

## Virtual Machines

## Virtual Machines

- What is a virtual machine?
- Examples?
- Benefits?

## Virtualization

- Creation of an isomorphism that maps a virtual guest system to a real host:
  - Maps guest state  $S$  to host state  $V(S)$
  - For any sequence of operations on the guest that changes guest state  $S1$  to  $S2$ , there is a sequence of operations on the host that maps state  $V(S1)$  to  $V(S2)$

## Important Interfaces

- Application programmer interface:
  - High-level language library such as `clib`
- Application binary interface (ABI):
  - User instruction (User ISA)
  - System calls
- Hardware-software interface:
  - Instruction set architecture (ISA)

## What's a Machine?

- Machine is an entity that provides an interface
  - Program view:
    - Machine = Entity that provides the API (language)
  - Process view:
    - Machine = Entity that provides the ABI
  - Operating system view:
    - Machine = Entity that provides the ISA

## What's a Virtual Machine?

- Virtual machine is an entity that emulates a guest interface on top of a host machine
  - Language view:
    - Virtual machine = Entity that emulates an API (e.g., JAVA) on top of another
    - Virtualizing software = compiler/interpreter
  - Process view:
    - Machine = Entity that emulates an ABI on top of another
    - Virtualizing software = runtime
  - Operating system view:
    - Machine = Entity that emulates an ISA
    - Virtualizing software = virtual machine monitor (VMM)

## Purpose of a Virtual Machine

- Emulation
  - Create the illusion of having one type of machine on top of another
- Replication
  - Create the illusion of multiple independent smaller guest machines on top of one host machine
- Optimization
  - Optimize a generic guest interface for one type of host

## Types of Virtual Machines

- Emulate (ISA/ABI/API) for purposes of (Emulation/Replication/Optimization) on top of (the same/different) one.

## Types of Virtual Machines

- Emulate (ISA/ABI/API) for purposes of (Emulation/Replication/Optimization) on top of (the same/different) one.
  - Process/language virtual machines (emulate ABI/API)
  - System virtual machines (emulate ISA)

## Types of Virtual Machines

- Emulate (ISA/ABI/API) for purposes of (Emulation/Replication/Optimization) on top of (the same/different) one.
  - Process/language virtual machines (emulate ABI/API)
  - System virtual machines (emulate ISA)

## Example 1 - Multiprogramming

- Emulate what interface?
- For what purpose?
- On top of what?

## Example 1 - Multiprogramming

- Emulate the ABI (user ISA + OS system calls)
- For purposes of replication (multiple processes share a physical machine)
- On top of the same ISA/OS

## Example 2: Emulation

- Emulate one ABI on top of another
  - Emulate a Intel IA-32 running WindowsXP on top of PowerPC running MacOS (i.e., run a process compiled for IA-32/Windows on PowerPC/MacOS)
    - Interpreters: Pick one guest instruction at a time, update (simulated) host state using a set of host instructions
    - Binary translation: Do the translation in one step, not one line at a time. Run the translated binary

## Example 3: Binary Optimization

- Emulate one ABI on top of itself for purposes of optimization
  - Run the process binary, collect profiling data, then implement it more efficiently on top of the same machine/OS interface.

## Example 4: Language Virtual Machines

- Emulate one API on top of a set of different ABIs
  - Compile guest API to intermediate form (e.g., JAVA source to JAVA bytecode)
  - Interpret the bytecode on top of different host ABIs
- Examples:
  - JAVA
  - Microsoft Common Language Infrastructure (CLI), the foundation of .NET

## Types of Virtual Machines

- Emulate (ISA/ABI/API) for purposes of (Emulation/Replication/Optimization) on top of (the same/different) one.
  - Process/language virtual machines (emulate ABI/API)
  - System virtual machines (emulate ISA)

## Types of Virtual Machines

- Emulate (ISA/ABI/API) for purposes of (Emulation/Replication/Optimization) on top of (the same/different) one.
  - Process/language virtual machines (emulate ABI/API)
  - System virtual machines (emulate ISA)

## System Virtual Machines (Conventional – Same ISA)

- Implement VMM (ISA emulation) on bare hardware
  - Efficient
  - Must wipe out current operating system to install
  - Must support drivers for VMM
- Implement VMM on top of a host OS (Hosted VM)
  - Less efficient
  - Easy to install on top of host OS
  - Leverages host OS drivers

## Whole System Virtual Machines

- Emulate one ISA on top of another
  - Typically runs on top of host OS (e.g., install Windows compiled for IA-32 on top of MacOS running on PowerPC)
  - Note: this is different from a process virtual machine that emulates the Windows interface and user IA-32 instructions on top of MacOS running on PowerPC)

## Taxonomy

- Process VMs
  - Same ISA
    - Multiprogrammed systems
  - Different ISA
    - Dynamic translators (Java)
- System VMs
  - Same ISA
    - Classic VMs
    - Hosted VMs (VMWare)
  - Different ISA
    - Whole system VMs