

Taking advantage of the SAS[®] System on Windows NT

Gary Mehler, SAS Institute Inc., Cary, NC

ABSTRACT

Windows NT is becoming the universal desktop client operating system as well as an important file and compute server for mission critical applications. This paper presents a discussion of the state of Windows NT and how the SAS System Version 7 for Windows takes advantage and integrates with the operating system. Data access and Microsoft BackOffice integration, and hardware considerations are also presented.

INTRODUCTION

Microsoft Windows NT sales grew dramatically in 1997, as many corporations which have been investigating Windows NT have now begun to deploy Windows NT for the client desktop. Microsoft claims that eleven million licenses for Windows NT Workstation have been sold. This is approximately 9% of the client market. The client platform with the most penetration (as of the end of 1997) is Windows 95, which is on 51 million systems (about 66% of client systems). Many of these 11 million deployments were replacing Windows 3.1 and to some extent Windows 95. Given that there are over 100,000 32-bit Windows applications available, the upgrade is safe for many organizations. Several major PC manufacturers (including Compaq, Dell, DEC, HP, IBM, and Micron) are preinstalling Windows NT Workstation on their business PCs, so the growth is likely to continue.

Also, Windows NT Server is making major inroads for mission critical file and compute servers. As hardware capabilities continue to increase on high-end NT Servers, more organizations are taking advantage of this platform as well. This is true in many areas, even for web-based applications. Microsoft estimates that about 360,000 web sites are now NT based (and also using Microsoft's IIS web server software).

This paper presents the current state of the SAS System for Windows NT. It introduces the basic system features of Windows NT as well as some

of the more advanced features of Windows NT. At each point, features and products of the SAS System that integrate well with Windows NT are presented. A brief overview of Windows NT Server is discussed within the context of SAS Institute's plan for continued support. The paper concludes with hardware considerations for the SAS System on Windows NT and a glimpse of future directions for Windows NT and the SAS System.

The Windows Family

Windows NT sales have grown steadily since the release of Windows NT Workstation 4.0 and Windows NT Server 4.0. NT Workstation and NT Server share the same microkernel, but are configured differently, and support different application types. Although operation on the DEC Alpha AXP is still supported (all other RISC platform support has been dropped), the majority of NT installations still run on the Intel IA32 class or processors: Pentium, Pentium II, and Pentium Pro. The SAS System Version 7 only supports the Intel platform, and the Pentium Pro or Pentium II processors are highly recommended. At this writing, there is no support on Windows NT for the DEC AXP, although internal research is ongoing for the support of this platform.

Although this paper focuses on Windows NT, it is important to note that the SAS System for Windows Version 7 supports the family of Windows 32-bit operating systems with a single executable image. There is a single CD that supports all three Windows 32-bit operating systems Windows NT Workstation, NT server, and Windows 95. However, the product has run time checks for Windows NT to take advantage of NT specific APIs (application programming interfaces). Development at SAS Institute focuses on Windows NT.

User Interface

Windows NT 4.0 incorporates the Windows 95 Explorer shell as the default user interface shell. This means that the user interface between the

highly popular Windows 95 and Windows NT, and Windows CE is largely the same.

It is important for applications to follow the user interface guidelines provided by Windows, and the SAS System is no exception. The SAS System is a native Windows NT application. Developers at SAS Institute have paid special attention to Windows 95 user interface guidelines outlined in Microsoft's *Windows User Interface Handbook* in designing the SAS System user interface.

Additionally, Windows NT supports messaging through Microsoft Exchange. NT provides for an email client named the Inbox, and Microsoft Office97 provides for a Personal Information Manager named Outlook that combines email inbox, scheduling, and contacts. Through support of Microsoft Messaging API (MAPI), the SAS System integrates nicely with the Inbox and Outlook, allowing email messages containing rich text and attachments to be created and mailed by the SAS System. The SAS System uses the email user interface consistent with the mail program installed. The messaging APIs supported are MAPI and VIM.

Windows NT also supports the Internet Explorer 4 user interface enhancement called "Web View". In this mode, the desktop or any folder can be portrayed as a web page, with hyperlinks and other interface enhancements.

Security

An essential part of an industry strength client workstation or mission critical server platform is security. NT provides for per-file and per-directory security within the NTFS file system, as well as local desktop security. A user id and password is required for logging into Windows NT. Across the network, APIs are provided for a single network login, and NT prevents use of resources by unauthorized users.

The SAS System takes advantage of the security features of Windows NT within the file system. In addition, SAS products such as SAS/CONNECT[®] and SAS/SHARE[®] only allow data access from authorized remote sessions.

File Systems

The file systems supported on Windows NT include FAT and NTFS. The FAT file system is provided for compatibility with MS-DOS. The NT File System provides for a secure file system that provides for 64-bit addressing, supporting files larger than 2 gigabytes. (While FAT32 on recent updates to Windows 95 supports filesystems greater than 2 gigabytes, individual files are still limited to 2GB on Windows 95 filesystems.) Version 7 of the SAS System properly supports NTFS allowing SAS Data Sets to be accessed that are larger than 2 Gigabytes. Data Sets around the size of 8 to 10 Gigabytes are common in this scenario. Windows NT and the SAS System also support long filenames up to 256 characters long.

Disk sharing services are provided for peer connections. And NT provides for file compression per file or per directory with the NTFS file system.

Data Access - ODBC Support

The SAS System provides for open data access. To provide data access to SAS Data Sets from other applications such as Microsoft's Excel or Microsoft Access, the SAS ODBC Driver provides for support for both local and remote access. Local access is supported by a local SAS ODBC Server provided in BASE SAS software. Remote access to SAS Data Sets is provided through the SAS/SHARE ODBC Server. And using SAS/SHARE*NET[®] software, clients can access data on remote SAS/SHARE servers without requiring the SAS System to be executing on the local PC. The SAS ODBC Driver supports data access through TCP/IP sockets, DDE and Network DDE communication access methods.

SAS/ACCESS[®] software supports accessing any ODBC compliant data source for both 16 and 32 bit drivers, thus making available data external from SAS, to the SAS System for processing. This provides for a large array of PC and data base data sources to be accessible from the SAS System.

Data Access – SAS Universal ODBC Driver

The Universal ODBC Driver provides stand-alone access to SAS data in read-only mode. In

this way, the Universal ODBC Driver allows client applications to directly access SAS data without requiring that the SAS System be present and running on the system providing the data.

Data Access - Microsoft SQL Server

The SAS System can access data in Microsoft's SQL Server data base through the SAS/ACCESS for SQL Server product. This provides for the ability to read, write or update a SQL data base. The data base can reside locally or across the network.

Multitasking, Memory Management, CPUs

Windows NT provides for preemptive multitasking. Applications execute within their own address space and are prevented from corrupting other applications. Windows NT runs in full 32 bit mode, using the protected mode of the Intel architecture. All of Windows NT is 32 bit, and all of the SAS System is 32 bit code. The SAS System is compiled for optimum performance on Intel Pentium and Pentium Pro microprocessors. This includes using a highly tuned floating point compiler optimizer for creating optimized floating point instruction sequences.

Differences between Pentium, Pentium II, and Pentium Pro processors are not major, but can have an impact on the performance of the SAS System running on them. The original Pentium (with or without MMX) is our baseline processor. The next chip to ship was the Pentium Pro, the chief benefit of which is its larger, closely coupled cache, which runs at full chip speed. Pentium II, which is internally the computing portion of a Pentium Pro, may also benefit from a larger cache, but access to the cache is not as direct as on the real Pentium Pro. As a result, applications that utilize large amounts of data may pay a penalty on the Pentium II. As of this writing, the Pentium II processor is available with the highest clock speeds (up to 300MHz), and as a result, can return the best performance on computations that don't require large amounts of data (such as iterated numeric calculations). When data throughput is the major issue, a slower Pentium Pro (200MHz) may outperform it.

In the original version of NT 4.0, memory is addressed linearly up to a limit of 2 gigabytes of address space per process. Applications execute in a robust tasking environment and the memory management architecture virtually removes the out-of-memory constraints of previous Windows 3.1 releases. This includes the removal of the 64 kilobyte limitations of the USER and GDI segments. NT provides for the ability to execute multiple concurrent SAS sessions. Multiple SAS sessions do not strain the system because these concurrent sessions share all the same code in memory, and only the data and stack sections of the SAS System are allocated in memory per executing session. This is often the best way to take advantage of multiple processors.

As of Windows NT Server Enterprise Edition (NTS/E), a feature called *4 Gigabyte RAM Tuning* allows specially built large applications to access up to 3 gigabytes of RAM. This 50% increase in RAM capacity (assuming you've purchased the RAM for your server already) is possible because as of NTS/E, the usual memory space split of 2GB for each application, and 2GB for the NT kernel is changed. The NT kernel can run fine in a 1GB address space, which frees an additional gigabyte of address space for application use. The SAS System Version 7 is built specially to take advantage of this feature on very large server systems.

This tuning can be useful for applications that use extremely large amounts of data. In some standard benchmarks (such as TPC/C), throughput improvements of 20% have been measured. The benefit comes because applications can utilize more memory to store data instead of re-accessing parts of the data through the I/O subsystem on disk.

Symmetric Multiprocessing Support

Out of the box, NT Server supports up to four processors (NT Workstation supports up to two processors) in which any thread of execution can execute on any processor, at any time. Windows NT Workstation and Server take advantage of multiple threads even down to the kernel level. This provides a significant value to applications, because even though applications are not full multi-threaded, applications may often gain

considerable performance from multiple processors because API requests tend to be multi-threaded as the API is completed down at the NT kernel level.

The SAS System uses multiple threads of execution, but SAS procedures do not take advantage of SMP. Where possible the SAS System does take advantage of multiple task execution such as in SAS/CONNECT[®] software. Research is underway so that future releases of the SAS System will utilize threads in a much smarter design.

It is recommended to install multiple processors especially in the case of using SAS/SHARE or SAS/CONNECT software on a Windows NT Server. Typically there are other server processes running on the NT Server system, and can benefit from multiple processors. Additionally, multiple SAS sessions running concurrently can significantly improve if multiple processors are available.

As of NTS/E, NT Server can take advantage of more than four processors, and high-end vendors are now shipping 8-way systems. These very high-end systems usually also sport multiple, fast buses for increased throughput.

Portability

From a strategic point of view, portability can be measured in two ways: 1) portability of the SAS applications, and 2) portability of the Windows operating system to be provided on a new and different processor and architecture. First, SAS applications are portable across SAS Institute's MVA[®] supported platforms. This includes SAS Data Step code, procedure code, and SCL code. Care needs to be taken not to use platform specific extensions typically found in advanced file I/O APIs and in the user interface.

Second, Windows NT itself is portable across differing hardware platforms, including Intel and DEC Alpha. In fact, the lightweight Windows CE machines are actually running a kernel based on Windows NT (not Windows 95) because the small systems are not running Intel processors. Portability was one of the design goals of Windows NT, and provides customers with hardware independence. This is accomplished with the use of a portable microkernel built on

top of modules named HALs (Hardware Abstraction Layers).

The combination of Windows NT and the SAS System provides one of the best portable or vendor independent solutions available. Windows NT can and most likely will be ported to the most successful workstation platforms in the future, and SAS Institute has already proved the portability of the SAS System, having ported to every major operating system in the industry. Also, SAS applications remain portable even across differing platforms.

The SAS System's approach to Client Server

The Institute's MVA architecture provides for the entire SAS System functionality to be able to run on either a client (Windows NT Workstation) or server (Windows NT Server). There is not a unique SAS System release for Windows NT Server, instead the same executable system runs on both NT Workstation and NT Server. This provides for maximum flexibility in designing a client server solution. Functionality on the client side or the server side is not limited. By providing the SAS System on both client and server, total flexibility and granularity is available in creating the optimal client server application. These are covered in the paper, "Integrating Windows Clients and the SAS System into the Enterprise", also in these proceedings.

Support for Microsoft BackOffice

Designed for Microsoft has created a logo certification program that certifies a level of functionality for products designed to run optimally on Windows NT Server. This logo certification program is called **Designed for Microsoft BackOffice**. Microsoft BackOffice is a family of products designed to run on Windows NT Server. Since the 6.12 Release, the SAS System for Windows has been certified as Designed for BackOffice. This certifies that the SAS System was designed and tested to integrate with the BackOffice family of server applications.



As a BackOffice client, the SAS System supports e-mail access through MAPI so that user can send e-mail messages, documents, and files from within the SAS software environment, either through standard interactive dialogs or programmatically through the SAS language.

Another BackOffice client feature is the ability to interact with and access SQL Server data. This access is available through DB-Library, the direct access provided through SAS/ACCESS to SQL Server and also through ODBC.

The bulk of the benefit is on the server side, where the following benefits are available:

- Network Independence. The SAS System supports multiple networking protocols, including TCP/IP and IPX/SPX for intersystem communication.
- Windows NT unified logon. After logging on to a client system and validating user rights and password, that information is used throughout the NT network. In this way, separate passwords are not required on remote systems, nor do passwords need to be sent across the network.
- Operation as a Windows NT Service. The SAS/CONNECT spawner is a Windows NT Service, meaning that it can be started at server boot time, independent of user logons. In addition, the service can be remotely administered by the site administrator on another Windows NT system. SAS/SHARE software is also run as an NT Service, allowing SAS/SHARE servers to be automatically started at server boot time. The SAS/SHARE service can also be administered remotely, meaning that SAS/SHARE can be started and stopped remotely, and in an orderly fashion.
- SMS support. The BackOffice certification requires that the application be installable through Microsoft's Systems Management Server (SMS). This is done through package information provided by SAS Institute, describing to SMS how the SAS System can be remotely installed in unattended mode.


Interoperation with SMS means that site administrators can more easily install and uninstall the SAS System (or updates) from a central location without having to run an interactive Setup session on each system. Managing installations through SMS also enables inventorying of installed software more easily. Use of SMS is a good way for administrators to easily deploy new software in large organizations, and the SAS System is fully supported within SMS.

Windows NT Support


In addition to enhancements made for the BackOffice logo, additions to the system have been made in the following areas:

- Performance Monitor support. Operation of the SAS System can be monitored in the Windows NT Performance Monitor. Performance Monitor is a powerful utility shipped with Windows NT that allows various performance metrics and counters to be displayed graphically or captured in log files for later analysis. The SAS System includes support for specific counters within the SAS System that can help administrators and site engineers determine whether optimal performance is being delivered by the SAS System during various operations. Performance Monitor allows knowledgeable users to do real-world performance analyses with their actual data and application workload. Performance Monitor can access data about the local system as well as from other NT systems on the network. Performance Monitor operation doesn't compromise the operation of the SAS System. During normal operation, the SAS System will update internal counters during various types of operations. These counters are stored in a separate memory location also accessible to Performance Monitor. When monitoring SAS System operation, these counters are interpreted by SAS System code that runs within Performance Monitor to decode and analyze the values.

Using Performance Monitor with the SAS System is straightforward. The first step is to select counters to monitor, as shown in the figure below.



Next, use the SAS System normally. While Performance Monitor is running, it will capture information about the operation of the SAS System. In the figure below, a sequence of data step operations were running.



- NT Event Log support. Various diagnostic messages, such as Setup program configuration changes, remote user connections, and runtime errors can be logged to the Windows NT Event Log. Event Log is another administrative utility that is shipped with Windows NT that acts as a central repository for system and application errors or configuration updates. The SAS System supports writing diagnostic and configuration changes messages to the event log. Event log messages can be reviewed locally or for other systems on the NT network.
- Four gigabyte tuning support. The SAS System has been built to allow access of up to 3 gigabytes of RAM when running on NTS/E.

- Microsoft Management Console Support. With NT 5, MMC replaces some functionality of the Windows Control Panel, and provides a way to generate application-specific consoles for administration, monitoring, and general information about an application. The SAS System supports MMC with a “snap-in” that provides a custom application for interacting with the SAS System installed on a particular system (local or remote).
- Microsoft Systems Management Server Support. To the administrator on an NT Server, SMS is a good way to distribute and inventory software packages on the systems in a large organization. The SAS System functions as such an SMS package.

Performance Improvements

Since the 6.12 Release, SAS Institute has worked to deliver increased performance for the SAS System on Windows NT. The Institute was assisted in this effort by Intel Corporation, through their ASC lab. Performance Engineers at Intel helped test the SAS System in a wide range of areas to determine ways in which performance could be improved. During this process Institute developers ran many SAS code streams on testing systems and analyzed results. When results pointed to a potential bottleneck, a performance drill-down utility called V-Tune (an Intel software package) was used to help the developers identify areas that could be made more efficient.

The compiler used to build the SAS System was tested and compared with the best commercially available, and found to be either as good or superior to it, depending on the area. Nonetheless, additional performance improvements were made to further boost performance on PCs.

In addition to targeted performance investigations, Version 7 also takes advantage of ways to make use of NT-specific I/O optimizations. In the areas of sequential I/O as well as random I/O, NT-specific enhancements were made to substantially boost throughput.

Finally, build-time “basing” is now applied to allow executable images to load from disk more efficiently.

Performance Considerations

Consider the following when determining how to obtain optimal performance using the SAS System on Windows NT:

- Memory Bottlenecks. Does your system have enough RAM to run the application(s) required? Extra memory allows NT to cache recently or frequently used files in its disk cache, greatly improving performance. Insufficient memory leads to swapping, which can greatly hurt performance.
- CPU bottlenecks. If your system runs several applications, ensure that the CPU(s) aren't overworked. Waiting for processor time means that your application(s) aren't running, but rather waiting for a chance to run.
- I/O bottlenecks. Since SAS processing is usually highly data intensive, these bottlenecks can be the worst for performance. Similar to CPU bottlenecks, when applications are waiting for I/O, they aren't running, and the longer they wait, the longer they take to complete. When several applications are running, I/O bottlenecks can become the single most important factor inhibiting optimal performance.
- Network bottlenecks. These come in the form of bandwidth limitations and latency. Bandwidth is exhausted when the network is moving data at its rated speed. Beyond that, some data has to wait. Latency is also an issue for applications that read data piecemeal, since each read has to be processed by the network and remote data source before it can be sent back to the application.
- Disk configuration. Striped disks (also known as RAID 0) are a good idea. Do the striping in hardware on a RAID disk controller. Fast/wide ultra SCSI gives maximum bandwidth per channel, and large systems incorporate multiple SCSI channels to prevent channel saturation. Striping across more drives (spindles) can improve performance, and the ideal is to stripe across channels on a large number of spindles.

- Disk fragmentation. It happens. As data are written and re-written, the file layout becomes fragmented. I/O that would normally be sequential becomes essentially random. Defragment disks from time to time. Even though NT doesn't include a defragmentation utility, it doesn't mean that defragmentation wouldn't help.
- Understand your system. Use Performance Monitor to look for bottlenecks and resource limitations.
- Understand data layout. A large number of disparate, large data files on a single disk volume probably means that nobody will get optimal access to their data. Partition data as much as possible across disk partitions.
- Understand application loads. A large number of processes sharing the same server can overload it in any bottleneck area. An NT Server acting as a SAS application server, a SQL Server, a web server, file server, domain logon controller, and an e-mail server is probably severely overloaded, and performance will likely suffer.

Comparing NT Server to NT Workstation, Release 4.0

NT Workstation and NT Server are built upon the same code base and execute the same programs. Typically, not even special server versions of programs are required to run on NT Server. What then are the basic differences in NT Workstation and NT Server? NT Workstation is optimized for the client workstation and for power users, while NT Server is optimized for system administration of NT based networks and for setting up large server applications.

NT Server is generally tuned for different memory allocations than NT Workstation. A typical configuration on NT Server allows a large percentage of RAM to be allocated to the disk cache, which keeps some disk contents in memory to speed access.

Multitasking behavior on NT Server favors longer runtimes before interruption. NT Workstation tends towards shorter runtimes (and higher priority for the foreground interactive application) so that interactive applications appear to be more responsive.

While NT Workstation supports a maximum of two processors, NT Server supports up to four (or eight on NTS/E).

Also on NTS/E, applications are allowed to access up to three gigabytes (GB) of application memory. This is not possible on NT Workstation.

For network administrators, NT Server is capable of controlling a network domain, the unit of security and management (such as validating domain logins) on NT networks. NT workstations can only be participants in a domain, not its controller.

Some applications and server packages are only available with NT Server. Many of them refuse to run on NT Workstation.

Hardware Considerations

In general, Windows NT Workstation is the preferred environment for SAS data analysis and end user application development. It is highly recommended for advanced development with the SAS System, to have at least 16 additional megabytes of memory available for use by only the SAS System, above what is required for the operating system. Typically, this would suggest a minimum of 32 Megabytes memory on the NT Workstation. 64 Megabytes performs exceptionally well. An Intel Pentium Pro 200 MHz processor is highly recommended, but a 200 MHz Pentium also is very acceptable. Multiple processors should be considered especially if performing concurrent SAS processing, or if substantial concurrent processing is required by other applications. (When purchasing new hardware, check to ensure the motherboard allows for additional processors to be added.)

When using Windows NT Server, the size of machine is dependent on the type of processing and serving being provided, as well as the other applications to be executed. NT Server machines can be configured at the low end (32-64 MB

RAM and 4 GB disk storage) to the high end (4GB RAM and >100 GB of disk storage). In general, for optimal performance using the SAS System on Windows NT Server, at least dual Intel Pentium Pro 200 Mhz processors and at least 64 Meg memory is recommended. Provide 16 Megabytes of memory for the SAS System and an additional 12-16 Megabytes of memory for SAS data processing.

The disk space required by the SAS System varies widely, but can range from under 100 Megabytes to 500 Megabytes for the entire system. The product can be custom installed for less disk utilization. On-line books and geographic map data are also available on separate CD-ROMs, and will require more space if installed to disk.

Care should be taken reviewing the hardware requirements for Windows NT. Windows NT is not supported across all Intel based PCs. *The Windows NT Hardware Compatibility Guide* should be consulted to determine compatibility with existing hardware

Future Directions

The future on the NT platform means Windows NT Version 5.0. A summary of expected improvements follows.

- Hydra. Out of the box, NT Server will support terminal emulation. Currently available technology known as WinFrame (from Citrix) provides similar functionality for NT 3.51. No support of this type is available on NT 4.0.
- Multiple Monitor Support. Each PC can support more than one monitor, intended for the use of a single user who needs more screen real estate. Monitors can be of varying resolution and color depth. Applications can run on any monitor, and can be moved across monitors.
- Clean Application requirements. NT 5 will support a configuration requiring that all installed applications be "clean". A clean application is one that doesn't copy files into

the Windows system directory, keeps the user's data separate from the program files, and generally doesn't "dirty" either the system area or the user's part of the system. A clean application is easier to install and manage, and reduces the total cost of ownership (TCO). TCO is a major concern in some organizations.

- Migration DLLs. Vendors of applications on pre-NT 5.0 operating systems are expected to provide a migration DLL that upgrades the application to one that runs properly on NT 5.0. Without a migration method, the user would have to reinstall any application after upgrading the operating system.
- Microsoft Management Console. MMC is a means for applications to develop administration modules (called MMC snap-ins) to manage all aspects of configuration and general operation of their application.
- Web-based Enterprise Management WBEM defines mechanisms to model managed systems. WBEM will provide a management infrastructure, a single representation of information, and integration of management data from multiple resources (i.e. disk drive space free, network card status, etc...). WBEM unifies management data administration and simplifies management applications by providing a common API to access information throughout a heterogeneous enterprise. WBEM also allows for SQL-like queries to be performed on system resources across the enterprise on-the-fly.
- Active Directory . The next generation Windows NT Directory Services is a directory service integrated with Windows NT Server that offers the hierarchical view, extensibility, scalability, and distributed security required by all customers, no matter the size. Active Directory will allow a single point of Administration for all published resources, which can include files, peripheral devices, host connections, databases, Web access, users, arbitrary other objects, services, and so forth. This next generation directory service can subsume and manage application-specific directories, as well as other NOS-based directories, to provide a

general purpose directory service that can reduce the administrative burdens and costs associated with maintaining multiple name spaces.

- Security. Microsoft has added new security mechanisms and unified its security model with NT 5.0. As NT 5.0 embraces enterprise-level computing, features such as Public Key Authentication(Kerberos support), single enterprise logon, and decentralized administration for large domains will be added. Security for all system objects, access verification, auditing, and cryptographic support has been enhanced. Windows NT Distributed Security services now interacts with the Directory Service for access rights across the enterprise. New security administration tools have also been added to support enterprise-wide administrative access to systems.
- Hierarchical Storage Management. HSM is an enhancement to NT's data storage capabilities. With HSM, NT creates virtual disk storage out of tape drives, optical drives, and other storage media when NT runs out of RAM. Also two other powerful features are added to allow for multi-level HSM systems, reparse points and volume-mounting. The reparse tool in HSM lets users create a new directory by gluing together two existing directories. NT 5 uses volume-mounting utilities to let users store data in whatever drive they choose, no longer dependent on the DOS/critical path method, single-drive-letter model.
- Clustering. Already available in a limited form, clustering provides for more reliability (through machines that step in when another one has failed) and more performance (through shared common computing resources that appear as one unit). Although currently the SAS System doesn't take advantage of clustering, this is an area we'll explore fully in the future.
- 64-bit support. Intel is working on the Merced (IA-64) processor to move into the 64-bit application space. SAS Institute is committed to delivering a 64-bit version of the SAS System on Windows NT when that processor becomes available. The 64-bit processor will boost performance by virtue

of increased parallelization, higher speed, and larger data path. The processor will boost scalability by virtue of eliminating 32-bit limits such as 4GB filesize issues and 4GB RAM limits.

- NTFS 5 enhancements. Support for a concept similar to UNIX filesystem mount points and symbolic links (called junction points in NT 5.0) allows NTFS 5 filesystems to join multiple drives into tree-oriented configurations.
- Encrypting Filesystem. Encryption on the filesystem ensures data security, even at the hardware level. This is available in NTFS 5.

CONCLUSION

Version 7 of the SAS System for Windows targets Windows NT Workstation and Windows NT Server Version 4.0. The Institute's MVA Architecture provides for the optimal client server application to be developed, because it provides for the possibility to have specific processing or data access, placed on either the client or server.

SAS Institute will continue investing in research and development for Windows NT and can expect to further embrace Windows NT Server development.

ACKNOWLEDGEMENTS

Robert Hamm and Tracy Carver of SAS Institute contributed to the preparation of this paper.

Trademarks

SAS, SAS/ACCESS, SAS/CONNECT, SAS/SHARE, and SAS/SHARE*NET are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand or product names are registered trademarks or trademarks of their respective companies.

REFERENCES

Cates, Mark (1997), "Taking Advantage of the SAS System on Windows NT", *Proceedings of*

the Twenty-second Annual SAS Users Group International Conference

Clegg, Jennifer and Mehler, Gary (1998), "The SAS System for Windows, Version 7", *Proceedings of the Twenty-third Annual SAS Users Group International Conference*

Mehler, Gary (1998), "Integrating Windows Clients and the SAS System into the Enterprise", *Proceedings of the Twenty-third Annual SAS Users Group International Conference*

Microsoft Corporation (1995), *The Windows Interface Guidelines for Software Design*, Microsoft Press, Redmond, WA, Microsoft Corporation

Author Information

Gary Mehler
SAS Institute Inc.
100 SAS Campus Drive
Cary, NC 27513
(919) 677-8000
gjm@wnt.sas.com