

# Windows 2000: A More Reliable OS Environment

By Eddie Ho

**The Windows 2000 operating system (OS) includes fundamental improvements to the kernel, addressing a number of issues that affect reliability and availability in prior versions of Windows NT. These new features help prevent system failures, improving both reliability and availability by decreasing maintenance and restart intervals after failures. This article describes the OS improvements found in Windows 2000.**

**P**roblematic kernel-mode software, memory conflicts, or corruption often cause system failures in Windows NT®. In the past, developers found it difficult to write and test software that reliably interacted with the Windows NT OS kernel while not interfering with memory being used by other software. To reduce the incidence of errant code in Windows® 2000, new kernel-mode code-testing tools let developers more easily create reliable drivers and other system components. In addition, architectural changes help protect system memory and core OS processes.

Availability also has been improved by decreasing the amount of time required for maintenance and for restarts after a failure. Windows 2000 introduces new administrative and maintenance features to address these post-failure issues. By reducing the number of tasks that require a system reboot, routine maintenance does not require as much downtime as in the past. And, in the event of a system failure, improved utilities make determining the cause of the problem and restarting the system faster.

The reliability and availability improvements in Windows 2000 mean that business users can rely on their systems to be up and running, resulting in improved satisfaction for system users and customers. For information technology (IT) users, these improvements provide a more robust

system architecture, fewer reboots, and more reliable application performance.

Major architectural changes have led to significant improvements in system stability, with a strong focus on OS protection and shared memory access. These changes include the following enhancements:

- **Kernel-mode write protection.** Helps prevent unsolicited write operations
- **OS file protection.** Prevents removal and replacement of system files
- **Driver signing.** Identifies drivers that have passed the Windows 2000 Hardware Quality Lab tests, and also submits a warning during an attempted installation of an unsigned driver

New tools also are available to help developers create high-quality and more reliable drivers. This includes pool tagging and guard pages to assist developers with debugging device drivers and tracking memory usage. The Driver Verifier and Device Path Exerciser utilities are test tools to better debug the driver in the kernel environment.

Another area of major improvement is availability. There has been a major reduction in the number of maintenance functions that require a system reboot. If

a reboot is needed, however, shutdown and restart times are also improved.

### Internal Architecture and Processes

The Windows 2000 OS provides the runtime environment in which applications execute. The OS contains a collection of small components that work together to perform different tasks. Each component provides a set of functions that act as an interface to the rest of the system. This collection of components provides the interface to access processor and all other hardware resources (see Figure 1). A process can be executed either in user mode or kernel mode.

### User Mode Provides Limited System Access

User mode is an application execution environment. Software in user mode operates in a nonprivileged state with limited access to system resources.

All applications run in protected subsystems in user mode. These subsystems can be either environmental systems that provide application programming interfaces (APIs) with different environments, or integral subsystems that provide system-level services.

### Kernel Mode Provides Unlimited System Access

Kernel mode processes can access all hardware and sensitive system resources. These processes are the building blocks for the OS and can be grouped as follows:

- **Executive.** This interface layer for all other components includes input/output (I/O), file management, virtual memory management, resource management, and inter-process functions.
- **Device drivers.** These support I/O and control the interface for a special hardware environment, such as a video adapter with building functions.
- **Hardware abstraction layer (HAL).** This layer protects the rest of the Windows 2000 executive from the specific hardware, making the OS compatible with multiple processor architectures.
- **Microkernel.** This provides the life-support for one or more processors. It includes scheduling, interrupt, exception dispatching, and CPU synchronization.

### Virtual Memory is Controlled by VMM

Windows 2000 supports 4 GB of virtual memory. The upper 2 GB are reserved for kernel-mode processes and the lower 2 GB are shared by kernel-mode and user-mode processes.

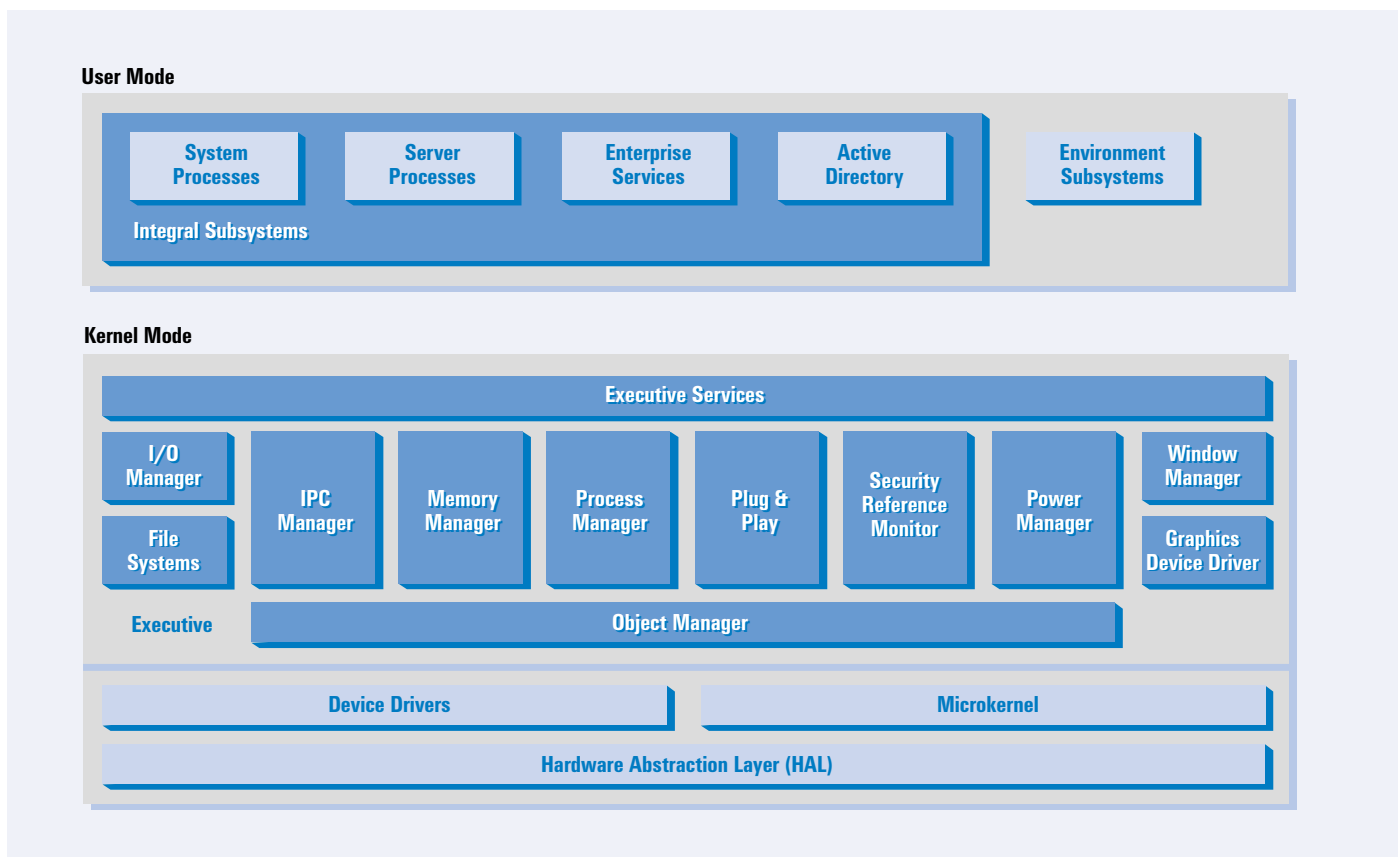


Figure 1. Windows 2000 Architecture

---

Addressability beyond 4 GB or 32 bit requires additional Intel processor architecture and API cooperation. This includes the physical addressing mode called Physical Address Extension (PAE) from Intel architecture and the additional operating system API.

Address Windowing Extensions (AWE) provide user applications with 32-bit virtual addressing to greater than 32-bit regions of physical memory. See <http://www.microsoft.com/HWDEV/NewPC/PAEdrv.htm> for a description of large memory addressability.

### **Reliability is Improved**

Windows 2000 improves reliability by providing early detection and prevention of improper memory-management practices in applications, kernel components, and device drivers. The OS is designed to gracefully manage application and system errors and exceptions without bringing down the server. Improvements are centered around developer debugging tools, with an emphasis on user-mode and kernel-mode protections.

### **Write Protection is Now Provided**

Windows 2000 adds write protection for code and read-only subsections of the kernel and device drivers. To provide this new protection, hardware memory mapping marks the memory pages containing code, assuring that they cannot be overwritten. This prevents kernel-mode software from corrupting other kernel areas inadvertently.

### **Windows File Protection Helps Prevent OS Corruption**

With Windows NT, installing software may overwrite shared system files such as dynamic link libraries (DLLs) and executable files. If not protected, the OS can fail due to corruption.

In Windows 2000, Windows File Protection verifies the source and version of a system file before it is installed. This prevents the replacement of protected system files. This service runs in the background and protects all future file installations.

Windows 2000 also implements digital signature technology to check whether the module to be installed is the correct version, and logs an entry into the event log.

### **Driver Signing Promotes Device Driver Quality**

Driver signing is designed to improve and promote the quality of device drivers in the Windows 2000 community. The digital signature of a device is associated with its individual driver package, and it identifies that the driver has been certified by Windows Hardware Quality Labs tests. This certification proves to users that the driver versions they installed are identical to the ones Microsoft® has tested. It notifies users if they attempt to install an

unsigned driver. This warning can be set to ignore or block within an enterprise. It also can be used to reduce total cost of ownership if unsigned driver installation is blocked.

### **Device Driver Code Development Tools Assist Debugging**

Memory management and driver debugging present a major challenge to the development community. The system might usually be partially up, and therefore error reporting and handling might not be operational. This can impact debugging productivity and result in poor code quality. With the overall focus on driver stability and reliability, Windows 2000 includes the following features and tools to assist debugging:

- **Pool tagging.** This debugging feature allows kernel developers to make all memory allocations to selected device drivers from a special pool, rather than from a shared system pool. To aid in debugging and to protect other code, the memory for the special pool is set to cause a system crash if a driver writes beyond its boundary. To assist in detecting memory leaks, pool tagging also lets developers put an extra tag on all allocations made from the shared pool to track tasks that make changes to memory.
- **Guard pages.** The guard pages tool creates boundaries for the special pool. These memory pages give the developer a means to find buffer-overwrite error. Using this feature, when a program requests a memory buffer, the OS puts the buffer at the edge of page memory. The OS then maps the next page as a guard page and denies access to the special pool. The system will generate a hardware error if the special pool attempts to write beyond its boundary.
- **Driver Verifier.** This set of checkpoints should be enabled during the debugging process. Driver Verifier can assist developers by exposing memory corruption from the allocation pool, writing to pageable resources while holding a spinlock or at a raised interrupt request level (IRQL), and handling memory-allocation errors.
- **Device Path Exerciser.** This layer of code exercises the device driver interface by calling the driver, synchronously or asynchronously, through various user-mode I/O interfaces and testing to observe driver reaction.

### **PageHeap Helps Uncover Memory-Access Errors**

Windows 2000 has a new tool, PageHeap, which can help developers find memory-access errors when they are working on user-mode code. Heap corruption is a common problem in software development. It typically occurs when an application allocates a block of heap memory of a given size, and then writes to memory addresses beyond the requested size of the heap block. When PageHeap is enabled for an application, all heap allocations in that application are placed in memory so that the end of the heap allocation is aligned with the end of a virtual page of memory. Any memory

---

access beyond this boundary will cause an immediate access violation within the application.

### **Availability Features Enhance Recovery and Diagnostics**

Windows 2000 reduces the amount of scheduled downtime for routine maintenance. It accomplishes this through a set of speedy recovery and diagnostic enhancements.

#### **Maintenance Downtime is Reduced**

Service Packs (SPs) can now be easily slipstreamed into the OS, which means customers do not have to reinstall SPs after installing new components. Also, many of the configuration changes that require Windows NT Server 4.0 to be rebooted no longer require rebooting in Windows 2000. These configuration changes include file system maintenance, hardware installation, network and communication protocol activation, memory management with PageFile, and installation of selected applications.

#### **Diagnostic Improvements Reduce Troubleshooting Time**

The diagnostic feature in Windows 2000 has undergone major improvements. This feature assists users with troubleshooting system errors in the following ways:

- **Kernel crash dumps.** In addition to full-memory crash dumps, the kernel dumps allow faster reboots for systems with large amounts of physical memory. When a Windows NT system fails, a snapshot of the memory is copied to disk; this can take as long as 45 minutes. Kernel dumps can decrease both the size and time required to perform the dump by up to 80 percent.
- **Faster CHKDSK.** The performance of this command, which checks the hard disk for errors, has been greatly improved.

#### **Recovery and Restart Time is Improved**

Several improvements in Windows 2000 help reduce the amount of time it takes to recover from a system failure and restart the OS. These improvements are found in the following features:

- **Recovery Console.** This is useful for repairing a system by copying a file from a floppy disk or CD-ROM to the hard drive, or for reconfiguring a start-up service. Using the console, users can start and stop services, format drives, and perform many administrative tasks.
- **Safe Mode boot.** Windows 2000 uses default hardware settings that allow administrators to troubleshoot systems or change the default settings of a driver.
- **Kill process tree.** This feature gives the Task Manager the power to stop not only a single process, but also any processes created by the parent process.
- **Recoverable file system.** A system log of Windows NT File System (NTFS) operations allows for fast recovery in the event of a disk failure. This can reduce downtime, since the file system can quickly return to a functioning state.

- **Automatic restart.** This can shorten unscheduled downtime and provide maximum unattended uptime.
- **Internet Information Services (IIS) reliable restart.** This one-step process starts all IIS services. It is started via Microsoft Management Console by using a command-line application.

#### **Storage Management is Enhanced**

Windows 2000 provides storage enhancement to help administrators maintain sufficient free disk space with minimal effort. For example, you can perform online tasks such as creating, extending, or mirroring a volume without shutting down the system. Key management features include the following:

- **Remote Storage Services (RSS).** This set of tools monitors the amount of space available on a local hard disk and usage level, and it can remove local data that has been copied to remote storage to make room for more free disk space.
- **Removable Storage Management (RSM).** This presents a common interface to robotic media and media libraries.
- **Disk quotas.** These are supported for monitoring and limiting disk space use on NTFS volumes.
- **Dynamic volume management.** This allows online administrative tasks to be performed without shutting down the system.

#### **Clustering Services are Standard**

Clustering is a collective environment to increase data and application availability. Windows 2000 products in the server family provide system services for server clustering as a standard part of Windows 2000 Server for a two-node environment. This can reduce IT costs by keeping systems running in the event of a single system failure. Clustering can address both planned downtime and unplanned system outages.

#### **Windows 2000 is Business Ready and Reliable in 2000**

Windows 2000 is ready for business and has become a more reliable OS than Windows NT. Its enhanced features ease driver development, which achieves a more reliable programming model, faster recovery during start-up, and easier administration with minimum interruptions. ◆

*Eddie Ho (eddie\_ho@dell.com) is a program strategist in the Dell Enterprise Systems group. He has worked in computer system design, architecture, and development in various OS environments with leading companies, including IBM and Dell. He is a frequent contributor to industry journals and is the author of a thin client book. He has a B.S. in Computer Science from the University of Wisconsin and an M.S. in Computer Science from North Dakota State University. He also is certified as a Microsoft Systems Engineer and Trainer.*