# Finding a platform, part 2

## The view from the top end

*If your business needs a high-end GIS, what operating system should you choose? Last month, we detailed what's on offer for desktop GIS. This month, we profile UNIX and Windows NT and predict the future of such high-end operators. ROBERT BUCKLEY reports.*

There are few offices left in the UK which have just one computer. Even small companies usually provide their employees, from the managing director down, with their own machines. And without a corporate network to share data—particularly geographic data—a company's efficiency can be substantially reduced and significant mistakes made. Usually a server, a machine dedicated to providing functions for other computers or "clients", is required to administer the network and store shared data.

A computer's operating system (OS) is the software required for it to actually work. Without it, the computer's hardware would be unusable and even the simplest task would be impossible. However, the OS required by a server, where speed and stability rather than ease of use and graphic capabilities are the priorities, is usually very different from that required for a desktop system—or at least it used to be.

### UNIX

UNIX was probably the first "modern" operating system. This means it includes pre-emptive multitasking, multithreading and protected memory (see the glossary on page 30 for an explanation of these terms). Its name is a pun on Multics, an OS from the same designers. It remains, 27 years after its inception, the most popular system for servers. It is extraordinarily stable and versatile, and it can cope with even the most powerful and varied hardware.

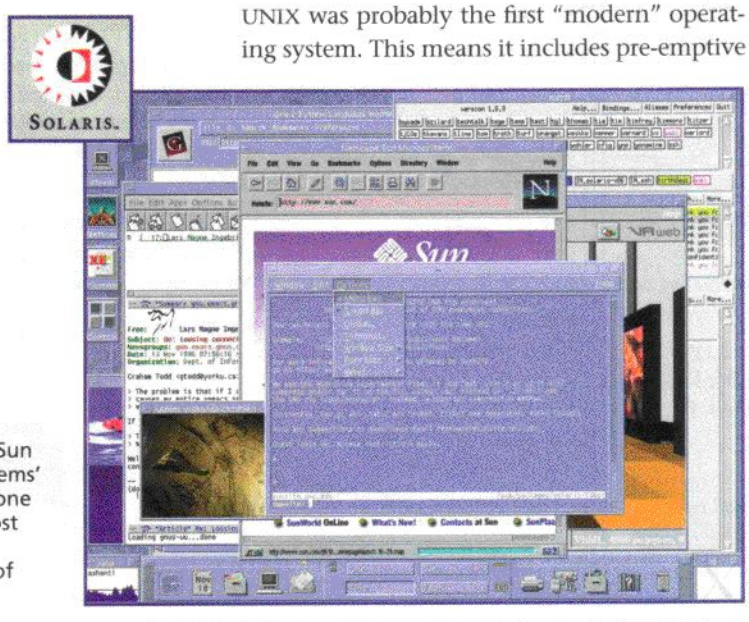


Figure 1: Sun Microsystems' Solaris is one of the most popular varieties of UNIX.

UNIX, however, is not a single entity. There are probably as many dialects of UNIX as there are programming languages. The most popular varieties include Solaris from Sun Microsystems (see figure 1), HP-UX from Hewlett-Packard, NEXTSTEP from NeXT software, IRIX from Silicon Graphics, the 64-bit Digital UNIX from Digital, and Mach from Carnegie Mellon University. As a result, you can't buy a "GIS for UNIX"—you have to buy a program that matches both the hardware and the software of your system (such as ERDAS IMAGINE for IRIX for Silicon Graphics workstations). Most varieties of UNIX have emulators for other operating systems; Solaris users can use Windows programs, for instance, if they buy something called WABI (see figure 2).

A dialect of UNIX can be found to run on almost any processor from Motorola's 68020 to Intel's Pentium Pro, from Digital's Alpha to Sun Microsystems' SPARC. There are 8-, 16-, 32- and 64-bit systems. 32-bit systems in particular have been around for a long time. As a result, for very high-end GIS use, there is still no competition to UNIX's dominance.

With a UNIX-based server, there is only one real option for GIS configuration—the software is run on the server with the client computers acting as little more than network computers. To run individual GIS packages on client machines with the server acting as a data storage system would be a waste of the server's power—why have a massively parallel, 64-bit computer acting as a disk drive when you can use it for processing instead? The result is that a UNIX-driven server almost necessitates that client machines use UNIX interfaces.

However, UNIX is a command-line based OS with no built-in graphics capabilities. For GIS use, a graphical user interface (GUI) is almost mandatory, and there are several X Windows systems which can be used to add this capability. Of course, with no standard, UNIX has a steep learning curve.

UNIX has always been expensive: a trained

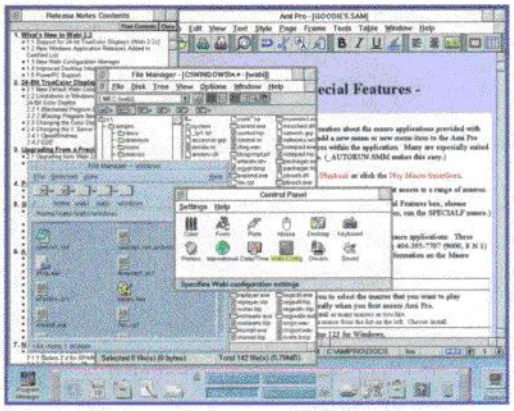

Figure 2: WABI is an emulator that enables Solaris users to use Windows programs.

system manager is vital, workstations have never been as cheap as modern PCs, and the costs of buying and maintaining hardware are high. There are few office applications for UNIX—Microsoft Office is unavailable, and only the forthcoming Java-based Corel Office might be able to provide some lifeline to isolated UNIX users. Indeed, while most of the major GIS packages are available for varieties of UNIX, small firms tend not to produce UNIX software. Large businesses tend to buy customised systems instead, specially tailored to their particular needs. This again pushes up costs and, as a result, such systems are mainly the preserve of government and ex-government sectors, research institutes and higher-education bodies.

So who should buy a UNIX system for GIS? If a company needs a high-end GIS, can afford the high prices associated with the system and doesn't mind dedicating technical staff and computers to GIS alone, then UNIX is worth the investment since it will outperform its rivals. The small or medium-sized business should only consider using UNIX as a network server and should look elsewhere for GIS.

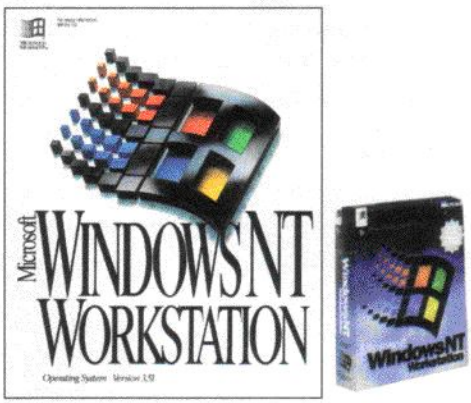


## Windows NT

Windows and Windows 95 are both based on DOS, a primitive operating system whose eccentric architecture was originally designed for the simple computers of the late 70s. As a result, no one with more than two users and no money would consider running a server on a Windows box. Realising this, Microsoft designed a new system called Windows NT ("new technology") which would correct these faults.

Windows NT has many advantages over UNIX. The vast majority of GIS and desktop mapping packages can run under it since it is capable of using most OS/2, DOS, Windows and Windows 95 programs (see figure 3). In fact, as long as a Windows 95 program is written correctly (technically, it has to obey the rules of the Win32 API) it will be a Windows NT program as well. Certain programs designed for UNIX can run under NT with minimal rewriting. It has a GUI, which, although by no means the best on the market, is considerably better than most varieties of X Windows. It can exchange data between applications in the same way as other Windows operating systems. It can understand more network protocols than UNIX, which confines itself almost exclusively to TCP/IP and related protocols. It can also run effectively on standard PCs (48Mb is recommended for an easy life) as both a workstation OS and a server OS. Instead of the great distinction between client and server that exists in the UNIX world, Windows NT offers a situation in which the client and the server are on almost equal terms, differentiated only by hardware capabilities. Processing is therefore split more or less equally and may be confined entirely to the client computers, the server only dictating access to centrally stored data.

Figure 3: The interfaces for Windows 95 and Windows NT are the same, although some of the individual features vary.

NT is slower than UNIX, less secure, and doesn't cope with more than a few user-connections at all well. Its functions don't scale to high-end systems easily, and the NT equivalent of multiprocessing at this level—"clustering"—is a poor substitute. Most IT managers therefore balk at using an NT server for large networks. There will be no 64-bit version of NT for quite some time, and Intel users will have to wait even longer.

A medium-sized business would do better to consider the more traditional NetWare server for its network and install NT on its workstations. The extra stability afforded by NT will be worth the investment in the long term.

At the moment, Windows NT is easily the best option for a small business's server. For PCs, it is also the best choice for an operating

## Glossary

**64-bit systems**—Being able to use a 64-bit processor's addressing capabilities is of greatest use in handling large files (of the order of 4Gb). It is therefore of greater interest to GIS and database users than to other IT sectors.

**Java**—a language designed to run on any operating system with an appropriate *virtual machine*. It is seen by many as the means to end the divides between operating systems in terms of program availability.

**Multiprocessing**—Modern computers often have more than one processor. To take advantage of this extra computational power, a computer must be capable of multiprocessing. The most efficient form of this is "symmetric multiprocessing", in which all the processors are used equally for calculations. Older systems which only have asymmetric multiprocessing capabilities use one chip to decide which processor, including itself, should be used for a calculation. While Windows NT has symmetric multiprocessing capabilities, these are not good enough for large numbers of powerful processors. The recommended configuration for greater processing ability is to link Windows NT servers together in a "cluster".

**Multitasking**—the ability to run more than one program at a time. This OS feature comes in two main varieties: *cooperative* and *pre-emptive.* With the former, programs are written so that they periodically allow other software to run. With the latter, the operating system is coded so that it allocates processing time to all the programs running. Windows 95 uses both. Programs written specifically for Windows 95 ("32-bit applications") can run pre-emptively with other such applications. When Windows 3.1 ("16-bit applications") or DOS programs are being run as well, only cooperative multitasking is used.

**Multithreading**—A "thread" is a portion of time allotted to a program by the operating system in which it can use the computer's processor or processors. If multithreading is available, a program can have more than one thread allocated to it, enabling it to perform several functions at the same time.

**Plug-and-play**—the ability of an operating system to recognise when a piece of hardware has been connected to it and to use it. A piece of software called a "driver" might have to be installed to use certain capabilities of the hardware.

**Protected memory**—When a program runs on a computer it requires a certain amount of RAM allocated to it by the operating system. If that memory isn't "protected", an incorrectly written application can use that program's memory, resulting in either the program or the whole machine crashing.

**RISC (Reduced Instruction Set Chip)**—These execute a smaller number of instructions than normal chips but do so more efficiently. All high-end processors are RISC chips because of the greater speed, smaller electricity requirements and greater processing power than standard chips. The newer Intel chips now incorporate RISC elements but maintain the older architecture overall in order to preserve compatibility with Windows operating systems and programs.

**Source code**—the original instructions used by programmers to create software. In order for source code to be turned into a program, it has to be compiled (translated into the computer's code), either before being run as in the case of C and C++ programs or while the program is being run—BASIC programs are an example of this.

**Virtual machine**—In order to run programs designed for other operating systems, an OS must have a *virtual machine* or an *emulator*. An emulator mimics the capabilities of a processor or operating system while the virtual machine is the actual operating system running as part of, or on top of, the computer's operating system. Windows NT has OS/2, Windows 3.1 and DOS virtual machines built-in. In general, emulators are slower than virtual machines.

**X-Windows**—A graphical user interface designed for UNIX systems. There are several such systems including Motif and OpenWIN. In order to qualify as an X Windows system, an interface has to be submitted to the X Open Consortium, the current owners of the UNIX trademark. This consortium also decides which systems qualify as UNIX systems.

system, although in common with UNIX, it has no plug-and-play capabilities.

There is also no guarantee that a piece of hardware that works under Windows 95 will work under Windows NT. In order to ensure the stability of the operating system, Windows NT's executive (the "brain" of the OS) stops programs from accessing hardware directly—all such calls must be made through the executive. This extra layer of operating system for programs to negotiate means that Windows NT is considerably slower than Windows 95. Windows NT therefore requires Pentium or Pentium Pro processors by necessity.

A small company should think twice about choosing Windows NT as its workstation operating system. The greater cost of the hardware, the operating system and the software means that the company will find it almost prohibitively expensive to install across anything more than a small user base.

## The future

Although the GIS market is going to be dominated almost entirely by Windows 95 (and its successor Windows 97, due out in 1998) well into the next century, Windows NT is already growing in popularity as a basis for GIS packages. With the interchangeability of Windows 95 and Windows NT applications increasing NT's software base daily, it's only a matter of time before NT becomes the favoured platform for the medium-sized business's GIS use. Microsoft's intended policy of moving people towards NT 5.0 from Windows 98 will eventually push the small business towards NT as

well. Judging by its early decision to port its UNIX software to NT, this is something that Intergraph, the major NT-based GIS vendor, appears to believe.

UNIX, while it will lose some custom to NT, is destined to remain as popular for large-scale GIS use as always. It will also receive new customers from central and local government as they upgrade from legacy systems such as VMS (Windows NT is basically a rewrite and updating of VMS for small-scale Intel-based computers and networks). This dichotomy between the small-scale Windows NT/Intel-based GIS and the large-scale UNIX/RISC-based GIS will continue to exist until Windows NT can be improved to such an extent that it can take on UNIX's greater capabilities.

One attempt, by Hewlett-Packard and Intel, to bridge the gap with a new chip, the Merced, stands some chance of working. The Merced will use portions of Intel chip and Hewlett-Packard's PA-RISC chip design, in conjunction with a new variety of UNIX and various developer tools, to allow applications designed to run on the new UNIX to run on NT as well. The Merced will not now be available until 1999.

In 1998, however, a fully object-oriented OS capable of running UNIX, MacOS and probably Windows/Windows 95/Windows NT software is to be launched. It will incorporate features of the best graphical user interface

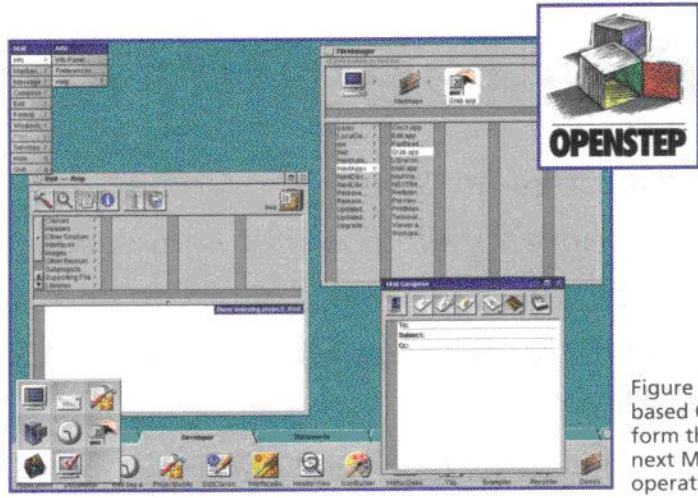


Figure 4: The UNIX-based OPENSTEP will form the basis for the next Macintosh operating system.

available with the power and stability of UNIX (see figure 4), and will be able to run on both Intel and PowerPC chips. This system is the next Macintosh operating system. Whether it will be able to bridge the gap between UNIX and NT, recapture Apple's old market share, and more importantly give greater power and ease of use to GIS packages, remains to be seen.

## Further reading

*Inside Windows NT* by Helen Custer (Microsoft Press, 1992), with a foreword by NT and VMS designer David Cutler, offers a good, if technical, introduction to modern operating system theory and both Windows NT and UNIX systems.

***ROBERT BUCKLEY*** *is associate editor for* Mapping Awareness *and IT manager for GeoInformation International. He can be contacted at: GeoInformation International, 307 Cambridge Science Park, Milton Road, Cambridge CB4 4ZD. Tel: 01223 423020; Fax: 01223 425787; Email: robertb@geoinfo.co.uk*