

10/3/22

DME

APP

user mode

kernel code

kernel mode

Types of Transfers

① Exception

→ Synchronous

→ hw switch mode (kernel)

save current state

exception handler

② interrupts

→ I/O completion, mouse or keyboard, network packet

→ asynchronous

→ timer interrupt

② System Calls

→ read, write

→ voluntary transfer to kernel

(syscall, system)

→ resumes at PC+1

→ synchronous

Kernel to User

→ upon syscall completion

→ new process starts

→ return from interrupts (exceptions)

Similarities btwn the 3 mode transfers

- ① mode switch
- ② save current state
- ③ run some kernel handler

OSDev Wiki

How kernel supports this.

→ Interrupt Vector Table (stores info for all 3 ^{handler} mode transfers)

- what to invoke
- triggers the handler + saves the current state.

→ Isolate with kernel stack.

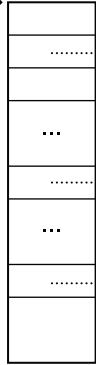
Processor Register

Interrupt Vector Table



Stores addr of the table Δ

array of functions



```
..... handleTimerInterrupt() {  
    ...  
}
```

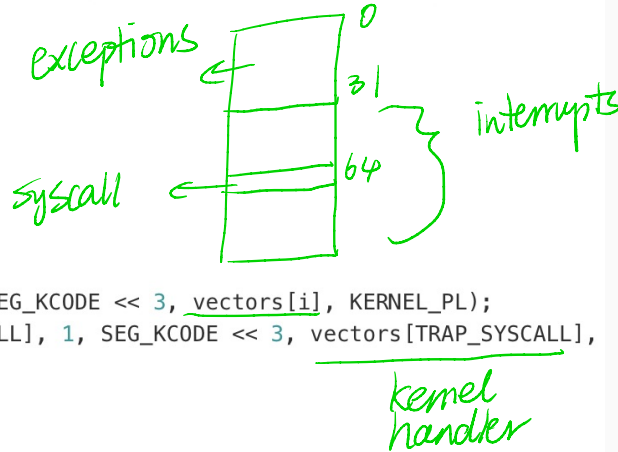
```
..... handleDivideByZero() {  
    ...  
}
```

```
..... handleSystemCall() {  
    ...  
}
```

```

11 // Interrupt descriptor table (shared by all CPUs).
12 struct gate_desc idt[256];
13 extern void *vectors[]; // in vectors.S: array of 256 entry pointers
14 struct spinlock tickslock;
15 uint ticks;
16
17 int num_page_faults = 0;
18
19 void tvinit(void) {
20     int i;
21
22     for (i = 0; i < 256; i++)
23         set_gate_desc(&idt[i], 0, SEG_KCODE << 3, vectors[i], KERNEL_PL);
24     set_gate_desc(&idt[TRAP_SYSCALL], 1, SEG_KCODE << 3, vectors[TRAP_SYSCALL],
25                 USER_PL);
26
27     initlock(&tickslock, "time");
28 }
29
30 void idtinit(void) { lidt((void *)idt, sizeof(idt)); }

```



trap.c

```

1 .globl alltraps
2 .globl vector0
3 vector0:
4     push $0 ← trap #
5     push $0 ← trap #
6     jmp alltraps
7 .globl vector1
8 vector1:
9     push $0
10    push $1
11    jmp alltraps
12 .globl vector2
13 vector2:
14    push $0
15    push $2
16    jmp alltraps
17 .globl vector3
18 vector3:
19    push $0
20    push $3
21    jmp alltraps
22 .globl vector4
23 vector4:
24    push $0
25    push $4
26    jmp alltraps
27 .globl vector5
28 vector5:
29    push $0
30    push $5

```

```

1  .globl alltraps
2  alltraps:
3      push %r15
4      push %r14
5      push %r13
6      push %r12
7      push %r11
8      push %r10
9      push %r9
10     push %r8
11     push %rdi
12     push %rsi
13     push %rbp
14     push %rdx
15     push %rcx
16     push %rbx
17     push %rax
18
19     mov %rsp, %rdi
20     call trap
21
22     .globl trapret
23     trapret:
24     pop %rax
25     pop %rbx
26     pop %rcx
27     pop %rdx
28     pop %rbp
29     pop %rsi
30     pop %rdi
31     pop %r8
32     pop %r9
33     pop %r10
34     pop %r11
35     pop %r12
36     pop %r13
37     pop %r14
38     pop %r15
39     add $16, %rsp
40     iretq
41

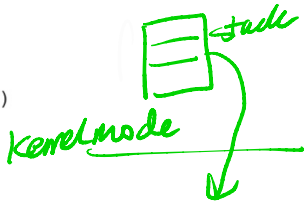
```

save state

```

32 void trap(struct trap_frame *tf) {
33     uint64_t addr;
34
35     if (tf->trapno == TRAP_SYSCALL) {
36         if (myproc()->killed)
37             exit();
38         myproc()->tf = tf;
39         syscall();
40         if (myproc()->killed)
41             exit();
42         return;
43     }
44
45     switch (tf->trapno) {
46     case TRAP_IRQ0 + IRQ_TIMER:
47         if (cpunum() == 0) {
48             acquire(&tickslock);
49             ticks++;
50             wakeup(&ticks);
51             release(&tickslock);
52         }
53         lapiceoi();
54         break;
55     case TRAP_IRQ0 + IRQ_IDE:
56         ideintr();
57         lapiceoi();
58         break;
59     case TRAP_IRQ0 + IRQ_IDE + 1:
60         // Bochs generates spurious IDE1 interrupts.
61         break;
62     case TRAP_IRQ0 + IRQ_KBD:
63         kbdtintr();
64         lapiceoi();
65         break;
66     case TRAP_IRQ0 + IRQ_COM1:
67         uartintr();
68         lapiceoi();
69         break;
70     case TRAP_IRQ0 + 7:
71         TRAP_IRQ0 + IRQ_SPURIOUS

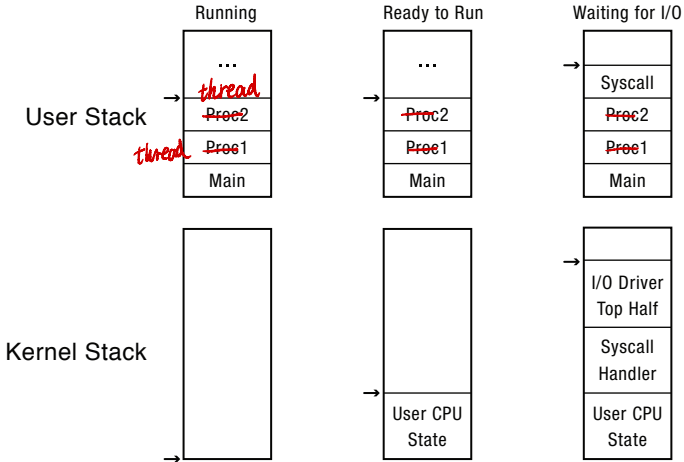
```



```

41 struct trap_frame {
42     uint64_t rax; // rax
43     uint64_t rbx;
44     uint64_t rcx;
45     uint64_t rdx;
46     uint64_t rbp;
47     uint64_t rsi;
48     uint64_t rdi;
49     uint64_t r8;
50     uint64_t r9;
51     uint64_t r10;
52     uint64_t r11;
53     uint64_t r12;
54     uint64_t r13;
55     uint64_t r14;
56     uint64_t r15;
57     uint64_t trapno;
58     /* error code, pushed by hardware or 0 */
59     uint64_t err;
60     uint64_t rip;
61     uint64_t cs;
62     uint64_t rflags;
63     /* ss:rsp is always pushed in long mode */
64     uint64_t rsp;
65     uint64_t ss;
66 } __packed;
67

```



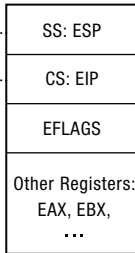
User-level Process

```
foo () {  
  while(...) {  
    x = x+1;  
    y = y-2;  
  }  
}
```

User Stack



Registers



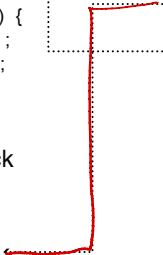
Kernel

```
handler() {  
  pushad  
  ...  
}
```

Interrupt Stack



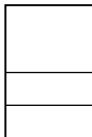
*using the
user stack*



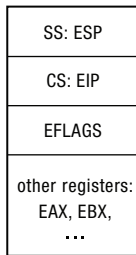
User-level Process

```
foo () {  
  while(...) {  
    x = x+1;  
    y = y-2;  
  }  
}
```

User Stack



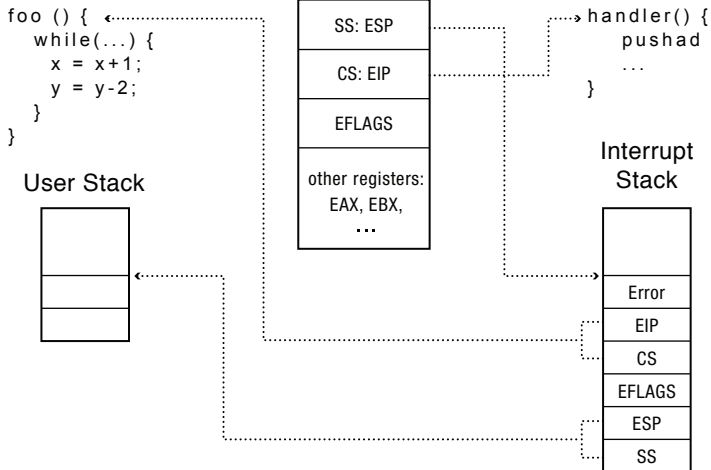
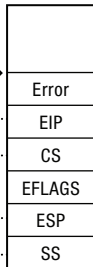
Registers



Kernel

```
handler() {  
  pushad  
  ...  
}
```

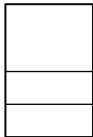
Interrupt Stack



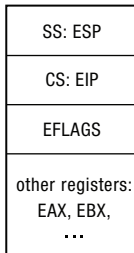
User-level Process

```
foo () {  
  while(...) {  
    x = x+1;  
    y = y-2;  
  }  
}
```

Stack



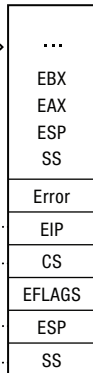
Registers



Kernel

```
handler() {  
  pushad  
  ...  
}
```

Interrupt Stack



All
Registers

