Olivier Gruber, Ph.D.

Full-time Professor Université Joseph Fourier Laboratoire d'Informatique de Grenoble

Senior Resarcher @ INRIA

Olivier.Gruber@inria.fr

Acknowledgments

• Reference Book

Virtual Machines Versatile Platforms for systems and processes

James E. Smith, Ravi Nair

Morgan Kaufmann

• Research Articles

- Cited on various slides

- Virtual Machines versus Real Machines
 - A virtual machine defines a machine (interface)
 - A virtual machine is a machine (implementation)



Olivier.Gruber@inria.fr

Operating System Basics

- Instruction Set Architecture (ISA)
 - Defines the instruction set
 - Defines other concepts such as memory, traps, interrupts, etc.
- Application Binary Interface (ABI)
 - Defines core concepts above the ISA
 - Example: Linux kernel system calls
 - Related to processes, threads, files, and devices





ISA

Operating System Basics

- Operating system design
 - Main goal is sharing hardware resources across processes
 - It provides each process with the illusion it runs alone on a real machine
- Operating systems are virtual machines
 - Subset of the processor instruction set (user-mode)
 - Concepts of the ABI



Dec Alpha Example

- Digital Alpha machine
 - Early provider of a 64bit RISC processor
 - Challenge: no existing software...
 - Support program binaries compiled to a different ISA / same ABI
 - Same ABI: ported the operating system
 - Different ISA: emulate one instruction set on a different instruction set



Hookway and Herdeg. *Digital FX!32: Combining Emulation and Binary Translation*. Digital Technical Journal, January 1997, pp 3-17 Zheng and Thompson. *PA-RISC to IA-64: Transparent Execution, No Recompilation*. IEEE Computer, March 2000, pp. 47-53

• Emulator Designs

- Interpretation:

- Interpretation of individual guest instructions (fetch, decode and emulate)
- Easy but slower

- Binary translation

- Binary translation of blocks of guest instructions to native instructions
- More complex but fast (close to native performance)
- Classical trade-off
 - Slow interpretation versus high overhead of binary translation



- System Vms (hypervisors)
 - Such as Vmware ESX or Xen
 - Goal: multiplex out-of-the-box operating systems
 - Often virtualize a similar hardware (but not always)



- System Vms (hypervisors)
 - Typical use in the Cloud
 - Provides an ubiquitous hardware
 - Provides remote management
 - Virtualize a similar hardware
 - Because performance is critical
 - Same instruction set
 - Similar devices, maybe less memory or less cores
 - Enables hardware sharing to reduce the costs
 - Energy-saving strategy
 - Use a few real machine as necessary
 - The VM has a long life
 - Until the underlying real machine is rebooted

- In-Process Virtualization (Process VMs)
 - Virtualizing a real machine within a process
 - Runs one out-of-the-box operating systems
 - Same hardware or not



- In-Process System Vms
 - Typical uses
 - Kernel development
 - Application availability
 - Help desks
 - Virtualize a similar hardware or not
 - Performance is often not as critical as it is for hypervisors
 - It is the availability of the platform that matters most
 - Similar devices, maybe less memory or less cores
 - The VM has a shorter life
 - Until the process is killed
 - More like an application than a real machine

- Virtualization at different layers...
 - Hypervisors, operating systems, and in-process system VMs



Olivier.Gruber@inria.fr

• Process-level Virtual Machines

- High-level language virtual machines..
- Examples:
 - Oracle Java or Microsoft C#
 - Eclipse Rich Client Platform (Java and Eclipse libraries)
 - Google Android (Java and Android's libraries)
 - Web applications (Flash or HTML5)



• Towards language independence

- Microsoft Common Language Infrastructure
 - Common Language Runtime (CLR) and Common Type System (CTS)
- The Java Virtual Machine is going in the same direction
 - Already a target for many languages (JavaScript, Scheme, Perl, Python, etc.)



• Virtualization at different layers...

- Hypervisors, operating systems, and in-process system VMs
- Adding high-level language VMs



Olivier.Gruber@inria.fr

Virtual Machine Taxonomy

