Cahners
MICROPROCESSOR REPORT
www.MPRonline.com
THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

# Microsoft Opens Windows Plans

## *Future Operating Systems to Be More Capable and Reliable*

***By Peter N. Glaskowsky {5/22/00-01}***

It is a truism in the PC industry that Microsoft always takes three tries to get anything right. The process mirrors the company's alleged business model—embrace, extend, and extinguish (the competition). This model is based in part on the history of the Windows operating system, which finally became usable in its Version 3.0 incarnation, and popular in the form of Windows 3.1. Microsoft's Windows NT operating system is in its third iteration, now called Windows 2000, and it seems to be following the rule. Windows 2000 is greatly superior to its predecessor, Windows NT 4.0, but it looks as if Microsoft will need to tweak Win2K a little before it can reach its full potential for consumers.

At WinHEC last month, Microsoft revealed additional details about its plans for future Windows operating systems. These details indicate that for most users, the Windows to wait for is the one code-named Whistler. Whistler will be the next major version of Windows 2000 and will include a configuration designed for consumers. Whistler will include some significant improvements to Windows 2000 while improving application and user-interface compatibility with Windows 98.

### Windows Me Offers Improvements for All

Microsoft says the Windows Millennium Edition operating system, aka Windows Me, will be the last Microsoft OS based on the Windows core. Windows Me is something of a transition product between Windows 98 and Whistler. Win98 still contains remnants of DOS to support software that must run in 8086 "real mode," while Whistler's WinNT core runs exclusively in 80386-style "protected mode."

Microsoft, seeking to eliminate the delays and diminished reliability associated with real-mode processing in Windows 98, has eliminated real-mode software support from Windows Me. The new OS boots directly into protected mode, bypassing the old DOS startup files CONFIG.SYS and AUTOEXEC.BAT. Device drivers and applications that rely on these startup files will not work under Windows Me.

The elimination of real-mode support from this consumer OS will help prepare software and hardware makers, as well as end users, for the arrival of Whistler. Though the change will surely inconvenience those who still rely on real-mode DOS software, Windows 98 will remain available for some time to support legacy hardware and software.

Windows Me borrows from Windows 2000 a feature that should benefit all users. Known as "system file protection" (SFP), the feature protects about 800 critical system files from being changed or removed by errant software-installation programs, misguided users, or viruses. When the operating system detects an attempt to modify such a file, it saves a copy of the original file and allows the modification, then checks the result against a database of authorized changes. If the change was not authorized, the original file is restored from the backup.

Though it might seem to be easier to write-protect critical system files, this solution would cause many software installers to fail. Microsoft believes that in many cases the offending installer is attempting to replace a newer file with an older one, even though the newer file would work equally well with the new application.

SFP does not solve all configuration problems, nor can it completely eliminate the common Windows condition known as "DLL Hell," caused by applications with conflicting requirements for specific versions of common dynamically linked libraries. SFP also does not protect against file corruption that happens outside the normal file-access mechanisms. It cannot protect files that change when programs are run, data files, or files that are not listed in its database. Drivers are also generally excluded from SFP, since they can be protected by digital signatures that show the driver has been certified by Microsoft. Ordinarily, only a signed driver can replace another signed driver.

Despite all these limitations, SFP is likely to yield a marked improvement in the overall reliability of Windows systems, especially those that use only new applications designed to be aware of SFP. The inclusion of SFP in Windows Me should also smooth the transition to Whistler, which has similar technology based on the "Windows file protection" (WFP) feature of Windows 2000.

### Other Features Enhance Ease of Use

For future operating systems, Microsoft plans a variety of other features that are meant to make PCs easier to use. The company is working with Adaptec, Compaq, Philips, and Sony to provide direct OS support for CD-RW (CD rewritable) drives. The group is developing an open standard, code-named Mt. Rainier, that should make CD-RW drives as easy to use as floppy drives. The standard, described at www.mt-rainier.org (this site was not online at the time of publication but may be available by the time you read this), will allow CD-RW drives to replace both floppy and CD-ROM drives, resulting in less expensive PCs. Version 1.0 of the new standard is planned for incorporation into the CD-ROM "Orange Book" by midyear.

The Mt. Rainier standard will require new CD-RW drives to manage media defects in hardware, rather than using software defect-management like that currently used for floppy drives and some other removable-media storage systems. Mt. Rainier CD-RW discs will be readable on standard CD-ROM drives, with the use of a software driver, but the need for new CD-RW drives to create these discs will surely annoy those who have already purchased drives that don't support the new standard. Data from Philips suggest the installed base of CD-RW drives will be more than 60 million units by 1H01, when the Mt. Rainier standard will first be incorporated into new CD-RW drives. Philips and Sony stand to sell quite a few new drives to replace those made obsolete by the standard they are helping to develop.

### Video and Graphics Support Enhanced

Future Windows OSs will also boast improved multimedia support, especially for professional users. Microsoft is developing DirectVA, an application-programming interface (API) that will provide a standard method for application software to access hardware video-codec accelerators such as those from C-Cube (see *MPR 11/16/98-msb*, "C-Cube Offers Affordable MPEG-2 Codec") and other companies.

Microsoft previously relied on hardware and software vendors to develop their own APIs for these chips; the result has been expensive proprietary add-in cards and inflexible applications. On the Macintosh platform, Apple's QuickTime API defines standard interfaces for software and hardware, making the Mac generally superior to Windows machines for video editing and playback. Apparently, Microsoft hopes DirectVA will make Windows more competitive in this market.

Though the first version of DirectVA supports only playback, Microsoft plans to add authoring support to the API in later versions. This enhancement will surely be welcomed by professional video editors as well as by consumers who would like to use their PCs to edit home movies. In the meantime, DirectVA will make it easier to design video-decoding chips and integrate them into Windows systems, producing an immediate benefit for consumers.

One feature, aimed directly at professionals, is a new RGB color model that is greatly superior to the sRGB model described at WinHEC last year (see *MPR 5/10/99-05*, "Microsoft Updates Windows Roadmap"). Where sRGB defines a smaller range of colors than can be recorded by cameras, displayed on monitors, or printed on color inkjet printers, the new sRGB64 defines a dramatically larger color space than any real-world device could ever support. Instead of defining a common subset to improve consistency among various devices, sRGB64 provides a way to identify every color the human eye can possibly see—along with some "colors" that can't possibly exist. The sRGB64 space is large and precise enough to distinguish among all Pantone and Munsell colors, which should simplify the use of these "spot" colors in computer applications.

As its name implies, the sRGB64 specification, proposed by Microsoft and HP to the International Electrotechnical Commission (IEC) last September, uses 64 bits of precision to define color values. Red, green, and blue each get 16 bits; the remaining 16 bits are used to define an alpha (transparency) value. Each of these four values includes a sign bit, allowing negative colors to be specified. Though negative colors cannot be used in real life, they are a useful abstraction when blending or comparing real colors; they also prevent clipping and saturation. Figure 1 shows the sRGB64 space plotted against the CIE standard model of human color perception and the smaller sRGB space.

Another key difference between sRGB and sRGB64 is the way color values are defined. The sRGB color space is a perceptual space, so called because it is based on the way humans perceive color. Numerically similar distances across the plane of the sRGB space represent visually similar changes. The sRGB64 space is a physical space, representing the physics of lighting. Perceptual spaces can represent more-subjective colors with a smaller dynamic range, allowing them to better model the effective behavior of human-interface devices such as monitors and printers, but physical spaces are better suited to the mathematical calculations involved in 3D rendering and similar tasks.

Microsoft and HP expect their proposal to be ratified by the end of 2000, after which Microsoft plans to spend about two years integrating the sRGB64 model (which will likely be renamed as part of the IEC standardization process) into future Windows operating systems. The sRGB64 color model will define the internal representation of colors in applications, and these values will be converted to sRGB or other perceptual spaces before images are displayed or printed.

One open question is how sRGB64 will be represented in image files. No current image-file format (such as JPEG or TIFF) is compatible with sRGB64. Few file formats are formally associated with any color space at all, a fact that contributes to the common color-matching problems experienced by consumers trying to view and print images from digital cameras and other sources.

### DirectX 8.0 Shows Microsoft's 3D Vision

The graphics chips introduced at WinHEC by ATI and NVIDIA (see *MPR 5/22/00-02*, "WinHEC Sees Great 3D") include features not supported by DirectX 7, the current version of Microsoft's multimedia API. DirectX 8, due out later this year, will be the first to give software developers access to these new features, as well as to several more that will be included in the next generation of graphics accelerators from these chip makers and others.

The 3D pipeline in DirectX 8, shown in Figure 2, is much more powerful and flexible than that found in DirectX 7. The new pipeline allows 3D models to be defined by curved surfaces; supports programmable transform, lighting, and rendering engines; and provides for an optional full-scene antialiasing step.

The curved-surface feature in the new pipeline allows applications to define shapes using three-dimensional polynomial splines (both Bezier and B-splines) as well as domain patches, which are triangular or quadrilateral shapes with a 3D profile. Objects composed of such shapes occupy less memory and disk space and are easier to transform in regular ways. Because all graphics chips must have conventional 2D polygons for the final rendering stage, curved surfaces are tessellated (subdivided) into polygons by the graphics chip, if it is equipped with a tessellation engine, or by the host processor.

DirectX 8 (DX8) supports several tessellation modes, selected by the application. Surfaces close to or facing the viewer can be tessellated more densely to produce a better visual result. The API also defines an adaptive tessellation mechanism that prevents sudden changes in the way surfaces are subdivided as the viewpoint changes.

The vertex-shading step introduces into the DX8 pipeline a feature formerly found only in high-end scientific-visualization systems such as those from SGI—programmable geometry transform and lighting calculations. Applications can load vertex-shader programs (up to 128 instructions long) into the graphics chip; these programs process each polygon vertex that passes through the pipeline.

The instruction set, defined by Microsoft in cooperation with major graphics-chip vendors, was developed specifically for this application. As a result of this optimization, some instructions are more complex than those commonly found in mainstream CPUs. It takes just one instruction, for

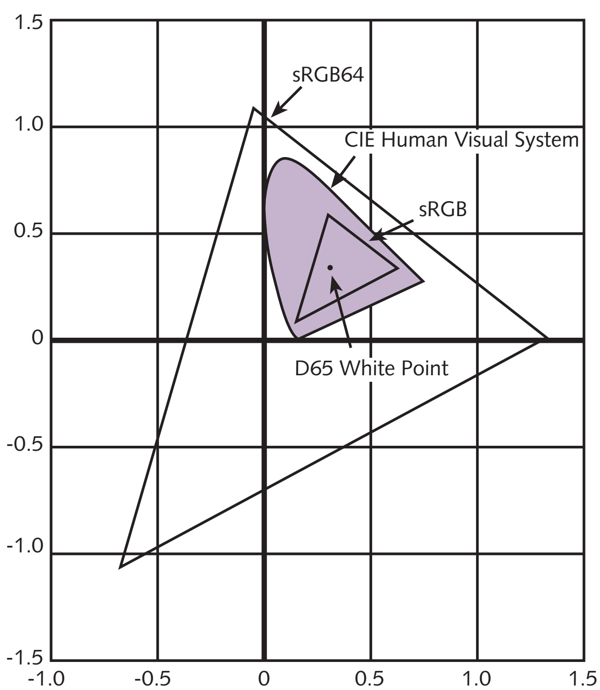


**Figure 1.** The sRGB64 color space encompasses a much larger range of possible (and impossible) colors than the human eye can see. The sRGB space, represented by the smaller triangle, cannot represent all visible colors. The D65 white point represents a color temperature of 6,500° K; it is commonly used to define the white point of computer monitors.

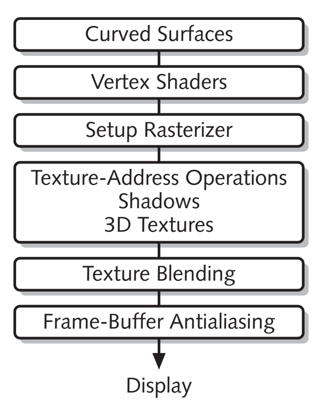


**Figure 2.** The new graphics pipeline in Microsoft's DirectX 8 supports high-level object geometry and programmable graphics chips such as those recently introduced by ATI and NVIDIA.

example, to perform a four-element vector multiply-accumulate operation. Values in this step are represented using single-precision floating-point numbers.

The new vertex-shading capability can offload complex object-manipulation tasks from the host processor. Rather than using the host processor to bend the joints of a complex model of a human body, DX8 applications will be able to use the same model from frame to frame, changing joint positions only as needed.

The setup-rasterizer portion of the DX8 pipeline shown in Figure 2 encompasses most of the functionality of PC graphics chips prior to 1999—an indication of how far PC 3D technology has come in the past two years.

Texture-address operations convert a set of texture coordinates (which define a three-dimensional position and a scale factor) into a color. Like geometry calculations, DX8 allows texture-address operations to be specified by small programs. These calculations are also performed using single-precision floating-point math.

The same stage of the pipeline is used to calculate the position of shadows and three-dimensional textures. Both of these scene elements are also new to DX8. Shadows modify the lighting of objects, based on the relative positions of other objects and light sources. The display of three-dimensional textures is conceptually related to volume rendering, found on Mitsubishi Electric's unique VolumePro (see *MPR 10/25/99-08*, "VolumePro 1000 Expands 3D Vision"). Unlike 2D textures that define only the surface appearance of 3D objects, 3D textures are defined over a 3D region. If a portion of an object is clipped off, a 3D texture can define the exposed surface. A program could use a 3D texture to define the internal grain of a block of wood, for example, then allow the user to carve the virtual block into a shape. (Perhaps more likely are 3D games that reveal internal organs when an adversary is shot.) DX8 defines a new format for compressed 3D textures.

Texture blending is used to combine the various results of previous stages of the pipeline. This process is also programmable. ATI's Radeon 256 and NVIDIA's GeForce2 GTS both implement this portion of the DX8 3D pipeline and can use it to create realistic visual effects, such as bumpy reflective surfaces and brushed-metallic finishes.

The final step in the DX8 pipeline is a flexible anti-aliasing mechanism. In addition to the basic function of smoothing edges between polygons, this stage is used to add special effects, such as motion blur and depth of field, as found in 3dfx's VSA-100 (see *MPR 12/6/99-02*, "Napalm Ignites Graphics Market").

## GDI+ Shows Microsoft's Goals for GUIs

Somewhat further in the future of Microsoft's operating systems lies GDI+, formerly known as GDI2K and also discussed at WinHEC '99. This year's update on GDI+ had little to say about 3D user interfaces, once the focus of the GDI+ work, other than to identify some basic goals for the 3D features that will eventually be included in GDI+.

Most of the progress Microsoft has made to date on GDI+ is related to 2D rendering and compositing—combining multiple independent hardware windows at the time the screen is drawn to the display. The goal of the new compositing architecture is to eliminate clipping tests for overlapping windows when background windows are updated, as well as to eliminate the need to redraw windows when they are moved around on the desktop. To achieve these goals, each open window will be assigned to a separate virtual screen with its own memory space on the graphics card. As far as each task is concerned, its window will be the only window on the screen and will always be up to date. Since cards will have some physical limit to the number of hardware windows they can support, GDI+ can fall back on software compositing—the same method used to create overlapping windows on PCs today.

Each window can contain 2D or 3D, or a mixture of the two kinds of graphics. DirectX already supports intermixed 2D and 3D graphics on the same display, but 2D and 3D graphics use different APIs, so it is difficult to synchronize the two types of graphics. In DX8, Microsoft will phase out the older DirectDraw 2D API and accept 2D-drawing requests through the same API stack used for 3D. The GDI+ hardware-windowing mechanism, however, will be required to realize the full potential of this simplification. Also, graphics subsystems will require more memory to accommodate multiple hardware windows—and more still if sRGB64 frame buffers are supported in hardware.

Microsoft is clearly working hard to make its operating systems more powerful, more capable, and more reliable. Whatever the outcome of the Microsoft antitrust litigation, we hope the company will continue this good work. ◇