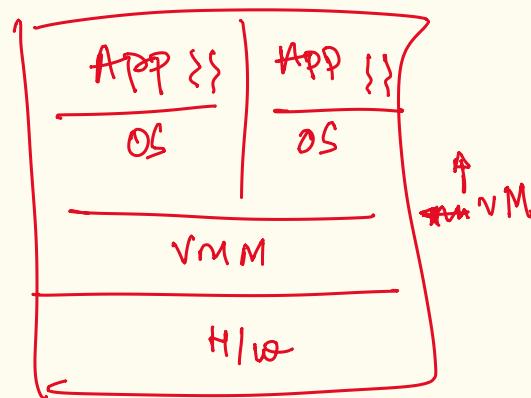


# Lecture #3

# CS 695



new abstraction — VM  
virtual m/c

that provides a system/m/c view.

Why VMs?

additional  
~ isolation sandboxes  
on a single m/c

~ server consolidation

(sharing /multiplexing PM with several VMs).

~ portability — tricky! (not easy  
from cost).

~ elastic resources

~ multi-OS systems  
in single m/cs

— testing & debugging

new  
— mechanisms: snapshot, record-replay,  
ballooning ...

Assignment #1

due: 29th Jan  
11.59 pm

3 parts

- LKMs
- ioctl
- procs

Assignment #2

— our own VM!

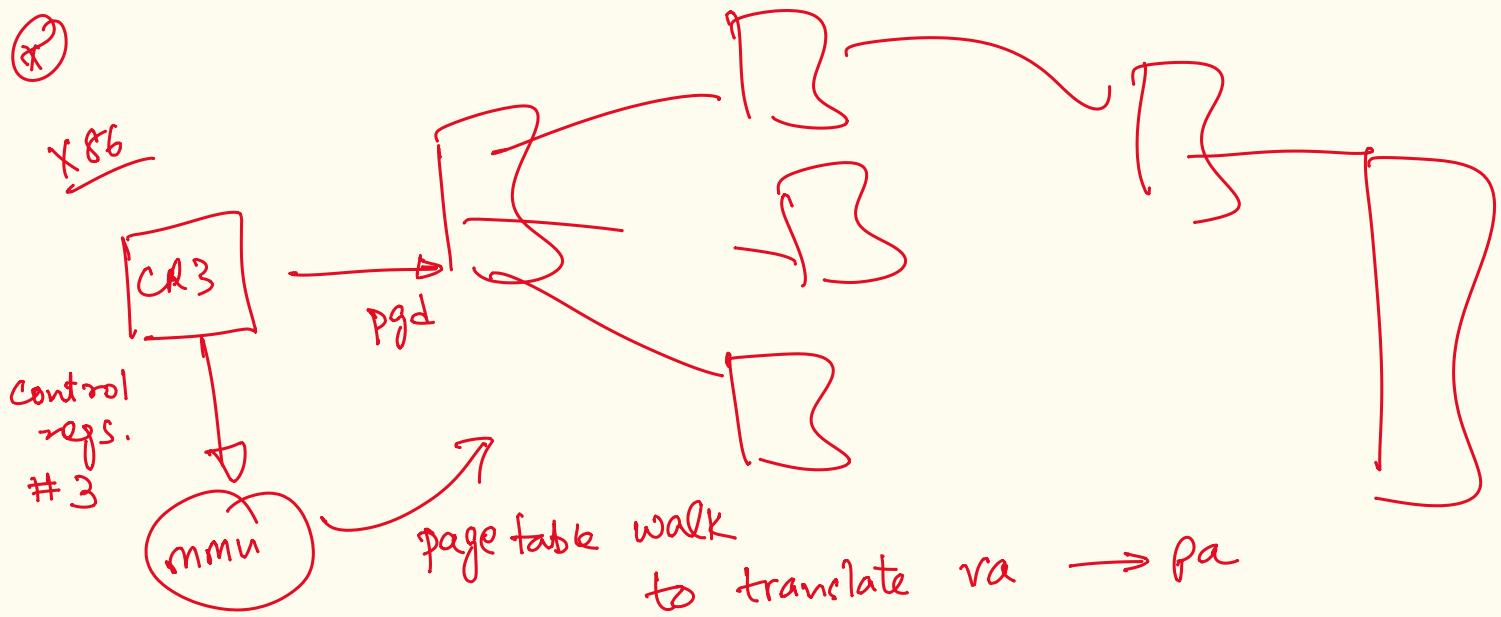
~ /dev/kvm

+ ioctl

+ io + mem

virtualization

mechanism  
+ policy



ld cr3 — load value to CR3 reg.  
instruction

⑤ OS building blocks  $\Rightarrow$  ensure that the OS is the sole owner / arbitrator of all resources

- (i) LDE - limited direct execution.
- (ii) interrupt/interrupt-driven execution
- (iii) system call interface

⑥ execute all non-OS instructions directly on the CPU, inject when critical / sensitive updates are to be made.

- privilege levels of exception.

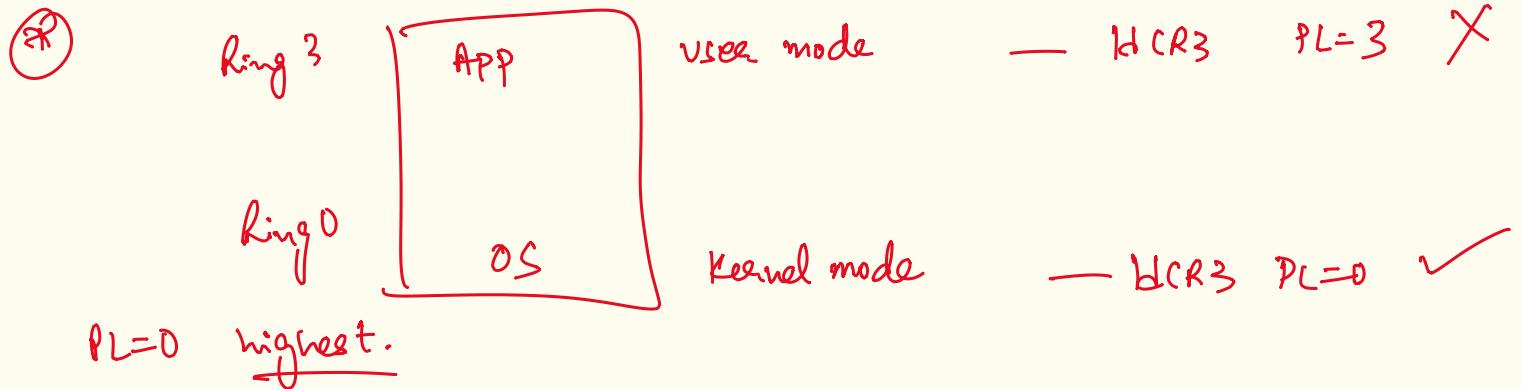
CPU executes in a

P<sub>L</sub>  
 $\times 86$  → 3  
2  
1  
0

every instruction of the ISA has min P<sub>L</sub>

for correct execution semantics.

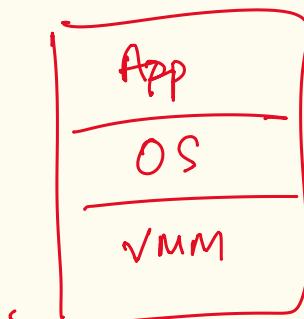
ldcr3



## # Challenges of hypervisor design!

① Who is the owner?  
of all resources

- Guest OS believes that it is the owner
- multiple guest OSes.
- CPL ~ bit mask of the CS / CR0



② System call handling

- SW modes from user to kernel
  - save user state
  - jump to interrupt handler
- %eax + 2      open's sys call no.  
int 0x80      interrupt 0x80
- APP      Lib (User态)      open      sys call  
↓      ↓      ↓      ↓  
Syscall handler      int. vector  
int. handler      Syscall number
- ↓      ↓      ↓      ↓  
IDTR

(iii) resource virtualization  
+ CPU, memory, IO  
- eg. CR3

(iv) device virtualization  
~ eg: NIC  
network interface card

Design: trap & emulate.