White Paper

VMware Virtual Machine Technology

September 2000

VMware, Inc.
3145 Porter Drive
Building F
Palo Alto, CA 94304 USA
www.vmware.com

Introducing VMware Virtual Machine Technology

With its unique virtual machine technology, VMware™ provides the foundation for a new, more efficient approach to computing. VMware virtual machines enable users to run two, three, or more operating systems on a single computer — all at the same time. Each virtual computer has its own private workspace, providing the same level of security as a physical computer. And, like a physical computer, each virtual machine can use network connections to share files, printers, and other resources, at the user's discretion.

This gives users unprecedented flexibility.

- A developer who writes code in Linux can easily read email in the Windows world.
- A quality assurance lab can test on Windows NT — plus Windows 3.1, Windows 98, and Windows 2000 — all on one computer, without rebooting.
- A technical support department can switch quickly to a configuration that matches the one a customer is using. Dozens of virtual machines, each with a different combination of operating system and applications, can be stored as files on a server and called up quickly on the support rep's desktop computer.

This technology has proved itself on the desktop with more than 350,000 registered users of VMware Workstation around the world. Our new server products will address some of the most critical issues facing providers of hosted applications and services over the Internet. These include ways to:

- Eliminate the need to use a separate physical server for each application, service or customer
- Guarantee security on shared servers
- Maintain the stability of existing servers when adding new applications and services
- Overcome problems with operating system compatibility
- Substantially reduce expenses and problems in systems administration and management

VMware Virtual Machine Technology

The notion of the virtual machine first emerged in the era of big-budget IBM VM/360 mainframes. Virtual machines made it possible to share extremely expensive hardware — protecting users from the actions of others and giving each the impression of being in full control of the machine. These early virtual machines used a thin software layer called a virtual machine monitor to manage the computer's resources and share them among multiple users.

The low cost of today's x86-based PCs has made it possible for each user to have an individual machine. But the typical user taps only a fraction of the power of a modern desktop computer. Why not take full advantage of that computing power by reviving the virtual machine?

VMware has succeeded in applying the concept of virtual machines to the Intel x86 architecture — a significant breakthrough, because x86 processors are not designed for virtualization. In addition, VMware has incorporated recent advances in operating system technology to minimize the performance overhead of virtualization and add advanced disk management features and multiprocessor scalability.

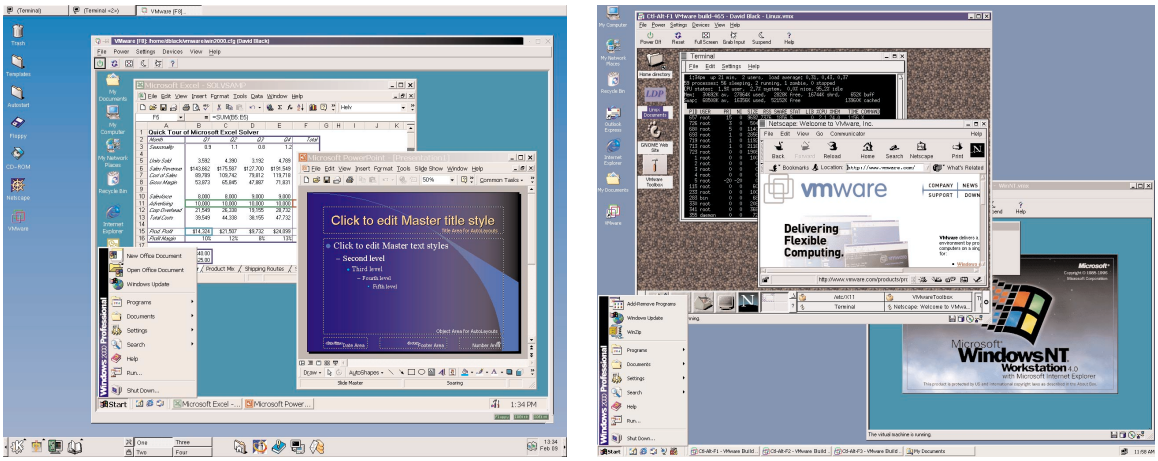

Figure 1. With VMware Workstation users can create one or more new computers — virtual computers — running at the same time on their Windows NT, Windows 2000, or Linux host computers

VMware Workstation is the first implementation of the VMware virtual machine technology. It uses software to create one or more virtual computers, each running its own operating system. As shown in Figure 1, these virtual machines typically appear as computers in a window — though a virtual machine can also be used in full-screen mode, taking full control of the display. Two VMware Workstation products are available. One runs on Windows NT 4.0 and Windows 2000 computers; the other runs on Linux computers.

VMware Workstation meets the key needs identified by users:

High performance — VMware overhead is impressively low. Applications running on virtual machines perform comparably to those running on real machines. Unlike PC simulators, which typically have high overhead for running their emulation processes, VMware technology enables virtual machines to make full use of the processor's power.

Portability — The VMware architecture runs on any x86-based computer, regardless of its manufacturer or its I/O devices.

Easy installation — VMware Workstation installs as easily as any other application. It installs without requiring changes to the existing operating system, and there is no need to repartition the host computer's hard disks.

Flexible disk management — VMware virtual machines can be encapsulated in a single file on the host called a virtual disk — or installed on a raw disk partition that can be booted natively, if desired. Disks can operate in any of three modes: persistent, nonpersistent, and undoable. *Persistent* disks behave like conventional disk drives on a computer, with changes made permanently. With *nonpersistent* disks, changes are discarded at the end of each session. *Undoable* disks allow the user to decide later whether to keep or discard changes.

Isolation — Each VMware virtual machine operates in its own protected space. A problem application in one virtual machine — even a system crash — will not affect other virtual machines or the host computer.

The fundamental technology developed by VMware has already proved useful in a large and diverse set of markets. VMware Workstation lets users:

- Run multiple operating systems and their respective applications — or multiple copies of the same operating system — concurrently, without restrictions.
- Encapsulate an entire computing environment and move it between computers as easily as copying a file.
- Test the same application concurrently using different operating system configurations — for example, with different amounts of memory, different operating system revisions, different Web browsers, or different system settings.
- Dedicate a virtual machine to run non-trusted applications downloaded from the Internet. The VMware technology guarantees that the virtual machines can be isolated from each other and from the host.
- Upgrade obsolete hardware and system software without losing compatibility. The legacy operating system and its applications can simply be transferred to a virtual machine.
- Rely on a hardware platform — defined by a VMware virtual machine — that is known to be stable. Virtual machines configured for this stable hardware platform will run correctly on any hardware that supports VMware Workstation.
- Suspend a virtual machine at any time. Its state is written to disk or saved in memory. Using Instant Restore, the suspended virtual machine can be restarted in seconds, avoiding the delay of the boot process and restoring any applications that were open.

VMware server software will expand this list of applications to include, for example, balancing loads among a collection of virtual machines on a scalable cluster of computer servers.

The VMware Architecture

How does this virtual machine technology work?

First, look at the way a standard PC operates, as shown in the illustration on the left in Figure 2. You have the hardware, complete with processor, memory, and disk drives. An operating system runs on the hardware. Applications — from word processors and Web browsers to databases and CAD programs — run on the operating system. And in this traditional model, each PC is limited to running one operating system at a time, and thus is limited to running applications designed for that operating system.

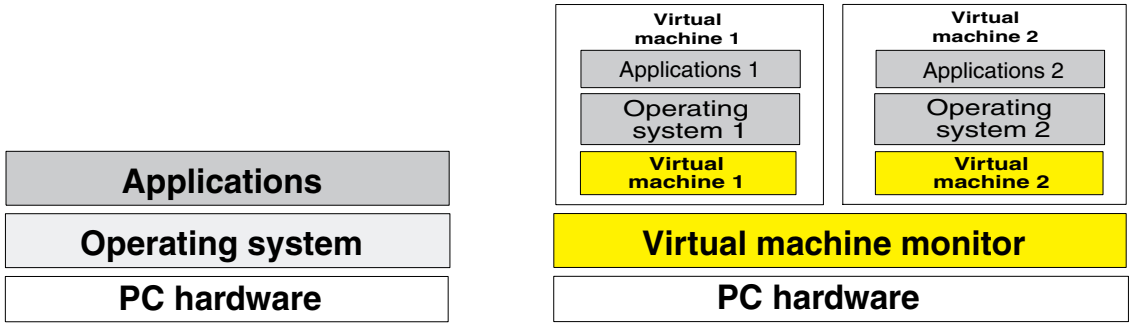

Figure 2. A standard PC (left) is limited to running one operating system and the applications it supports. Virtual machines (right) make it possible to run multiple operating systems at the same time.

A virtual machine monitor runs directly on the hardware and creates virtual machines, each of which behaves like a physical computer and can run its own operating system, as shown in simplified form in the illustration on the right in Figure 2.

VMware Workstation operates in conjunction with an existing, or host, operating system and runs guest operating systems inside virtual machines.

VMware is installed on a physical computer that is already running an operating system — the host operating system — and may be running applications on the host operating system. It creates a virtual machine, running its own operating system — the guest operating system. That guest operating system and any applications running in the virtual machine behave as if they were running on a separate physical computer.

VMware itself is a virtual machine monitor running in kernel mode directly on the hardware. At the same time, it is a normal application running in user mode on top of the host operating system, as shown in Figure 3 on the next page.

For some functions, such as access to memory and the CPU, the virtual machine has direct access to the hardware of the host PC. For other functions, such as disk access, the virtual machine uses services provided by the host operating system.

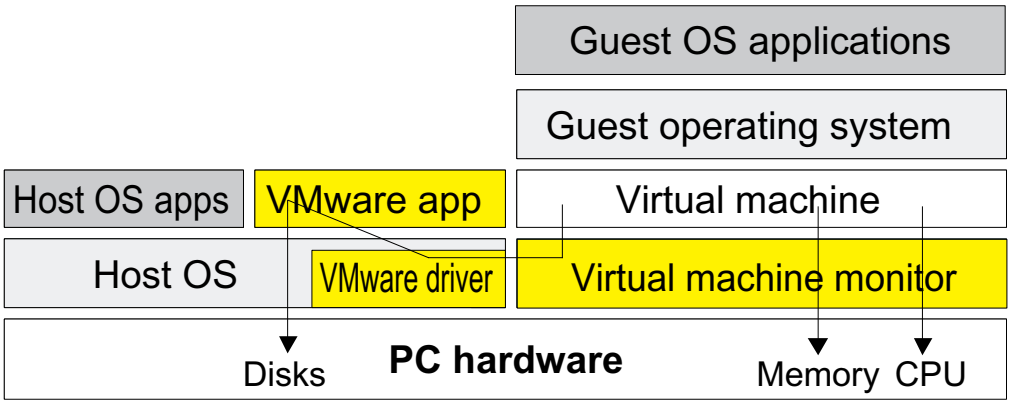

Figure 3. The VMware Workstation architecture

VMware Workstation uses three software components to implement its special dual abilities.

VMware application — The application portion of VMware Workstation installs, starts, and runs like a normal application on top of an existing operating system and uses the graphical user interface of the host operating system to configure, launch, and administer virtual machines. To launch a virtual machine, the VMware application loads and runs the monitor. During execution, the monitor calls back to the VMware application when it needs access to resources controlled by the host operating system. The application then calls the host operating system to access these resources.

Virtual machine monitor — The monitor portion of VMware Workstation runs in privileged mode directly on top of the hardware. The virtual machine monitor eliminates overhead through a virtualization technique known as direct execution. This technique allows the virtual machine to run directly in an environment where it is isolated both from the host operating system and from other virtual machines.

The virtual machine monitor is specifically structured to virtualize the x86. For maximum flexibility, it suspends some of the system state set up by the host operating system. For example, the monitor operates in an address space that is distinct from that used by the host operating system. However, the virtual machine monitor defers to the host operating system for the management of all system resources, including the physical memory of the machine and the scheduling of the processor.

VMware driver — VMware Workstation also includes a device driver that runs under the host operating system. The VMware driver is specific to the host operating system and acts primarily as a communication gateway between the application and monitor portions of VMware Workstation. The VMware driver also effectively hides the virtual machine monitor from the host operating system. As a result, VMware appears like a normal process to the host operating system.

VMware Workstation always uses its application and monitor portions. Without its application portion, VMware Workstation would need its own drivers to handle the diversity found in I/O devices. Without its monitor portion, VMware Workstation would resemble a simulator, and would have the limited performance found in simulators. The monitor gives VMware the ability to arbitrarily take over the processor so it can minimize the overhead of running the virtual machine.

For maximum performance, the core components of VMware Workstation use the virtual machine monitor and configure the hardware to match the requirements of the virtual machine. Thus the virtual machine has direct access to the host computer's CPU and a portion of its memory. For maximum portability, the device-dependent portions of the system use the VMware application and rely on the host operating system and its device-specific drivers. Thus I/O calls, for example, are routed through the VMware application and take advantage of host operating system services.

Virtual Machines with the Features You Need

What this technology creates for users is an amazingly capable working environment. Choose your preferred host (a typical strategy is to run most applications on the host — so the choice of host depends on the operating system needed by those applications). Then add VMware for Windows NT and Windows 2000 or VMware for Linux, and you get virtual machines that can run the most popular x86-based operating systems and their applications at near-native speeds.

Take a quick look at the key capabilities a virtual machine provides.

CPU — The processor seen by the virtual machine corresponds to the physical processor in the host computer. So, for example, Intel MMX capabilities are available to the virtual machine if they are present in the physical processor. In symmetric multiprocessor systems, each virtual machine has access to one of the processors on the host.

Memory — A virtual machine can use as much as 512MB of RAM if there is enough RAM installed in the host computer. The memory used by a VMware virtual machine is subject to the standard allocation and paging policies of the host operating system. The VMware virtual machine system does need to lock a subset of the pages in memory at any given time. The size of this subset can be configured using controls available in the VMware interface.

The amount of memory required for a virtual machine varies with the operating system and the application software to be run there. A good starting point for planning purposes is to allow 64MB of memory for each virtual machine that will be running at a given time plus 64MB for the host.

CD-ROM and floppy drives — Virtual machines can use up to two 1.44MB floppy devices. These can be mapped to physical drives on the host computer or connected to floppy image files on the host file system. CD-ROM drives on the host can be mapped to corresponding devices in the virtual machine. Both IDE and SCSI CD-ROM drives on the host appear as IDE CD-ROM drives in the virtual machine.

Hard disks — The virtual machine's virtual hard disks can be IDE, SCSI, or both. The virtual IDE controller supports up to four devices, which can include both hard disks and CD-ROM drives. The virtual SCSI controller can support up to seven virtual hard disks. If the host operating system has access to a storage device, VMware can build and use the virtual hard disks of the virtual machine on that device. This includes raw disks, disk partitions, RAID arrays, and files on a local or network file system. The latter option allows users to create virtual machines without having to repartition or add disk resources.

A good starting point for rough planning purposes is to allow about 500MB of disk space for each virtual machine. For more precise estimates, it is necessary to know how much disk space is required for the operating system, applications and data in each virtual machine. These space needs will be the same in a virtual machine as they are on a physical computer. If the virtual machine's drives are used in undoable or nonpersistent mode, additional space is needed for the files that store information on changes made during a working session.

Network — VMware takes advantage of host operating system drivers for access to network interface cards. This means that if the host can use the device, so can the virtual machine.

The VMware software emulates an Ethernet card in the virtual machine and an Ethernet hub on the host, effectively giving the virtual machine a network identity indistinguishable from that of a real machine. As a result, virtual machines can share files, printers, and other resources with other virtual machines, the host operating system, and machines connected to the LAN or the Internet. In fact, a single PC can host multiple VMware virtual machines networked together to create a self-contained client/server development and test environment. The use of standard distributed file system protocols such as NFS and SMB allows the tight integration of the virtual machines.

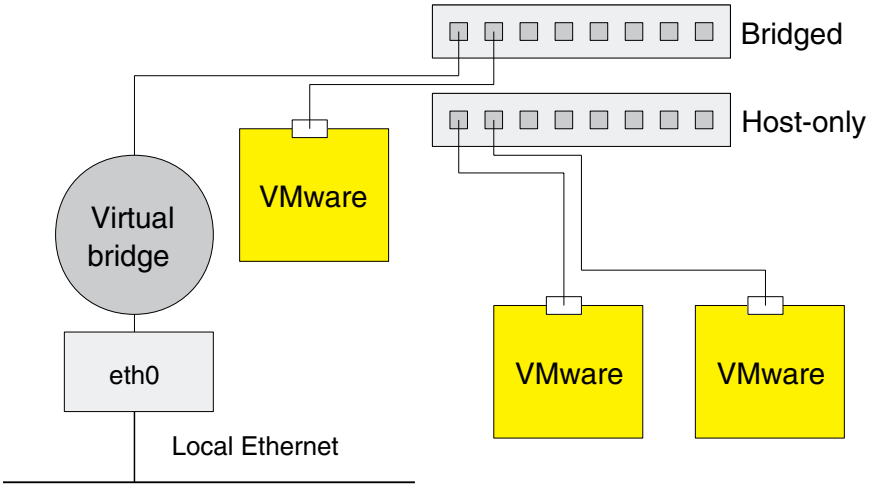

Figure 4. Bridged networking and host-only networking in VMware

VMware virtual machines can operate in one of three networking modes:

- **Bridged** — The virtual machine connects to the host, other virtual machines, and the external real network through a virtual network interface card and a virtual Ethernet hub running on the host machine. To the outside world, the virtual machine appears as a totally independent machine on the same Ethernet as the host.
- **Host-only** — The virtual machine is limited to networking only with the host machine through the virtual network interface card. The host can then be configured to connect the virtual machine to the outside world, using routing or masquerading.
- **No networking** — The virtual machine is isolated from all network activity.

Input devices — VMware Workstation multiplexes user input devices such as the mouse and keyboard. User events are either sent to a specific virtual machine or left to the host operating system for default processing. Positioning the mouse cursor in the window that runs the virtual machine gives it exclusive user input. Moving the mouse cursor to the host desktop or pressing a hot-key sequence releases the display input to the host operating system.

Guest operating systems — VMware virtual machines provide full support for Windows 95, Windows 98, Windows NT, Windows 2000, popular distributions of Linux, and FreeBSD. Other x86-based operating systems — including MS-DOS and Windows 3.1 — will run in a virtual machine, although there is no performance-enhancing VMware Tools package available for these operating systems.

Other devices — VMware software also virtualizes the other devices that complete a personal computer, such as sound cards, serial drawing tablets, parallel ports, and serial ports.

A Better Approach to Computing

VMware virtual machine technology provides the foundation for a new approach to computing. Users can mix and match multiple operating systems and applications — running otherwise impossible combinations on a single computer. They can exploit the full power of their computers, work more flexibly and efficiently, and dramatically reduce the cost of computing throughout their enterprises. This technology has proven itself on the desktop and will soon demonstrate its potential as a platform for delivering applications and services on corporate intranets and over the Internet.