Technical White Paper

February 1999



VMware, Inc. 3145 Porter Drive Building F Palo Alto, CA 94304 USA Tel: 650-475-5000 Fax: 650-475-5001 www.vmware.com

 $^{\odot}$ 1999 VMware, Inc. All rights reserved. VMware and Virtual Platform are trademarks of VMware, Inc. All other trademarks are the property of their respective holders.

Introducing VMware[™] Virtual Platform[™]

VMware offers a dramatically new approach to personal computing that enables users and organizations to use their PCs more flexibly and productively than ever before. VMware shatters the limits of personal computing and frees developers and end users to be more innovative and experimental—without risk. VMware's approach is based on its patent-pending *Virtual Platform*[™] technology.

VMware Virtual Platform is a thin software layer that allows multiple operating system environments to run concurrently using the same hardware resources. It accomplishes this by transparently multiplexing all hardware resources into multiple virtual machines, each resembling the underlying machine. Figure 1 shows two virtual x86-based machines running on the same x86-based hardware.



Figure 1. VMware Virtual Platform

VMware Virtual Platform allows virtual machines to work in concert with each other, sharing files and devices. This is possible because each virtual machine has its own unique network address. VMware Virtual Platform supports the integration of multiple environments so that these environments perform like multiple applications on a single computer.

VMware Virtual Platform also provides additional features not offered by standard operating systems. For example, VMware Virtual Platform can encapsulate a virtual machine and enable it to be moved freely among different physical machines. VMware Virtual Platform also provides each virtual machine with fault isolation and containment

🗊 **vm**ware

capabilities, meaning that application crashes or data corruption occurring in a given virtual machine will not affect data or applications outside of that virtual machine.

VMware Virtual Platform meets three key requisites demanded by users and organizations:

High performance – VMware Virtual Platform overhead is impressively low. Applications running on virtual machines perform comparably to those running on real machines. Unlike PC simulators, VMware Virtual Platform uses all the components of the processor and eliminates the overheads typically associated with PC emulators.

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

Easy installation – VMware Virtual Platform installs as easily as an application. On an existing system, it installs without requiring changes to the operating system, and without requiring the addition or repartitioning of disk resources.

These features make VMware Virtual Platform a general-purpose technology with applications for a large and diverse set of markets. VMware Virtual Platform lets users:

- Run multiple operating systems and their respective applications concurrently, without restrictions.
- Encapsulate an entire computing environment and move it between computers as easily as copying a file.
- Test the same application concurrently on different operating system configurations. For example, with different amounts of memory, different operating system revisions, or different system settings.
- Dedicate a virtual machine to run untrusted applications downloaded from the Internet. The VMware Virtual Platform technology guarantees that the virtual machines are isolated.
- Upgrade obsolete hardware and system software without losing compatibility. The legacy computer system and its applications are simply transferred to a virtual machine.
- Rely on a known stable hardware platform, defined by VMware Virtual Platform. Virtual machines configured for this stable hardware platform will correctly execute on any hardware that supports the virtual platform.
- Load-balance a collection of virtual machines on a scalable cluster of computer servers.

The VMware Virtual Platform Architecture

This section provides a brief description of the architecture and the components of the VMware Virtual Platform system.





One unique feature of VMware Virtual Platform is its ability to work in conjunction with an existing or *host* operating system, and yet retain the ability to arbitrarily configure the processor to minimize the virtualization overhead. VMware Virtual Platform solves this apparent contradiction through its *dual-mode personality*. As shown in Figure 2, VMware Virtual Platform is both a virtual machine monitor running directly on the hardware (left), and a normal application running on top of the host operating system (right). For maximum performance, the core components of VMware Virtual Platform use the virtual machine monitor personality and configure the hardware to match the requirements of the virtual machine. For maximum portability, the device-dependent portions of the system use the application personality and rely on the host operating system and its device-specific drivers.



Figure 3. The VMware Virtual Platform Architecture

🗊 **vm**ware

VMware Virtual Platform uses three software components to implement its special dual abilities. The following text describes each component, and Figure 3 shows how each one is related to the host operating system:

Virtual Platform (VP) Application – The application portion of Virtual Platform 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, it loads and runs the monitor. During execution, the monitor calls back to the application to access system resources. The application then calls the host operating system to access these resources.

Virtual Platform (VP) Monitor – The monitor portion of Virtual Platform runs in privileged mode directly on top of the hardware. The VP Monitor eliminates overheads through a virtualization technique known as *direct execution*. This technique allows the virtual machine to run directly in an environment where it is isolated from both the host operating system and from other virtual machines.

The VP 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 disjoint from the host operating system. The VP Monitor, however, 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.

Virtual Platform (VP) Driver – VMware Virtual Platform also includes a device driver downloaded into the host operating system. The VP Driver is specific to the host operating system and acts primarily as a communication gateway between the application and monitor portions of the system. The VP Driver also effectively hides the VP Monitor from the host operating system. As a result, VMware Virtual Platform appears like a normal process to the host operating system — a process that relies heavily on a specific device driver, but a normal process nevertheless.

VMware Virtual Platform always uses its application and monitor portions. Without its applications portion, VMware Virtual Platform would need to handle the diversity found in I/O devices, and without its monitor portion, VMware Virtual Platform would resemble a simulator, and would have the limited performance found in simulators. The third component (the driver portion) is necessary only if the host operating system is protected from user-level applications; it enables such systems to communicate with the VP Monitor. This division of labor into various components allows VMware Virtual Platform to achieve its twin goals of performance and portability.



Capabilities Enabled by VMware Virtual Platform

VMware Virtual Platform not only allows users to share the same hardware concurrently among many virtual machines; it also enables a number of features that are otherwise difficult or impossible to implement without the use of this technology:

Stable hardware platform - VMware Virtual Platform defines the virtual hardware interface on top of which the virtual machines run. As such, it is largely independent of the underlying hardware. For example, VMware Virtual Platform can include IDE disks, even if the underlying system only has SCSI disk controllers. The predictability of the hardware platform eliminates most operating system installation problems and allows virtual machines to be transported from system to system.

Only certain components of the virtual hardware platform, most notably the processor, are determined by the underlying hardware platform. For example, the x86 instructions introduced by the Pentium III are available to virtual machines only if the underlying processor is compatible with a Pentium III.

Encapsulation - VMware Virtual Platform totally encapsulates the execution of the virtual machine. For example, VMware Virtual Platform can encapsulate an entire session of a virtual machine within a transaction. At the end of the session, the user can commit the transaction and make all changes to the virtual disks permanent, or abort the transaction to revert the virtual machine to its original state.

This encapsulation will also enable users to checkpoint the state of a running virtual machine, so that they can later restore it multiple times, either on the same or different underlying computer.

Isolation - The Virtual Platform uses hardware protection mechanisms to isolate virtual machines from each other, and virtual machines from the host operating system. This isolation does not make any assumptions concerning the software that runs within the virtual machine. Even a rogue application or operating system is confined to the VMware Virtual Platform sandbox.

Resource Management

Users can flexibly configure their virtual machines by assigning them a collection of I/O devices. VMware Virtual Platform then relies on the host operating system to access the underlying devices. VMware Virtual Platform defers to the host operating system for the management of all system resources and appears as a normal application to the host operating system. The most common devices are listed here, along with a description of how VMware Virtual Platform interacts with them:

CPU – The virtualization of the x86 processor is handled entirely by the VP Monitor portion of the system. However, the execution of each virtual machine is subject to the scheduling policies of the host operating system. Each virtual machine uses one main process that runs its VP Application component (see the VP Application section of Figure 3). Users can control the priorities of the different virtual machines and of the applications running directly on the host operating system by controlling the priorities of the VMware Virtual Platform processes.

🗊 **vm**ware

VMware Virtual Platform takes advantage of the powerful management features of the x86 architecture to reduce the load on the underlying processor. For example, the halt instruction, which is found in the idle loop of most operating systems, suspends the execution of the VMware Virtual Platform process.

Memory – The memory needed to run the virtual machine is allocated through the address space of the application portion of VMware Virtual Platform. As such, the memory used by VMware Virtual Platform is subject to the standard allocation and paging policies of the host operating system. VMware Virtual Platform system does need to lock a subset of the pages in memory at any given time. The size of this subset is part of the configuration parameters of the virtual machine.

CD-ROM and floppy disks – Virtual machines can be configured to include removable media devices such as CD-ROM and floppy disks. VMware Virtual Platform system allows these devices to be mapped to the corresponding device of the underlying platform. It uses the host operating system to control access to the devices among the virtual machines.

Hard disks –VMware Virtual Platform can build and use the virtual hard disks of the virtual machine on any storage device accessible to the host operating system. This includes raw disks, disk partitions, and files in a local or network file system. This latter option allows users to create virtual machines without having to repartition or add disk resources. The system can even find free disk space on the system or strip a virtual disk image across multiple file systems to improve performance.

In addition, the disk storage manager forms the basis of the transactional features offered by VMware Virtual Platform. Users can run a virtual machine session as part of a transaction; this allows them to later commit or undo all the changes made to the virtual machine. A variant of this mechanism allows multiple virtual machines to share the same set of disks, thus greatly reducing the disk space demands of multiple virtual machines.

Network – VMware Virtual Platform emulates an Ethernet card, effectively giving the virtual machine an identity that is equal to the network identity 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. The use of standard distributed file system protocols such as NFS or SMB allows the tight integration of the virtual machines.

Display – VMware Virtual Platform includes two different mechanisms to present the virtual display. In the first mechanism, the virtual display appears as a window of the host graphical user interface. In the second mechanism, VMware Virtual Platform works in conjunction with the host operating system or its display server to give a single virtual machine exclusive access to the entire screen. A hot-key sequence allows the user to switch between modes of operation.

Input devices –VMware Virtual Platform 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. Clicking the window that runs the virtual machine gives it exclusive user input. A hot-key sequence releases the display input to the host operating system.

Miscellaneous – VMware Virtual Platform also virtualizes the other devices that complete a personal computer, such as sound cards, parallel ports and serial ports.



Compatibility and Performance

One of the main goals of VMware Virtual Platform is to be as transparent as possible to the virtual machines. With total transparency achieved, all operating systems and applications execute as correctly on VMware Virtual Platform as they would on a real hardware platform. VMware Virtual Platform currently supports the most common types of I/O devices, such as IDE disks, standard floppy drives, as well as an Ethernet card and soundcard. VMware Virtual Platform has been tested on a large number of different operating systems, including all Microsoft operating systems from DOS to Windows to NT, and a number of Unix x86 operating systems.

One potential incompatibility for VMware Virtual Platform is any operating system that relies on undocumented or undefined features of the PC hardware. VMware Virtual Platform can handle these incompatibilities if it knows about them. Untested operating systems that rely on unknown or undocumented features of the architecture may not work correctly.

Overall, our experience in developing VMware Virtual Platform has been that each operating system uses different features of the processor and of the system devices. However, our experience has also been that different applications running on the same operating system use the same architectural features, and that applications run correctly once the operating system has successfully booted.

Unlike API emulators such as WINE and WABI, VMware Virtual Platform runs the applications with their intended operating system. By design, VMware Virtual Platform guarantees the compatibility of all applications.

Another essential goal of VMware Virtual Platform is to eliminate performance overheads. Our use of direct-execution techniques to multiplex hardware resources has allowed us to achieve this goal. For example, the overhead of VMware Virtual Platform can be as low as 3%-5% for certain computation-intensive benchmarks.

A Better Approach to Computing

VMware Virtual Platform is an innovative and powerful approach to computing. With VMware Virtual Platform, operating systems, applications, and hardware can be "mixed and matched" in a variety of ways. This strategy frees users and organizations to exploit the full potential of their computing platforms, to work more flexibly and efficiently, and to dramatically reduce the cost of computing throughout their enterprise.