Linux Rootkit

Adrien 'schischi' Schildknecht

July 17, 2015

Section 1

IDT hooking

Why?

Main interface between the kernel and the world (userland, hardware...)

Modifying the IDT

- The Address of the IDT is stored in a register;
- Changing an entries:
  - Modify the table (RO);
  - Create a new table;

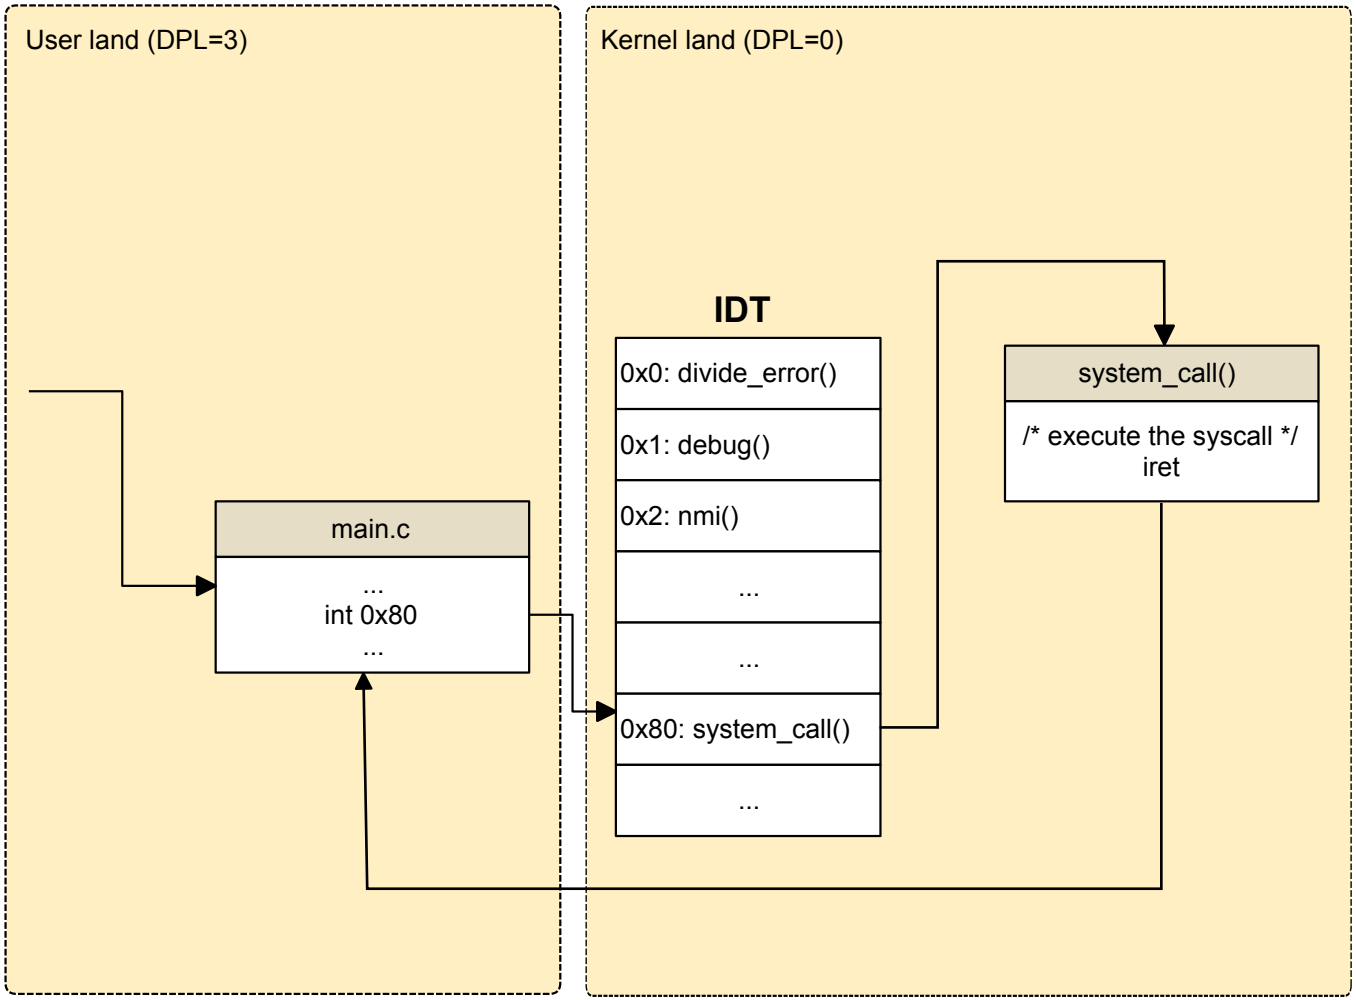

Interrupt stack frame

Pre hook

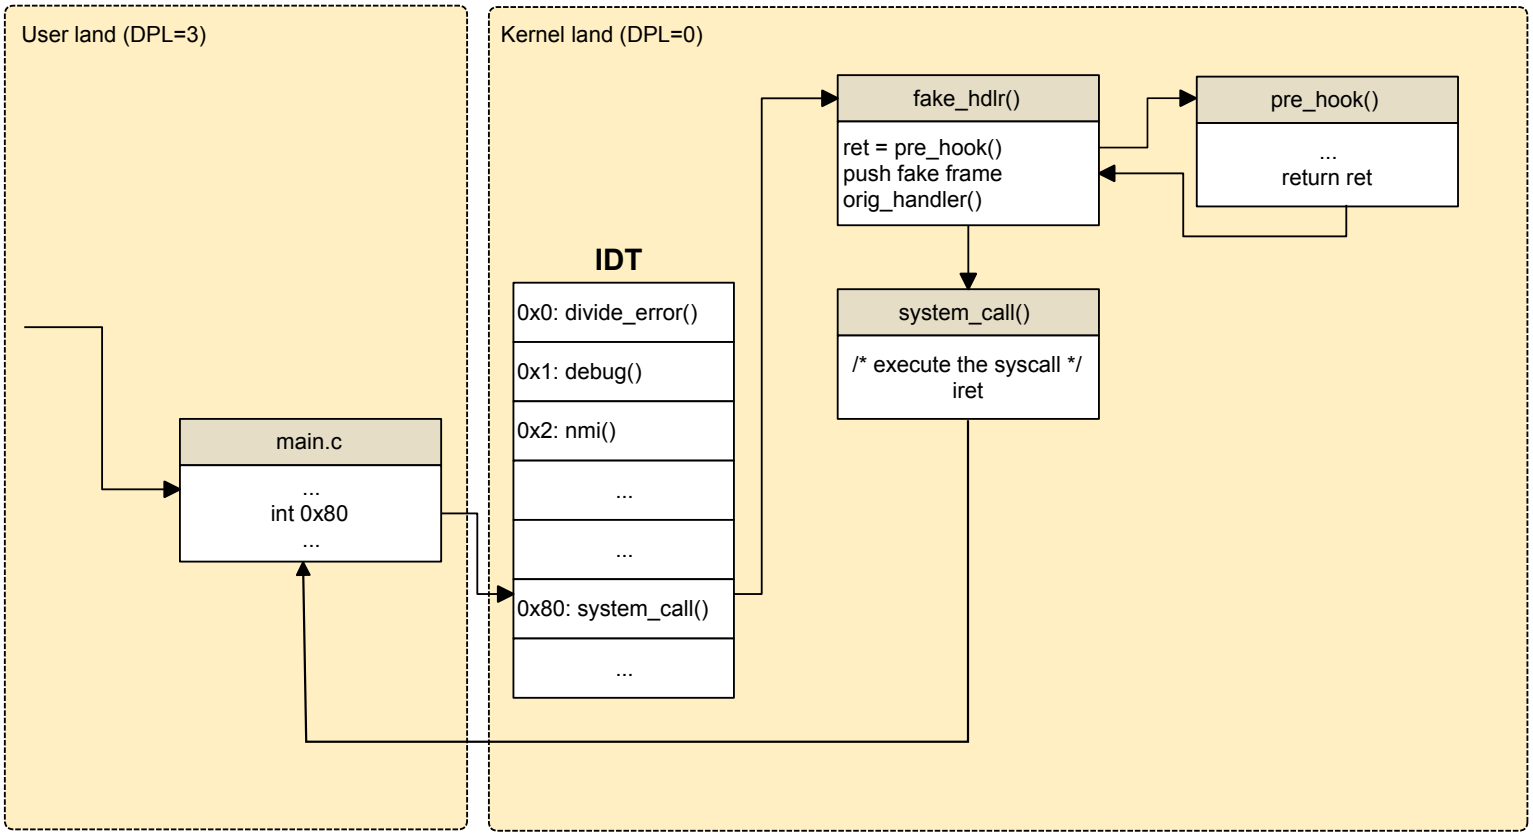

Post hook

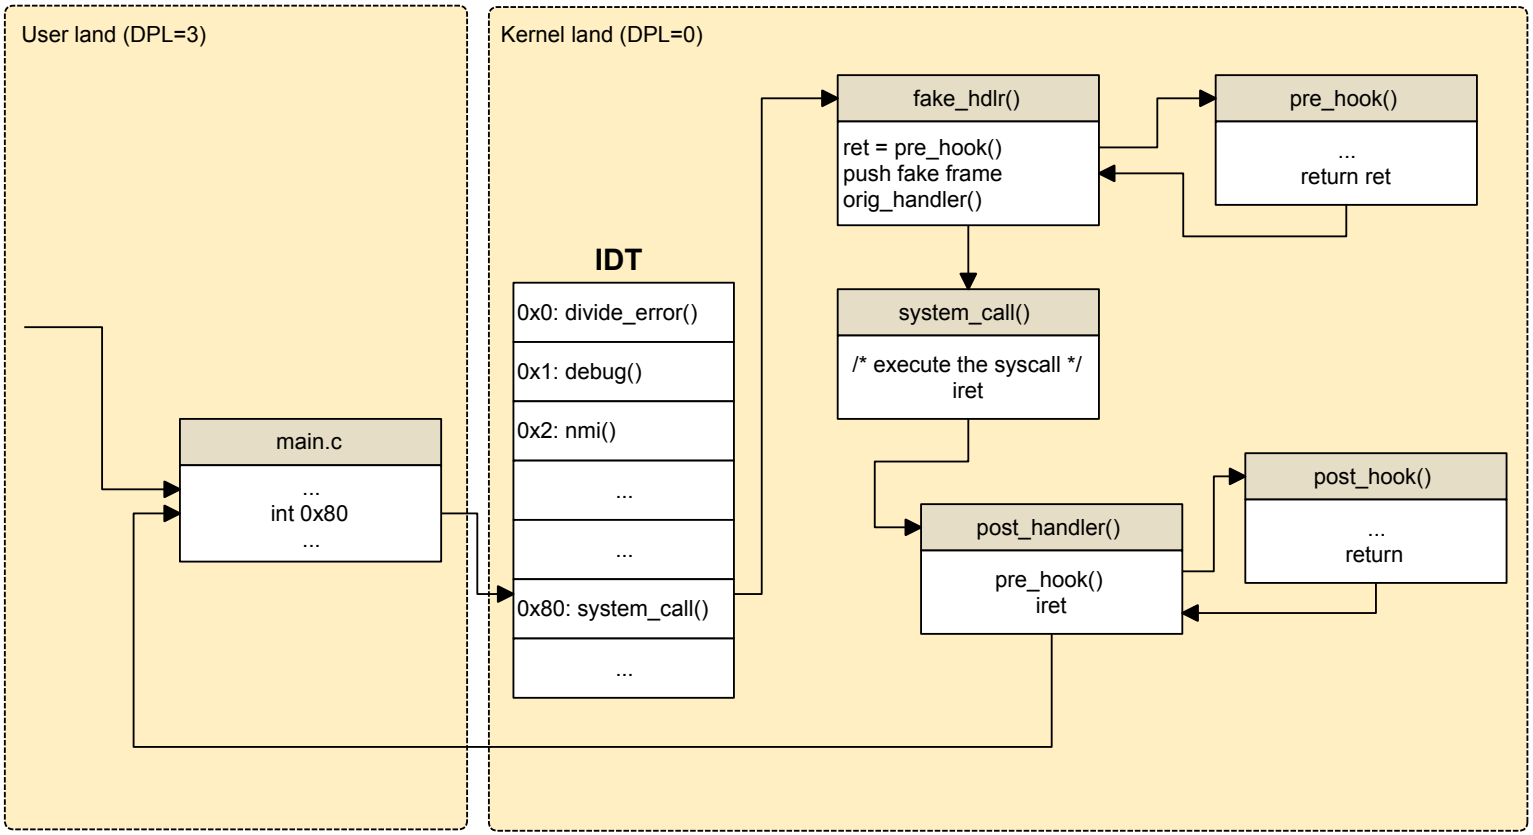

Summary

Section 2

Syscall hooking

How to make a syscall

3 ways:

- 32bits: int 0x80, sysenter (Intel), syscall (AMD);
- 64bits: syscall;

Int 0x80

Sysenter


```
1   /* Obtain a valid pointer to per cpu data*/
2   swapgs
3   /* Setup a stack */
4   mov $stack_sysenter, %rsp
5   add %gs:this_cpu_off, %rsp
6   /* Save registers on the stack */
7   sub $0x28, %rsp  /* Skip exception frame */
8   SAVE_REGS
9   /* Fill exception frame */
10  movl 12(%rbp),  %eax        /* RIP */
11  movq %rax,      0x80(%rsp)
12  movq $0x23,     0x88(%rsp)  /* CS */
13  movq $0x0,      0x90(%rsp)  /* RFLAGS */
14  movl 0x0(%rbp), %eax        /* RSP */
15  movq %rax,      0x98(%rsp)
16  movq $0x2b,     0xa0(%rsp)  /* SS */
17  mov %rsp, %rdi
18  /* Set an invalid esp as return addr */
19  movl $__stringify(0x42cafe42), 12(%rbp)
20  /* Pre-hook ! */
21  call *sysenter_pre_hook
22  RESTORE_REGS
23  /* Call the original handler without swapgs */
24  jmp *(sysenter_orig_hdlr + 3)
25
```

Syscall

Section 3

Conclusion

Conclusion

```c
#define MEGA(S) ((S) * 1024 * 1024)

int main(int argc, char *argv[]) {
    char buf[4096];
    int fd = open("/home/schischi/foo", O_CREAT | O_WRONLY,
    0660);

    if (argc == 2 && !strcmp(argv[1], "-f"))
        if (fallocate(fd, 0, 0, MEGA(700)) != 0)
            return 1;
    for (int i = 0; i < MEGA(700) / sizeof (buf); ++i)
        write(fd, buf, 4096);
    write(fd, buf, MEGA(700) % sizeof (buf));

    unlink("/home/schischi/foo");
    return 0;
}

```

```
$ repeat 100; ./a.out
    ./a.out  0.01s user 1.46s system 18% cpu 8.018 total

$ repeat 100; ./a.out -f
```

Conclusion

Questions ?

schischi@lse.epita.fr

schischi - irc.rezosup.org

References

- FS design
  - Book "Practical File System Design" by Dominic Giampaolo
- VFS
  - http://git.kernel.org/cgit/linux/kernel/git/torvalds/linux.git
  - http://lwn.net/Kernel/Index/
- Journaling, logging
  - http://pages.cs.wisc.edu/~remzi/OSTEP/file-lfs.pdf
  - http://research.cs.wisc.edu/wind/Publications/sba-usenix05.pdf
- Ext4
  - https://ext4.wiki.kernel.org/index.php/Ext4_Design
  - http://www.ibm.com/developerworks/library/l-anatomy-ext4/
- Btrfs
  - http://video.linux.com/videos/chris-mason-btrfs-file-system
  - http://atrey.karlin.mff.cuni.cz/~jack/papers/lk2009-ext4-btrfs.pdf