The Soul of a Virtual Machine: VM Save, Restore, & Rewind

Patrick F. Wilbur

Department of Mathematics & Computer Science,
Clarkson University

November 14, 2008

Overview

- Virtualization basics
- Virtual machine state components
- Approaches to saving states
- State rewind uses and problems

Types of Virtualization

Two main categories:

- Process virtual machine
  - Runs a single program
  - Program can only access resources and abstractions provided by VM
- System virtual machine
  - Provides an entire platform for running an OS
  - May provide shared direct access to some hardware, or emulate hardware
  - OS can only access resources provided by VM

Definitions

- Virtual Machine Monitor, Hypervisor
- Guest
  - Guest program
  - Guest operating system

State Components

- Process virtual machine state:
  - Individual program stack and memory
  - Open files/devices and locks
  - (Filesystem state)
- System virtual machine state:
  - Entire memory allocation
  - Virtual processor state
  - I/O and device states
  - (Block device state)

Pause and Save

- Trivial in virtualized environment (like hibernation)
- After pausing operation of guest:
  - Save virtual processor states
  - Save contents of guest memory
- A paused-and-saved guest is easily portable

Live Save (Checkpointing)

Building on the backs of giants:

- Live Migration
- LVM Logical Volume Snapshotting

State Rewind

- Useful in several contexts:
  - Rapid recovery
  - Debugging (analyzing problem progression)
  - Testing (rewind and traverse other branches of possible-states tree)

State Rewind

- Problems:
  - Difficult to snapshot file system or disk states at precisely the same time as memory
  - Difficult to optimize frequency of snapshots
  - Some tasks are time or timing dependent
  - Difficult to determine how many snapshots are needed for recovery

Questions?

References

- Matthews, Jeanna, et al. *Running Xen: A Hands-On Guide to the Art of Virtualization* . Prentice Hall, 2008.
- <http://www.cl.cam.ac.uk/research/srg/netos/papers/2005-migration-nsdi-pre.pdf>
- <http://www.usenix.org/event/nsdi08/tech/full_papers/cully/cully_html/index.html>
- Lewis, AJ. "LVM HOWTO." November 27, 2006. November 13, 2008. <http://tldp.org/HOWTO/LVM-HOWTO/>
- <http://en.wikipedia.org/wiki/Logical_Volume_Manager_(Linux)>
- <http://en.wikipedia.org/wiki/Virtual_machine>
- <http://en.wikipedia.org/wiki/Virtualization>

License Information

Last modified: November 15, 2008 9:52 PM EST.
Presented: Patrick F. Wilbur, Clarkson University, Potsdam, NY USA, November 14, 2008.

==

LICENSE:

Patrick F. Wilbur
Department of Mathematics and Computer Science
Clarkson University
Potsdam, NY USA

http://pdub.net