


CS 3214 lecture # 27

"VMS & containers"

SPOT ←



Cloud computing characteristics

- elasticity ← granularity is shrinking
- multi-tenant ← statistical multiplexing
- need for isolation - "private view" of system

How do we get private view?

Process: address space
CPU

what isn't private?

- filesystem
- sockets: port numbers
- pids are shared

rm /lib
/etc/config
ps

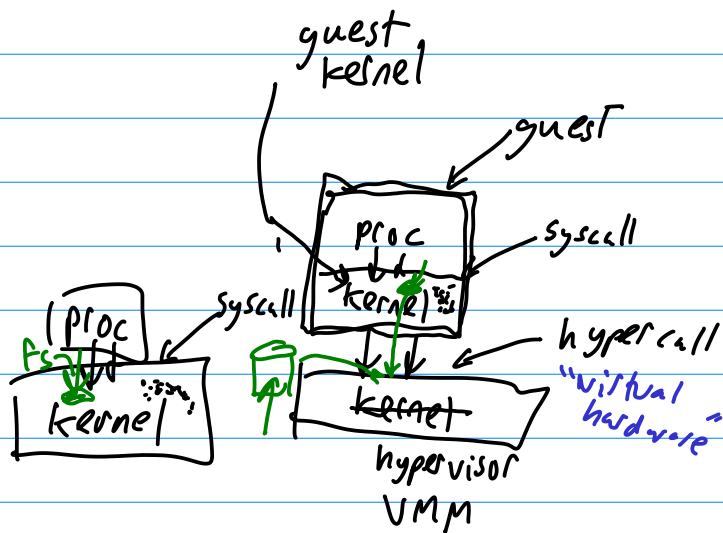
process is not a great abstraction for cloud

Alternatives

- bare metal
- virtual machines <
- containers <

Virtual machine

- virtual disk has everything app needs
- guest kernel (nothing shared)
- hyp/vmm multiplexing



1970s: Goldberg 1972, 1974

- improve test OS
- developing H/W diagnostics
- running different OSes / versions

Popok / Goldberg requirements

- equivalence / fidelity
- resource control
- efficiency (most inst. need to execute directly on H/W)

How to run VM?

- basic idea: "deprivilege" (run in user mode)
- trap and emulate

g: will all ^{privileged} instructions trap

IBM/360 yes


x86 (prior to VT-x) << no

Early 2000s

- modify guest: Xen para-virtualization (2002)
- dynamic binary translation: VMware (1999)
- fix the hardware: VT-x (2005/2006) AMD-V

lots of classic mechanisms to rethink

Memory mgmt
fixed RAM, CPU
device I/O
:



"Heavyweight" - image



- 2 kernels are running
- I/O

Containers: the OS already does this!

kernel mechanisms

- chroot "jail"

- use some other dir as / (for process + children)

Namespaces

- UTS - hostname

- PID

- mount

- user

- IPC

- network

:

cgroups

- limit resource consumption

- CPU

- mem

clone()

unshare()

container = set of processes sharing - chroot

- ns

- cgroups

layered filesystems

kubernetes CNCF

microVMs