals.cat_buf);
}
let main(int argc, char** argv)
{
    if(argc != 1)
    {
        printf("Usage: %s\n", argv[0]);
        return EXIT_FAILURE;
    }
    concatenate_first_chars();
    printf("Roll Number: %s and %s", rollnumber, rollnumber2);
    printf("Value to glib is %d", glib);
}
mpit_0x@ubuntu:~/Desktop/0x00000000$
```

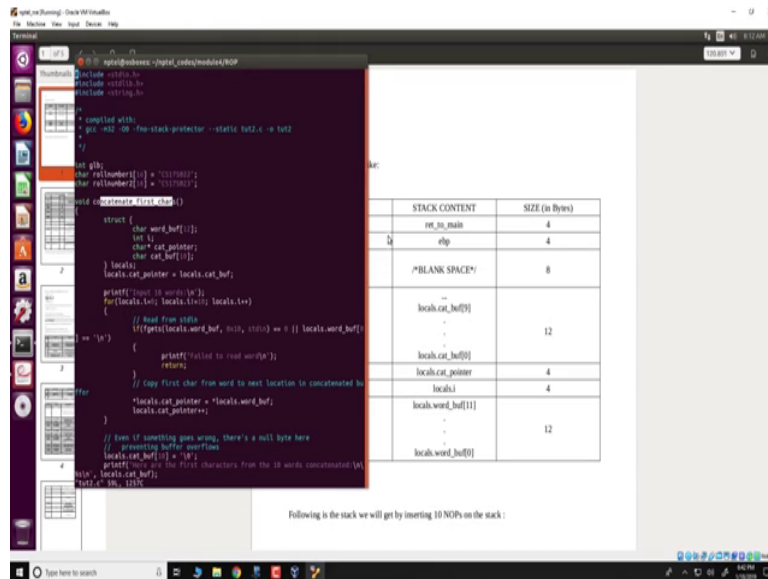
```
mpit_0x@ubuntu:~/Desktop/0x00000000
Terminal
mpit_0x@ubuntu:~/Desktop/0x00000000$ gcc -std=c99 -g -fstack-protector -static test2.c -o test2
mpit_0x@ubuntu:~/Desktop/0x00000000$ ./test2
[
]
let glib;
char rollnumber[10] = "1111111111";
char rollnumber2[10] = "1111111111";
void concatenate_first_chars()
{
    struct {
        char word_buf[10];
        int i;
        char* cat_pointer;
        char cat_buf[10];
    } locals;
    locals.cat_pointer = locals.cat_buf;
    printf("Input 10 words:\n");
    for(locals.i=0; locals.i<10; locals.i++)
    {
        // Read from stdin
        if(fgets(locals.word_buf, 1024, stdin) == 0) { locals.word_buf[0] = '\0';
        printf("Failed to read words");
        return;
        }
        // Copy first char from word to next location to concatenated by
        *locals.cat_pointer = *locals.word_buf;
        locals.cat_pointer++;
    }
    // Even if something goes wrong, there's a null byte here
    // preventing buffer overflow
    locals.cat_buf[10] = '\0';
    printf("Now see the first characters from the 10 words concatenated:\n");
    printf("%s", locals.cat_buf);
}
let main(int argc, char** argv)
{
    if(argc != 1)
    {
        printf("Usage: %s\n", argv[0]);
        return EXIT_FAILURE;
    }
    concatenate_first_chars();
    printf("Roll Number: %s and %s", rollnumber, rollnumber2);
    printf("Value to glib is %d", glib);
}
mpit_0x@ubuntu:~/Desktop/0x00000000$
```

```
mpit_0x@ubuntu:~/Desktop/0x00000000
Terminal
mpit_0x@ubuntu:~/Desktop/0x00000000$ gcc -std=c99 -g -fstack-protector -static test2.c -o test2
mpit_0x@ubuntu:~/Desktop/0x00000000$ ./test2
[
]
let glib;
char rollnumber[10] = "1111111111";
char rollnumber2[10] = "1111111111";
void concatenate_first_chars()
{
    struct {
        char word_buf[10];
        int i;
        char* cat_pointer;
        char cat_buf[10];
    } locals;
    locals.cat_pointer = locals.cat_buf;
    printf("Input 10 words:\n");
    for(locals.i=0; locals.i<10; locals.i++)
    {
        // Read from stdin
        if(fgets(locals.word_buf, 1024, stdin) == 0) { locals.word_buf[0] = '\0';
        printf("Failed to read words");
        return;
        }
        // Copy first char from word to next location to concatenated by
        *locals.cat_pointer = *locals.word_buf;
        locals.cat_pointer++;
    }
    // Even if something goes wrong, there's a null byte here
    // preventing buffer overflow
    locals.cat_buf[10] = '\0';
    printf("Now see the first characters from the 10 words concatenated:\n");
    printf("%s", locals.cat_buf);
}
let main(int argc, char** argv)
{
    if(argc != 1)
    {
        printf("Usage: %s\n", argv[0]);
        return EXIT_FAILURE;
    }
    concatenate_first_chars();
    printf("Roll Number: %s and %s", rollnumber, rollnumber2);
    printf("Value to glib is %d", glib);
}
mpit_0x@ubuntu:~/Desktop/0x00000000$
```

So this assignment essentially it work with a file in this directory called ROP and you can actually download this file, it is part of VM which comes along with this course and in the VM,, if the directory NPTEL codes module 4 ROP, you could actually find this required code, so the file that was attacked was TUT2.C or it actually tutorial 2.C and there is essentially did a lot of things, there was a main and there was a concatenate first chares the function and I have to give credit to the students who would this codes.

So their roll numbers are present here, the right working of this code is as follows, so you, this is the executable tutorial 2, it inputs ten words as like this and the output is the first characters of every input word, so you see over here that the first characters ABCD to J are actually printed out on the screen, so this is what this program actually does and there is a wonderbilty in this particular program, the wonderbilty is present in this function, concatenate first chares and we will all go into what the wonderbilty is but believe it as an assignment and as a challenge to any of the viewers to actually try to find out these wonderbilty in this function.

But what we will see over here is, we will try to build an ROP attack which uses that wonderbilty, so the aim of the ROP attack is to somehow compute 10 factorial, as you know 10 factorial is the product of numbers from 1 to 10, so the aim of these attack is to determine or computer 10 factor using ROP gadgets and fill in the value of the result in this global variable GLB, so this GLB value would be printed somewhere, over here print GLB and you can actually look at this file and note that nowhere in the file we are actually using GLB anywhere else except in this print F and also note that we are not computing 10 factorial or any other factorial anywhere in this particular program.



Okay, so starting with the ROP attack, the first thing to do with the ROP attack is to identify the gadgets which are present in the program, so the gadgets can be obtained using this particular, using a particular tool which is present in this directory ROP gadget master, it can go to this ROP gadget master and use the tool python ROP gadget.PY specify a binary as the input and provide the path to the executable, which you want to evaluate, so what this gadget tool would do?

The ROP gadget tool would do was that it would pass through the entire executable, tutorial 2 and identify all possible gadgets that it can find, so we can just redirectory output into this file E1 okay and will open E1 and see that these are the gadgets that it will found okay, so note that the each of this gadgets has a few instructions and then has a return or a return F or call instruction at the end, this particular address over here is the address of the particular gadget.

So as you see the tool has found a large number of gadgets present in the executable tutorial 2 and in fact there are like 13,276 gadgets that are present, now a coming back to the payload that we want to write, we want to compute 10 factorial using these gadgets, so essentially what we need to do is pick out a gadgets from these 13,276, arrange them in a manner on the stack, so that in the end 10 factorial is computed.

So will not go into details about how we are actually picking up these gadgets but we will just see a report on this, so report is present over here okay, so we just go look at this report for the entire attack and what we see over here is the contains of the stack and this is the contains of the stack when the function concatenate first chares is been executed.

So as we know that when this function is executed there is a stack frame which is active and this frame look something like this will, so there is return address which is stored, which is of 4 bytes, then there is a previous frame pointer essentially the frame pointer corresponding to main, then there is some other data and locals which are present, so for our specific compilation the locations of the stack and the stack frame are present as follows over here but when you run it or when you actually try to execute this program, you may get a different value for these addresses, however the structure of the stack and the relativeness of wed, the various data are present on the stack would be identical.

(Refer Slide Time: 7:54)

(i)

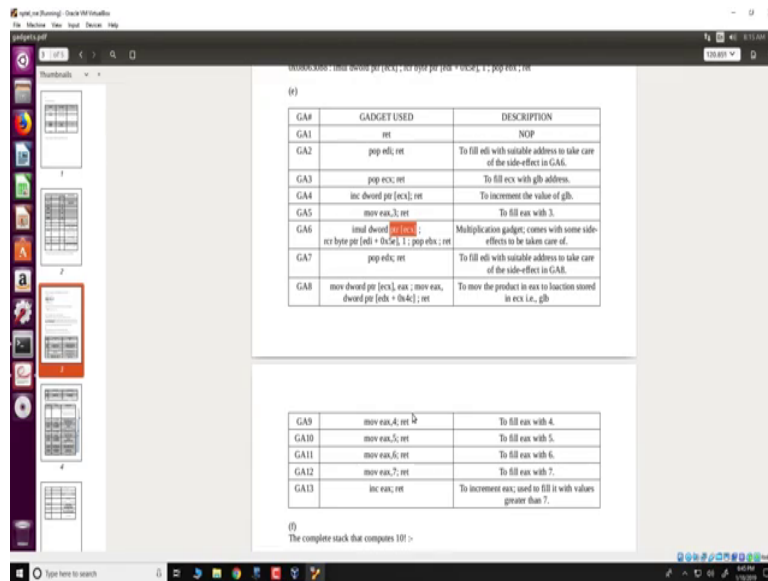
The original stack looks like:

LOCATION	STACK CONTENT	SIZE (in Bytes)
0xffffcc	ret_0x_main	4
0xffffc8	ebp	4
0xffffe0	/*BLANK SPACE*/	8
0xffffd4	...	12
...	locals.cat_buf[0]	
...	...	
0xffffb4	locals.cat_buf[0]	4
0xffffb0	locals.cat_pointer	
0xffffac	locals.i	4
0xffffa8	locals.word_buf[11]	12
...	...	
0xffffa0	locals.word_buf[0]	

Following is the stack we will get by inserting 10 NOPS on the stack :

So what an ROP attack actually does, is that it replaces this return to main and starts filling in gadgets over here, so starting from the first gadgets which is placed at this point, you would have gadgets going upwards like this, such that each of these gadgets have some operation in computing the 10 factorial.

(Refer Slide Time: 8:14)



The gadgets that we use are actually there are 14 of them, pick from this space of that 13,000 odd gadgets which are found in tutorial 2 using the gadget tool and these are the 13 gadgets which were used, so these gadgets varied in functionality from very simple things like return which is equal an to just doing no-op operation, to something which is much more complicated like this gadget number 6, which is essentially integer multiplication, it multiplies some contents pointed to by these ECX register with the EX register, the description of each of this gadgets is present over here, so what we do in order to compute 10 factorial is we select gadgets and start to build a sequence of operations which he place on the stack such that after all of these gadgets complete its execution, we would be able to complete or obtain 10 factorial.

(Refer Slide Time: 9:24)

gdbgui

Address	Disassembly	Comment
0x00000000	GA1 <0x0040112>	GA1: ret
0x00000001	GA6 <0x0003008>	GA6: imul dword ptr [ecx]; rct byte ptr [edi + 0x5c], 1; pop ebx; ret
0x00000002	GA5 <0x0006f00>	GA5: mov eax, 3; ret
0x00000003	GA4 <0x000446d>	GA4: inc dword ptr [ecx]; ret
0x00000004	GA4 <0x000446d>	GA4: inc dword ptr [ecx]; ret
0x00000005	<0x0004b20>	gh add in ecx
0x00000006	GA3 <0x00040b1>	GA3: pop ecx; ret
0x00000007	<0x000e50>	So that [edi+0x5c] is valid.
0x00000008	GA2 <0x0004480>	GA2: pop ecx; ret

Address	Disassembly	Comment
0x00000009	GA1 <0x0040112>	Originally this had the return to main. GA1: ret
0x0000000a		We just replaced the contents of ebp by its previous value itself.
0x0000000b	a	Blank space filled with some random stuff.
0x0000000c	a	
0x0000000d	a	
0x0000000e	locals_cat_buf	
0x0000000f	a	
0x00000010	a	The first byte from word_buf[] is copied at each iteration.
0x00000011	locals_cat_pointer	
0x00000012		Its value starts from 0x00000010 and increments in each iteration.
0x00000013	0 or 9	locals i

gdbgui

Address	Disassembly	Comment
0x00000014	GA6 <0x0003008>	GA6: imul dword ptr [ecx]; rct byte ptr [edi + 0x5c], 1; pop ebx; ret
0x00000015	GA5 <0x0006f00>	GA5: mov eax, 4; ret
0x00000016	GA8 <0x00077be>	GA8: mov dword ptr [ecx]; mov eax, dword ptr [edi + 0x5c]; ret
0x00000017	<0x000e00>	[edi + 0x5c] is valid
0x00000018	GA7 <0x000403a>	GA7: pop ebx; ret
0x00000019	GA1 <0x0040112>	GA1: ret
0x0000001a	GA6 <0x0003008>	GA6: imul dword ptr [ecx]; rct byte ptr [edi + 0x5c], 1; pop ebx; ret
0x0000001b	GA5 <0x0006f00>	GA5: mov eax, 3; ret
0x0000001c	GA4 <0x000446d>	GA4: inc dword ptr [ecx]; ret
0x0000001d	GA4 <0x000446d>	GA4: inc dword ptr [ecx]; ret
0x0000001e	<0x0004b20>	gh add in ecx
0x0000001f	GA3 <0x00040b1>	GA3: pop ecx; ret
0x00000020	<0x000e50>	So that [edi+0x5c] is valid.
0x00000021	GA2 <0x0004480>	GA2: pop ecx; ret

Address	Disassembly	Comment
0x00000022	GA1 <0x0040112>	Originally this had the return to main. GA1: ret
0x00000023		We just replaced the contents of ebp by its previous value itself.
0x00000024	a	Blank space filled with some random stuff.
0x00000025	a	
0x00000026	locals_cat_buf	

gdbgui

f) The complete stack that computes 10! :-

LOCATION	CONTENT	EXPLANATIONS
	<0x0040953>	return to main, the original return address.
	..	
	..	Repeating the same set of gadgets with a few minor tweaks to increment eax, we finally get the value of 10! in gh.
	..	
	..	
0x00000014	GA8 <0x00077be>	GA8: mov dword ptr [ecx]; mov eax, dword ptr [edi + 0x5c]; ret
0x00000015	<0x000e00>	[edi + 0x5c] is valid
0x00000016	GA7 <0x000403a>	GA7: pop ebx; ret
0x00000017	GA1 <0x0040112>	GA1: ret
0x00000018	GA6 <0x0003008>	GA6: imul dword ptr [ecx]; rct byte ptr [edi + 0x5c], 1; pop ebx; ret
0x00000019	GA5 <0x0006f00>	GA5: mov eax, 4; ret
0x0000001a	GA8 <0x00077be>	GA8: mov dword ptr [ecx]; mov eax, dword ptr [edi + 0x5c]; ret
0x0000001b	<0x000e00>	[edi + 0x5c] is valid
0x0000001c	GA7 <0x000403a>	GA7: pop ebx; ret
0x0000001d	GA1 <0x0040112>	GA1: ret
0x0000001e	GA6 <0x0003008>	GA6: imul dword ptr [ecx]; rct byte ptr [edi + 0x5c], 1; pop ebx; ret
0x0000001f	GA5 <0x0006f00>	GA5: mov eax, 3; ret
0x00000020	GA4 <0x000446d>	GA4: inc dword ptr [ecx]; ret
0x00000021	GA4 <0x000446d>	GA4: inc dword ptr [ecx]; ret
0x00000022	<0x0004b20>	gh add in ecx
0x00000023	GA3 <0x00040b1>	GA3: pop ecx; ret

The modified stack with all gadgets placed look something like this, so it starts from gadget number 1 placed in this location, this location originally had the return to main which we had seen and now we overflow a buffer and replace this return to main this sequence of gadgets, so what happens during executions, is first gadget 1 executes and the return in gadget 1 would pick this address from the stack which corresponds to gadget 2 and execute, than gadget 3, gadget 4 and so on executes in such a way, the sequence of placing the gadgets in such a way so that at the end of the sequence 10 factorial is computed.

(Refer Slide Time: 10:14)

<0x0040953>	return to main, the original return address.	
..	..	Repeating the same set of gadgets with a few minor tweaks to increment eax, we finally get the value of 10 in glb.
..	..	
..	..	
..	..	
..	..	
0x0040114	GA8 <0x00077be>	GA8: mov dword ptr [ecx], eax ; mov eax, dword ptr [edx + 0x4c] ; ret
0x0040110	<0x00e960>	[edx + 0x4c] is valid
0x00400fc	GA7 <0x000403a>	GA7: pop ecx; ret
0x00400f8	GA1 <0x0040132>	GA1: ret
0x00400f4	GA6 <0x000308b>	GA6: imul dword ptr [ecx] ; rcr byte ptr [edx + 0x5e], 1 ; pop ebx ; ret
0x00400f0	GA9 <0x0006f7d>	GA9: mov eax, 4 ; ret
0x00400ec	GA8 <0x00077be>	GA8: mov dword ptr [ecx], eax ; mov eax, dword ptr [edx + 0x4c] ; ret
0x00400e8	<0x00e960>	[edx + 0x4c] is valid
0x00400e4	GA7 <0x000403a>	GA7: pop ecx; ret
0x00400e0	GA1 <0x0040132>	GA1: ret
0x00400dc	GA6 <0x000308b>	GA6: imul dword ptr [ecx] ; rcr byte ptr [edx + 0x5e], 1 ; pop ebx ; ret
0x00400d8	GA5 <0x0006f0d>	GA5: mov eax, 3 ; ret
0x00400d4	GA4 <0x000446f>	GA4: inc dword ptr [ecx] ; ret
0x00400d0	GA4 <0x000446f>	GA4: inc dword ptr [ecx] ; ret
0x00400cc	<0x0040132>	glb add in ecx
0x00400c8	GA3 <0x0004083>	GA3: pop ecx; ret
0x00400c4	<0x00e975>	So that [edx+0x5e] is valid.
0x00400c0	GA2 <0x0040140>	GA2: pop ebx; ret

GA12: mov eax, 7 ; ret To fill eax with 7.

GA13: inc eax; ret To increment eax; used to fill it with values greater than 7.

(f) The complete stack that computes 10! :-


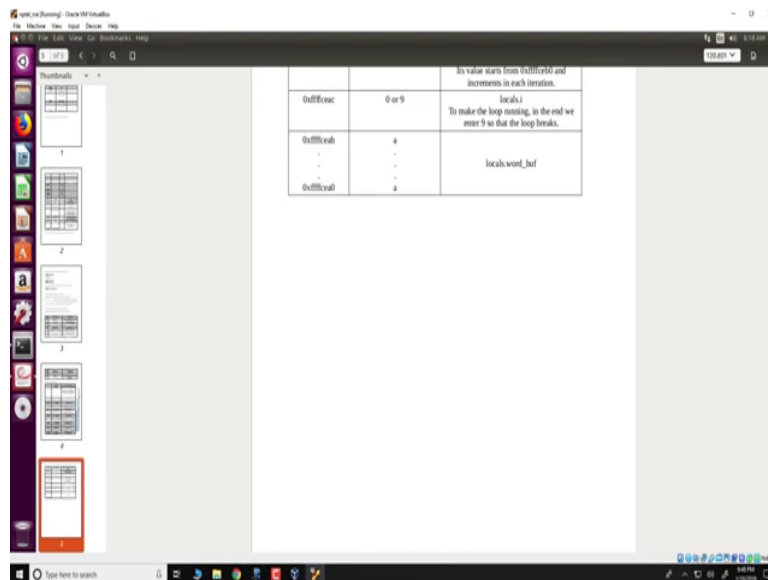
LOCATION	CONTENT	EXPLANATIONS
<0x0040953>	return to main, the original return address.	
..	..	Repeating the same set of gadgets with a few minor tweaks to increment eax, we finally get the value of 10 in glb.
..	..	
..	..	
..	..	
..	..	
0x0040114	GA8 <0x00077be>	GA8: mov dword ptr [ecx], eax ; mov eax, dword ptr [edx + 0x4c] ; ret
0x0040110	<0x00e960>	[edx + 0x4c] is valid
0x00400fc	GA7 <0x000403a>	GA7: pop ecx; ret
0x00400f8	GA1 <0x0040132>	GA1: ret
0x00400f4	GA6 <0x000308b>	GA6: imul dword ptr [ecx] ; rcr byte ptr [edx + 0x5e], 1 ; pop ebx ; ret
0x00400f0	GA9 <0x0006f7d>	GA9: mov eax, 4 ; ret
0x00400ec	GA8 <0x00077be>	GA8: mov dword ptr [ecx], eax ; mov eax, dword ptr [edx + 0x4c] ; ret
0x00400e8	<0x00e960>	[edx + 0x4c] is valid
0x00400e4	GA7 <0x000403a>	GA7: pop ecx; ret
0x00400e0	GA1 <0x0040132>	GA1: ret
0x00400dc	GA6 <0x000308b>	GA6: imul dword ptr [ecx] ; rcr byte ptr [edx + 0x5e], 1 ; pop ebx ; ret
0x00400d8	GA5 <0x0006f0d>	GA5: mov eax, 3 ; ret
0x00400d4	GA4 <0x000446f>	GA4: inc dword ptr [ecx] ; ret

To be more specific for example each of these sequences for example at this position we would have computed 3factorial and the result for 3 factorial would be placed in the global variable GLB, here to here from the gadgets GAN9 present in this location to the gadget GA8

over here we will compute 4factorial, so this is evident from this move EX to 4 and then to the contents of GLB, we multiply 4 and therefore after a few more gadgets we would have stored a 4factorial in GLB, now this is repeated several times in a sequence are like 5, 6 and so on until the entire 10 factorial is computed, the one hurdle which one would probably face while trying to go this is that the stack segment may actually overflow for instant because this is from the second function which is invoked soon after main.

Therefore, the content of the stack is relatively small and if we keep adding gadgets onto this stack, the stack may quite easily overflow and thing not work, so it is quite tricky way to choose gadgets and place gadgets in such a way so that the require payload gets executed.


(Refer Slide Time: 11:52)



The image contains two screenshots of a terminal window. The top screenshot shows a list of ROP gadgets with their addresses and assembly instructions, such as `0007100 0000 0000 e883 0000 0000 0000 0000`. The bottom screenshot shows a block of assembly code starting with `CALL MOV_ESI_5_1_POC` and `CALL MOV_ESI_5_2_POC`, which are typical gadgets used in ROP chains to construct a payload.

So let us see how this thing works, will not go into details about the internals of this because it is extremely complex and essentially out of scope for this particular course but people just run this and the input is, the input where the gadgets are placed is in this file, import VM, so it is a binary file, so we could open it with octal dumb profile and this is how, this is what the input is actually used to the tutorial program, so let us see how this ROP gadgets actually executes, so there is a payload_VM.PY, will just have a look at that and see that, this particular code arranges the gadgets and forms the required input which is placed in Y, in the string Y over here and finally Y gets printed.

(Refer Slide Time: 13:02)



So this particular python script or you could actually go through it to see how these gadgets are arranged and eventually the file that gets created is, this file called input._VM, so we run this tutorial 2, give it this malicious input, which is input_VM and we see that the value in GLB is 3628800, this essentially is the value for 10 factorial are which you can actually verify, so one problem in this code is that, there is a segmentation fault and after the 10 factorial gets printed and this is because the program was not terminated in a safe way or with a nice exit, so I would leave an assignment or a very hard challenge, one of the most difficult challenges in this course is to modify this particular code in such a way that the ROP gadget terminates in a grace hole way. Thank you.