Bartłomiej Świercz

System calls

Operating system

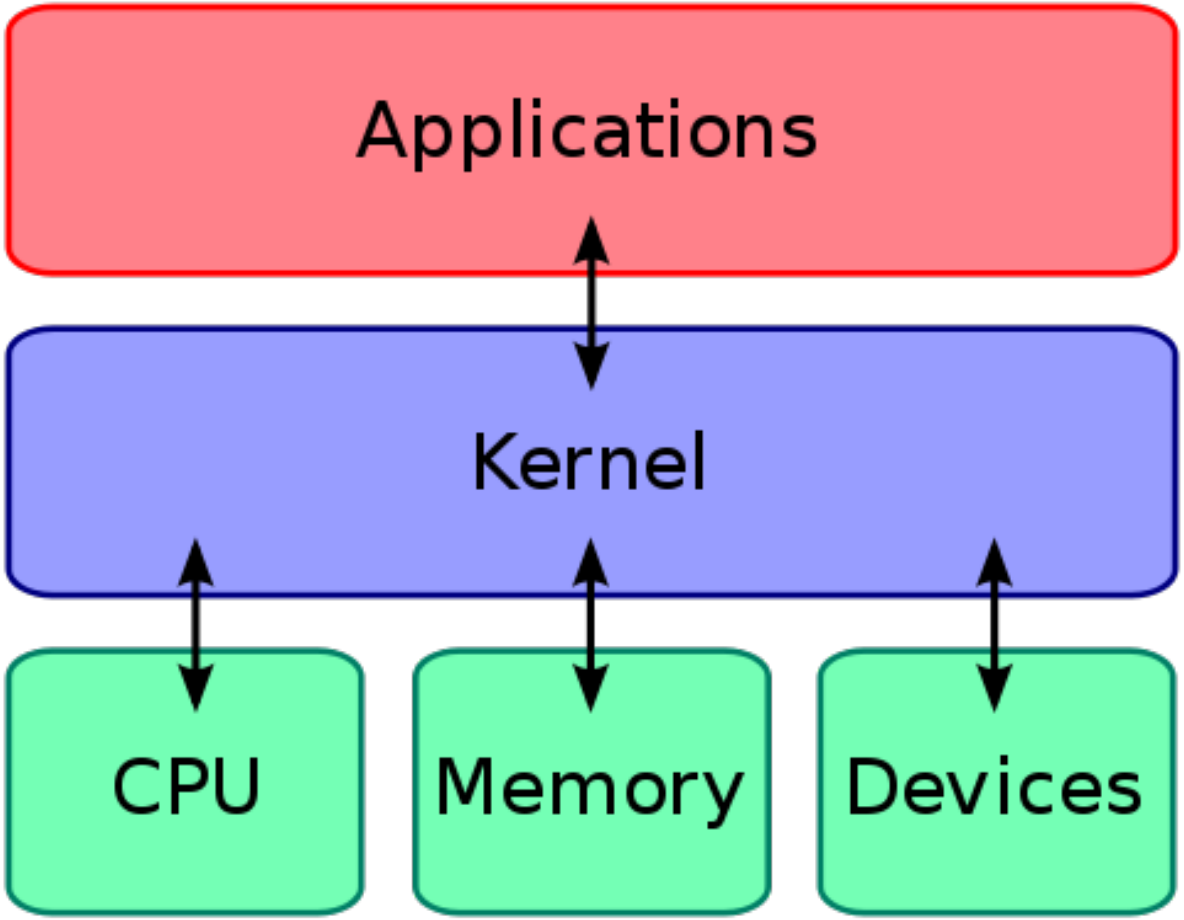

Go to kernel ...

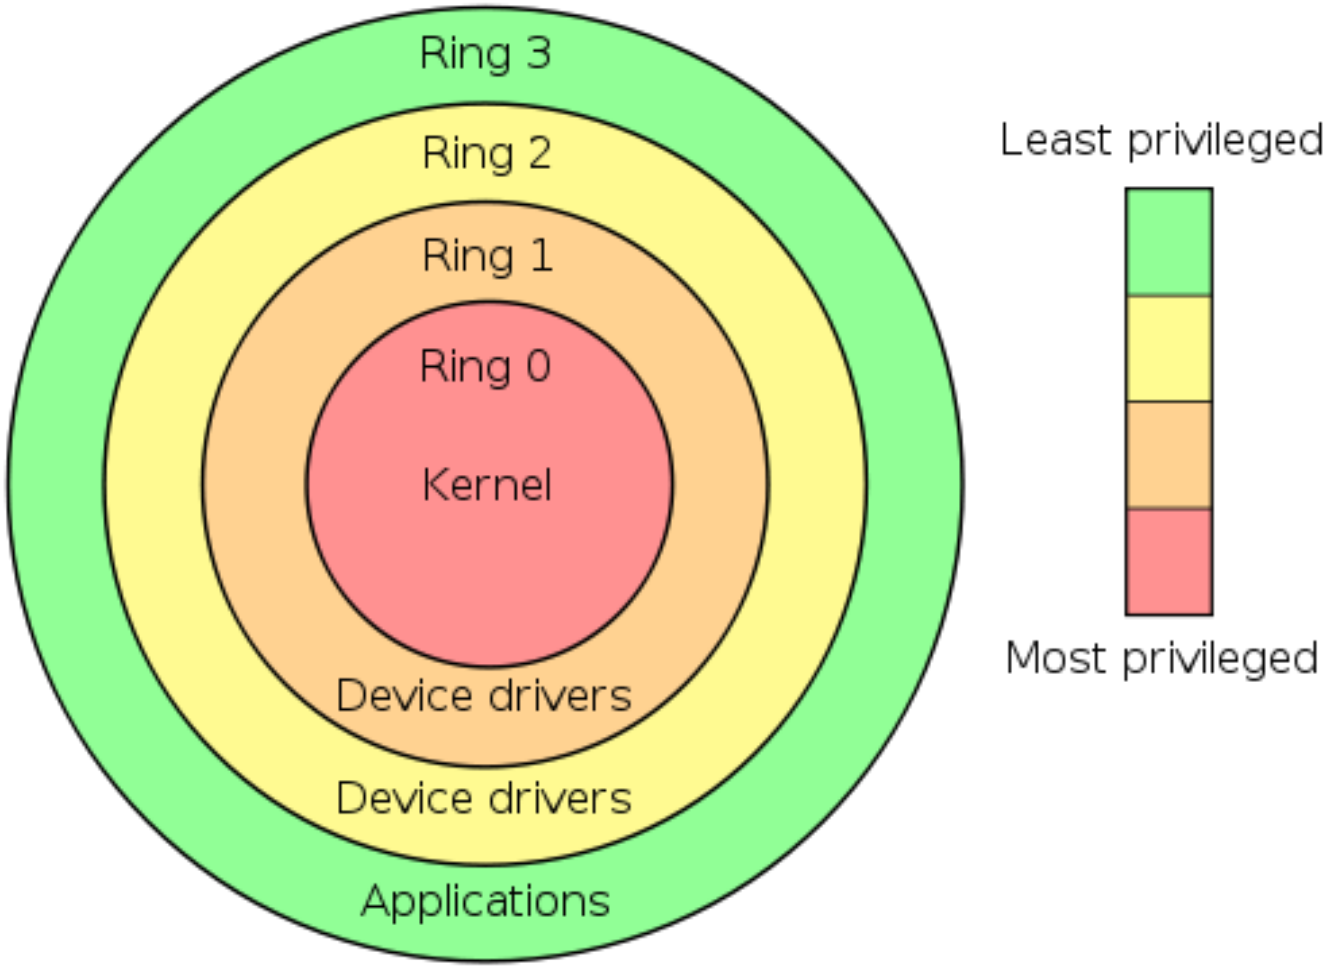
Ring 3
Ring 2
Ring 1
Ring 0
Kernel
Device drivers
Device drivers
Applications
Least privileged
Most privileged

Interrupts

- Hardware interrupts
- Software interrupts
  - System calls
- Slim top half
- Fat bottom half
  - Softirq
  - Tasklets
  - Task queue

System call

- Mechanism to request operation with increased privileges
- System call is provided by OS
  - Kernel code snippet run by user process
- The mechanism is wrapped by library
  - Transmission is hidden
  - Context switching

Syscall path

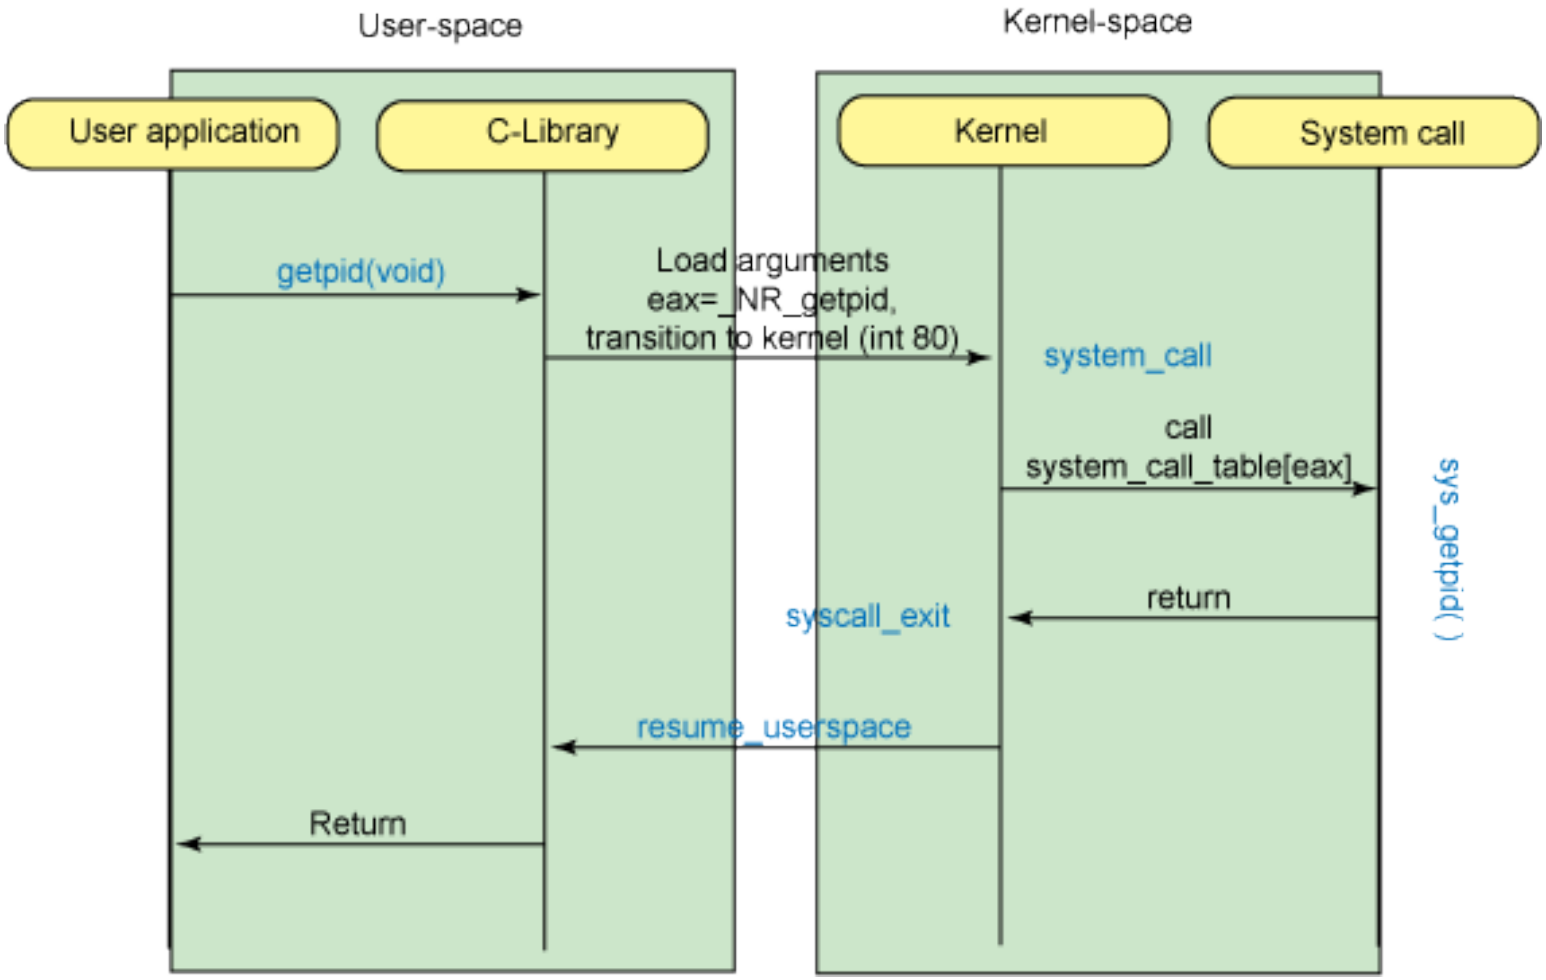

Syscall table

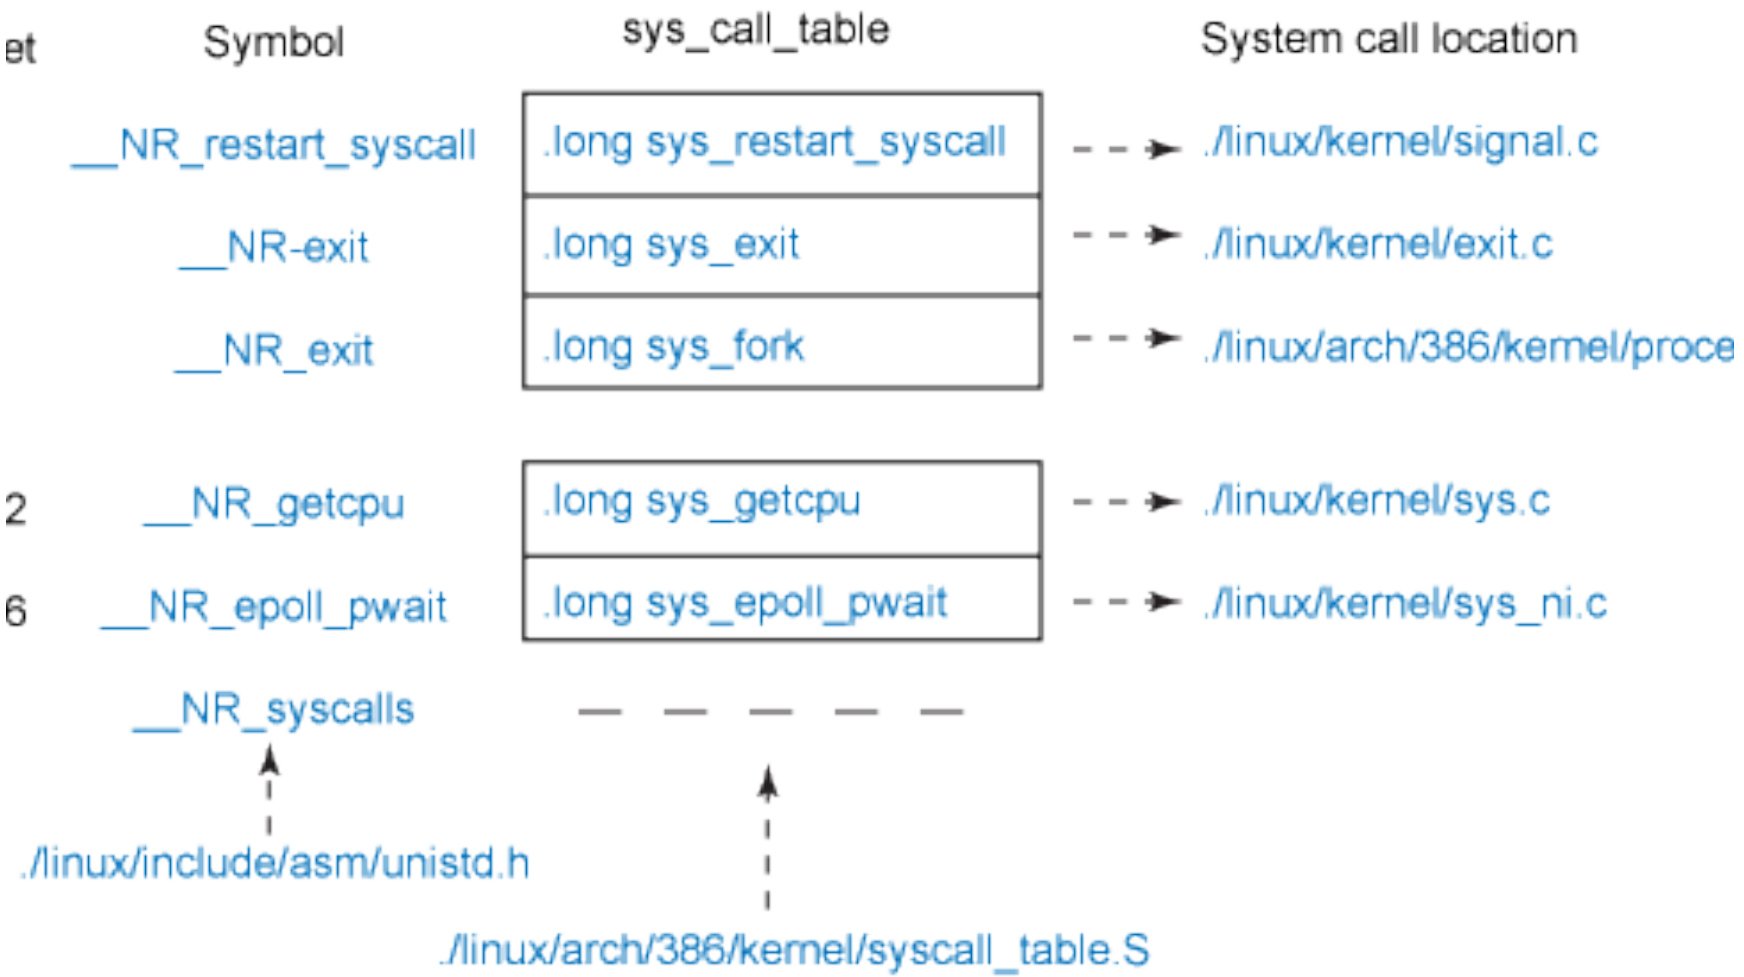

How to enter to kernel space?

- int 0x80
  - Syscall number read from %eax
- SYSCALL/SYSENTER
- SYSRET/SYSEXIT

- http://lkml.org/lkml/2002/12/9/13