

lec 3

Interrupts + System Calls

Warm up:

How do function calls
work in assembly?

- args?
- return value?
- return address?
- registers?
- how to make the call? return?

Brief reminder about function calls

main:

0x1008

callq

foo

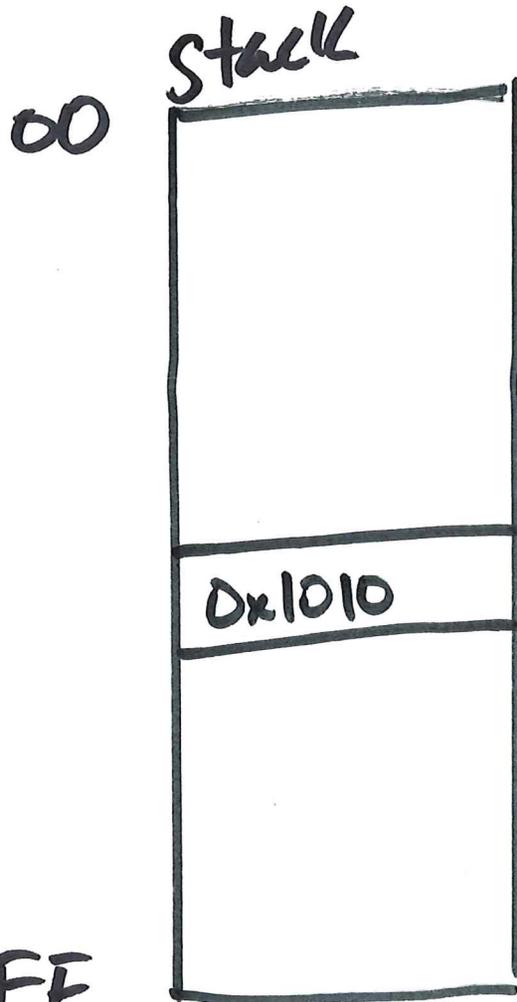
0x1010

:

foo:

:

retq

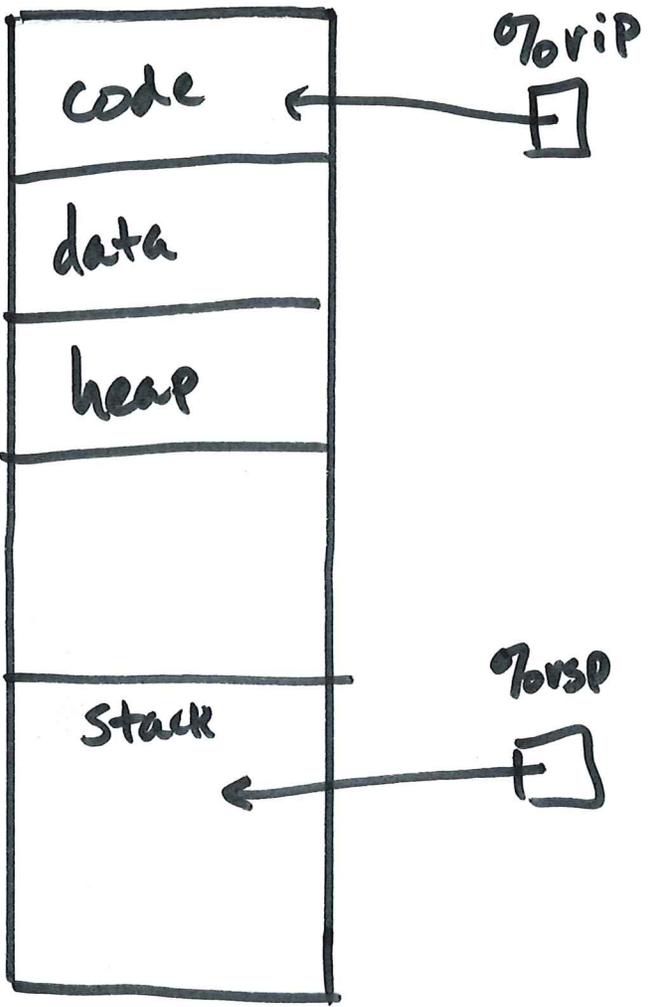


FF

x86-64
calling convention:

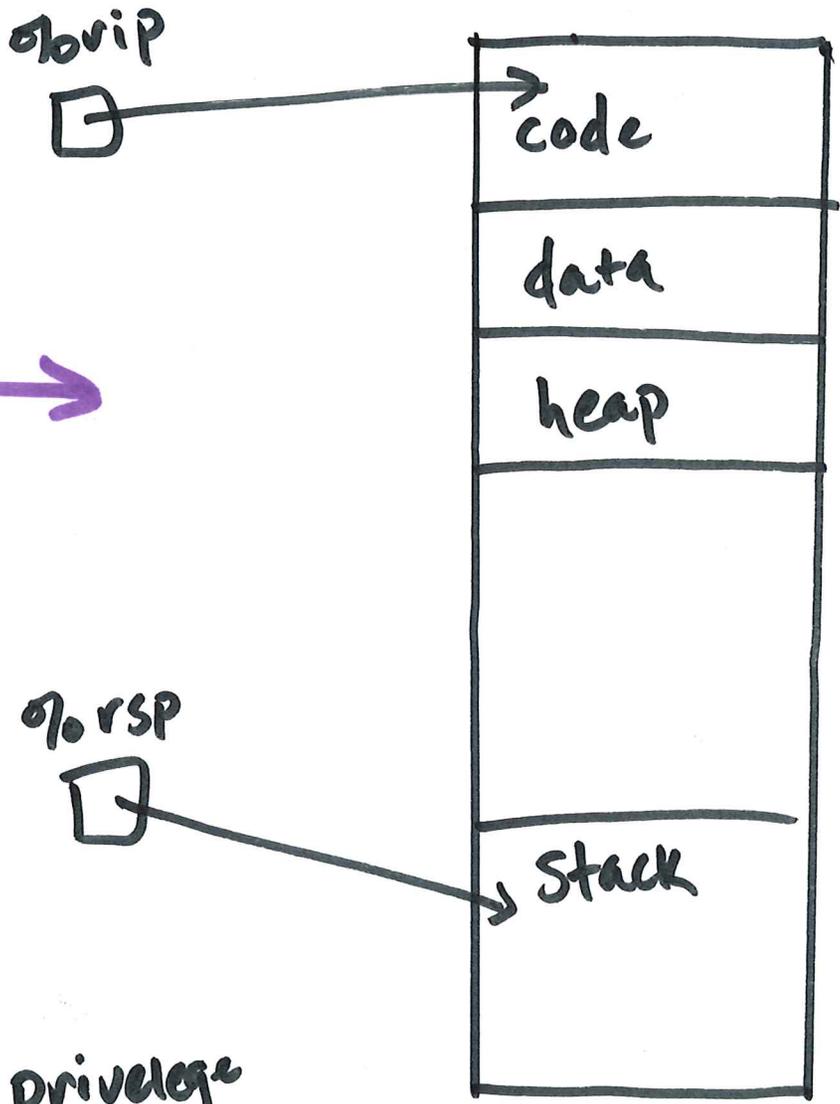
callee-saved: rbx, rsp, rbp,
r12, r13, r14, r15

proc's virtual memory



privilege = user

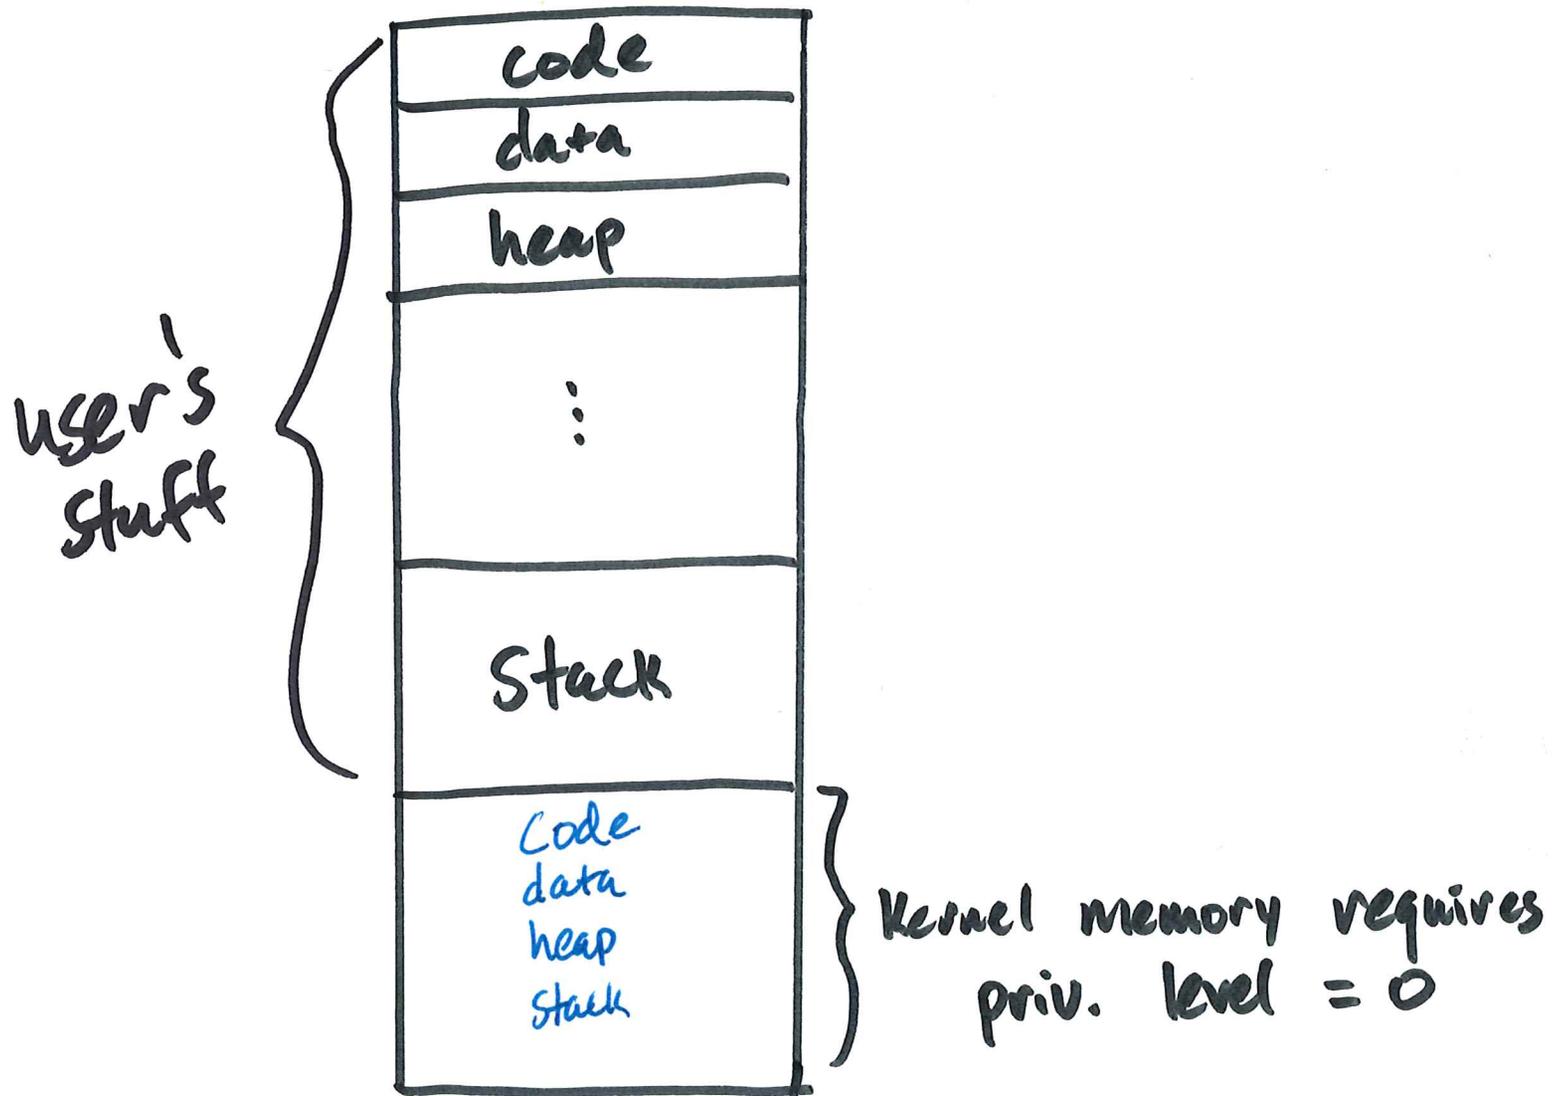
Kernel



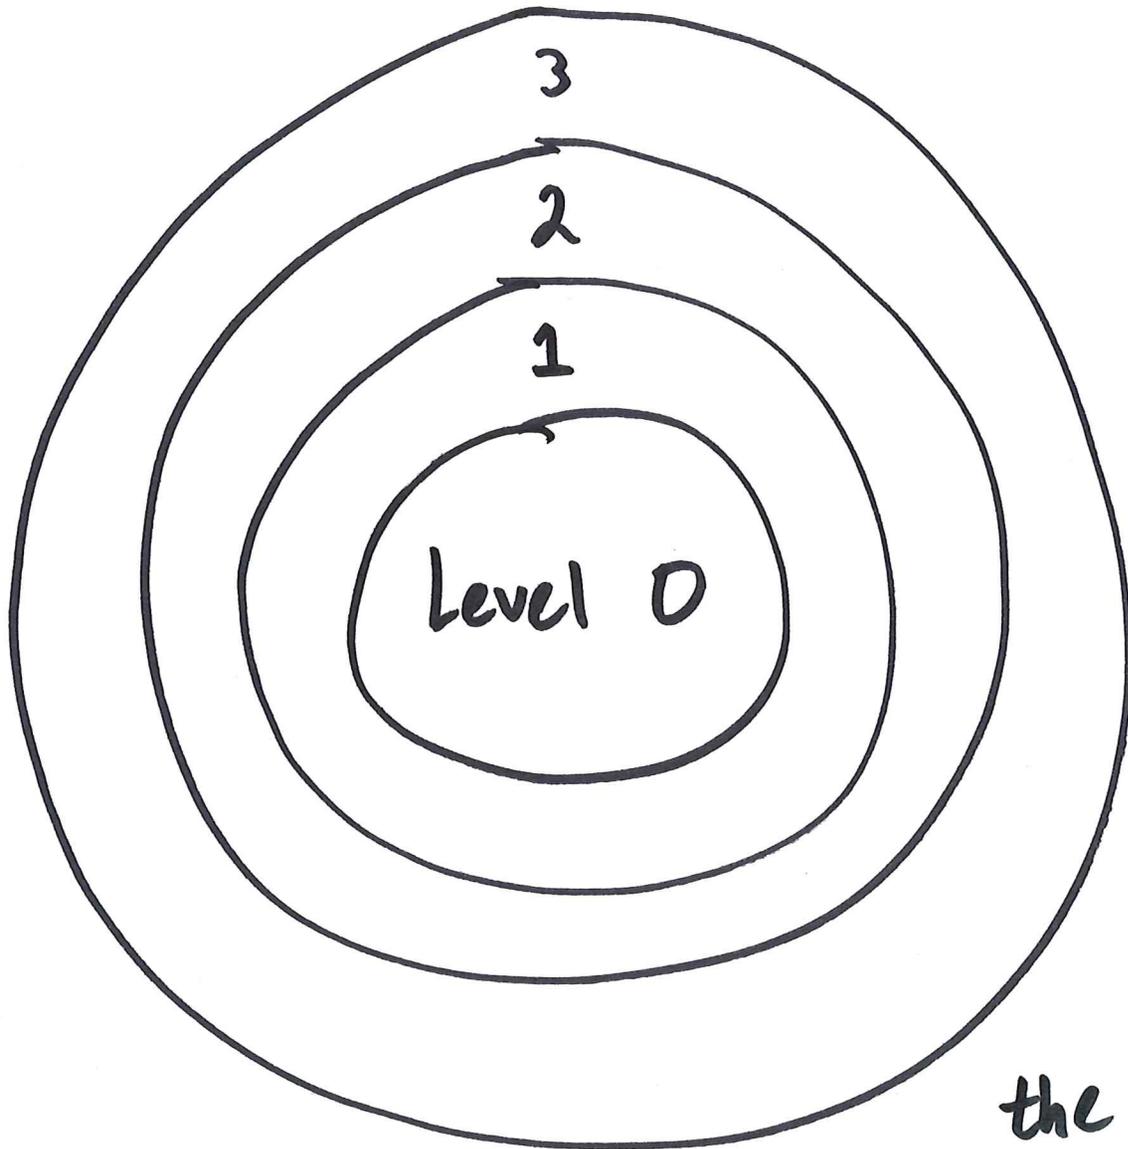
privilege = kernel



Proc's virtual address space
also maps kernel data, etc.



x86 privilege levels

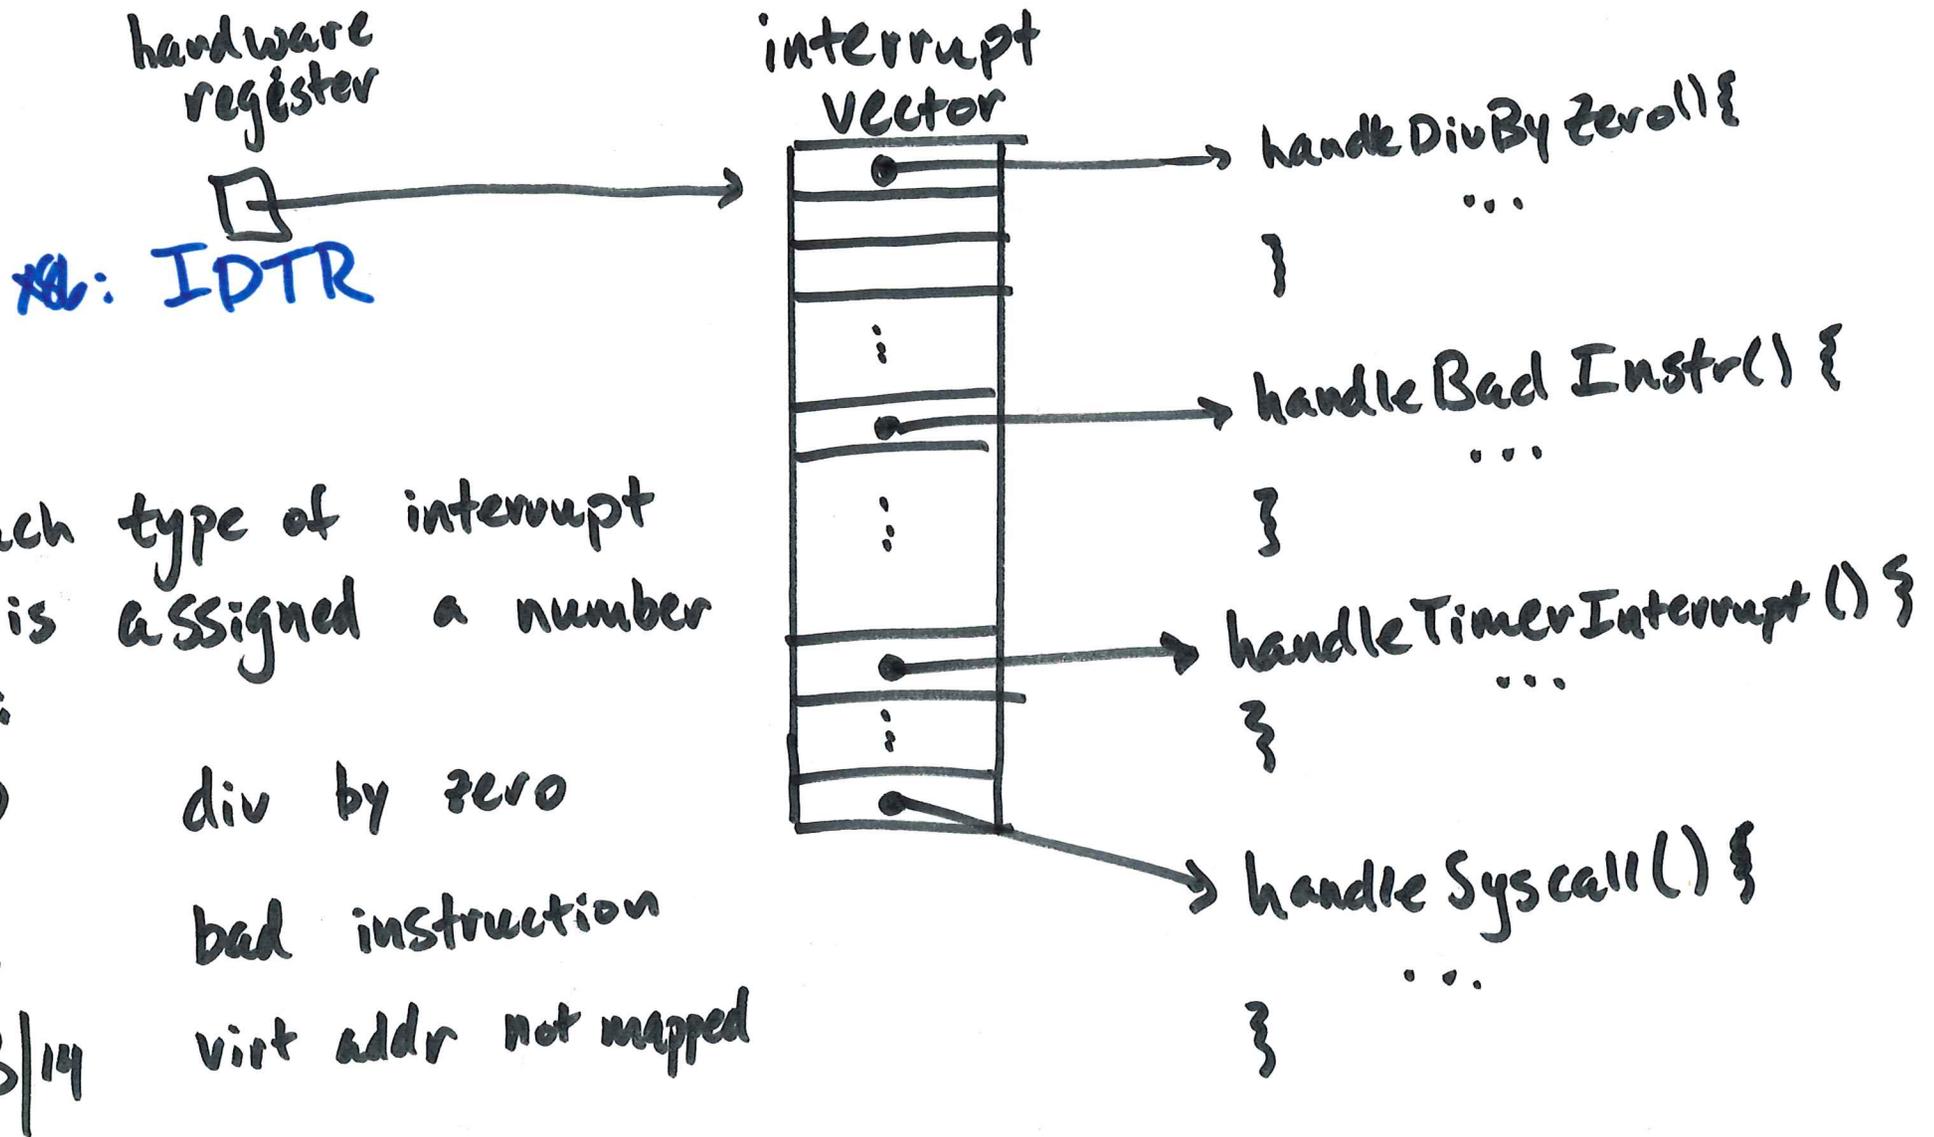


"ring 0"
= kernel mode

"ring 3"
= user mode

the current privilege level (CPL) is stored in low 2 bits of CS segment reg 3.5

Hardware Interrupts



- each type of interrupt is assigned a number

x86:

- 0 div by zero
- 6 bad instruction
- 13/14 virt addr not mapped

32-255 available

x86: syscall = 64

What must be true about mode transfer?

- Kernel must control entry points

 - malicious app must not jump to arbitrary kernel code

- program counter, stack pointer, priv. level all change

 - what would happen if we used the proc's stack?

 - could be invalid

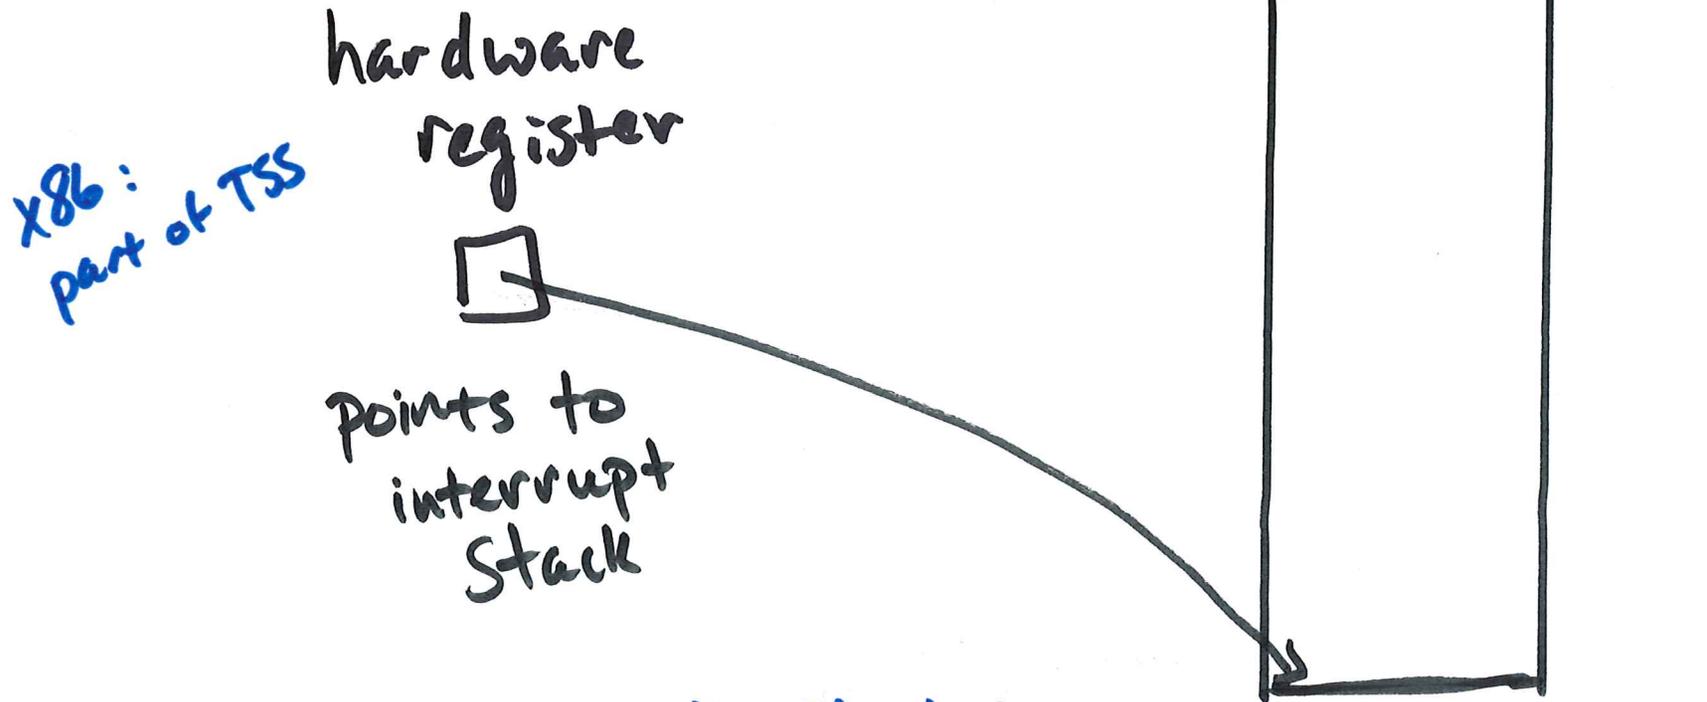
 - could point to kernel memory

- old values of PC, SP, CPL must be saved

 - transparently restart process

 - mode transfers must happen at known place in instruction stream

Interrupt Stack



- in user mode, kernel stack empty
- on interrupt, set $\%rsp$ to this register
- convenient to allocate separate kernel stack for each process

Hardware support for mode transfer (x86-64)

- before calling handler
set up kernel stack
like this

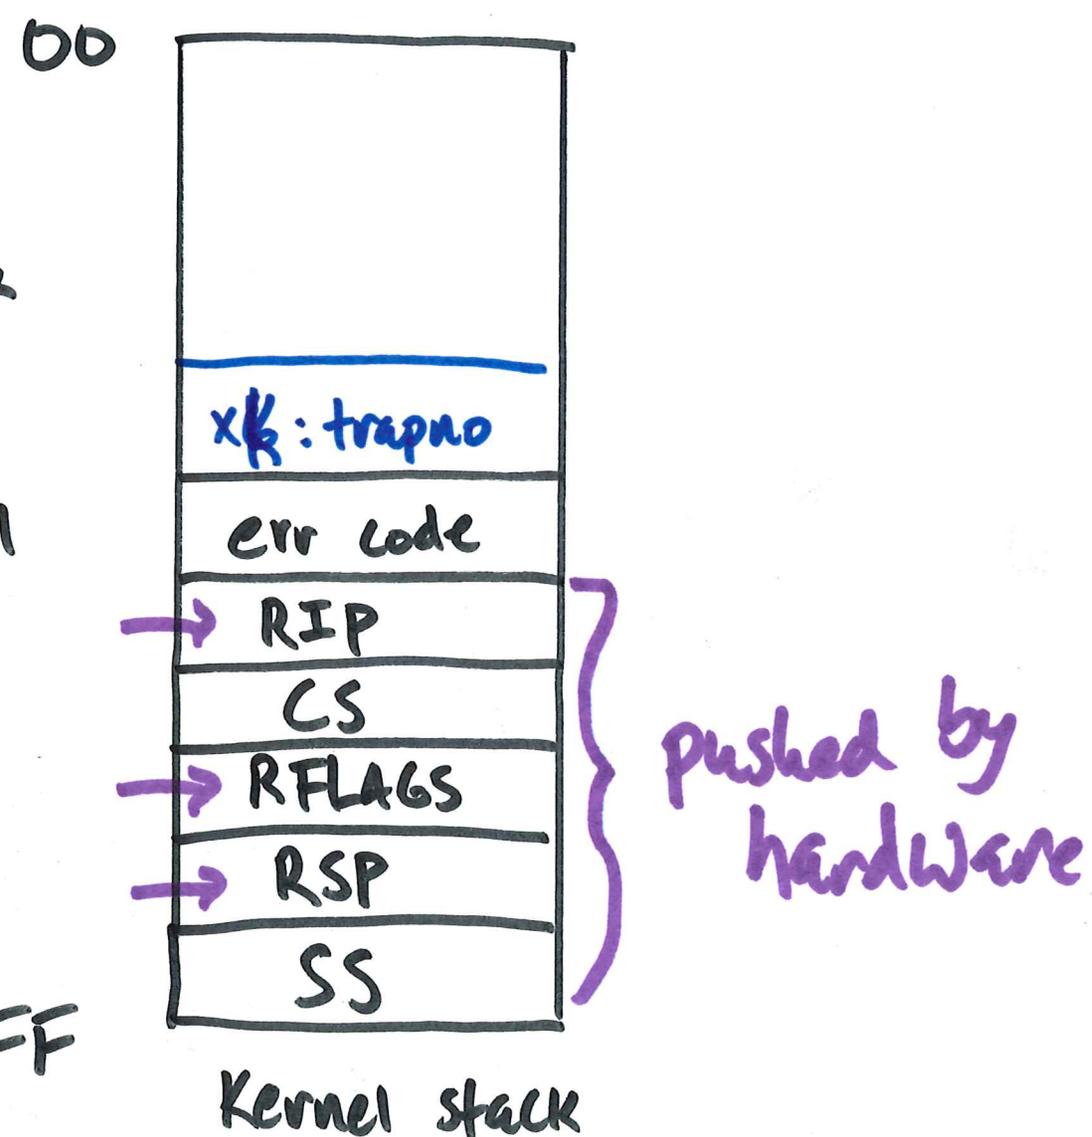
- handler should save all
other registers!

- to restore:

→ iretq

(after restoring
all other regs

+ popping trapno+err code) (see inc/trap.h)

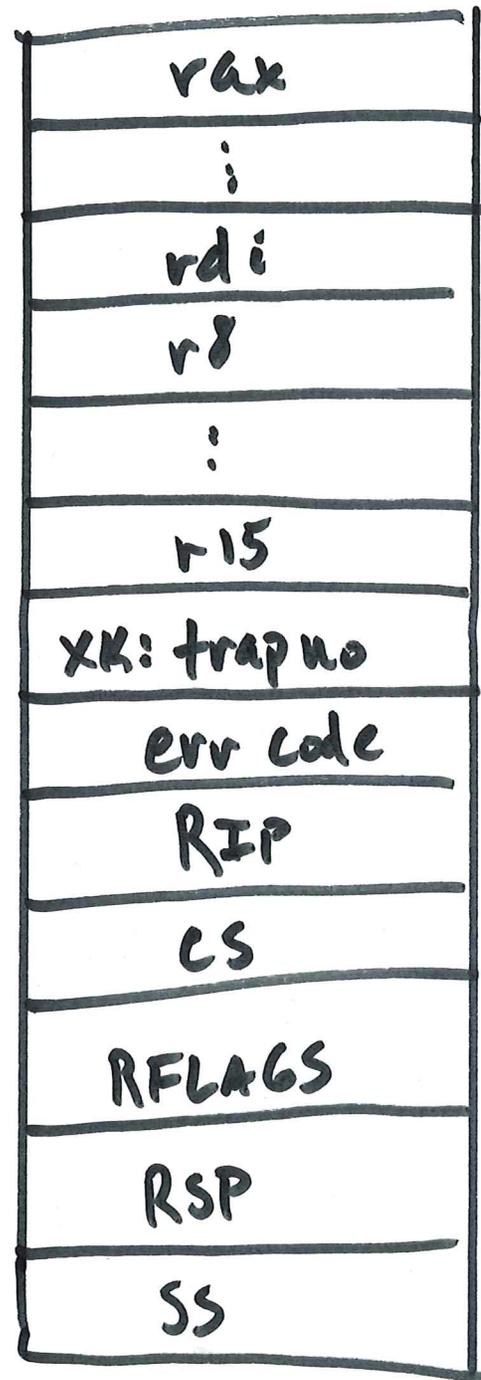


After saving all registers,
stack looks like this

Only now is it safe to
call C functions in kernel.
(Why?)

- this data structure is called
a trap frame

struct trap_frame



Kernel stack

Finishing interrupts

- timer goes off
- hardware calls handler
- handler saves all regs + calls C
- C decides what to do
 - run other process?
 - go back to same process?
- C returns
- handler restores regs + executes `iretq`

System Calls

- application wants to request kernel service
 - read a file, for example
- can do this via "software interrupts"
 - "int" instruction on x86
 - takes interrupt num. as arg
 - invokes that handler
- assign an interrupt number to system calls
 - xk: 64

```
main() {  
    file_open(arg1, arg2);  
}
```



```
file_open(arg1, arg2) {  
    movq $SYS_OPEN, %rax  
    int $64  
    retq  
}
```



```
file_open(arg1, arg2) {  
    // do operation  
}
```



```
file_open_handler() {  
    // checks --  
    // copy args  
    file_open(arg1, arg2)  
    // copy ret val  
    return;  
}
```

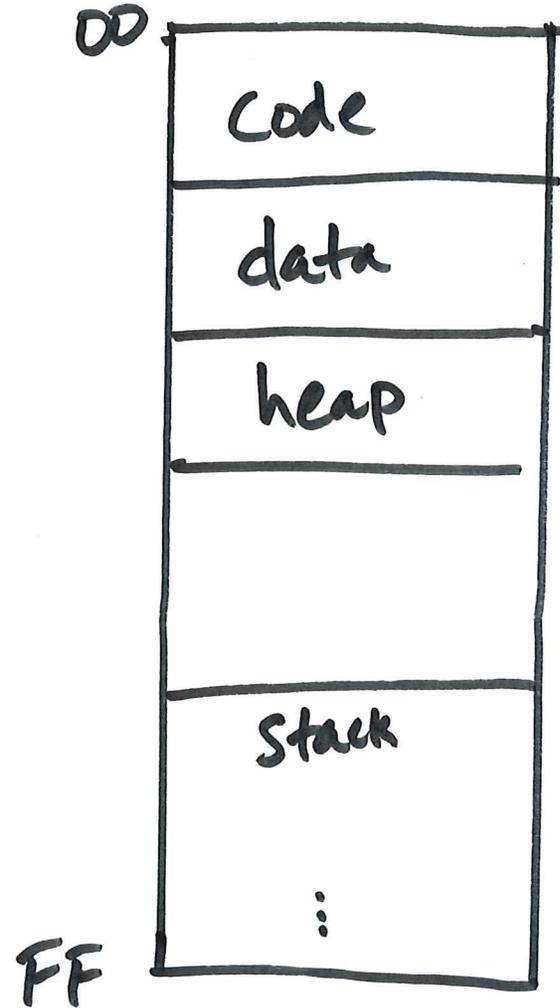
PROG



main:

⋮
addq %rdi, %rax
⋮
callq foo
⋮
pushq %rax
⋮
system call ...
⋮

Virtual mem



3. extra