

## Discussion Plan:

### 0) Background on Virtualization

- simulation of computer

- uses

- run multiple OS

- kernel development

- accuracy:

- reproduce quirks malicious

- maintain isolation w/ a screen

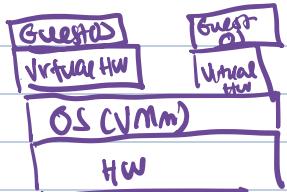
- cannot tell you are running a guest

- History:

- 1960s, 70s: IBM

- 90s: VMware repopularized VMs  
for x86 hardware

- 2000s: AMD / Intel built  
specialized CPUs for  
virtualization



- General goals of virtualization:

- 1) Fidelity = identical execution

- 2) Performance

- 3) Safety

simulation?

execute guest instructions

- lacks perf

on real CPU?

perf?

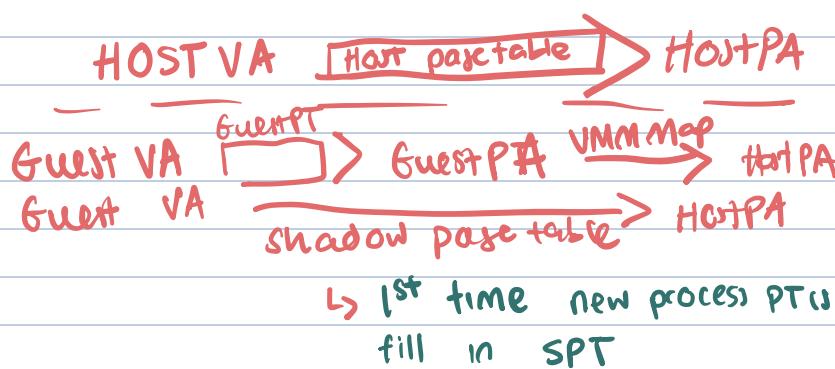
run at  
CPL3 (protection  
level)

Trap-and-emulate?

↳ only traps and  
emulates PRIVILEGED  
instructions

- How does paging work?

- but some instructions behave  
differently - no fidelity



Issue: VMM  
can't modify PT  
in place

- Problem: x86 not virtualizable w/ trap-and-emulate

- main problem: CPL3 vs. CPL0 → you can tell

you're in CPL3

- sometimes privileged bits masked out

• Two solutions:

BT (Binary Translation) - rewrite offending instructions

HW (Hardware Virtualization) - CPU maintains  
privileged state,  
executes privileged  
guest instructions

- CPU maintains guest copy of privileged  
state in VMCS

- hardware saves and restores

privileged reg. state to / from  
VMCS as it switches from NMX root  
to VMX non-root

\* MMU: EPT



- EPT better when PT contexts  
change frequently

- shadow PT better when PT are static

• Now: what is VT-x?

↳ hardware extensions that make x86 virtualizable  
w/ fidelity, perf, and safety

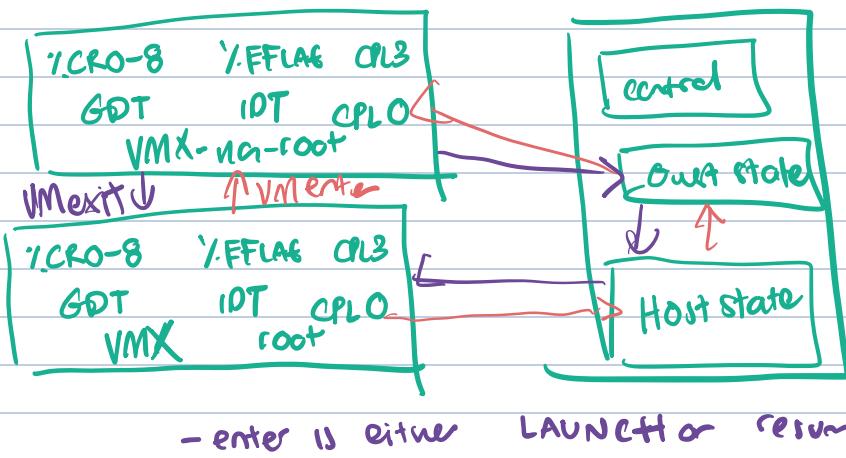
• goals: DIRECT execution of privileged  
instructions

- two CPU modes:

- VMX root mode

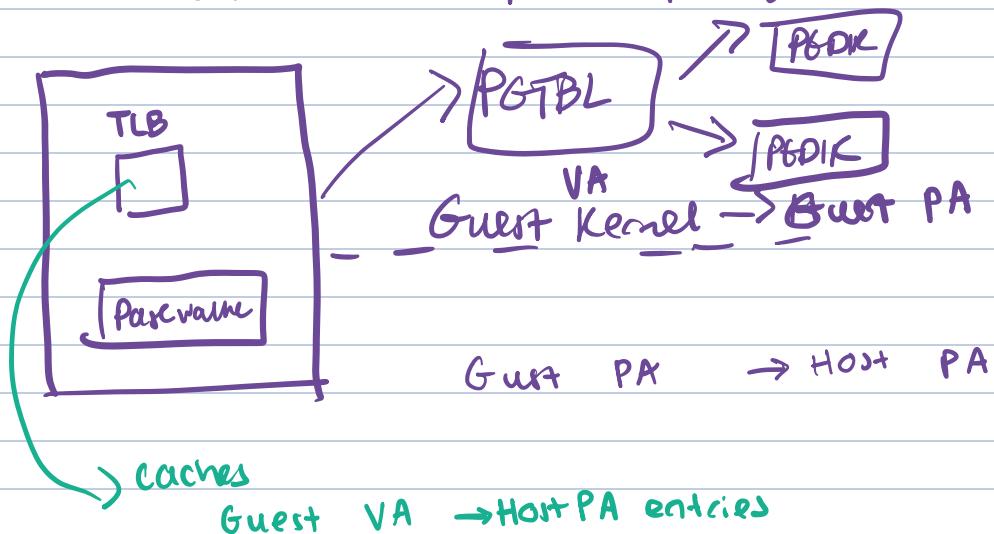
- VMX non-root mode (Guest)

- each has its own protection ring



EPT? enables DIRECT execution of guest PT interactions

- reads and writes to PT in memory
- maintain two layers of paging translation



### Dune:

- use EPT, VT-x to support Linux Processes
- dune mode=isolated w/ VT-x non-root mode

↳ not w/ CPL3 + PT protections  
(trap + emulate)

- configure EPT so dune process can only access PAs it can normally have access to

\* process can manipulate its own PT! (using %CR3)

- fast exceptions due to shadow IDT
- can run sandboxed code  
(where most opt-in comes)

\* GC app that is sped up by Dune:

- GC: find live data, reclaim non live data

- concurrent GC (Boehm):
  - mutator runs in parallel w/ tracer
  - at some points, mutator might modify pointers in already traced objects
  - we need to figure out which ones were modified

- Dunes:

- clear dirty bits at beginning of GC
- scan for ret PTE
- dirty bits to find written pdes

Sources: [https://abelay.github.io/6828seminar/notes/65810\\_lec\\_dune.pdf](https://abelay.github.io/6828seminar/notes/65810_lec_dune.pdf)