

## Feeling Anxious about Transitioning from Desktop to Server? Key Considerations to Diminish Your Administrators' and Users' Jitters

Kate Schwarz, Margaret Crevar, SAS Institute Inc.

### ABSTRACT

As organizations strive to do more with fewer resources, many modernize their disparate PC operations to centralized server deployments. Administrators and users share many concerns about using SAS® on a Microsoft Windows server. This paper outlines key guidelines, plus architecture and performance considerations, that are essential to making a successful transition from PC to server. This paper outlines the five key considerations for SAS customers who will change their configuration from PC-based SAS to using SAS on a Windows server: 1) Data and directory references; 2) Interactive and surrounding applications; 3) Usability; 4) Performance; 5) SAS® Metadata Server.

### INTRODUCTION

Many organizations that have successfully transitioned to using SAS® on a server initially encountered strong resistance from experienced SAS programmers who want to have personal control of everything, particularly copies of data and programs on their own local drives. What are the reason these organizations are "forcing" their SAS users off their beloved desktops? With the current federal security demands on an organization's data, most heavily regulated industries such as financial services, communications, and health care discourage this practice because its inherent security exposures and redundancies do not conform to government regulatory compliance plus other industry practices and legal constraints.

The solution to the regulatory restrictions is deploying a centralized server environment that prevents users from having local copies of programs, data, reports, and intelligence on their laptops or desktops that might be redundant, unsecured, outdated, customized, error-prone, or lost.

When proficient SAS programmers and administrators hear that PC-based SAS or the SAS® windowing environment will no longer be available, their most common fear is that they will completely have to change the way they work with SAS. They worry that their skills, programs, reports, models, and other analytics resources and assets from PC-based SAS will no longer be viable. *This is a misconception; in a server environment, these assets are even more valuable.*

While it is true that interfaces such as SAS® Studio and SAS® Enterprise Guide® provide robust point-and-click capabilities to generate SAS code, it is important to realize that generated code cannot duplicate everything that is possible within the SAS programming language. SAS Studio and SAS Enterprise Guide are simply different interfaces with distinct advantages over the SAS windowing environment that many PC users are accustomed to using.

In addition to offering a robust programming environment, SAS Studio and SAS Enterprise Guide provide tasks and wizards that can increase productivity. *All SAS code, whether submitted from PC-based SAS (either interactively or in batch), SAS Studio, or SAS Enterprise Guide, is executed by the same sas.exe process behind the scenes.* Figures 1 and 2 depict the Program Editors of these interfaces.

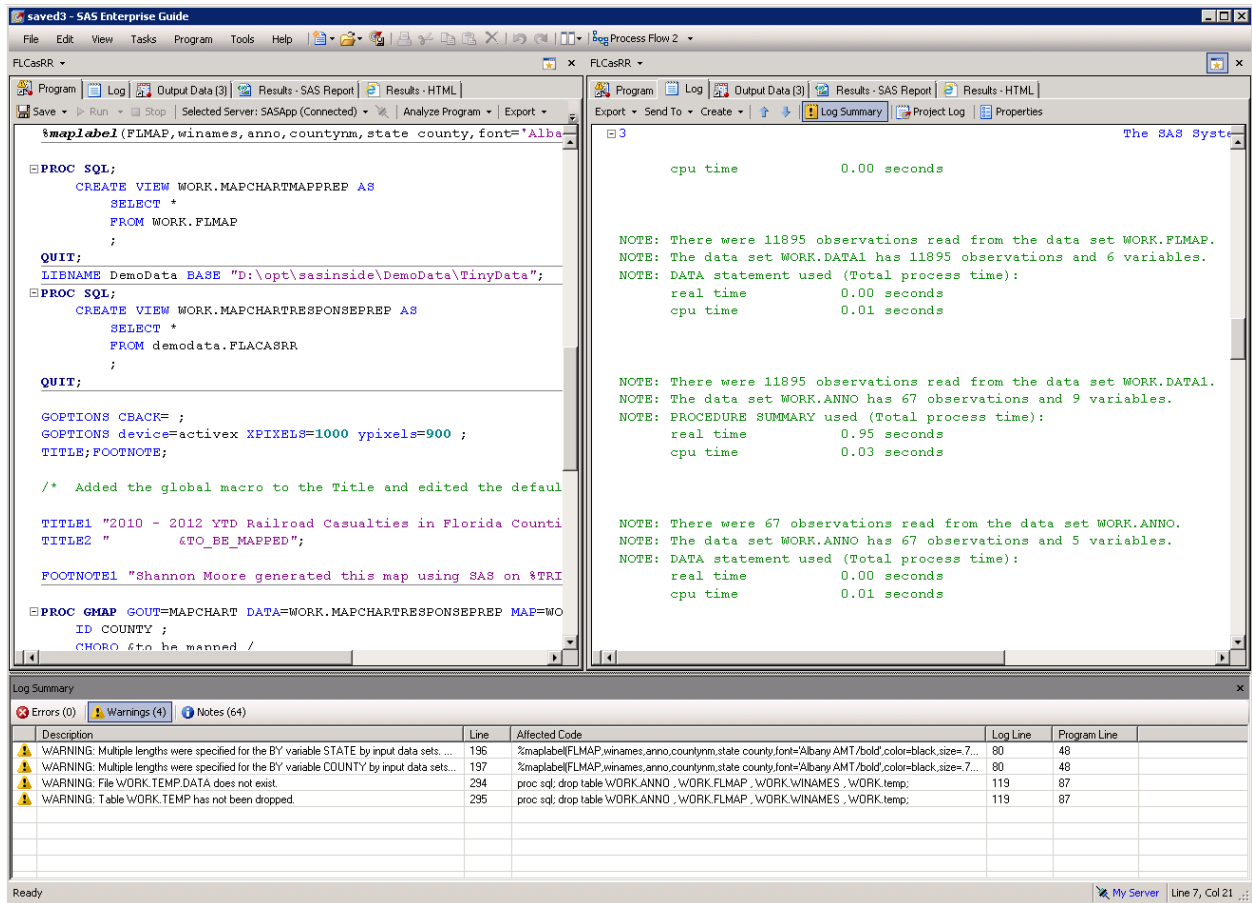


Figure 1: SAS Enterprise Guide Program Editor

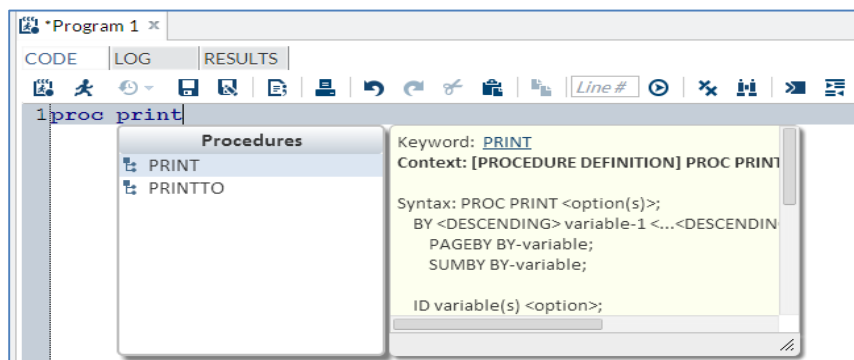


Figure 2: SAS Studio Program Editor

Because this transition involves moving from a decentralized environment to a centralized one in which all processing is done on the server, users and administrators need to keep in mind that all operations are conducted from the server's perspective.

A comprehensive review of the current SAS environment and the new server configuration is necessary, with special attention to the following five key considerations:

1. Data and directory references

2. Interactive and surrounding applications
3. Usability
4. Performance
5. SAS® Metadata Server

## WHAT IS THE SERVER'S PERSPECTIVE?

Before focusing on the five key considerations, this section outlines the server's view of its resources.

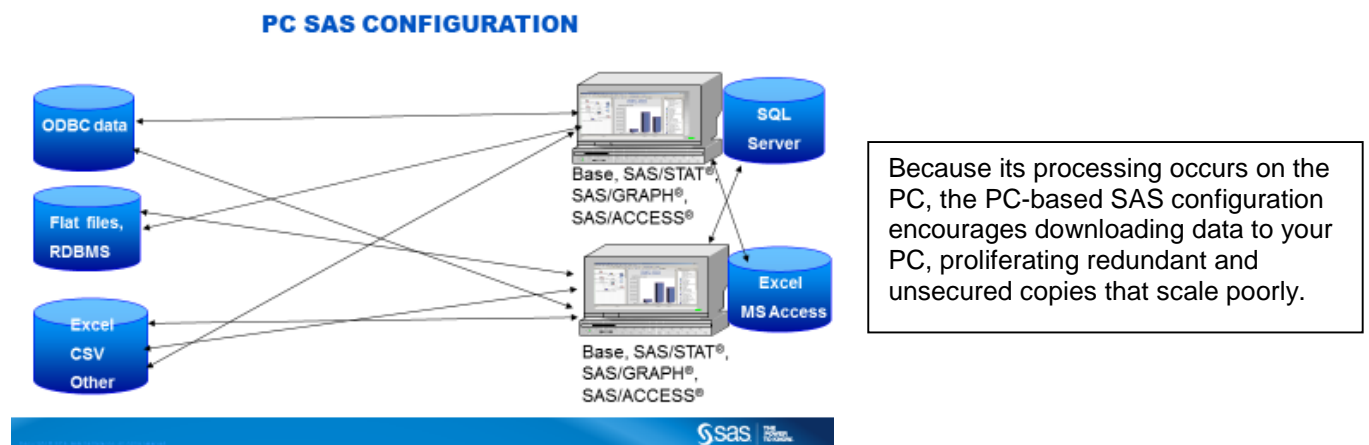
In order to understand how a server "sees" or references its resources, we consider how PCs "see" or reference their resources, and then compare and contrast the two. Going from a PC to a server is a lot like changing from being a driver of a private car (decentralized PC) to becoming a passenger on a commercial aircraft (centralized servers).

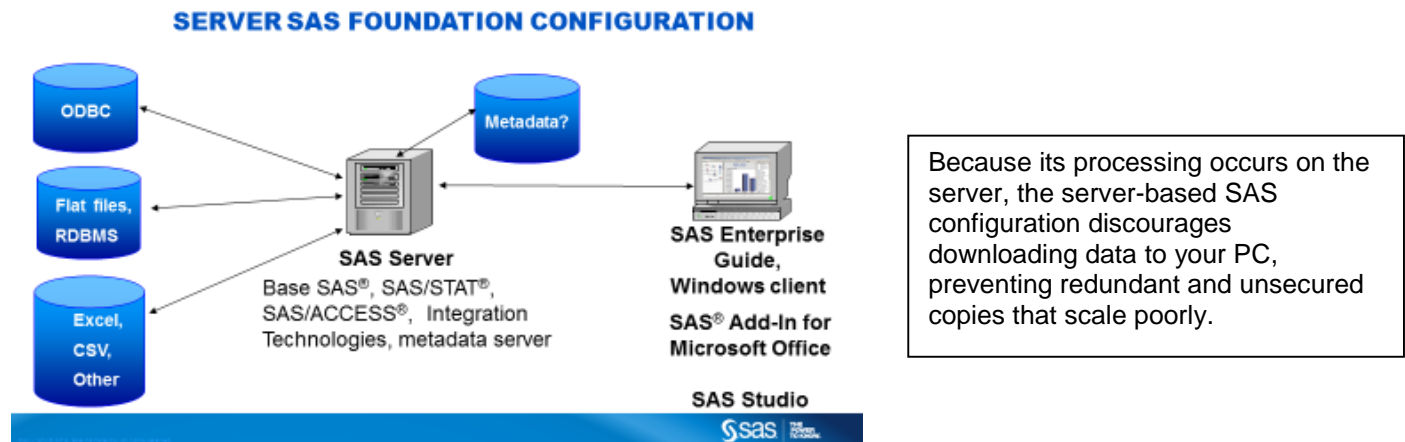
Driving a car is akin to using PC-based SAS because they have similar individual controls and constraints; being a commercial aircraft passenger is akin to using server SAS because of their centralized control, administration, and scalability.

Commercial aircraft and servers' advantages and consolidated resources become apparent when we consider how they are both accessed, coupled with how fast and far they let us, our belongings, and data travel:

- Centrally controlled airports allow public access to the aircrafts' great distance, speed, and cargo capacities.
- Centrally controlled servers allow broad team and enterprise access to vast computational power, speed, and storage capacities.

Note the PC- and server-based SAS configuration topologies' differences in Figure 3. In the PC-based SAS configuration, the SAS modules are installed on each PC, with data sources and targets stored in many places—this is a highly decentralized arrangement. In the server SAS configuration, the SAS modules are installed on a server, and a metadata server typically (and should, as a best practice) accompany this arrangement for central control of all SAS resources. The server SAS configuration resembles a hub-and-spoke, like many airports' terminal and gate arrangements.





**Figure 3: PC-Based SAS and Server-Based SAS Topologies**

The PC-based SAS configuration in Figure 3 resembles much of our personal transportation, especially when you deploy a private car to and from multiple starting points and destinations. When you are driving a car, you control your route, destination, departure time, number of passengers, cargo, maintenance, and so on. This individual level of control is like the PC-based SAS configuration; you control how and where you store and reference your data and programs. You conduct input and output operations based on the resources available on your PC.

The server SAS configuration in Figure 3 resembles commercial aircraft transportation, where your position is referenced according to your position in the aircraft (for example, seat 32A). In this scenario, you are sharing the centralized passenger cabin, engines, crew, cargo space, security, and mechanical systems with dozens or hundreds of other passengers. The aircraft's centralized resources are much like a server's resources—they are both regulated and controlled by other parties on behalf of the users, or passengers.

Just as our global economy cannot thrive by using private land transportation alone, so too do most organizations require the greater scalability, security, and speed of public air transportation and servers for their physical and digital logistics. Because most organizations experience rapid growth (especially in terms of data, users, employees, partners, contractors, customers, and sales), coupled with demands for more real-time intelligence delivery and self-service dissemination, the limited scalability available with PC-based SAS quickly becomes a barrier to transport analytics further and faster. As a result, a decentralized PC configuration is ill-equipped to deliver insight and foresight in a timely manner for most domestic and global organizations.

How then does the server (and its SAS users) see or reference its resources? It's a lot like swapping your driver seat in a private car to become a passenger on a commercial aircraft. A server's perspective is a centralized view of data sources and targets. *A server-based SAS configuration views its resources centrally much like airplane passengers are referenced via their seat's position in the cabin. Servers' centralization helps us achieve speed, scalability, security, and simplicity.*

Table 1 compares and contrasts PC and server key characteristics.

	PC	Server
<b>Purpose</b>	Individual computational processing	Group/team/enterprise computational processing
<b>Constraints</b>	Hardware storage, memory, and computational resources	Can house and harness more hardware storage, memory, and computational resources than a PC Can scale to add more storage, memory, and computational resources than a PC Processor speeds can exceed and trail PC processor speeds
<b>Key Activities</b>	Owner controls how/where/when/what Owner controls frequency and extent of ongoing maintenance Owner supplies locations; laptops provide mobility and portability Storage and computational capacities are shared with none or a few others	Administrators control how/where/when/what Administrators control frequency and extent of ongoing maintenance Centralized hardware and peripheral locations Storage, capacity, and utilization are shared with dozens, hundreds, or thousands of others
<b>Security</b>	Some physical and electronic barriers and deterrents, which are easily breached	Extensive physical and electronic barriers and deterrents, which are more difficult to breach
<b>Common capabilities and characteristics</b>	House and harness storage, memory, and computational resources Support and conduct transactional and analytics systems	
<b>Unique Capabilities</b>	Laptops' portability and ability to run for hours without electricity Can operate without connectivity or special indoor temperature-controlled facilities	Provide mission-critical resources for most private- and public- sector organizations' operations and digital assets, transaction, and analytics systems Support cloud and in-house computing Extensive business continuity, backup, and restores processes and resources.

**Table 1: PC and Server Key Characteristics**

## WHAT ARE THE FIVE KEY CONSIDERATIONS?

Now that we understand that transitioning from using PCs to using a server is a bit like swapping a private car for a commercial airplane, we have to plan for our trip. What will we bring along? How and where will our belongings be secured, stored, and shared with other passengers and traveling companions? What can stay at home? These same questions apply to our SAS resources; we need a way to categorize the aspects of our PC environment that will transition to a server.

The following five key considerations give structure to our planning process. These guide the migration from the current PC-based SAS environment to a new server configuration.

1. Data and directory references
2. Interactive and surrounding applications
3. Usability
4. Performance

## 5. SAS® Metadata Server

To relieve your anxiety about these five aspects, please note that in most cases, well over 90% of SAS code is platform-independent; far less than 10% of the code depends on the host environment, such as FILENAME and LIBNAME statements.

### DATA AND DIRECTORY REFERENCES

Because moving to a Microsoft Windows server includes a different architecture with new directories, storage locations, permissions, and possibly a newer SAS release, the server operations must be able to recognize its resources in the new operating environment.

In most cases, users will access centrally stored and maintained resources instead of owning and maintaining their own copy of files. The following guidelines accommodate operations such as reading, writing, and updating data that was created elsewhere or that is remotely mounted on your network.

- Conduct a comprehensive code review for all external file references. For example, files with external references need to be altered, because a PC's reference to C:\[local directory] becomes [server drive letter] \[drive directory], and users' mapped drives might not exist on the server, so Universal Naming Convention (UNC) paths might need to be used.
- Host-specific references such as INFILE, FILE, FILENAME, LIBNAME, ODS, %INCLUDE, libraries, formats, catalogs, and so on, need to be changed so that the server can refer to them.
- SAS data sets that were created in Windows and that will be stored on UNIX or Linux MUST be accommodated by one of the following remote access methods:
  - Cross-Environment Data Access (CEDA) (default method)
  - CPORT and CIMPORT procedures
  - XPORT engine with DATA step or PROC COPY
  - XML engine with DATA step or PROC COPY
  - Data Transfer Services (DTS) in SAS/CONNECT®  
For more information about these methods, see the topic "Strategies for Moving and Accessing SAS Files" in *Moving and Accessing SAS® 9.4 Files*.
- The storage location for confidential data that is secured only on a laptop or desktop will need to be reconsidered. Alternatively, keeping a few copies of PC-based SAS might be the best option if this change will be too disruptive or costly. See "Appendix A: When Should a PC Be Deployed Instead of a Server?".
- For database connections, LIBNAME statements might need to be re-specified into metadata server definitions that take advantage of the database's credentials. If metadata database libraries are defined, remove these LIBNAME statements from your programs so that the pre-defined librefs in metadata will prevail.
- FTP calls on the filename can fail if the server allows only secured FTP connections.

### SAS/ACCESS® Interface to PC File

If SAS/ACCESS Interface to PC Files was licensed alongside PC-based SAS foundation technologies, and it is no longer available when moving to a server, SAS Enterprise Guide can read Microsoft Excel and other PC file formats as long as LIBNAME and the IMPORT and EXPORT procedures are adjusted accordingly. We suggest using SAS Enterprise Guide to import and export data tasks.

If SAS/ACCESS Interface to PC Files is licensed on the server, its code might need to be modified because many PCs are 32-bit and most servers are 64-bit.

### 32-bit PC File Formats

If there are a lot of 32-bit PC files, a fast way to handle them is to install PC Files Server on the server.

If there are few 32-bit PC files, installing the PC Files Server and SA® Enterprise Guide on a workstation can be more cost-effective. This separate workstation can be used to ensure that SAS Enterprise Guide

scheduled jobs will be correctly executed. This approach can also ensure that the PC Files Server is not on a machine that might be powered down.

Once you have identified all of the files and libraries used by your SAS applications, remove them from the SAS code or jobs themselves and place them in either the SAS Metadata Server or the autoexec.sas file. Either way, they are in a single location that will make changing their location much simpler in the future.

## **INTERACTIVE AND SURROUNDING APPLICATIONS**

PC-based SAS capabilities that require user interaction will need to be accommodated on the server. Anything that pops up a window or a prompt, such as DEBUG, DM statements, PROC FSLIST, PROC FSEDIT, SAS/AF® applications, are some of the main areas for assessment and adjustment. Keep in mind that these situations will be capturing and using information about the server, not on the user's PC.

Frequently encountered situations involving user interaction are outlined in this section.

- Operations that invoke DOS windows to run DOS commands but that require no user interaction can be converted to equivalent SAS Enterprise Guide functionality, such as prompts.
- The DATA step debugger does not run in SAS Enterprise Guide. The debugger is often used at a cursory level in isolated incidents, such as in unit testing. Weigh whether there are sufficient scenarios to justify keeping a few copies of PC-based SAS or allow the SAS windowing environment to handle these instances.
- Code that includes automatic email might need to be altered depending on the SAS server's ability to harness email services. SAS Enterprise Guide and SAS Studio have exporting and "send to" features that can provide equivalent capabilities.
- BAT files might need to be modified because the .bat files will execute on the server and will be using server data, environment information, or both, rather than local data and environment information.
- Custom scripts, environmental controls, and system controls surrounding PC-based SAS operations will also need to be accommodated. Some operations such as scheduling might need to be redeveloped and tested, especially if third-party technologies are involved.
- Be aware that -NOXCMD is often the default setting with a server. As a result, the following operations might fail unless this setting is altered:
  - %SYSEXEC
  - CALL SYSTEM
  - X
  - PIPE
- Users will likely require shell access on the server for the following scenarios:
  - Using SAS/CONNECT with RSUBMIT code
  - Updating or executing custom scripts, such as entering passwords
  - XCMD permissions

### **Dynamic Data Exchange**

Dynamic Data Exchange (DDE) operations can often be redirected to workarounds such as PROC REPORT, ODS, and other modern ways to work with Microsoft Excel in SAS Enterprise Guide.

- For some scenarios in which PC-based SAS involves interactive DDE, there can be a fair amount of change, adjustment, or adaptation. In these cases, a hybrid configuration of using a client/server environment and keeping some licenses of PC-based SAS for heavy DDE usage can be viable.
- Review the following SAS note and blog entry:
  - <http://support.sas.com/kb/20/767.html>

- <http://blogs.sas.com/content/sasdummys/?s=DDE>

## **SAS/AF applications**

SAS/AF applications used in PC-based SAS will not run in SAS Enterprise Guide and should be converted to stored processes or other applications. Using SAS Enterprise Guide as a development environment for stored process is recommended because SAS Enterprise Guide contains specific guidance for creating and maintaining stored processes.

Creating custom tasks in SAS Enterprise Guide or SAS Studio is another option that allows high interactivity and more dynamic operations. The SAS/AF user interfaces will be modernized and they will often gain a much better look and feel once they become stored processes. SAS/AF is also functionally stable, so converting these applications also aligns with ongoing support for stored processes provided by SAS.

## **USABILITY**

Because moving to a Windows server includes a different architecture that relies heavily on organizational standards and administrators' deployment decisions, administrators must clearly communicate and set users' expectations regarding how these server settings will affect users. Administrators can best support SAS users by allowing a constant connection to the SAS workspace server. Assessing and altering firewall settings are also involved here.

## **Firewalls and idle time-out effects**

The SAS server or a firewall might have short idle timeouts. These settings can cause the following effects on how users become accustomed to working in SAS Enterprise Guide:

- Long-running jobs and your SAS session on the server can stop, and your WORK data can be deleted.
- Servers that have short idle times might cause users to lose their WORK tables. This situation often occurs when users manipulate WORK tables or spend long amounts of time coding in SAS Enterprise Guide. The server's short idle time can cause your WORK area be deleted every few minutes, and you might lose your SAS Enterprise Guide connection to the server.
- The server can be running in a different time zone. Thus, the date/time could be different from the user's workstation (where date="&sysdate9."d). The server's sort sequencing and support for international options can also differ from individual PC installations and therefore affect results. Clear communications regarding these server settings must be shared with the SAS users.
- Code that calls X commands or SYSTASK might not work unless administrators grant appropriate permissions. Assess the needs of users who are capable of handling shell commands and accommodate them accordingly. Please review this blog entry for more details:  
<http://blogs.sas.com/content/sasdummys/2012/08/31/the-case-for-xcmd-privileges-in-sas-enterprise-guide/>
- Because ABORT or ENDSAS calls terminate the connection between SAS Enterprise Guide and the server, consider removing or commenting out.
- The server's permissions might prevent SAS Enterprise Guide users from having a server location to write to other than WORK. Use the WORK path in the external file references and transfer those files before disconnecting SAS Enterprise Guide from the server.
- Alternative means of providing a scrolling log in SAS Enterprise Guide include:
  - submit each step separately and review its log
  - turn on the SAS Enterprise Guide project log
  - use the SYSECHO statement in programs to highlight your progress



## PERFORMANCE

This section outlines the main areas for planning and implementing the proper server resources that ensure happy SAS users. Understanding how SAS operates on a server is also covered in this section, with special emphasis on consumption of the I/O subsystems, file systems, and memory. SAS users and administrators will need to review and characterize their SAS tasks and SAS interactions in order to set up a Windows server that viably supports SAS jobs.

Server SAS often accompanies many new ways of executing SAS, including web browser and/or Java front-end applications to make it easier for business users to leverage the power of SAS. This change leads to needing more computer resources to support these new SAS sessions.

Each site should evaluate its overall system performance to determine which architecture aspects are candidates to run on separate instances of an operating system. This approach ensures having sufficient computer resources that avoid performance degradation. More details are provided in this paper "How to Maintain Happy SAS®9 Users," available at <http://support.sas.com/resources/papers/proceedings09/310-2009.pdf>.

With a PC, I/O is controlled individually. In RDBMS, it is typical to have only a single instance of the relational database management system (RDBMS) running at any given time. With server SAS, it is common to have many SAS instances running at any given time. Generally, SAS users start their own SAS session for each SAS job and application they are running. SAS creates a high volume of I/O, mostly as sequential reads and writes, so SAS data interactions are significantly different from RDBMS.

The most important task is a detailed assessment of how the PC users are using SAS. For example, if you have 50 PC users executing SAS jobs on their desktop that run for hours and consume at least one full core on their desktop, then you will need to make sure that the server can support 50 concurrent SAS analytical jobs at the same time.

In addition to making sure you have enough cores to support all the users in this new shared environment, you need to make sure that you have enough I/O throughput on the new server to support the concurrent SAS sessions. Again, understanding the I/O throughput requirements of each SAS analytic desktop user will make this task easier.

SAS has hardware sizing and performance experts on staff who can guide your efforts to properly size your server environment. Contact your SAS account team or the authors if you want to discuss your specific scenario in more detail.

You are also encouraged to review the following materials:

- Troubleshooting system performance information, available at <http://support.sas.com/kb/42/197.html>.
- SAS system option settings for best performance in UNIX and Windows environments, available at <http://support.sas.com/kb/46/954.html>
- Because SAS uses the operating system's file cache to read and write data, the maximum I/O throughput rate is restricted by how fast the operating system's file cache can process the data. The maximum I/O throughput rate we have seen via file cache has been near 1GB per second; more commonly it is around 600 MB to 800 MB per second per operating system instance. Discussions regarding how to enable SAS to do direct I/O are beyond the scope of this paper, but there are ways to make this happen.
- A large number of temporary files, ranging from very small (under 100 MB) to very large (tens or hundreds of GB) can be created during long-running SAS jobs. The location of these files is referred to as the SAS WORK area. *SAS WORK must point to a file system that has enough I/O throughput to support your users' I/O demands, as well as enough disk space to support all the temporary files that will be created during each SAS session.*

- You will have to decide between a single, large file system for SAS WORK that all SAS users will share or multiple, smaller physical file systems (not different directories on the same file system) for SAS WORK that you round-robin your SAS users between.
- SAS does not pre-allocate storage when initializing or when performing writes to a file. When SAS creates a file, it allocates a small amount of storage, but as the file grows during a SAS task, SAS extends the amount of storage needed. Since file extensions are limited to the amount of available space within the file system currently being used, you have to closely monitor it to prevent consuming the file system.
- Analytics procedures such as PROC REG, MIXED, GLM, NL MIXED, GENMOD, and HPF procedures tend to be more CPU-intensive in nature. A single job that contains these procedures can consume an entire core within your computer until it has finished its processing.
- You will need to determine how many long-running analytical tasks will occur simultaneously and make sure you have sufficient CPU cycles to support them along with your other SAS users.
- Matrix manipulation tasks such as IML, GENMOD, and MIXED tend to be memory-intensive. You will need to have enough physical memory to support all of these tasks plus sufficient file cache to hold the data of the other SAS tasks that might be running in memory.
- Because the SAS Metadata Server stores its specifications in an in-memory table, administrators must ensure that there is enough physical RAM in the computer to keep this in-memory database resident. When users run queries and updates, their requested data is written to memory. With these demands, administrators have to ensure that there is enough physical memory available. If physical RAM is exceeded, paging will ensue, which degrades performance. You might want to consider dedicating a separate server to host the SAS Metadata Server.
- For the logical SAS servers (SAS® Workspace Server, SAS® Stored Process Server, SAS® OLAP Server), you should set up the hardware for these servers the same way you set up the hardware for the legacy batch jobs discussed above because the SAS clients are just submitting SAS code to these back-end SAS servers in a similar fashion to the way the legacy batch SAS jobs submitted code. The only difference is that you might have a SAS server (SAS session) servicing multiple SAS clients instead of a single SAS user the way a legacy batch SAS job did.
- The shared nature of client/server operations can involve some latency, depending on network traffic and server sizing, plus other factors and settings. Knowing how much data will be shared between the multiple tiers of a SAS implementation will dictate the speed of the network between the various components of the SAS infrastructure. In general, 100 MB NIC cards in the desktop systems and 10 GB NIX cards in the back-end systems should be fine.
- SAS Enterprise Guide performance might be adversely affected when you connect to a SAS server over a wide area network (WAN). Even the fastest WAN cannot match the speed of SAS Enterprise Guide on a LAN. Consider deploying Citrix or Remote Desktop Services. If you have SAS Enterprise Guide installed on the same LAN or even on the same server, you can use a Citrix client or the remote login for Remote Desktop Services to access SAS Enterprise Guide.
- SAS Enterprise Guide conducts more automatic operations such as adding ODS statements to create a richer set of results, such as SAS Report outputs. Because these outputs are larger, there can be some longer response times. Review the following three links for more details:
  - <http://blogs.sas.com/content/sasdummy/2011/05/25/sas-enterprise-guide-options-my-favorite-5/>
  - <http://blogs.sas.com/content/sasdummy/2011/12/16/a-sas-options-viewer-for-sas-enterprise-guide/>
  - <http://blogs.sas.com/content/sasdummy/2011/11/22/inspecting-sas-macro-variables-in-sas-enterprise-guide/>

## SAS METADATA SERVER

For some experienced SAS users, this will be your first exposure to using SAS Metadata Server.

Because moving to a server often involves new SAS platform administration responsibilities such as managing metadata, this section outlines some key areas of defining source and target data libraries, structuring folders, and assigning permissions that administrators must conceive, implement, and maintain. See Figure 3 to review the SAS Metadata Server topology.

Because all of the client applications and all SAS servers, such as SAS Workspace Server or SAS Stored Process Server, depend on the SAS Metadata Server, properly configuring and administering the SAS Metadata Server is crucial—your SAS analytics cannot deliver insight and foresight without it.

SAS Metadata Server maintains all the information about the underlying data files, users, stored processes, and so on in a centralized location. The SAS Metadata Server keeps its specifications (referred to as a repository) registered by the SAS applications in an in-memory table.

- Administrators' deployment decisions can allow more data sources and targets to be centrally defined, minimizing the need for individual libname definitions and redundant copies of data.
- The strength of your SAS administration and communication skills greatly influences successful metadata environments.
- Permissions and centrally-defined metadata libraries are suggested as ways to control access to the server resources.
- Outputs can be written to local or remote drives.

## CONCLUSION

A successful transition from running SAS on PCs to running SAS on a server requires much more than properly configured SAS software or programming proficiency. Proper planning and accommodations for the five areas discussed here, coupled with communicating the new server's deployment nuances, will smooth and ease this transition.

Especially for analytics and personally identifiable information, organizations put themselves at unnecessary risk by conducting analytics on personal computers instead of a centrally managed server environment. For the sustained health and wellness of the organization, servers are favored, even if they are small servers. The key is having the computational environment appropriately sized for its workload. SAS has hardware sizing and performance experts on staff who can guide your efforts, so you can diminish or eliminate server-induced anxieties.

Although the five considerations outlined here were derived from numerous customer meetings and advisors, perhaps you have encountered others. Contact your SAS account team or the authors if you want to discuss your specific scenario in more detail. We welcome and encourage your anecdotes of how you overcame those challenges.

## APPENDIX A: WHEN SHOULD A PC BE DEPLOYED INSTEAD OF A SERVER?

A PC-based SAS environment is often a better fit than a server under some or all of the following circumstances:

- You do not have to share the source or target data, reports, and other outputs.
- Individual users work purely independently with just their own data (no sharing).
- Your data volumes are small (up to tens or hundreds of MBs) and they do not grow very much, or they are shrinking.
- You have highly mobile or disconnected analytics staff who operate across disparate locations—something like Peace Corps, with independent auditors, field researchers.

- You conduct analytics in an area with poor energy, cooling, or connectivity resources, such as research projects or military operations in developing countries or remote wilderness areas.
- You are running a very small business or one that lacks heterogeneous data and operational systems. For example, a sole proprietor or small mom-and-pop business with few operational systems and more homogenous data, such as one or few point-of-sale cash registers, or a sole customer database or invoicing system.
- Your data and analytics are purely one-time, fleeting, or temporary, and don't need to repeat, scale, or be secured with much vigor or rigor, such as a once-a-century event.
- Your organization has no centralized data center or is a legacy PC environment, and there is no one to manage the operations of a server.
- If single-job performance is paramount, then a single job will often run faster on a PC because PCs typically have faster processors.

## REFERENCES

Crevar, Margaret. 2015. "How to Maintain Happy SAS®9 Users". *Proceedings of the SAS Global Forum 2015 Conference*. Cary, NC: SAS Institute Inc. Available at <http://support.sas.com/resources/papers/proceedings15/>.

Hemedinger, Chris. 2014. "Why Your DDE Programs Don't Work Anymore". SAS® Blogs. Cary, NC: SAS Institute Inc. Available at <http://blogs.sas.com/content/sasdummy/2014/10/14/dde-is-doa/>.

SAS Institute Inc. 2008. SAS Note 20767. "Use of DDE in SAS® Enterprise Guide® or SAS/CONNECT® software is not supported." Available at <http://support.sas.com/kb/20/767.html>.

SAS Institute Inc. 2012. "Strategies for Moving and Accessing SAS Files." *Moving and Accessing SAS 9.4® Files, Second Edition*. Cary, NC: SAS Institute Inc.

SAS Institute Inc. 2014. SAS Note 46954. "SAS® system option settings for best performance in UNIX and Windows environments." Available at <http://support.sas.com/kb/46/954.html>.

## RECOMMENDED READING

- "Best Practices for Configuring Your I/O Subsystem for SAS®9 Applications"  
Available at <http://support.sas.com/resources/papers/proceedings15/>
- "Ensuring you Have the Proper Resources for your SAS®9 Applications"  
Available at [www2.sas.com/proceedings/forum2007/203-2007.pdf](http://www2.sas.com/proceedings/forum2007/203-2007.pdf)
- "Frequently Asked Questions Regarding Storage Configurations"  
Available at <http://support.sas.com/resources/papers/proceedings15/>
- "Important Things to Know when Deploying SAS® Grid Manager"  
Available at <http://support.sas.com/resources/papers/proceedings15/>
- "A Practical Approach to Solving Performance Problems with SAS"  
Available at [support.sas.com/rnd/scalability/papers/solve\\_perf.pdf](http://support.sas.com/rnd/scalability/papers/solve_perf.pdf)
- "SAS® Performance Monitoring – A Deeper Discussion"  
Available at [www2.sas.com/proceedings/forum2008/387-2008.pdf](http://www2.sas.com/proceedings/forum2008/387-2008.pdf)
- "Solving SAS Performance Problems: Employing Host Based Tools"  
Available at [support.sas.com/rnd/scalability/papers/TonySUGI31\\_20060403.pdf](http://support.sas.com/rnd/scalability/papers/TonySUGI31_20060403.pdf)
- "The Successful SAS® Shop: 10 Ideas, Suggestions, and Radical Notions"  
Available at <http://support.sas.com/resources/papers/proceedings13/192-2013.pdf>
- <http://blogs.sas.com/content/sasdummy/>

## ACKNOWLEDGMENTS

The following reviewers and trusted advisors made key contributions to this paper. The authors heartily thank them for their generous efforts.

Chris Hemedinger

M-E Eddlestone

Mike Ferraro

Anja Fischer

Maggie Marcum

Rick McElroy

Ferrell Shaffer

Jeff Simpson

## CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the authors at:

Margaret Crevar  
Kate Schwarz  
100 Campus Drive  
Cary, NC 27513  
SAS Institute Inc.  
margaret.crevar@sas.com  
kate.schwarz@sas.com  
<http://www.sas.com>

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.