More on Virtual Machines

Kai Shen

Dept. of Computer Science, University of Rochester

Live Migration

- Migrating a VM from one physical machine to another
  - minimal freeze time
- Migration approaches
  - stop the VM; move the VM state to the new machine; start it
  - stop the VM on the old machine; set up the skeleton on the new machine (all, or most, page table entries invalid) and then start it
  - keep the VM running on the old machine; move state over on the background; then repeatedly move dirty state until it is small; stop the VM on the old machine; move the final dirty state; start it on the new machine [Clark et al. NSDI'05]

Virtual Machine I/O

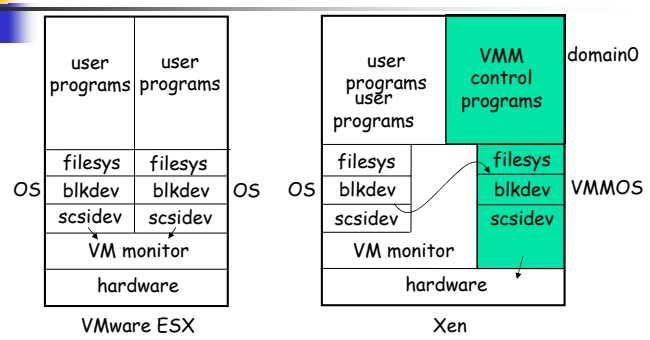

- two-level I/O scheduling
- anomaly with anticipatory I/O scheduling
  - process context tracking [Jones et al. USENIX'06]

VM-based Intrusion Diagnosis

- In a conventional system
  - a superuser has full trust of the machine
  - when an intruder assumes the superuser identity, he/she can erase all traces of the intrusion
- In a VM platform
  - even if an intruder assumes the superuser identity of a VM, he/she cannot erase information recorded by VM monitor
  - what information is recordable at the VM monitor?
- VM-based backtracking intrusions [King&Chen SOSP'03]

VM-based Trusted Computing

- Modern OSes are complex and can be compromised
  - trusted computing systems using private/thin OSes with customized security/trust policies
- How can virtual machines help?
- Terra [Garfinkel et al. SOSP'03]
  - support multiple systems of different security/trust policies on one piece of hardware
  - customize hardware platforms with different security models

VM-based Trusted Computing (cont.)

- What if a single system has parts with different security constraints?
  - secure browsing without leaking user behavior information
  - secure SSH server without leaking configuration information
- Proxos [Ta-Min et al. OSDI'06]
  - partition the system into two parts: secure part runs in a private with private OS; the rest runs in commodity OS
  - use system call routing to partition flexibly

VM-based Trusted Computing (cont.)

- Can we run on commodity OS without trusting it?
- Overshadow [Chen et al. ASPLOS'08]
  - protecting application data from OS peeking (two sets of shadow page tables: one for application, the other for OS)
- but OS does need to access application data
  - swapping
    - encrypt/decrypt dynamically
  - execute a system call with a pointer parameter referencing a memory address