The Solaris 8 Operating Environment and Microsoft Windows 2000 Race for Control of Web Infrastructures

January 2000

Prepared for Sun by:

D.H. Brown Associates, Inc.
222 Grace Church Street
Port Chester, NY 10573
(914) 937-4302
Fax: (914) 937-2485

TABLE OF CONTENTS

The Solaris 8 Operating Environment and Microsoft Windows 2000 Race for Control of Web Infrastructures

OVERVIEW

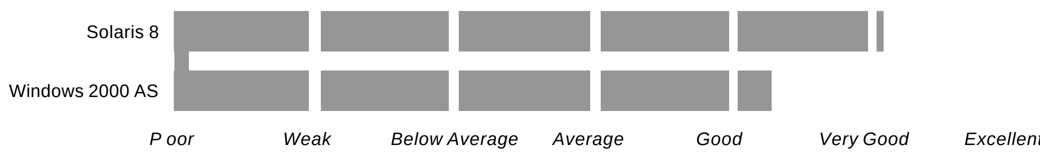

FIGURE 1: Solaris 8 Functions vs. Windows 2000 Advanced Server

After decades of proclaiming that the network was in fact the computer, Sun now focuses on expanding its definition to the web itself. Sun has taken the lead in promoting an architecture that is totally oriented around Internet protocols and that stresses the role of websites for providing a vast and diverse array of services which follow a utility model.

The Solaris 8 Operating Environment[1] represents the latest iteration in the system-software component of Sun's effort to achieve those goals. Having faced down Microsoft and held its ground so far, Sun now redoubles its efforts to use its operating-system technology as a competitive differentiator. With the Solaris 8 platform, Sun maintains its lead over the first version of Microsoft's Windows 2000, which still falls short of key Solaris software functions related to reliability, scalability, and manageability.

Sun will use Solaris 8 to promote its vision of Internet computing, even as Microsoft steps up its own assault on high-end computing with Windows 2000, the latest release of its Windows NT technology. From an operational standpoint, Sun's vision represents a mighty argument in favor of outsourcing, encouraging IT design to rethink all computational services as commodities that are simply tapped from Internet points-of-presence under optimal economic terms.

The idea of centralized infrastructures may recall some unpleasant memories of an earlier era, in which users were at the mercy of unresponsive IT departments. It also resonates with contemporary circumstances, however, since web browsers have become the primary entry point for a growing number of day-to-day computing activities. Developers have increasingly begun to explore possibilities for segmenting application designs along web boundaries – i.e., shifting application logic from clients to web servers and implementing user interfaces with HTML-based presentation layers.

[1] For the purposes of this document, the Solaris 8 Operating Environment will be referred to as Solaris 8.

This approach delivers several immediate benefits, including client independence – since web access is supported by a wide variety of platforms – and geographic independence – i.e., the ability to access both applications and data from any location. Further, organizations that have long struggled to maintain huge networks of PCs now sense they can potentially use a web-based application approach to ease their management burden. By centralizing applications and simplifying clients as much as possible, they can enable greater efficiency through increased economies of scale, managed either by in-house IT operations or a new breed of Application Service Providers (ASPs).

Simultaneously, the role of e-commerce has been growing, and a profusion of other services publicly available on the Internet has been emerging. These developments have created a need for vastly more complex applications that are deployed on servers and can be accessed reliably by huge numbers of globally dispersed web clients.

All these trends demand system software that can maximize reliability and scalability while providing as many network entry points as possible to native functions and services. Sun has put itself into a unique position to respond by relentlessly maintaining its focus on UNIX and extending its role as a leading UNIX visionary at a time when virtually all other UNIX vendors had adopted strategic roles for Windows NT.

Some observers had criticized Sun for its position, predicting that Windows NT would sweep UNIX out of many key low-end server markets targeted by Sun's Solaris Operating Environment. However, functional gains for Windows NT and Intel x86 servers materialized more slowly than expected. Microsoft encountered significant delays shipping the next release of its Windows NT technology, while the limitations in previous versions began to be felt ever more severely. Intel x86 technology, on which Windows NT is based, also progressed more slowly than expected up the scalability curve, with delays in standardized eight-way server technology and slippage in the schedule for IA64, Intel's next-generation 64-bit technology. Meanwhile, in an adroit and remarkably successful response to the encroachment of commodity technology, Sun embraced and extended the characteristics of PC servers and Windows NT networks in its SPARC technology-based workgroup server hardware and Solaris software products.

Traditionally, Sun's vision resonates most appealingly at the high end of the enterprise, where datacenter culture has shown no sign of diminishing, and administrators are comfortable with the concept of strongly centralized services akin to utilities. By delivering the superior reliability and scalability of Solaris software at competitive price-performance with acceptable levels of interoperability, however, Sun has extended its presence in the datacenter into outlying departments and workgroups. Sun's conflict with Microsoft has historically been most heated in the departmental/workgroup space, as Microsoft's efforts to grow infrastructures from the desktop up collide with Sun's more traditional datacenter approach. Sun now moves the contest to another level, seeking to define the infrastructure underlying local points of presence for the web and farms of web servers.

Solaris currently maintains the upper hand in terms of providing the functional capabilities it has identified as key factors for enabling web-oriented infrastructures. However, Microsoft is strengthening its arsenal with another Windows 2000 package, which specifically optimizes for datacenter requirements. This datacenter version of Windows 2000 is scheduled to ship later this year and will drive yet more function into the commodity space, potentially weakening Sun's differentiation. While both Sun and Microsoft offer compelling visions for the future of the web, the race to implement their views shows no sign of abating.

SOLARIS BACKGROUND

Sun announced Solaris 8 Early Access Edition Software for SPARC and Intel platforms in November 1999 and plans to ship the new system in volume in Q100. In addition to the base Solaris 8 Operating Environment, Sun offers a number of layered software products to enhance the operating system's affinity for specific environments.

Currently, Sun's available feature sets for Solaris include the following packages:

- *Solaris Operating Environment* represents the foundation of the operating-system platform, comprising the Solaris kernel and networking components. Sun positions the Solaris operating environment as an e-business platform that enables service providers to support consumer e-commerce, business-to-business e-commerce, and corporate Intranets. Sun packages a comprehensive suite of software from the Sun-Netscape Alliance with Solaris 8 software, including:
  - Netscape Application Server 4.0,
  - iPlanet Web Server, Enterprise Edition 4.0,
  - iPlanet Directory Server 4.11 and Solaris Directory Extensions 1.0,
  - iPlanet Certificate Management Server (CMS) 4.1,
  - iPlanet SunScreen EFS 3.0.1, and
  - iPlanet Webtop 2.0.1.

  With the exception of Directory Server, all these products are fully functional but licensed for evaluation and development only; deployment licenses are required for production.
- *Solaris Enterprise Software* is an integrated suite of products that provides mainframe-class capabilities in layered packages for the Solaris platform. Sun targets this feature set at enterprise datacenters and service providers that support Enterprise Resource Planning (ERP), e-commerce, data warehouses, and emerging High Performance Computing (HPC) applications. Success in this environment is defined by the ability to support mission-critical applications on highly resilient and scalable clustered hardware. Solaris

Enterprise Software consists of Solaris Resource Manager, Solaris Bandwidth Manager, and Sun Cluster software.

Sun plans to introduce major updates to the base Solaris Operating Environment about every 18 to 24 months, with quarterly releases for minor updates. The company will also update Solaris layered software services more regularly, providing Sun with greater flexibility to respond to customer needs for particular features without risking major modifications to the Solaris kernel.

WINDOWS 2000 BACKGROUND

Microsoft announced its shipment of Windows 2000, the long-awaited update to its Windows NT technology, in December 1999. Volume shipment will begin in February 2000. Windows 2000 ships in three packages:

- *Windows 2000 Professional*: A desktop version supporting two processors and 4 GB of memory
- *Windows 2000 Server*: A low-end server version supporting four processors and 4 GB of memory
- *Windows 2000 Advanced Server*: A midrange server version supporting eight processors and 8 GB of memory

As noted above, Microsoft plans to ship a high-end version called Windows 2000 Datacenter Server later this year. The Datacenter Server version has been designed to support 32 processors and will raise maximum memory support to 64 GB. It will also add four-node HA clustering capabilities, as well as other scalability enhancements.

METHODOLOGY

This study examines Solaris 8 and Windows 2000 Advanced Server (AS) in terms of several key criteria for meeting the computing requirements of web-oriented enterprises and organizations, including:

1. Reliability and Availability Functions
2. Scalability Functions
3. Manageability Functions
4. Applicability Issues

Within each section, callouts note the functional advantages and/or disadvantages for each competitor, relative to the average competitive entry in the server-operating-system space. Each section begins with a summary chart that shows the average rating of all functional details for that particular area using the same metric. D.H. Brown Associates, Inc. (DHBA) rates capabilities using the following nomenclature: Excellent, Very Good, Good, Average, Below Average, Weak, and Poor.

Both Solaris and Windows 2000 products draw significant strength from a large number of third-party add-on products. However, this study assessed only vendors' base product offerings. In addition, both Windows 2000 and Solaris systems receive frequent updates and enhancements, including some that have been announced but not yet delivered.

RELIABILITY & AVAILABILITY

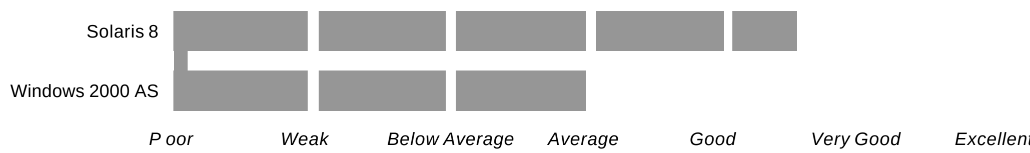

Businesses are becoming more globally oriented and Internet-based, and electronic commerce is continuing to enter the mainstream, trends that pose new demands for reliability, since more and more systems must be able to respond 24 hours a day and 365 days a year.

The reliability of a system depends on a variety of factors involving both hardware and software. On average though, hardware has become more reliable over time:

- Server designs increasingly build on highly integrated components, reducing complexity and hence the points of failure
- Environmental issues such as power failures can be protected against using Uninterruptible Power Supplies (UPS)
- Hardware areas that are vulnerable to mechanical failure, such as storage, can be protected through techniques such as RAID
- Components such as fans are now routinely replicated for reliability

These improvements have meant that system failures increasingly derive from non-hardware-related factors such as software.

Operating systems rank among the most complex pieces of software in existence, and reaching enterprise-grade reliability levels remains a notoriously challenging task for developers. Operating systems serve as the middle tier between hardware and applications. This means they must respond to a variety of unpredictable hardware conditions, while maximizing the performance of equally unpredictable software layered on top of the operating system.

The need for stability and scalability creates another challenge. The ability of a system to handle harsh conditions without failing or degenerating into unpredictable behavior describes its stability. Its ability to add resources and expand its workload capacity determines its scalability. Both qualities are critical for enterprise-class systems, but stability and scalability conflict fundamentally with each other, posing a significant problem for developers.

It is relatively simple to build a scalable system that is unstable, or to stabilize a system that does not have to scale, but delivering both remains an elusive achievement. Today's most venerable enterprise systems – including IBM's OS/390, IBM's OS/400, and Compaq's VMS – all reached this goal only through a long and tortuous process, during which scalability and reliability continuously leapfrogged each other. Along the way, failure became a rite of passage for all, sometimes very publicly. Products that have survived this process earn the most precious distinction that can be imparted on an operating system – *maturity*.

In spite of the growing robustness of hardware and software, some failures remain virtually unavoidable. Hence, operating systems try to protect against failures in both hardware and software through high availability (HA) techniques. These techniques tend to fall into two classes:

- *Resiliency* functions, which protect individual systems
- *Clustering* functions, which increase uptime by coupling the operations of multiple systems

Single-system resiliency functions allow an operating system to contain failures or planned outages within a server, in some cases drawing on technology long available in mainframe environments. Such “self-healing” features potentially may include the ability to:

- Adapt to processor failure by isolating failed CPU components
- Dynamically cordon off memory that has suffered single-bit errors, so that software no longer risks using potentially unreliable areas
- Support dynamic addition and removal of I/O adapters, CPUs, and memory modules for purposes of repair or upgrade

Clustering functions can be used to maintain the availability of system services by failing over to a backup system in the event of system outage. Clustering enables processing to continue by allowing one or more servers to take over for a server that has crashed due to hardware or software failure. Sometimes, clusters can also respond to the failure of individual components, such as disks, adapters, or individual applications. By isolating faults on a failed node, the remaining nodes can continue functioning. The overall clustered system, therefore, keeps functioning, albeit at reduced capacity.

The availability features discussed so far have all involved unplanned downtime related to some type of failure. System availability also depends heavily, however, on the ability to minimize planned downtime related to system management, upgrades, and other regular activities. Operating systems can reduce planned downtime by enabling *online maintenance*, i.e., allowing as many administration functions as possible to be performed without requiring reboots or otherwise disrupting system operations.

PRODUCT MATURITY

Solaris 8	*Very Good*
Windows 2000 AS	*Below Average*

Solaris can now be considered a relatively mature product, having shipped in production for over eight years (see Table 1). Moreover, the Solaris kernel is based on a combination of SunOS software, which had shipped in production for 10 years, and UNIX System V Release 4 (SVR4), which itself represents the outgrowth of a technology base that has shipped far longer. Solaris kernel functions have been stable for a number of years, and most enhancements related to the Solaris 8 feature set were implemented at the user level, involving relatively few modifications to sensitive kernel structures.

TABLE 1: Solaris and Windows NT Product History

	Sun	Microsoft
Q392	Solaris 2.0	
Q492	Solaris 2.1	
Q193		
Q293	Solaris 2.2	
Q393		Windows NT 3.1
Q493	Solaris 2.3	
Q194		
Q294		
Q394		Windows NT 3.5
Q494	Solaris 2.4	
Q195		
Q295		
Q395		Windows NT 3.51
Q495	Solaris 2.5	
Q196		
Q296	Solaris 2.5.1	
Q396		
Q496		Windows NT 4.0
Q197		
Q297		
Q397	Solaris 2.6	
Q497		
Q198		
Q298		
Q398		
Q498	Solaris 7	
Q199		
Q299		
Q399		
Q499		
Q100	Solaris 8	Windows 2000

Windows 2000 is based on the Windows NT kernel, which has been shipping in production for almost seven years. However, after several years without a major update, Windows 2000 introduces a profusion of new features that span a variety of operating system functional areas, some of which involve kernel modifications. Windows 2000's Active Directory and multi-user services, for example, reach deeply into the core of Windows NT's operations. While the exact amount of new code added can only be estimated, it is clearly substantial. Even if "show stoppers" – i.e., production-stopping bugs – have been eliminated by the time Windows 2000 ships, it will clearly take some time before the rough edges get smoothed out.

RESILIENCY FUNCTIONS

Solaris 8	*Good*
Windows 2000 AS	*Average*

Solaris 8 remains at the forefront of dynamic reconfiguration functions. It supports Dynamic I/O Reconfiguration and Alternate Pathing on most of Sun's server product line, giving it a unique differentiation from Windows 2000 and UNIX competitors. These functions enable online repair and reconfiguration of CPUs, memory, and I/O as follows:

- Dynamic Reconfiguration (DR) enables HA by allowing a system administrator to "dry up" defective server components – such as CPUs, memory, and I/O – without application interruption by off-loading processes. Sun's hot-plug hardware capability then allows the defective component to be replaced without creating any electrical problems. This reduces both planned downtime (e.g., for upgrades) and unplanned downtime (e.g., for component failures).
- Dynamic Reconfiguration Coordination: Solaris 8 software provides the tools and APIs to allow Dynamic Reconfiguration to occur automatically, based on administrative policies. Using the Dynamic Reconfiguration Coordination API, developers can adapt applications so that they will be notified of pending and requested configuration changes.
- Alternate Pathing (AP) allows an I/O path to be redirected transparently to applications, enabling a server to adapt to I/O device failure.

Windows 2000 matches Solaris 8 for support of most common hardware-based failure-protection mechanisms, such as UPS equipment and RAID storage. In addition, some developers of high-end Windows 2000 servers have introduced resiliency features that match those of Sun's SPARC servers. However, at the operating-system level, Windows 2000 does not yet support the advanced operating-system functions of the Solaris Platform– such as DR or AP – for protecting unclustered systems from hardware failure.

CLUSTERING OPTIONS

Solaris 8	*Good*
Windows 2000 AS	*Average*

Sun's clustering option for the Solaris 8 platform, called Sun Cluster 2.2, offers a competitive set of HA functions. While not industry-leading by UNIX standards, Sun Clusters 2.2 supports up to four nodes, enabling important HA cluster functions such as multi-directional and cascading failover. Multi-directional failover allows a failed node's workload to be split and failed over to multiple backup nodes. Cascading failover allows failover to continue when a backup node fails. Sun provides monitoring agents for both its own and third-party applications.

Windows 2000 Advanced Server builds in native clustering technology based on the layered technology that had been available for Windows NT 4.0 called Microsoft Cluster Server (MSCS – formerly code-named "Wolfpack"). Compared with UNIX-based HA clustering products such as Sun Clusters 2.2, Windows 2000 AS clustering delivers minimum HA capabilities. Cluster sizes remain limited to two nodes, precluding multi-directional and cascading failover.

ONLINE MAINTENANCE FUNCTIONS

Solaris 8	*Very Good*
Windows 2000 AS	*Good*

For several years, Sun has focused its Solaris software development efforts on emulating the feature set of mainframes. Mainframes originally introduced many of the technical concepts that allow systems to minimize planned downtime. Several Solaris 8 features reflect this effort, including:

- *Live Upgrades* – Live upgrades allow administrators to build a new Solaris software release using an existing production system as the driver. The process involves building the updated release in a new directory structure and rebooting the system using the new directory as the root, which allows immediate use of the new system. Falling back to the previous release is simple, since the root need only be reassigned to the old root directory, and changeovers are largely transparent to most applications.
- *Hot Relief* – Sun has designed a formal process for Solaris 8 that allows fixes, including kernel modifications, to be applied to a running system. Sun plans to apply the system to its patches for use by its Enterprise Services personnel.

Other Solaris 8 platform resiliency functions such as Dynamic Reconfiguration and Alternate Pathing can also be used for reducing downtime by enabling upgrades to processors, memory, and peripherals without rebooting the system.

Windows 2000 significantly lowers planned downtime from earlier versions of Windows NT, principally by sharply reducing the number of system-management operations that involve a system reboot. Windows 2000 also introduces some features that simplify patch management, for example service pack "slipstreaming," which eliminates the need to re-apply service packs every time new components are installed. Otherwise though, Windows 2000's upgrade and patch procedures remain fairly traditional.

SCALABILITY

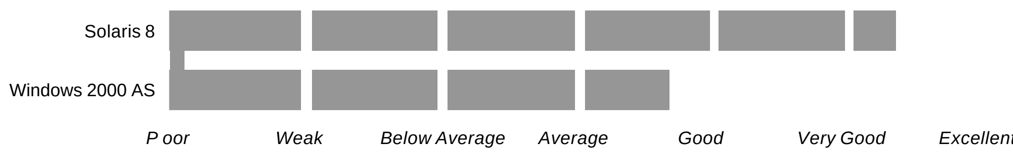

Web-based applications run in an environment where wildly fluctuating workloads and low tolerance for failure are the norm. Hence, developers of web-oriented systems must treat scalability as a primary optimization target. Three functions fundamentally determine the scalability of an operating system:

- *64-bit support* – The ability to exploit processing, memory, and storage beyond the 4 GB limitation imposed by 32-bit systems. Several levels of 64-bit capabilities exist, including 64-bit processor support, large file systems, large files, large physical memories, and large process address spaces (where “large” means greater than 4 GB).
- *Shared-memory multiprocessing (SMP)* – The ability to take advantage of multiple processors in a server. Criteria include kernel locking granularity, kernel thread mechanisms, and evidence of scalability based on industry-standard benchmarks.
- *Performance clustering options* – The ability to grow system capacity, including performance and storage, by lashing together multiple servers using high-speed interconnects. Typically, a system’s ability to handle technical applications and commercial applications (e.g., database or web) classifies its performance clustering capabilities.

64-BIT CAPABILITIES

Solaris 8	*Excellent*
Windows 2000 AS	*Good*

Full 64-bit compliance requires that an operating system be capable of scaling its code and data support beyond the 2 or 4 GB limits common to 32-bit systems. The ability to handle larger disk file systems, file sizes, physical RAM, and process address space forms a key part of the technology required. In addition to adding 64-bit capabilities in both hardware and software, vendors must also make the migration from 32- to 64-bit systems as painless as possible.

The biggest payoff for 64-bit capabilities has come for large databases that can cache complete database indexes (or the database contents themselves) in physical memory, offering a roughly 10x improvement in access time over disk. Performance improvements in real-world situations with real workloads are substantially more modest – as demonstrated by TPC-C results for various 64-bit vendors that are moderately higher, closer to a factor of 10%-100%.

Solaris 8 software is a fully 64-bit environment, providing applications with a full 64-bit virtual address space, supporting up to 64 GB of physical memory, and supporting up to 1 TB of storage. Although Sun's servers have now turned over completely to the 64-bit UltraSPARC CPU, Solaris 8 still comes in both 32-bit and 64-bit flavors, as chosen at install time.

Although Windows NT supported 64-bit files and file systems since it first shipped in 1993, Windows 2000 is still making the transition to full 64-bit capabilities. Windows 2000 Advanced Server supports up to 8 GB of physical memory using Intel's PAE extensions, but user applications still run in a 32-bit address space. Full 64-bit support awaits a future update to Windows 2000.

SHARED-MEMORY MULTIPROCESSING

Solaris 8	*Excellent*
Windows 2000 AS	*Average*

The ability of an operating system to exploit shared-memory multi-processing (SMP) architectures continues to represent a critical differentiator in server environments. Relevant factors include:

- The degree to which the kernel has been optimized to exploit multiple processors. This influences the absolute range of processors that can effectively be supported – ranging from two processors up to more than 100 processors in advanced NUMA architectures.
- The availability of mechanisms to support SMP-optimized applications, such as threads.
- The availability of industry-standard benchmark evidence on high-end systems. Results should be based on tests such as TPC-C and TPC-D, which stress I/O as well as computation.

Solaris 8 provides excellent SMP scalability. The Solaris kernel has been designed to support up to 256 processors, and currently runs on SPARC servers with up to 64 processors. Indeed, Sun's server product line is virtually unmatched in terms of scalability. It offers one of the broadest performance ranges in the industry – from uniprocessor systems to very large SMPs – without introducing gaps in terms of binary compatibility or operational requirements.

As one of the first UNIX environments to optimize for kernel threads, Solaris pioneered the MxN thread model. This model has proven highly effective in allowing scalable network applications. In addition, Solaris software has demonstrated strong performance on a broad range of benchmarks from 16- to 64-processor systems, including results for both TPC-C and TPC-D that employ all 64 processors in its highest-end SPARC servers.

TPC-C tests have been run on Windows NT systems with up to 12 processors, putting it in the same league as some UNIX systems. The results showed that Windows NT 4.0 could achieve moderate performance gains, but with somewhat limited linearity. The Windows NT kernel provides a very efficient thread mechanism for development of SMP applications, but its relatively limited SMP ranges had represented a significant limitation for Windows NT's scalability in the past.

Windows 2000 introduces a number of other kernel-level optimizations that directly target improved SMP scalability. This may mean that better results will begin to appear in the future, but Windows 2000 Advanced Server caps the number of supported processors at eight, so true next-generation performance benefits await the shipment of Windows 2000 Datacenter Server, which has been designed to support 32 processors.

PERFORMANCE CLUSTERING OPTIONS

Solaris 8	*Good*
Windows 2000 AS	*Good*

Clusters can help increase a system's capacity, including performance and storage. To scale performance on a cluster, applications work in concert with clustering software to partition their workloads into subtasks, which the clustering software then distributes across the clustered servers. Since even the fastest cluster interconnects usually have lower bandwidth and greater latency than the bus in an SMP (in some cases by several orders of magnitude), synchronization among the subtasks becomes a critical bottleneck that systems must minimize. Identifying opportunities for coarse-grained parallelism, therefore, proves key to effective scalability on clusters.

A variety of parallel-programming tools and techniques have emerged to assist in partitioning applications for clusters. Their use requires considerable expertise, however, and some classes of applications fundamentally cannot be adapted at all. If sufficiently partitioned, applications can exploit clustered systems containing hundreds or even thousands of nodes, resulting in monumental levels of performance.

Several commercial database systems, including Oracle Parallel Server (OPS), IBM DB2 Universal Database (UDB), and Informix XPS, have been extended to work on clusters of servers connected by high-speed interconnects. All these systems are available for both Solaris and Windows 2000 platforms. The credibility of Sun's commercial clustering performance has been validated with four-node, 24-way TPC-C results and a TPC-D result on a four-node, 24-way cluster/SMP combination. Although Windows NT had proven its ability to support scalable data warehousing on large clusters with TPC-D results, few clustering performance results have yet emerged for Windows 2000 clusters. However, Windows 2000 Advanced Server does earn an advantage from its built-in ability to support web server farms, a useful feature that requires add-on tools in Solaris.

MANAGEABILITY

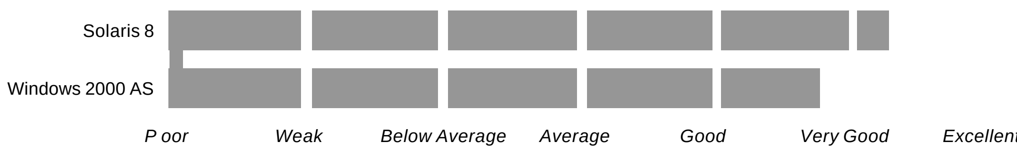

Like any complex technology, the performance of an operating system will depend a great deal on its ability to be managed and tuned. Administrators need the ability to carefully monitor the system's behavior and to easily configure critical functions in response to runtime conditions. Further, in the Internet era, enterprises routinely disperse servers geographically, in some cases across different continents and time zones.

Thus, effective *remote management* of operating systems becomes increasingly important. If an enterprise depends on a thousand servers, it is simply not feasible to maintain a thousand system administrators locally. A more methodical procedure is required. Several techniques have emerged to help manage servers remotely, including:

- *Server hardware features* – A number of server hardware vendors offer service processors for their systems. These independent processing units monitor server activity and allow remote diagnostic access in the event of failures that disable the main system processor.
- *Remote operating-system access* – Since the operating system controls all server functions, administrators must be able to communicate with it remotely. Ideally, a remote administrator should be able to use the system as if he or she were physically next to the hardware. Remote interaction might occur over character-oriented sessions – as if the administrator were using a local ASCII terminal – or via a distributed GUI – with graphics and keyboard/mouse events being passed back and forth based on the native look-and-feel of the environment being managed.
- *Web-based system management* – By using a Java technology-based system management GUI, systems can be managed remotely across networks from any Java technology-enabled web browser. Java software's user interface widgets closely match those of mainstream Windows widgets, enabling management tools that are relatively intuitive to inexperienced users.
- *Template-based installation* – The template approach involves a "cookie cutter" method, in which a template server is created and tested, then replicated across multiple servers using a distribution mechanism. The cost of this technique is that the template gets installed on a server that is not actually used. The benefit of this approach, however, is that if a critical server crashes, administrators can take an idle standard server – or one being used for low-priority tasks – and change its configuration to make it a replacement for the critical server. This approach provides tremendous flexibility for managing

systems. A further step automates the updating of common parts so that administrators can ensure all servers remain identical. This guarantees the backup server can substitute for the critical server once reconfigured. In addition, if a critical server behaves in a problematic fashion, administrators can use an identical server to replicate the problem rather than having to take the critical server out of service.

As part of the move to web-based infrastructures, server consolidation has become a key part of IT management. The primary driver toward consolidation is widespread recognition of the significant operating cost savings that can result from server consolidation. The client-server model of distributed computing resulted in a proliferation of servers that was accelerated by the seductively low acquisition costs of Windows NT servers. The large numbers of servers has created a nightmare of IT administration and support, resulting in increased costs that consolidation can help overcome.

Consolidation also represents the only cost-effective way that companies can meet the level-of-service requirements associated with e-commerce, which call for scalability and resource optimization. By aggregating multiple applications on a single machine, server consolidation attempts to reap such benefits as simpler management and greater flexibility in allocating computing resources to particular tasks.

Since critical applications are often designed to dominate the resources of a dedicated server, re-deploying such applications on SMP servers requires specialized management functions. Software-based *resource-management* tools efficiently allocate system resources such as CPU, memory, and network bandwidth to different applications. In effect, these tools override the operations of the default UNIX scheduler, taking customized policies into consideration.

In large networks, it becomes increasingly tedious for users and administrators to securely and reliably track user IDs, passwords, server host IDs, and printers throughout organizations. *Distributed systems management* itself becomes a database problem. Operating systems supporting web-based infrastructures must therefore provide effective distributed *network security* functions. They must also provide a special-purpose distributed database called a *directory service* that provides users and administrators with an up-to-date and global reference to all network resources. With fully distributed security and directory services, administrators can provide access to data and applications from anywhere on the network. Users can then log in from a client system regardless of its geographic location or the server through which it connects.

UNIX has historically had relatively complex system management procedures. Most UNIX systems require administrators to hand-edit a large and dispersed set of cryptic configuration files – a crude and error-prone process. This administrative nightmare resulted from the operating system's ad hoc development at the hands of those largely unconcerned with the problems of large-scale, production-computing environments. As networks have expanded across organizations and the web has become a growing part of day-to-day activities, lowly administrative chores such as restarting a printer queue or handling a backup procedure often fall to less-experienced (and less expensive) system administrators.

Some UNIX vendors have responded by developing more user-friendly *GUI-based system management tools*. Administrator tools have begun the transition from "remember and type" to "recognize and point."

REMOTE MANAGEMENT

Solaris 8	*Excellent*
Windows 2000 AS	*Very Good*

Solaris 8 natively supports all classes of remote management. A number of Sun's SPARC servers offer service processor options, which support the ability for administrators to dial in remotely and perform diagnostics without involvement of the host processor. Through this capability, maintenance can occur whether or not the system is up and running. Sun's SPARC servers derive an additional advantage from OpenBoot, a low-level PROM-based mechanism for managing information related to devices and boot sequences. Since OpenBoot relies on an independent processor, it can be accessed even when a server is disabled. Also, an OpenBoot session can be set up over serial lines, allowing a server to be accessed even when the network is not operational.

Because Solaris software has a multi-user design – like most UNIX implementations – it fully supports remote interaction via network or serial-line connections. In other words, the Solaris kernel can manage the processes associated with remote system administrators, while simultaneously running processes associated with server tasks and local users, appropriately tracking security privileges and processor usage. Based on this basic scheduling ability, remote administrators have several choices for communicating with an active Solaris Operating Environment, including:

- Logging in via an ASCII terminal over a serial port: Solaris software inherently has the ability to manage many serial ports to which ASCII terminals are attached. This option is typically useful in campus environments, where infrastructures may include RS-232 serial line configurations. It provides access to Solaris via the UNIX command shell. In turn, the shell provides access to all files and system configuration settings, which are stored by UNIX in traditional files that can be edited in character-based sessions.
- Logging in over telnet: Telnet, the traditional UNIX network terminal tool, allows administrators to access a Solaris system over a network as if they were connected locally through an ASCII terminal. As with any character-based session, administrators using telnet have total access to all files and UNIX system settings.
- Using X Window System distributed graphics to run Solaris GUI system management utilities: Administrators can use utilities such as admintool

remotely, as the X Window System used by Solaris software inherently has the ability to support remote interaction. Users on remote systems that are also configured with the X Window System can therefore run graphical applications as if they were local. Administrators can use admintool to graphically configure any server on the network.

- Using Solaris WebStart: WebStart, which is Sun's Java technology-based system management tool, has a Java implementation that allows administrators to perform software installation and configuration from any Java technology-enabled web browser.
- Using Solaris JumpStart for template-based installation: This function allows operators to create a master of a Solaris environment, including all necessary patch updates, that can then be rolled out to many distributed systems.
- Using Solaris Management Console: Sun began shipping a management console with Solaris 7 that allows administrators to take advantage of remote access functions supported in its underlying Web-Based Enterprise Management (WBEM) and the Distributed Management Task Force (DMTF) Common Information Model (CIM) standards.

Some Windows 2000 hardware suppliers have introduced service processors akin to those found in Sun's SPARC systems, allowing remote management at the hardware level. Windows 2000 also offers a significantly stronger ability to be managed remotely at the software level through several mechanisms. This occurs principally with the "remote administration mode" included in its new Terminal Services, which gives system administrators a GUI-based method to remotely administer any Windows 2000 server over the network. Windows 2000 introduces Microsoft Management Console (MMC), a new management framework that supports remote instrumentation based on WBEM and CIM standards. MMC can be installed remotely without local intervention on systems configured with specific hardware, such as network interface cards containing a Pre-boot Execution Environment (PXE) Boot PROM. Windows 2000 also includes some limited template-based installation capabilities based on its "disk image preparation" mechanism, which allows a fully-configured Windows 2000 system to be duplicated.

RESOURCE MANAGEMENT FUNCTIONS

Solaris 8	*Excellent*
Windows 2000 AS	*Average*

Solaris 8 addresses resource management comprehensively, including the Solaris Resource Manager function and the Solaris Bandwidth Manager. Solaris Resource Manager enables system administrators to guarantee minimum processor cycles and memory allotments to users or applications. Solaris Bandwidth Manager does the same for network connections.

In addition, the dynamic domains feature in Solaris 8 allows servers to be partitioned into logical divisions on a hardware basis. Each partition runs its own copy of the operating system and acts as a virtual system that is completely isolated from others. The boundaries between domains can be adjusted dynamically via operator commands or scripting, enabling large servers to accommodate different business cycles. For example, multiple LAN servers might be combined into one large environment during off-hours to process heavy batch jobs. This capability provides an additional level of robustness, since a catastrophic failure in one domain leaves applications running in other domains unaffected. While dynamic domains are currently only supported on Sun's Enterprise 10000 server, Sun plans to introduce the capability for a much broader range of its server platforms later this year as UltraSPARC-III processors appears.

Windows 2000 introduces basic resource-management functions that have long been available in Solaris software, such as disk quotas which regulate storage available to users or applications. Windows 2000 also includes some network bandwidth throttling functions, but it otherwise falls short of Solaris 8 software's abilities to partition processor and memory cycles by specific users or applications.

DISTRIBUTED SYSTEMS MANAGEMENT

Solaris 8	*Good*
Windows 2000 AS	*Very Good*

Both Solaris 8 and Windows 2000 software approach the directory service and network security problem with derivatives of the same mechanisms, Lightweight Directory Access Protocol (LDAP) and Kerberos network security.

LDAP provides an open standard for directory services based on a subset of X.500 – the vast, comprehensive, and formal standard for information exchange. In addition to comprehensively addressing data interoperability, the LDAP functional specification is also broad enough to leave room for implementation of a variety of enhancements for boosting scalability and reliability in very large networks. As LDAP continues to gain implementations on a wide variety of operating systems, it promises to become both a de jure and de facto standard for managing resource information on enterprise networks.

Kerberos is a distributed security system developed at MIT that authenticates users across large networks. When integrated with applications, Kerberos allows verification of the identity of an application user without sending any data across the network that might allow an attacker to impersonate that user. Kerberos establishes a user's identity at logon and then attaches a "ticket" containing a session key to the user's client. This ticket can then be used to continuously verify the identity the client by applications throughout the network. Because the ticket is encrypted in the Kerberos server's key, a client cannot modify the ticket without detection. Solaris 8 ships with the Netscape Directory Services 4.1.1, a strong LDAP V3-compliant implementation that has been tested with up to one million entries. It also includes Kerberos v5 support in the base operating system.

Windows 2000 introduces active directory, a new set of directory and security mechanisms that have been reimplemented from the ground up and integrated directly into virtually all its management aspects. Active directory is a directory service based on LDAP that provides a single point of administration for Windows 2000 network resources such as files, users, and other objects. Like Solaris 8, Windows 2000 uses Kerberos V5 for network authentication. Microsoft also provides a comprehensive API for accessing directory and security services from applications and scripting functions.

GUI MANAGEMENT TOOLS

Solaris 8	*Very Good*
Windows 2000 AS	*Excellent*

Historically serving highly technical users on its workstations, Sun began to address ease-of-use criteria in system management relatively recently. Its new emphasis represents part of its move to focus more on the requirements of commercial users. Although previous versions of Solaris software included a GUI-based system management tool, admintool, this tool lacks the breadth of SMIT in AIX, or SAM in HP-UX. It only covers the basics of adding and deleting user profiles, printers, host names, serial ports, and software.

Solaris 8 enables comprehensive, GUI-based system management with the Solaris Management Console software. Sun introduced this update in Solaris 7 as a point-and-click administration tool based on WBEM and CIM standards. The Solaris Management Console product provides a centralized integration point for Solaris system administration and management tools. The console is configurable and extensible, allowing integration of system-management applications based on a variety of development methods, including the X Window System, scripts, Java, and HTML. Solaris AdminSuite software, a tool running within Solaris Management Console software, allows GUI-driven management of users and groups.

Another Solaris tool, Solaris Web Start, provides a Java-based GUI for installing system software and software add-ons, offering both ease of use and remote manageability. Solaris 8 software also includes configuration wizards based on the popular InstallShield tool for Windows, which extend simplified installation to third-party Solaris applications.

Windows NT has long enjoyed an intuitive user interface for managing single systems, largely benefiting from the exceptional familiarity of the Windows look-and-feel adopted by the Windows NT GUI. Windows 2000 introduces a new GUI for managing system resources called the Microsoft Management Console (MMC). Like the Solaris Management Console product, MMC is compliant with WBEM and provides a general-purpose management display framework for hosting administration tools built as MMC Snap-Ins by Microsoft and third parties. MMC can be used from within an existing enterprise console, or it can be used to launch enterprise consoles.

APPLICABILITY

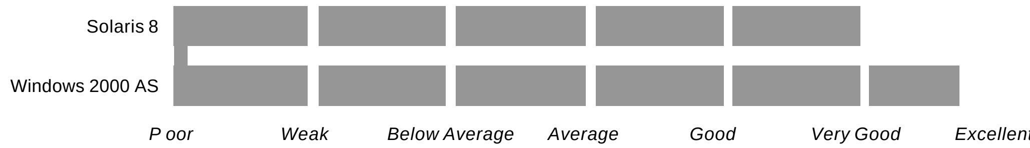

Aside from the functional tradeoffs associated with particular technologies, users need to consider the business and operational conditions under which a technology will be deployed. To arrive at a complete profile of an operating-system product, a number of factors come into play, which relate to what DHBA classifies as "applicability" issues. These include the following:

- *Platform availability* – Since a system consists of both an operating system and its underlying hardware, the server platform choices available for running the operating system become very important.
- *Application availability* – A platform is only as useful as the applications that run on it. While there are usually few technical barriers to porting an application to a given operating system, market perception plays an enormous role in determining which environments third-party developers decide to target first. Exact application counts for an operating system are difficult to come by reliably, due to the many gradations in terms of release levels and support of optional components. An informal scan of the product literature for various applications, however, quickly reveals the absolute leaders.
- *Development considerations* – Operating systems need to provide strong development environments in order to assure the flow of new applications and allow customers to customize an environment for their purposes. These environments should include a broad set of programming tools and a rich set of services that developers can exploit via Application Programming Interfaces (APIs).

PLATFORM CHOICES

Solaris 8	*Good*
Windows 2000 AS	*Very Good*

Although Solaris 8 runs on both Intel x86 and SPARC processor architectures, the vast majority of Solaris deployments occur on Sun's SPARC product line. Some other vendors, including NCR and Siemens, have committed to turning over their product lines to Intel systems running Solaris software. The majority of these plans have been timed to begin in earnest when Intel IA64 systems start shipping later this year.

In spite of Solaris 8 software's dependence on hardware that comes predominantly from a single vendor, Solaris systems nonetheless present a compelling value proposition. Sun's SPARC server hardware offers scalability that is virtually unmatched in the industry. Sun machines support binary and administrative compatibility ranging from the low-end Ultra 5S uniprocessors to 64-processor Sun Enterprise 10000 servers. Virtually all other server vendors either fragment their product lines across Windows NT, multiple UNIX offerings, and proprietary systems, which sacrifices binary compatibility, or deploy different system architectures at the high-end, such as clustering or MPPs, which require different administrative procedures.

Windows 2000 has an advantage in platform choices from the large number of Intel x86-oriented vendors that sell and support Windows 2000 on their products. These span a wide variety of system types, ranging from laptops up to mainframe-class systems from vendors such as Unisys. Although Windows NT had supported Compaq's 64-bit Alpha architecture, Microsoft and Compaq recently decided to cease Windows NT development for that platform. Hence, Windows 2000 runs exclusively on Intel x86 hardware.

APPLICATION AVAILABILITY

Solaris 8	*Excellent*
Windows 2000 AS	*Excellent*

Both Solaris and Windows 2000 platform are perceived as market leaders and thus offer very good application availability. Virtually all the leading third-party providers of the key components for web-based infrastructures – such as databases, web servers, middleware, and web application servers – provide optimized versions of their products for both Solaris and Windows 2000 software.

DEVELOPMENT CONSIDERATIONS

Solaris 8	*Very Good*
Windows 2000 AS	*Excellent*

Microsoft and Sun have each built very strong relationships with the development community over the years. Both vendors offer strong development tool product lines, which include compilers and Integrated Development Environments (IDEs) preferred by corporate developers. Sun and Microsoft have also defined overarching development architectures that provide a foundation for users to custom-build distributed infrastructures that can span the web and other network models.

Sun's approach to distributed computing is based on Java, its platform-independent programming environment. Sun and its partners in the Java community have defined standards for a broad range of distributed services that can be implemented with the Java language, ranging from Enterprise JavaBeans transaction processing to messaging. Solaris 8 includes a highly optimized Java Virtual Machine, which helps to boost the performance of Java software when deployed on the Solaris platform. Microsoft offers its Distributed Network Architecture (DNA), which is based on its COM+ distributed component model and provides a similar set of services for transaction processing and messaging.

Windows 2000 earns a slight advantage over Solaris 8 software for its high degree of integration, which comes at the expense of openness. Sun requires key infrastructure components such as web servers, messaging systems, and web application servers to be purchased and installed separately. Windows 2000 builds all of these features into the base operating system. It should be noted, however, that since COM+ runs only on Windows 2000, the growth path for COM-based applications is inherently constrained by Windows 2000's capabilities. Developers who have chosen COM+ and encounter scalability or reliability barriers must either wait for new releases of Windows 2000 or develop their own workarounds to overcome operating system limitations. By contrast, Java technology derives much of its value from the broad range of platforms it supports, from low-end PCs all the way up to clusters of IBM mainframes.