The Future of Virtual Machines: A VMware Perspective

Ed Bugnion
Co-founder, VMware Inc.

JUGS
September 27, 2001

Outline

- Historical Perspective
- MultipleWorlds™ Technology
  - Technology and Products
- Technology
  - Hosted and Host-less architectures
  - Performance
- 4 Usage scenarios

The Problem (1960's)

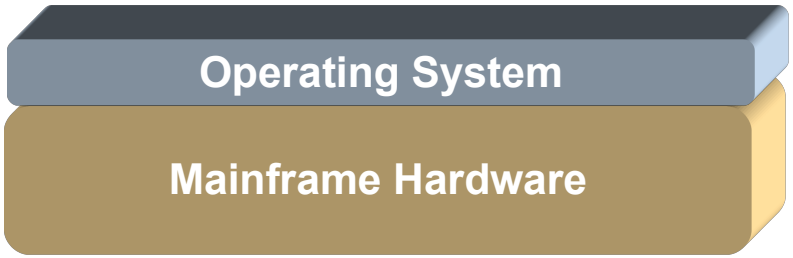

The Solution (1960's)

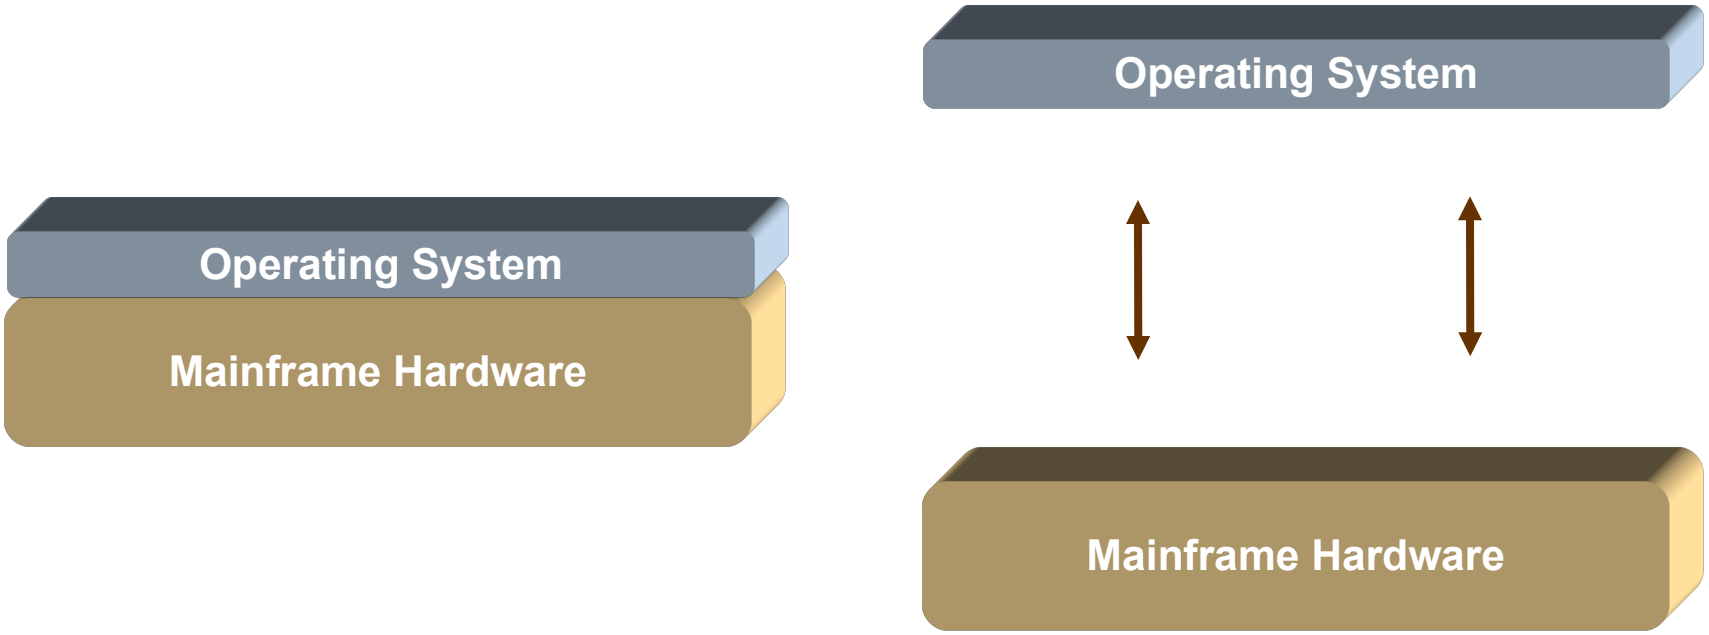

Virtual Machine Monitors

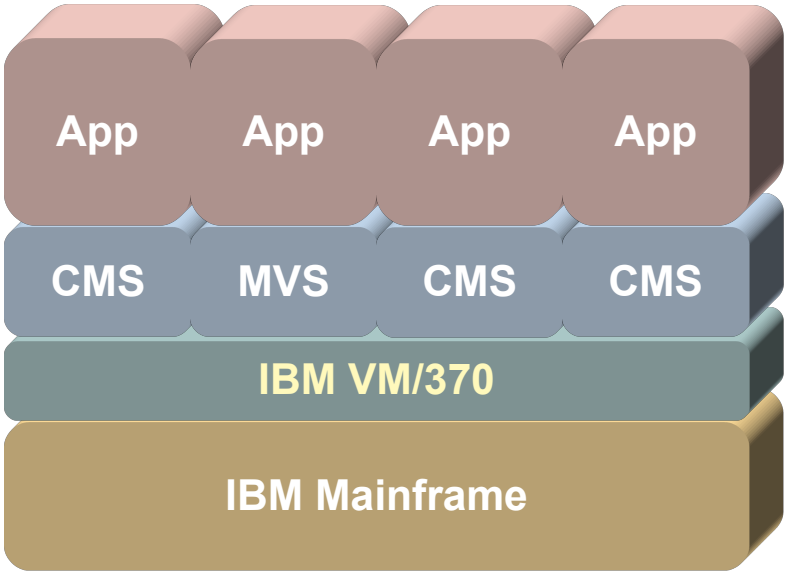

A thin software layer that sits between hardware and the operating system— virtualizing and managing all hardware resources

Old idea from the 1960s

- IBM VM/370 – A VMM for IBM mainframe
  - Multiple OS environments on expensive hardware
  - Desirable when few machine around
- Popular research idea in 1960s and 1970s
  - Entire conferences on virtual machine monitor
  - Hardware/VMM/OS designed together
- Interest died out in the 1980s and 1990s.
  - Hardware got cheap
  - Operating systems got more more powerful (e.g multi-user)

A return to Virtual Machines

- Disco: Stanford research project (1996-):
  - Run commodity OSes on scalable multiprocessors
  - Focus on high-end: NUMA, MIPS, IRIX
- Hardware has changed:
  - Cheap, diverse, graphical user interface
  - Designed without virtualization in mind
- System Software has changed:
  - Extremely complex
  - Advanced networking protocols
  - But even today :
    - *Not always multi-user*
    - *With limitations, incompatibilities, …*

The Problem Today

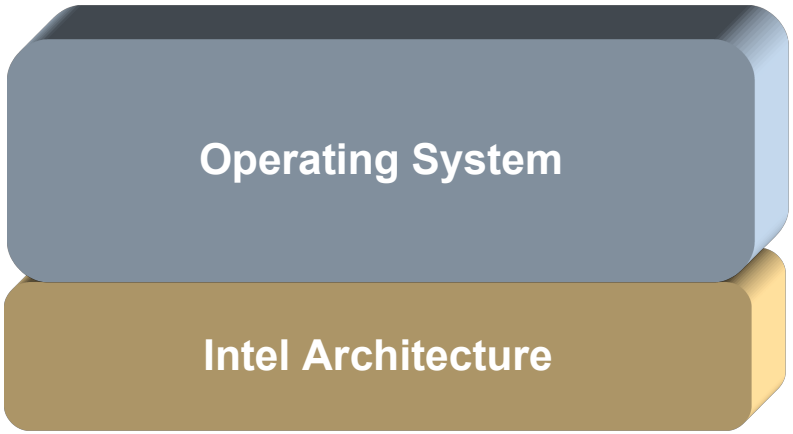

The VMware Solution

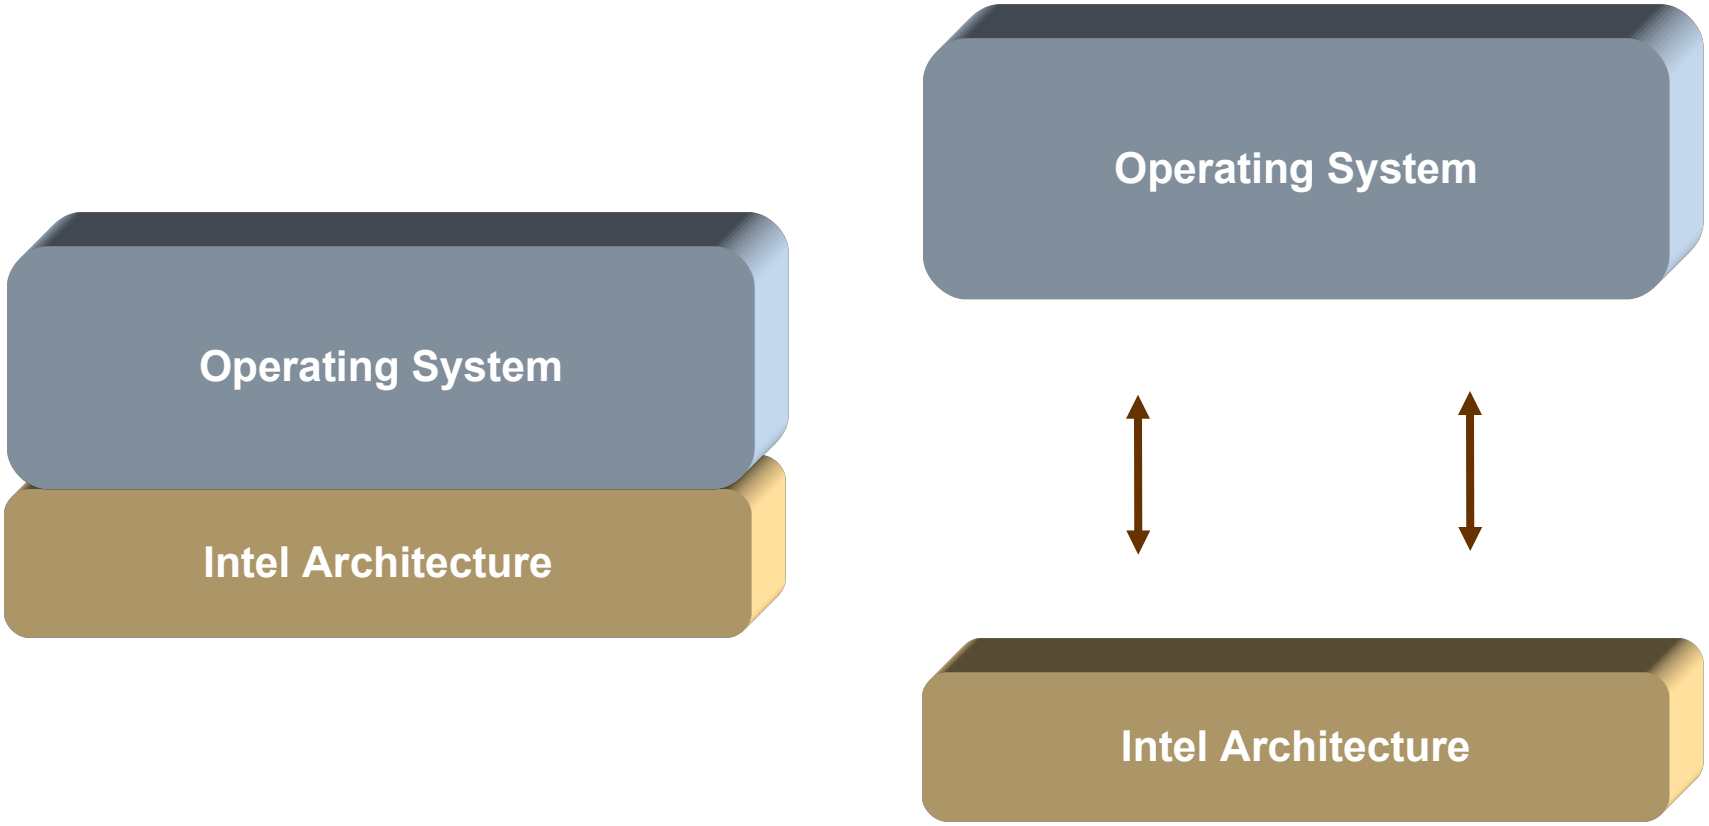

VMware™ MultipleWorlds™ Technology

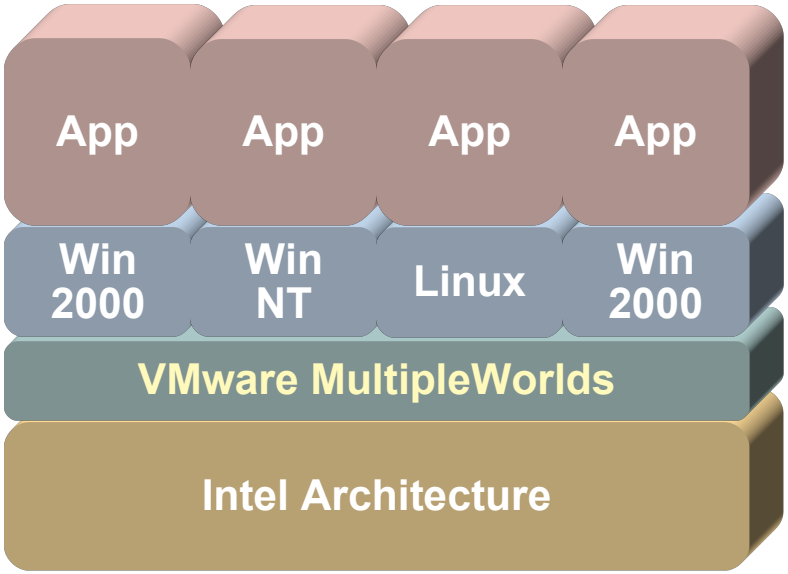

A thin software layer that sits between Intel hardware and the operating system— virtualizing and managing all hardware resources

MultipleWorlds Technology

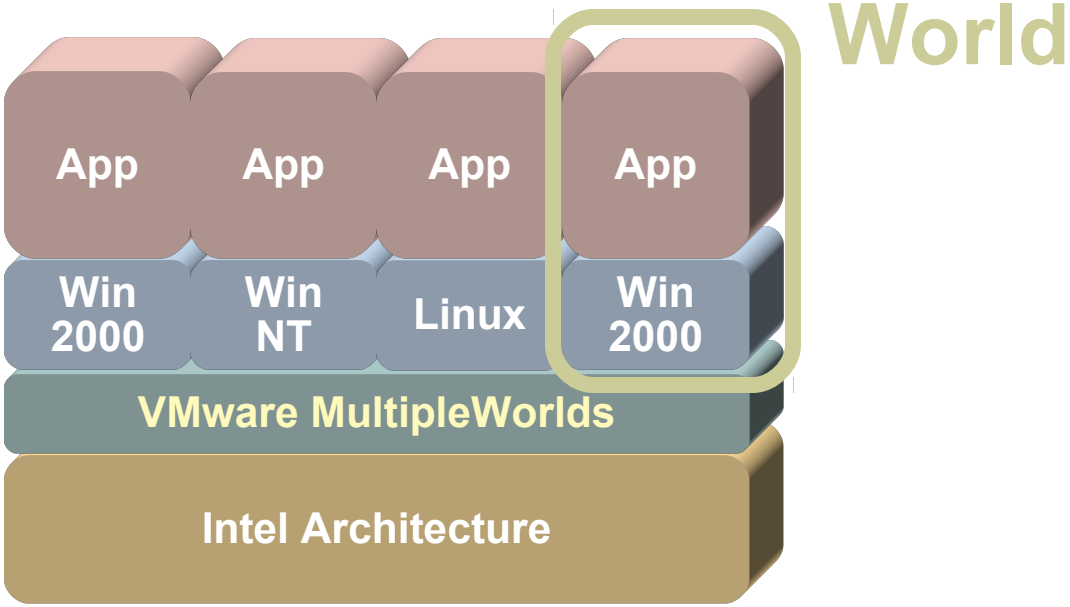

A world is an application execution environment with its own operating system

MultipleWorlds Technology

A world is an app ion environment with its own operating system

Challenges

- Virtualization of IA-32
- Hardware Diversity
- Acceptance

VMware Workstation– Screen shot

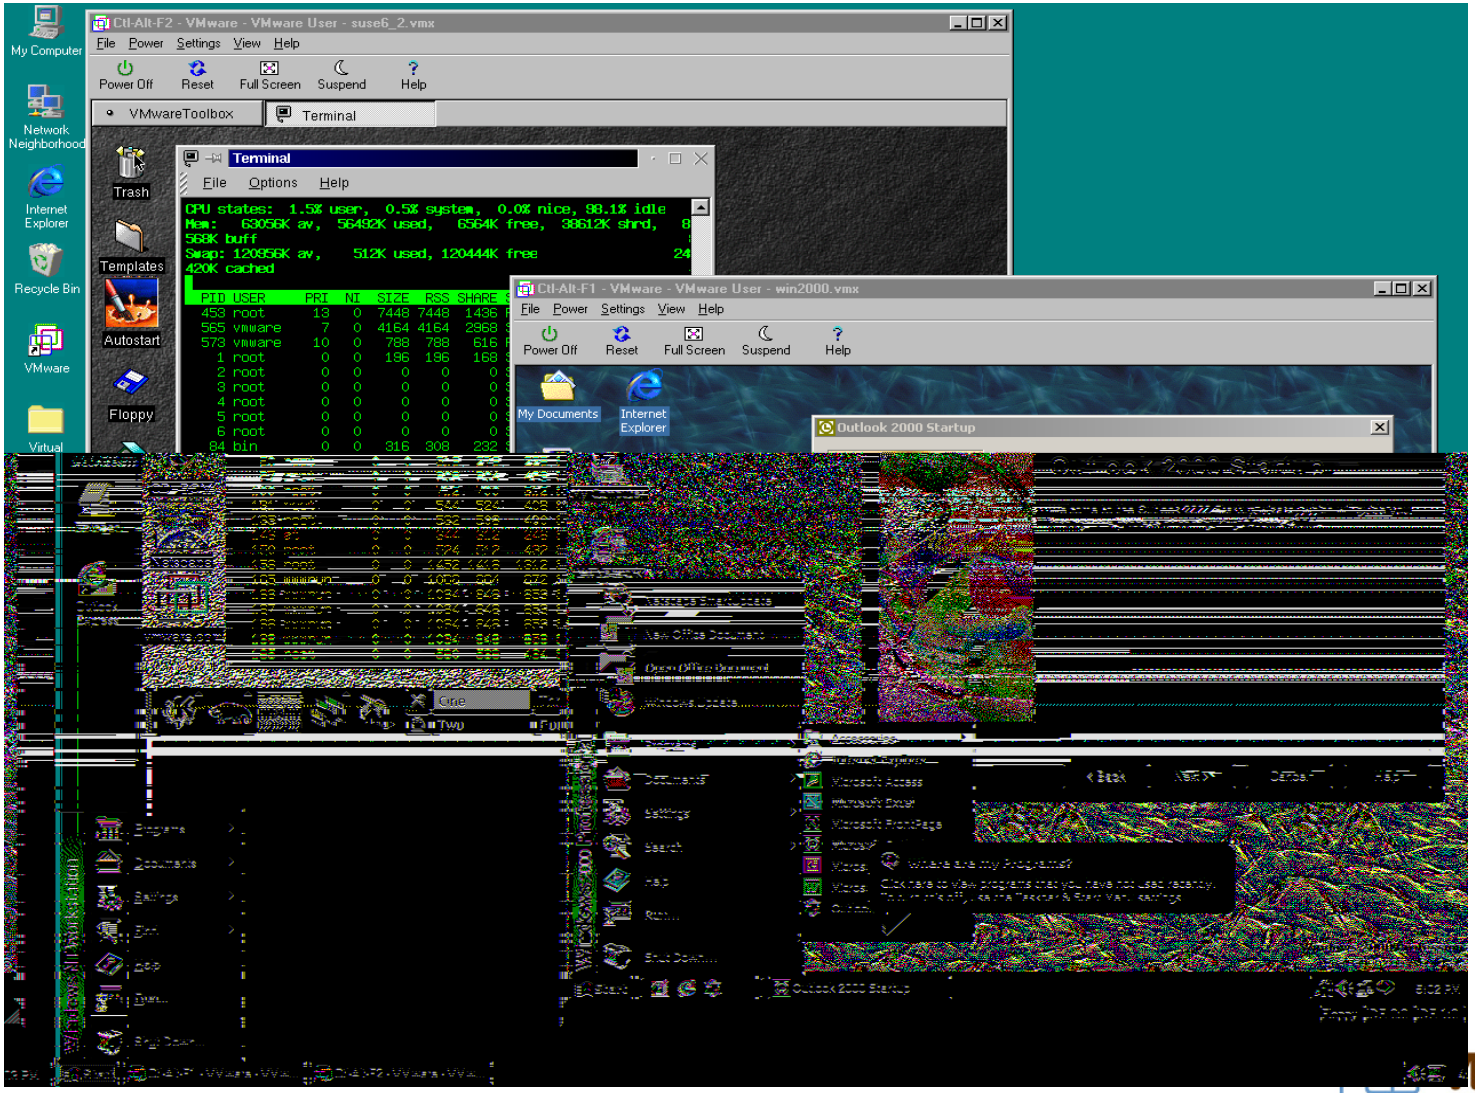

VMware Server – Screen Shot

- Web-based management interface
  - Stop, start, suspend/resume virtual machines
  - Monitor CPU usage
  - Run scripts
  - Secure user authentication
- Remote Console
  - Windows and Linux versions
  - Full desktop display
  - Full mouse and keyboard support
  - Secure user authentication
  - Access VMware configuration editor

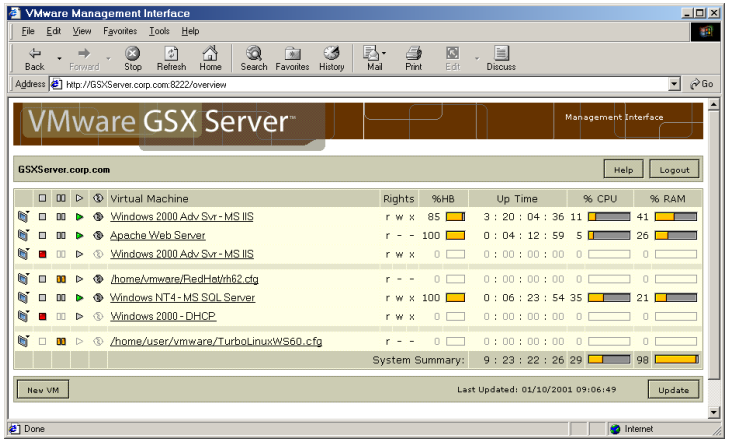

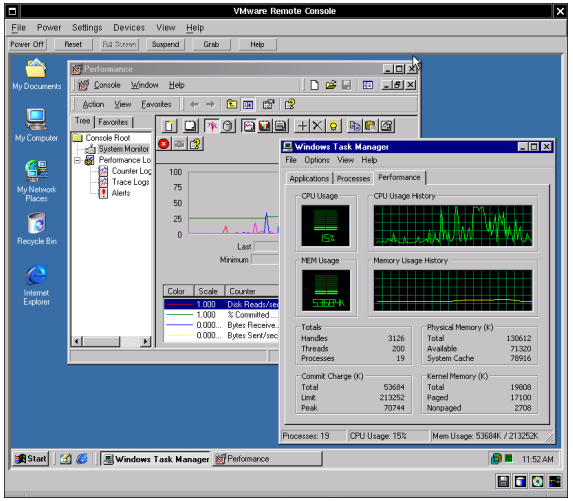

VMware Products

- VMware Workstation
  - Run Multiple Operating Systems on your workstation
  - Hosted Architecture
  - Available for Linux and Windows hosts
- VMware GSX Server
  - Run multiple servers on your server
  - Hosted Architecture
  - Available for Linux hosts and soon Windows hosts
- VMware ESX Server
  - + Quality of Service
  - + High-performance I/O
  - Host-less Architecture

Virtual Hardware

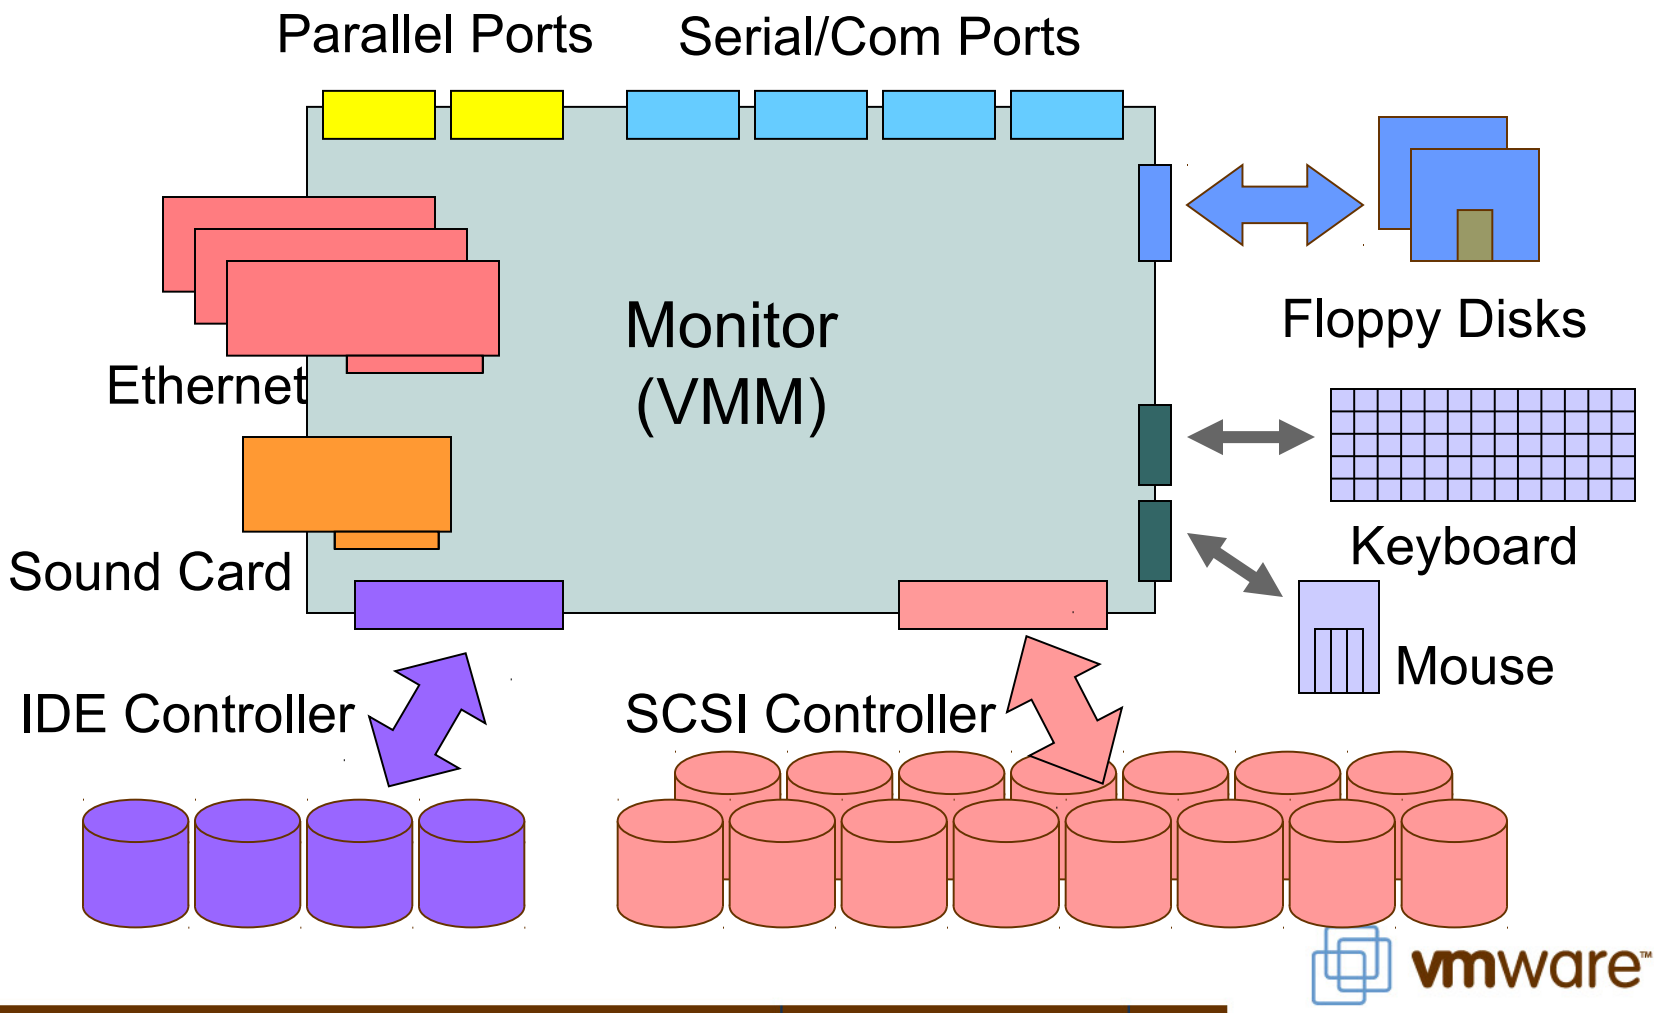

Attributes of MultipleWorlds Technology

- Software compatibility
  - Runs pretty much all software
- Low overheads/High performance
  - Near “raw” machine performance
- Complete isolation
  - Total data isolation between virtual machines
- Encapsulation
  - Virtual machines are not tied to physical machines
- Resource management

VMware Core Technology

The present

Virtualization through Ring Compression

Virtual Machine Monitor (VMM) runs at ring 0

Kernel(s) run at ring 1

Requires that CPU is virtualizable

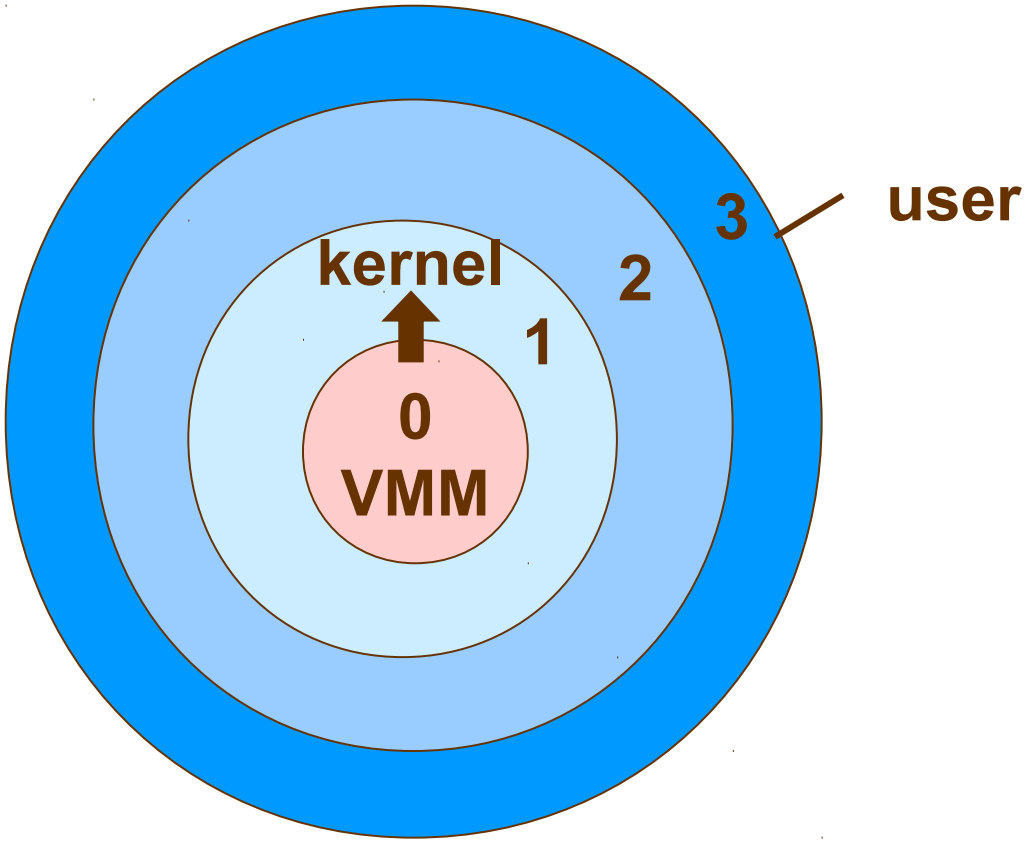

vmware™

Classification of processor architectures

- **Strictly virtualizable** processor architectures
  - Can build a VMM based on trap emulation exclusively
    - *No software running inside the VM cannot determine the presence of the VMM (short of timing attacks)*
  - Examples: IBM S/390, ~~DEC Compaq~~ Intel Alpha, PowerPC
- **(Non-strictly) virtualizable** processor architectures
  - Trap emulation alone is not sufficient and/or not complete
    - *E.g. instructions have different semantics at various levels (sufficient)*
    - *E.g Some software sequences can determine the presence of the VMM (complete)*
  - Examples: IA-32, IA-64
- Non virtualizable processor architectures
  - Basic component missing (e.g. MMU, …)

Hosted VMware Architecture

Host Mode

VMware, acting as an application, uses the host to access other devices such as the hard disk, floppy, or network card

VMM Mode

The VMware Virtual machine monitor allows each guest OS to directly access the processor (direct execution)

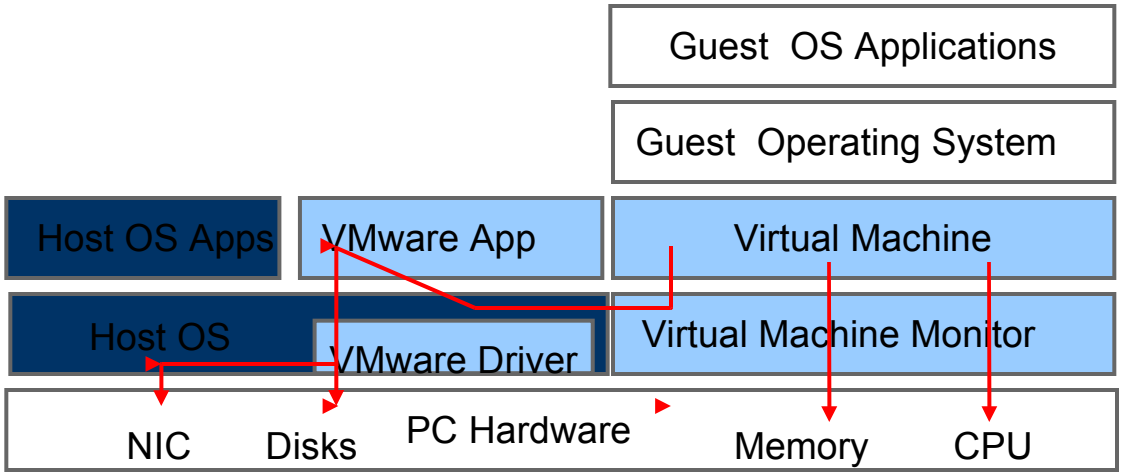

VMware achieves both near-native execution speed and broad device support by transparently switching* between Host Mode and VMM Mode.

*VMware typically switches modes 1000 times per second

Hosted VMM Architecture

- Advantages:
  - Installs and runs like an application
  - Portable – host OS does I/O access
  - Coexists with applications running on the host
- Limits:
  - Subject to Host OS:
    - *Scheduling Decisions*
    - *Resource management decisions*
    - *OS failures*
  - Performance overheads:
    - *World Switch*
    - *I/O access*
- Usenix 2001 paper:

  J. Sugerman, G. Venkitachalam and B.-H. Lim, “Virtualizing I/O on VMware Workstation’s Hosted Architecture”.

Virtualizing a Network Interface

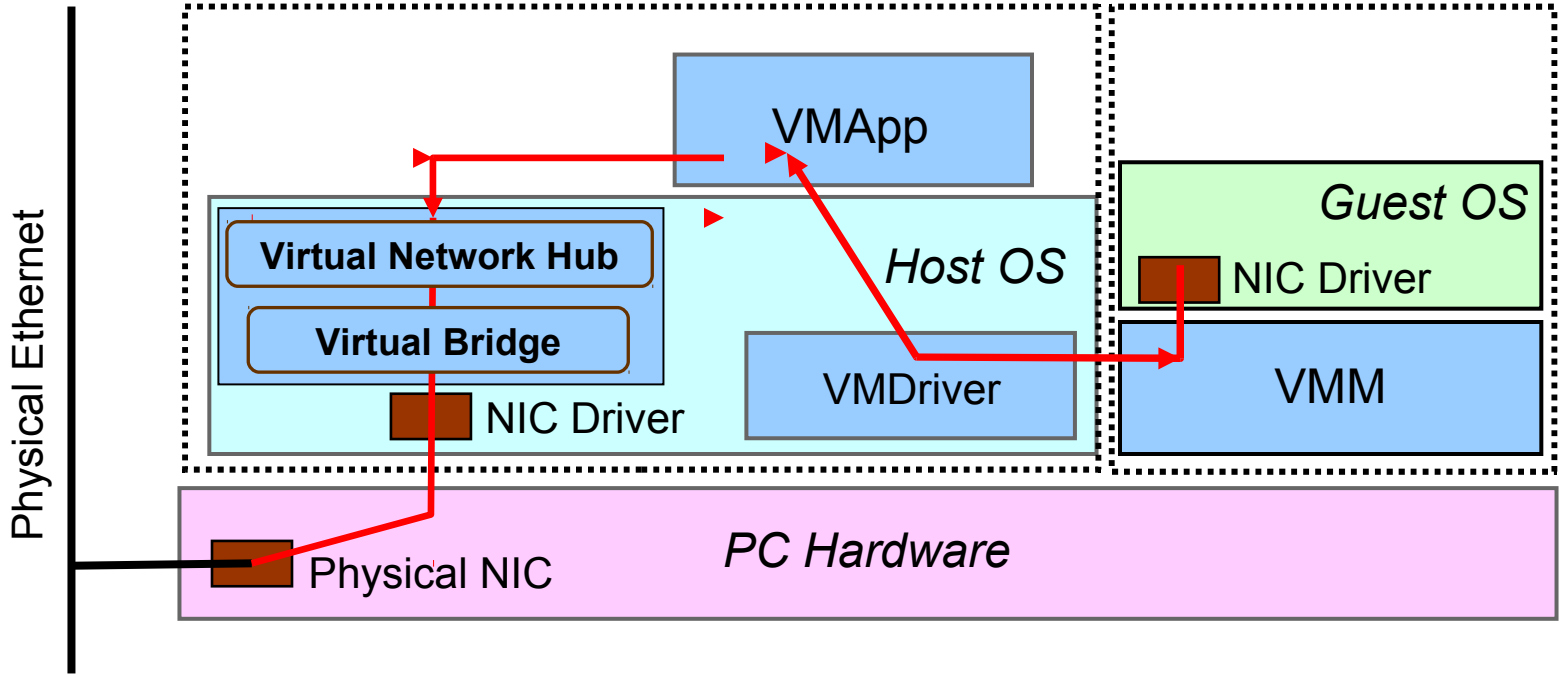

Experiment – TCP Throughput

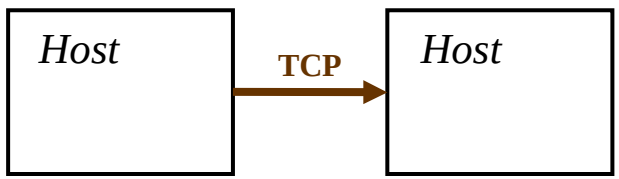

Native

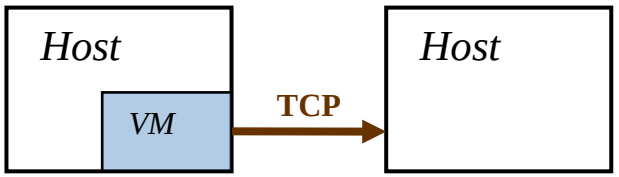

Virtual Machine

- Two speed of host:
  - Standard -- 733 MHz Pentium III
  - Slower -- 350 MHz Pentium II
- 100 megabit Ethernet connected via crossover cable
- Host and Guest OSes are Linux 2.2.x kernels
- 3 optimizations that reduce number of World switches

Optimized Performance– 733 MHz

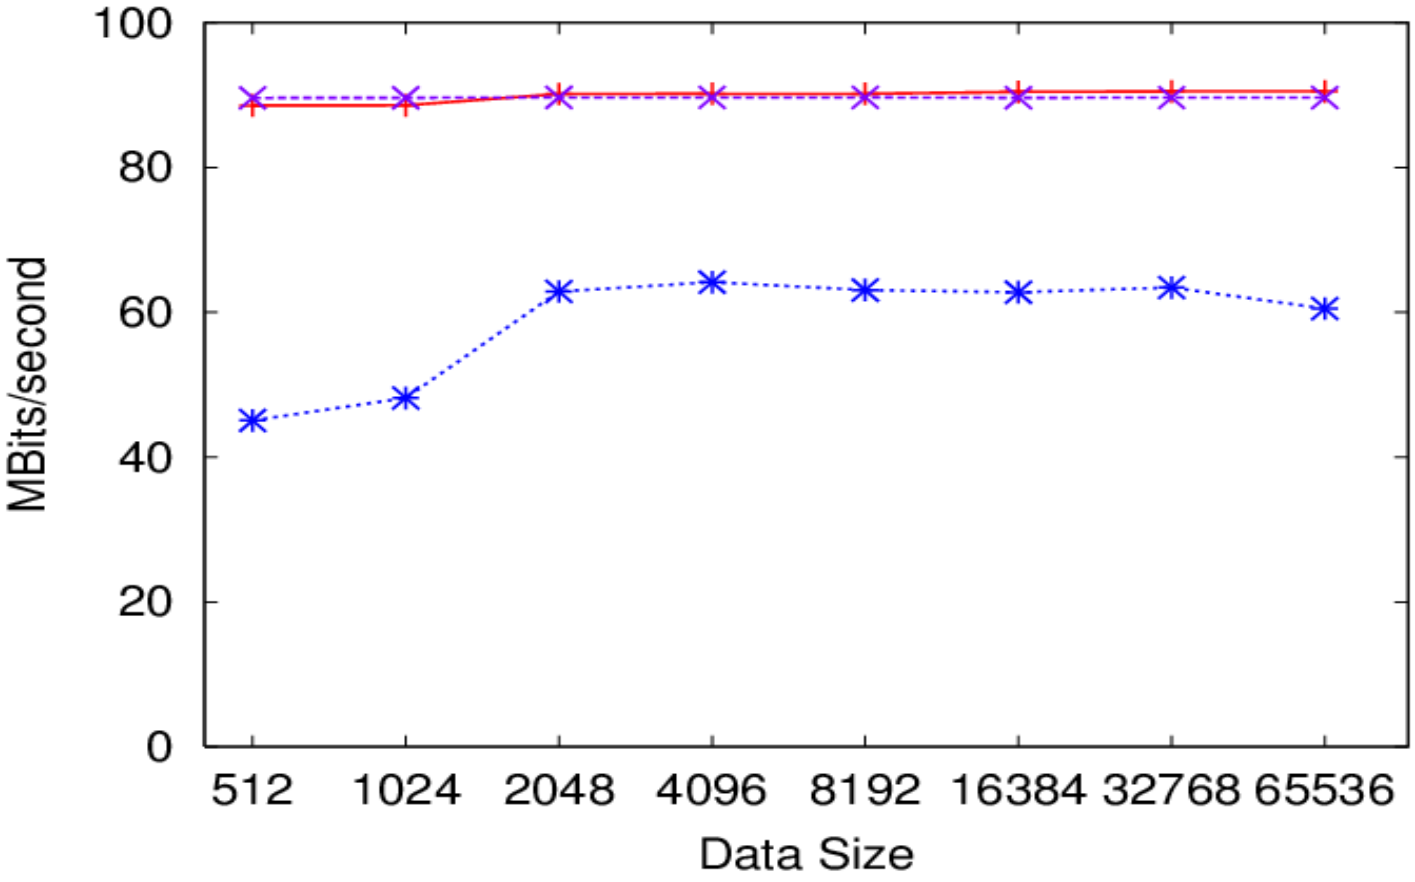

- Native
- VM/733 MHz Optimized
- VM/733 MHz Version 2.0

Optimized Performance– 350MHz

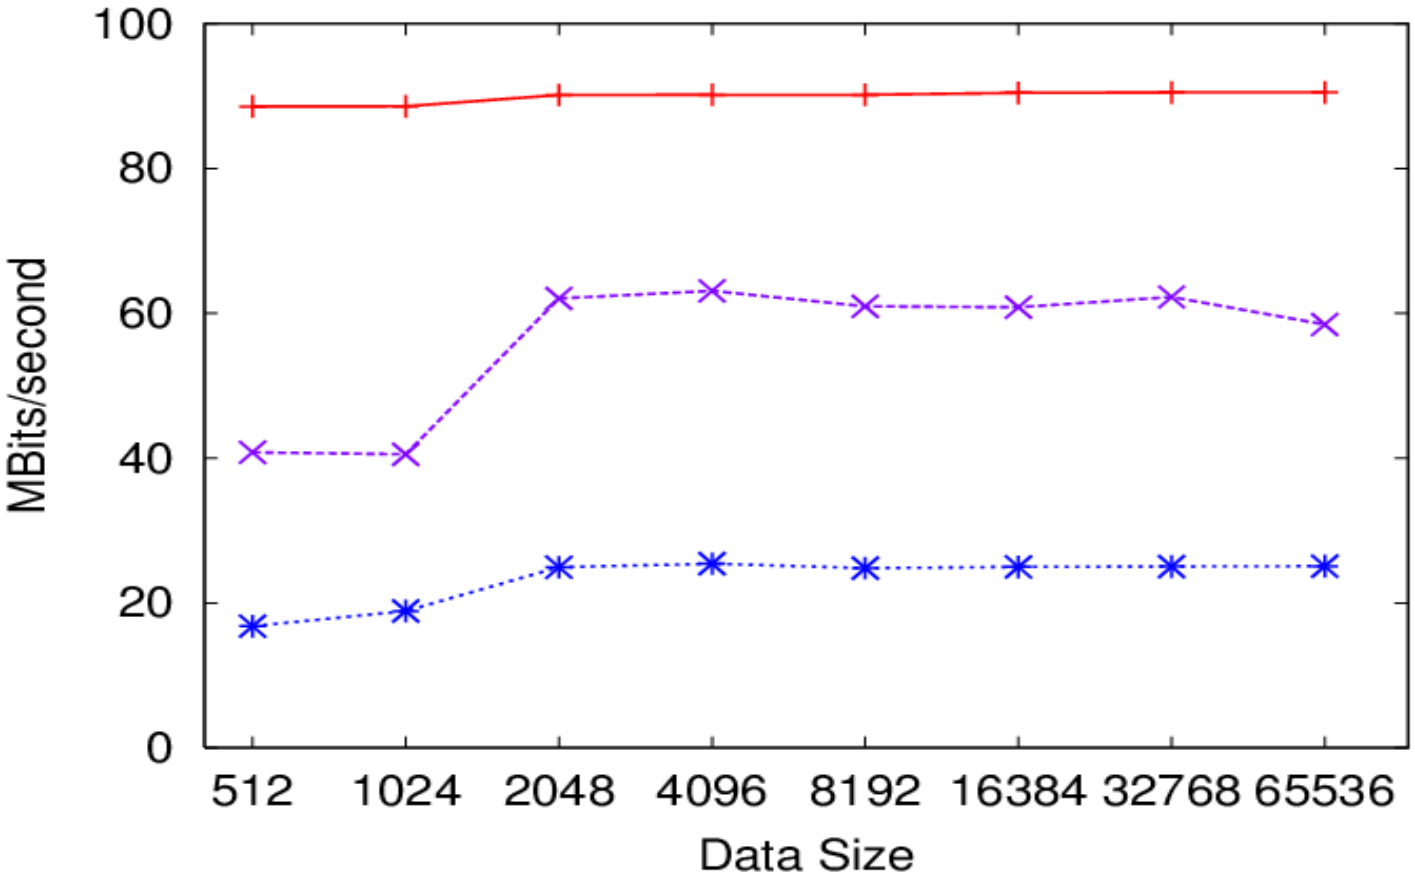

- **Native**
- **VM/350 MHz Optimized**
- **VM/350 MHz Version 2.0**

CPU Utilization – **VM/PC-733**

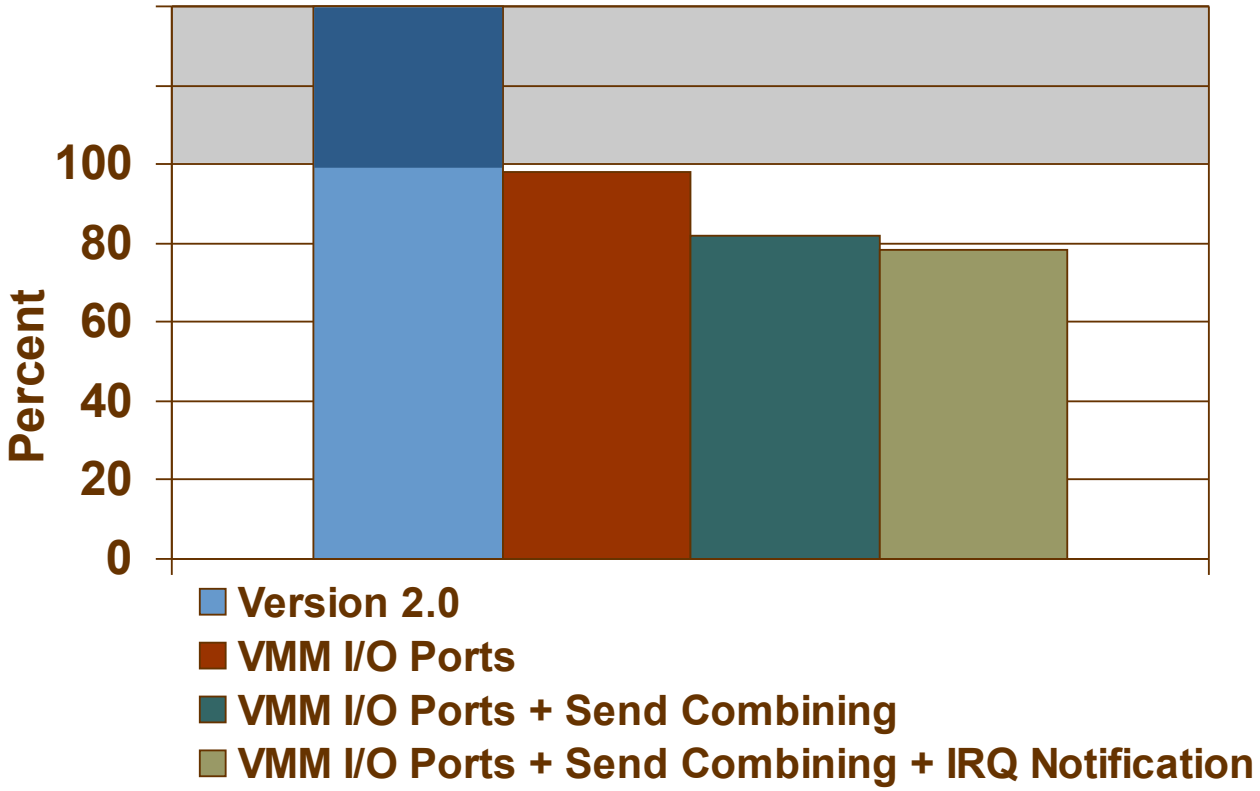

- Native **PC-733** is I/O bound with under 20% CPU utilization

Beyond the Hosted Architecture

- Limits of the Hosted Architecture:
  - World switch overhead – especially I/O
  - Hard to make QoS guarantees
  - Depend on the Host

- ESX Server Architecture:
  - Eliminate the host
  - All applications run in a VM
  - Looks closer to a traditional VMM system

ESX Server Architecture

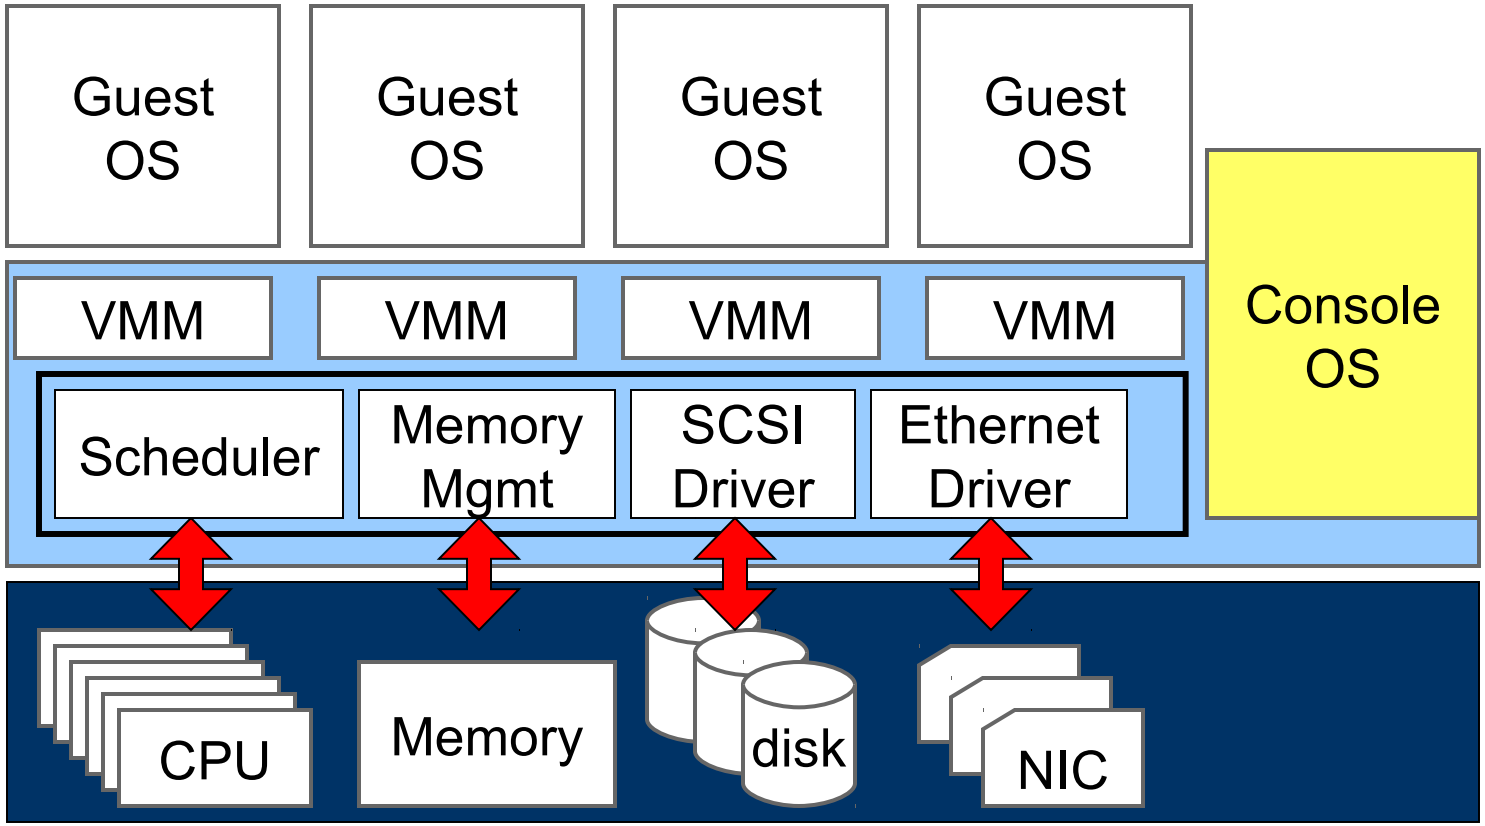

High Performance Network

- Ethernet and Gigabit Ethernet
- Each virtual adapter has its own MAC address
- No world switch !

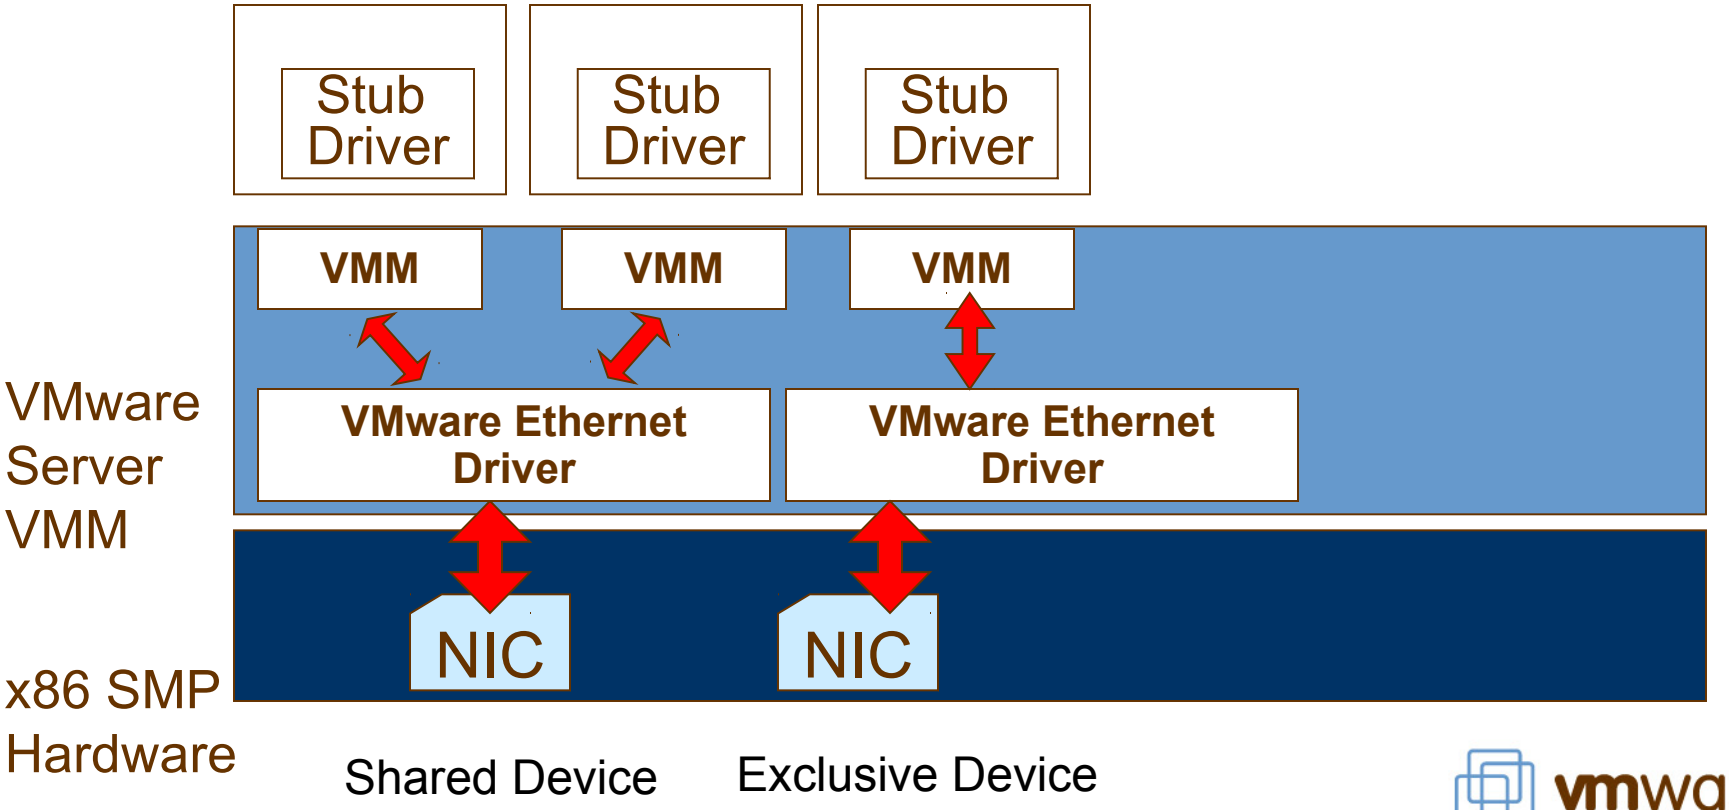

Intra-system networking

- Executes at memory speed

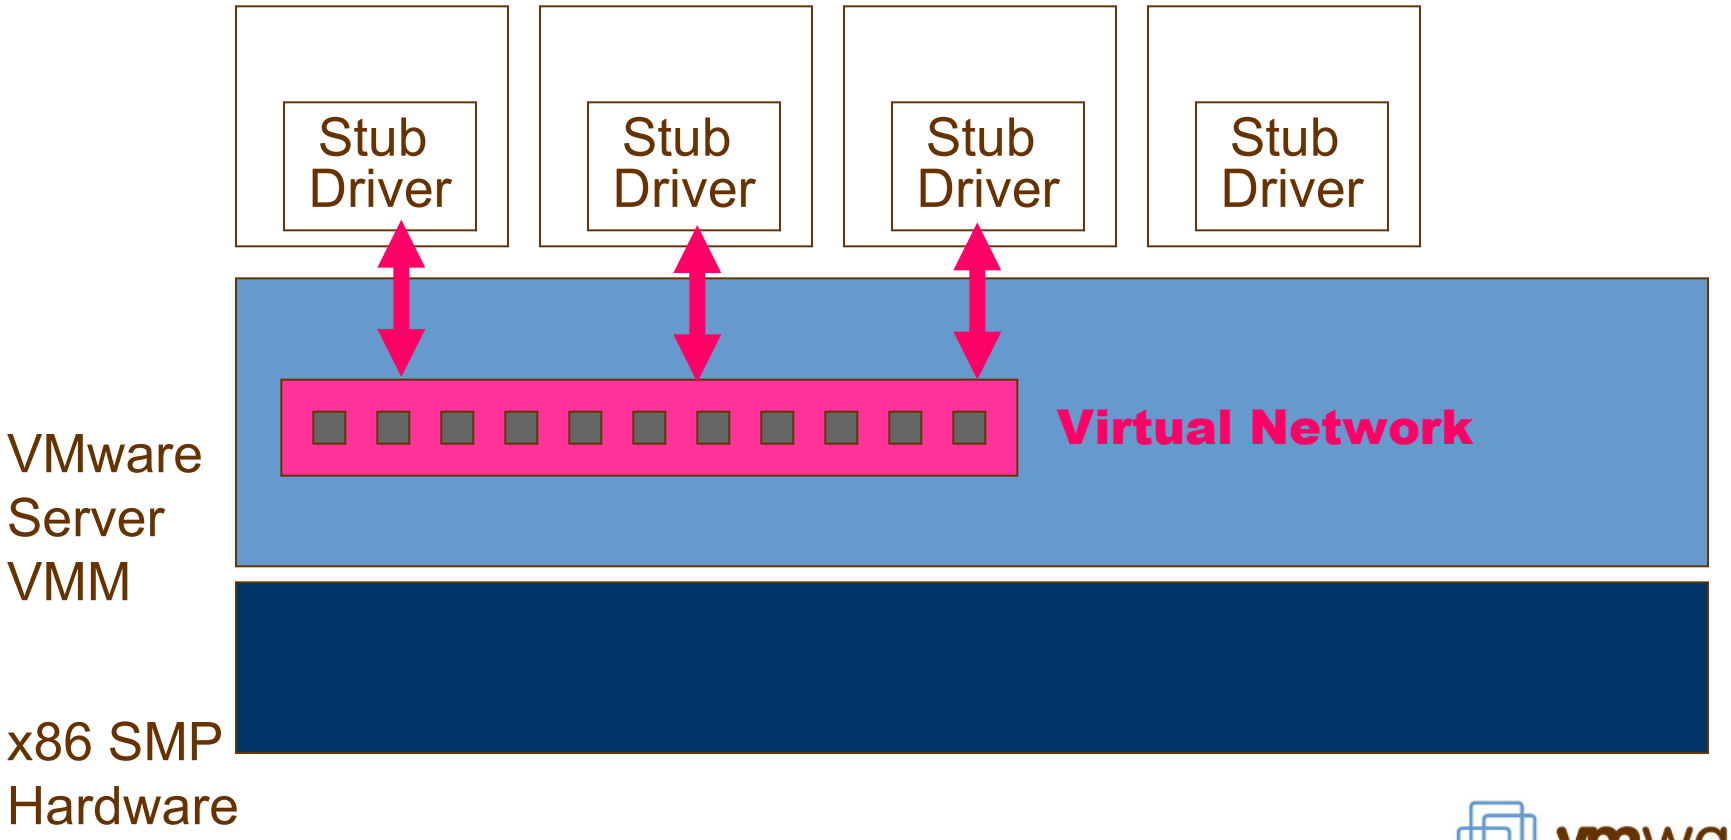

Usage Scenarios

4 Examples on Desktops and Servers

Scenario #1: Testing and Deployment

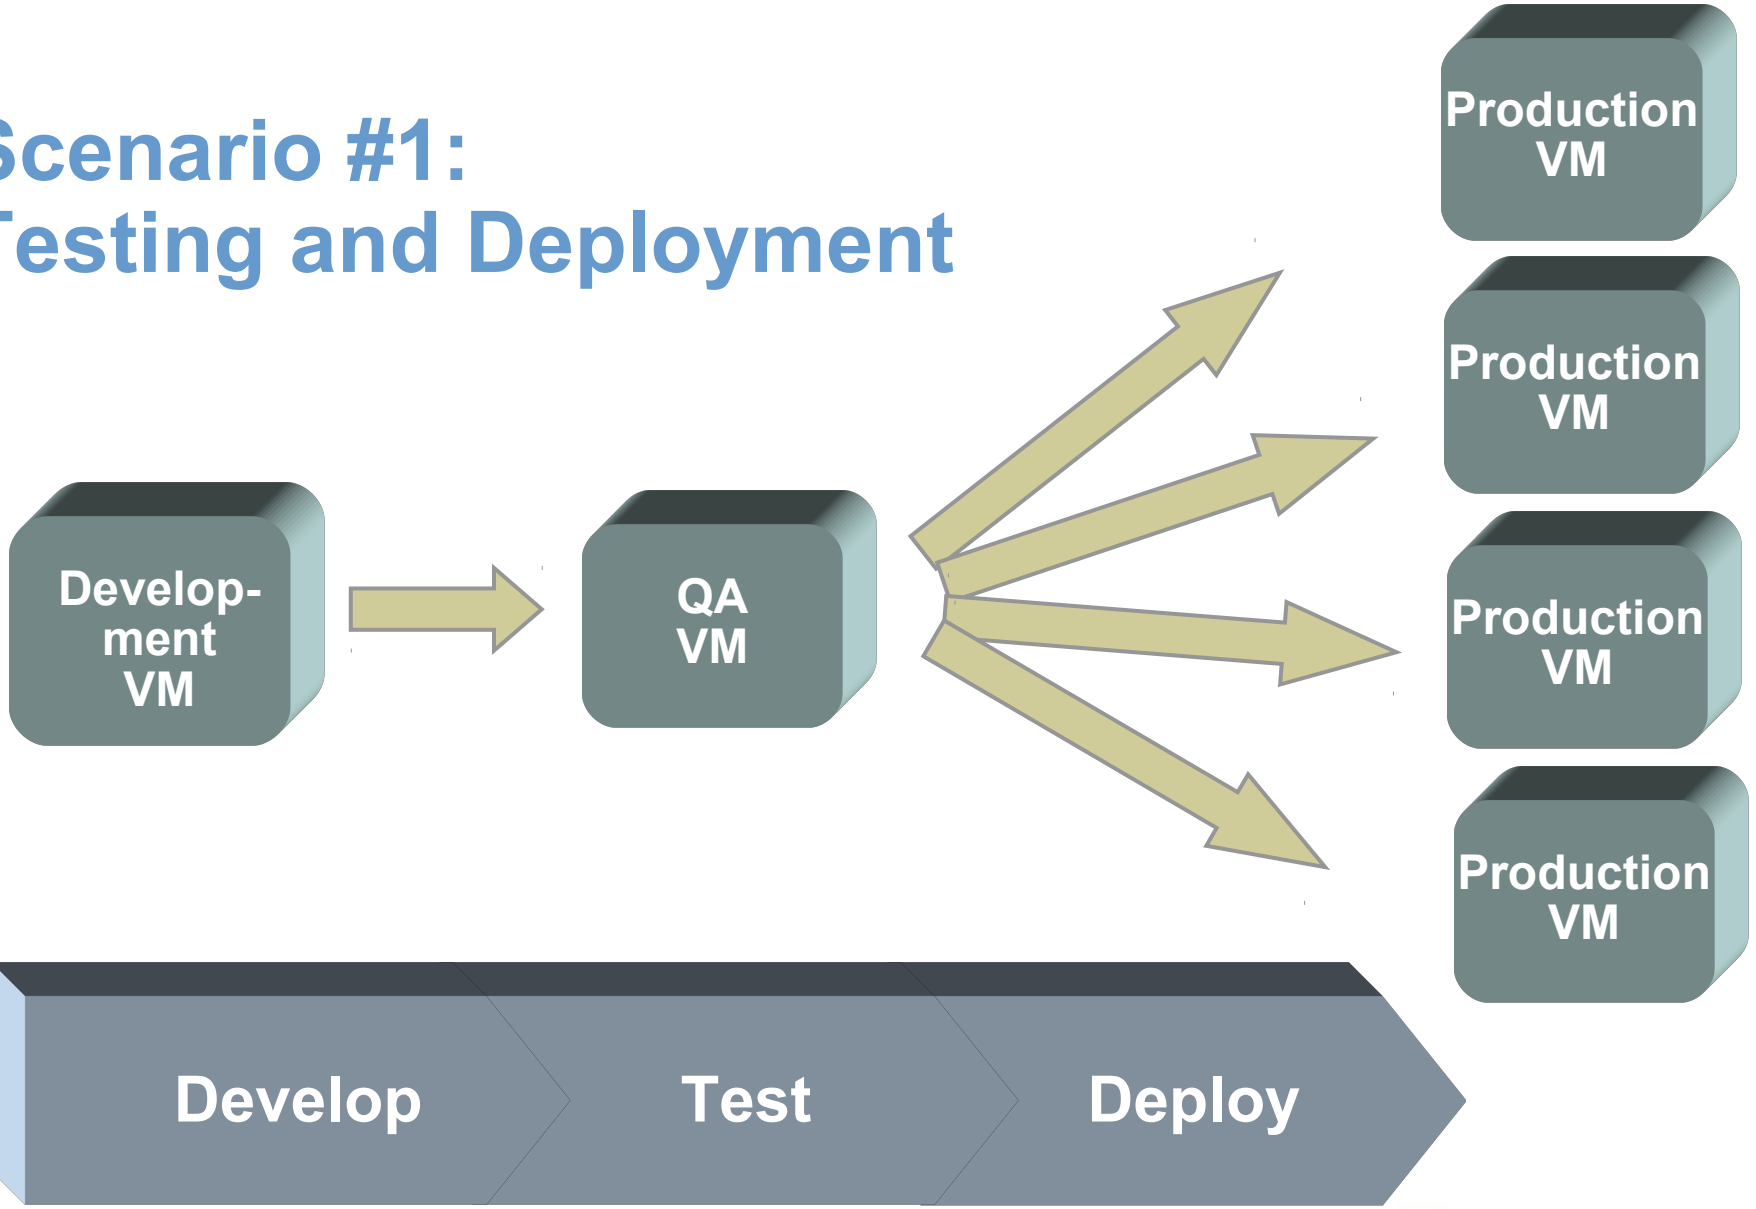

Major Wall Street Investment Banking Firm

Testing and Deployment

Challenge

Testing & deployment was error-prone and expensive

Solution

Test and deploy in VMware worlds

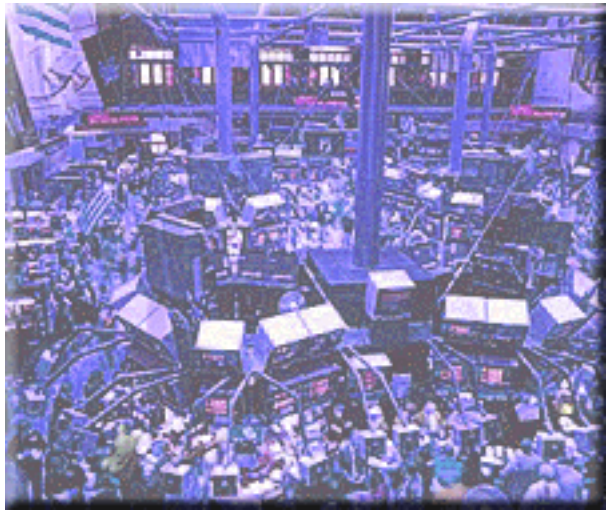

"VMware allows us to deliver well-tested and more reliable solutions in a shorter time frame at substantially lower costs."

vmware™

Scenario # 2:
Server Consolidation

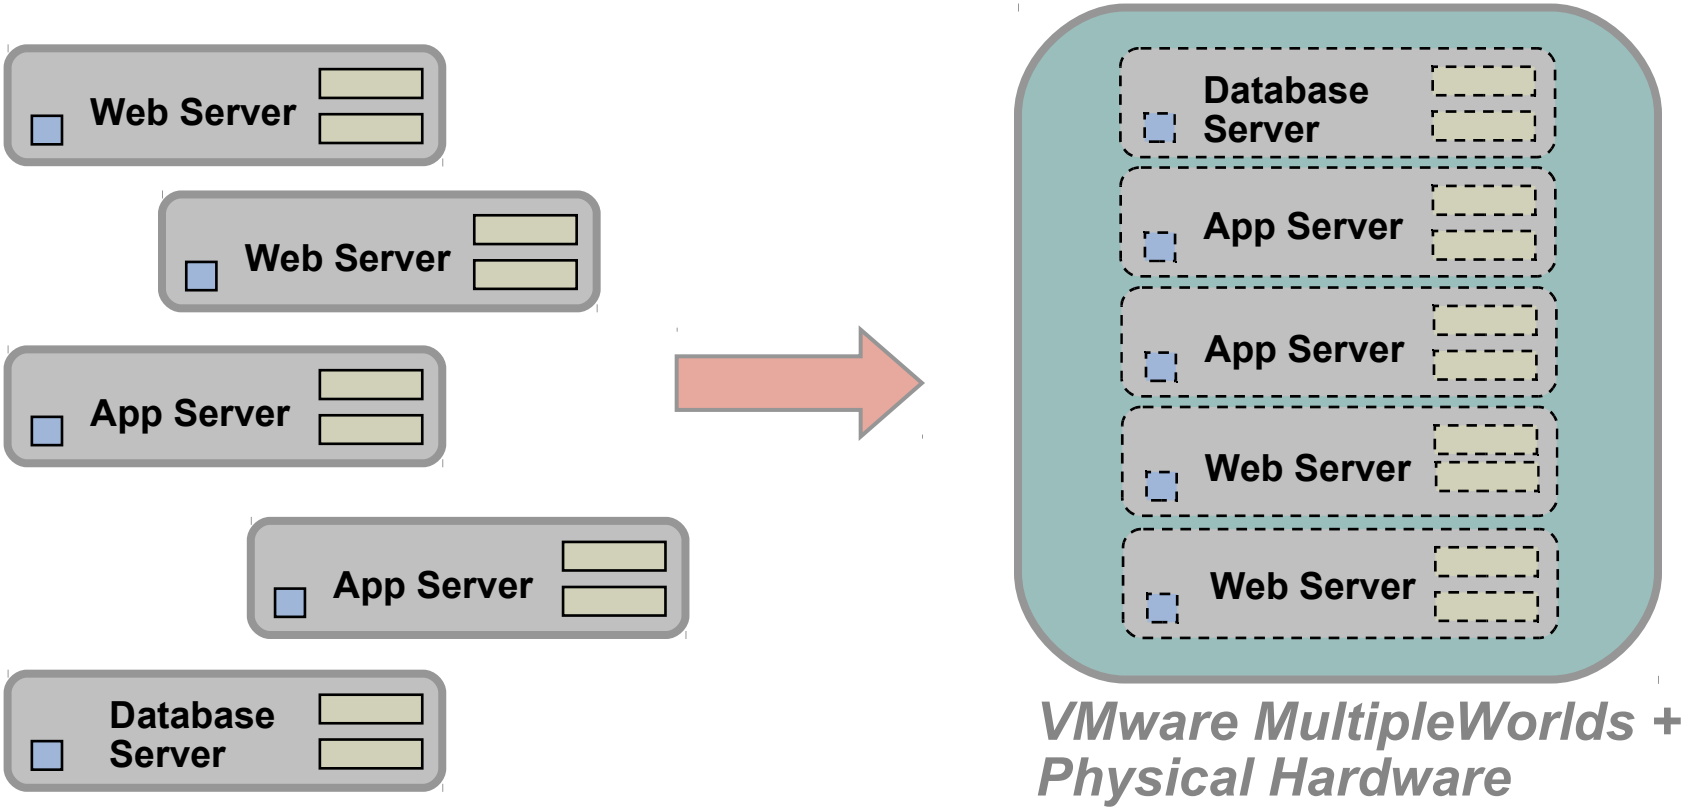

VMware MultipleWorlds + Physical Hardware

Halliburton Company

Server Consolidation

The Challenge

One database per oil well, one server per database

The Solution

Run each database in a VMware world

"We're able to run up to 10 database servers on a single server, which allows us to provide mainframe levels of reliability and data security at much lower cost."

vmware™

Scenario #3:
Application Compatibility

- Some applications require their OS
- Some solutions require multiple applications
- Appliances provide solutions

→ VMware in Appliances

Cisco Content Engine 590

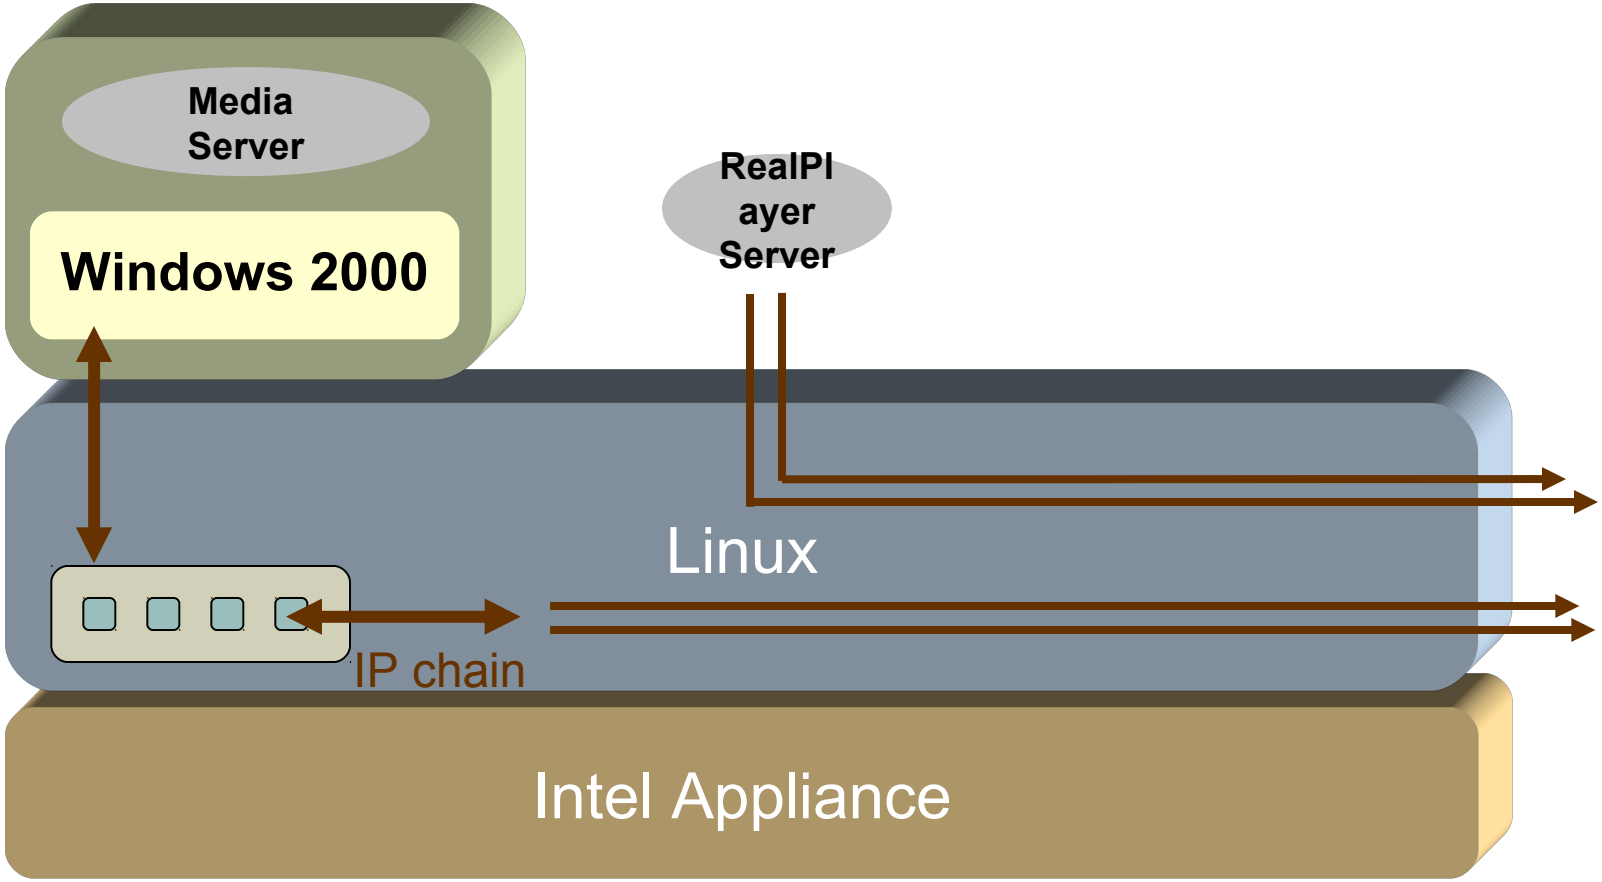

Scenario #4: Security Solutions

- Traditional tension : Security vs. Usability
  - Secure systems are not that usable
    - *E.g: require some particular OS setups*
  - Flexible systems are not that secure
    - *Many documented examples*
- Virtual Machines allow:
  - Secure Host
    - *that ensures the security of the whole system*
  - Flexible, Usable Virtual Machines
    - *that play no role in the security of the whole system*

National Security Agency NetTop

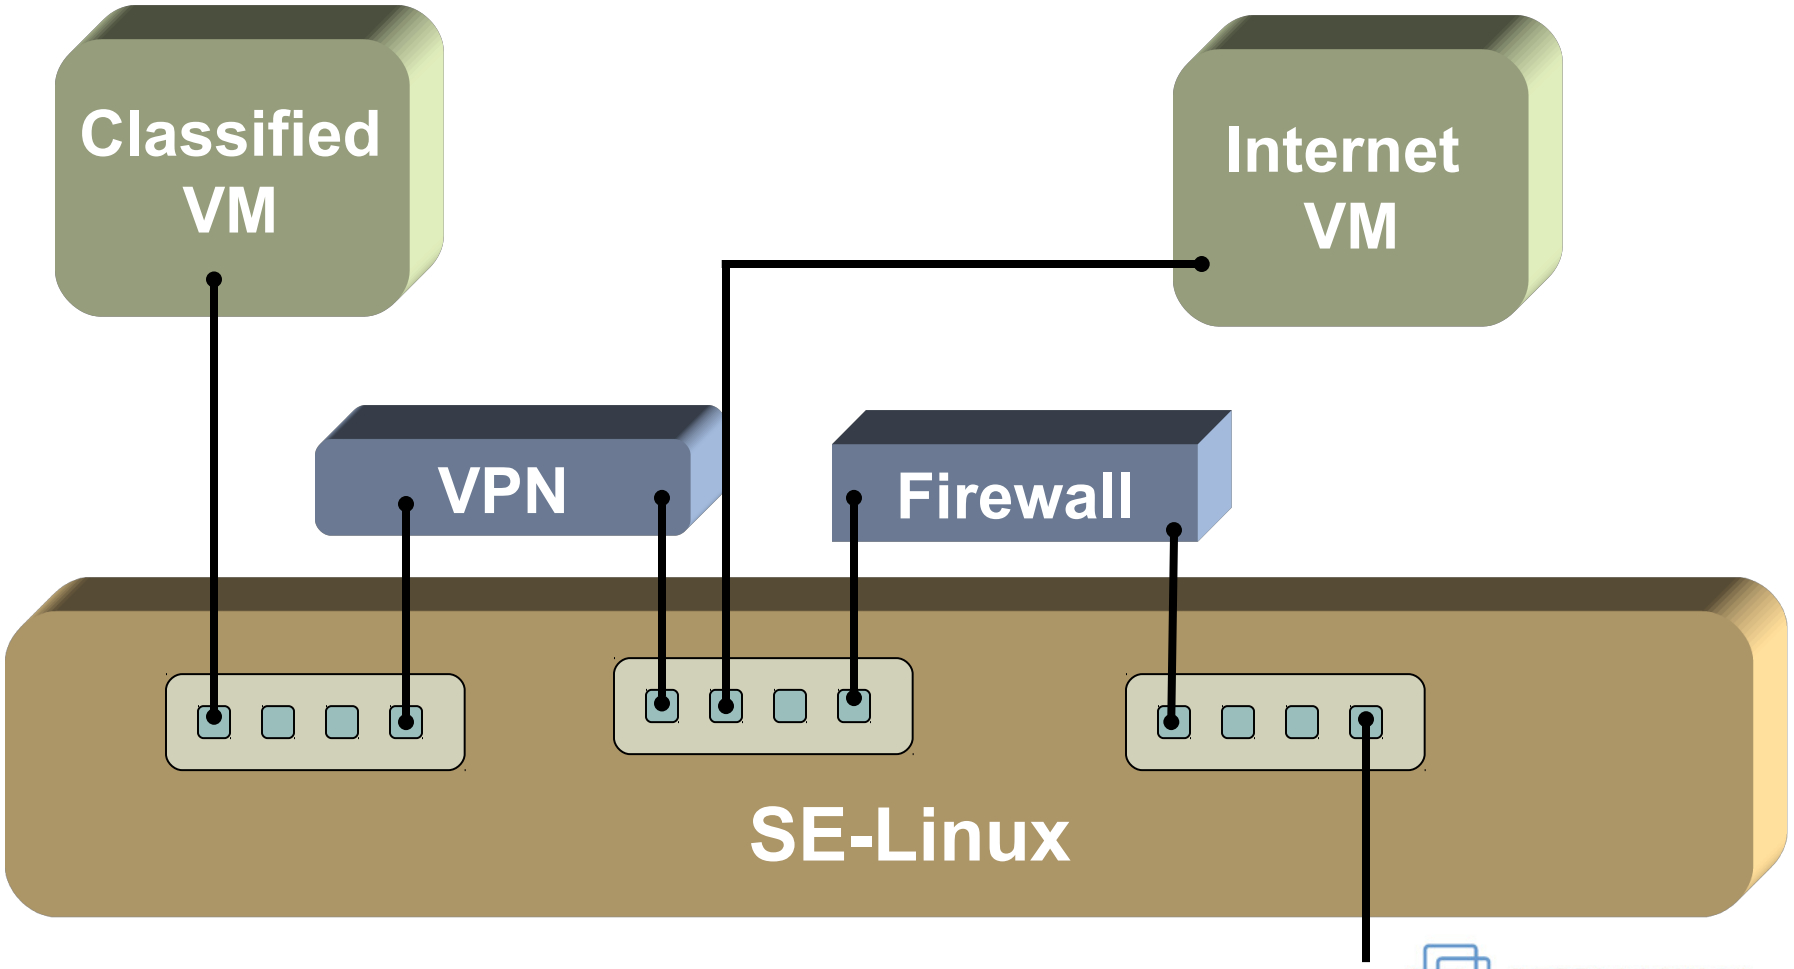

vmware™