Paper 220-2008

The Sum of All Fears

Or: So, You're the SAS Administrator?

Richard Phillips & Danny Grasse, ThotWave Technologies, Chapel Hill, NC

ABSTRACT

The SAS® Intelligence Platform is a complex software environment, with many moving parts. Server side configuration and optimization, links to databases and authentication directories, web-based java applications, and securing access to these resources all play a part. The care and feeding of such a complex environment is far from trivial. We aim to share our experiences and perspectives of what it takes to administer the SAS 9 architecture along with real world examples of successful and unsuccessful efforts to do the same. Attendees will learn what skills need to be available to a SAS Administrator, as well as which skills may be redundant to a person in this role.

INTRODUCTION

Historically, SAS® Software has been something of a niche tool. Expertise with the software lay within the business – which is to say those people who used the software on a day to day basis were most often statisticians, analysts and data managers involved with the normal operations of the organization. These folks did not answer to the IT department – except when a server went down.

Since the introduction of the SAS® 9 platform – Enterprise BI Server, DI Server and the like – things have got somewhat more complicated. SAS now provides software which has a more complex set of requirements and often involves technologies which 20 years' of SAS experience doesn't necessarily prepare one for. Also, the emphasis of these recent 'Enterprise' products means that SAS functionality is exposed very widely within the organization, and this introduces the postulant SAS Administrator to a new set of challenges.

In the end you may find what we have come to understand – that the SAS Administrator is not so much a *person* as it is a *role*. The spectrum of technological, business, and interpersonal skills is wider than any typical person in an enterprise environment has.

SOME DEFINITIONS

The verb 'administer' is defined by the *Oxford English Dictionary* as "to manage as a steward, to carry on, or execute (an office, affairs, *etc*.); to manage the affairs of (an institution, town, *etc*.)". We'd like to define the role of SAS Administrator as the responsibility to "manage the affairs" of a SAS environment.

Further, let's examine the "SAS environment" itself. By definition, SAS software is a technical environment – consisting of tools, interfaces, data sources and the like. This could extend to the hardware on which the software runs. But what's possibly more interesting is the "soft" environment – SAS is great software, but isn't much use without a set of people who understand how to use that software to produce insights. Also, those insights aren't much use if they don't tell the decision makers the right things at the right time.

Organizational Components of the SAS Environment

Figure 1 illustrates the general areas we believe a SAS Administrator should turn their attention to. Experience in a variety of SAS® Intelligence Platform deployments has taught us that success for the designated SAS Administrator is usually measured in terms of "satisfaction" within some combination of all these high-level areas. Each of these areas, however, has its own peculiarities - so we will deal with each separately.

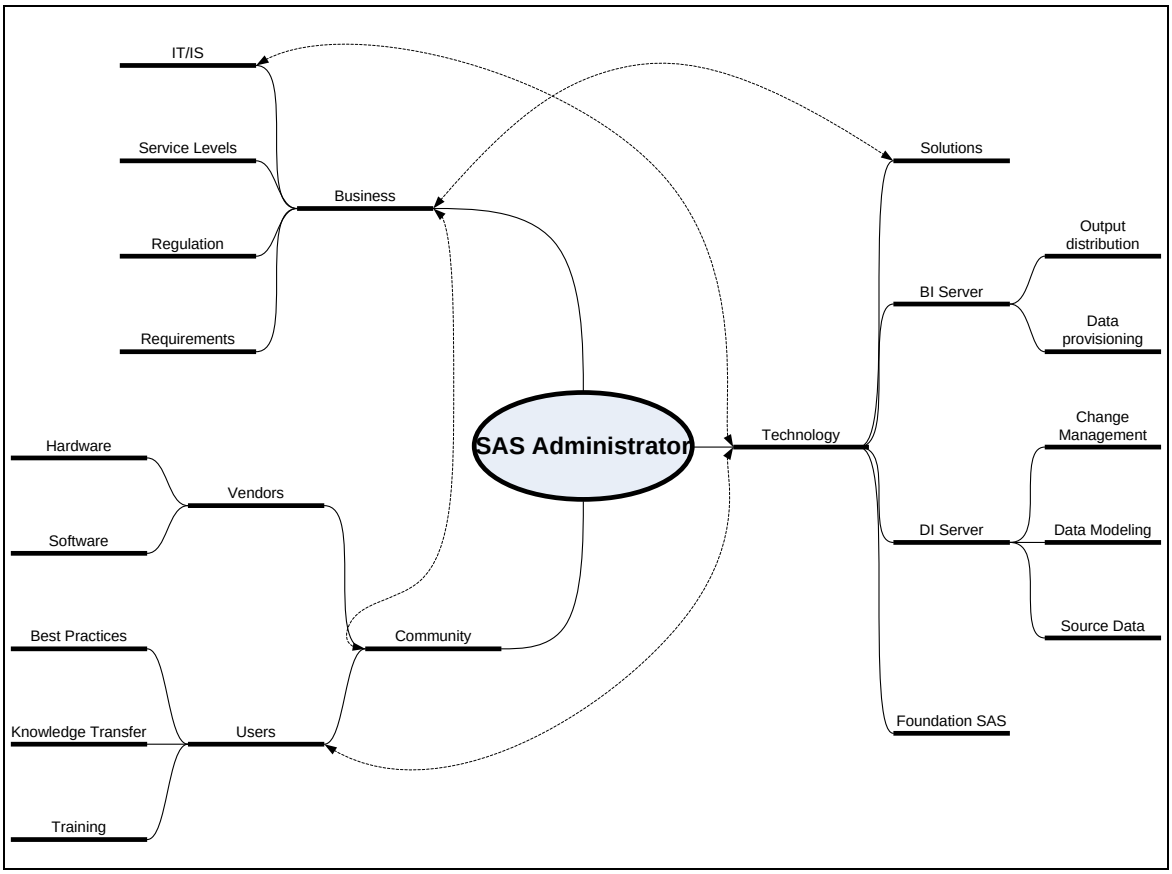

Figure 1 - Administrative topics in the SAS Environment

TECHNOLOGY

The Technology topic is possibly the easiest to define. We could inventory in detail each software component in every possible SAS environment, and detail the tasks involved in the ongoing maintenance of those separate pieces. This is something of a Sisyphean task, but it is an interesting exercise to take one product set and examine it in this way to illustrate the point. This list then represents the activities the SAS Administrator needs to be 'good at' in order to run the technology.

Taking SAS Data Integration Server as an example:

Component	Technology	Administrative tasks	Level of activity
Foundation SAS	SAS Modules	Hot Fixes Licensing	Continuous, proactive
SAS Services	SAS Metadata Server SAS Batch Server	Scheduling Backups Configuration Monitoring	Continuous, proactive
Scheduling Services	Platform LSF Grid Manager Platform Process Manager	User/Password synchronization Scheduling Flows Monitoring batch runs	Sporadic
Database Connectivity	DBMS clients – Oracle, ODBC, OLE DB, DB2 Connect etc.	Monitoring Configuration	Continuous, proactive
Change Management	SAS Metadata Server Source Code repository	Change control, testing configuration changes	Sporadic
Data Storage	DASD, SAN, NAS	Monitoring, provisioning, capacity	Continuous, proactive

Component	Technology	Administrative tasks	Level of activity
		planning	
Metadata Storage	DASD, SAN, NAS	Monitoring, Backups, Project Repositories	Continuous, proactive
Security	Authentication – LDAP, MSAD, Database, Host etc Authorization – SAS Metadata, Database roles and privileges	Shared credentials User training Metadata privileges	Sporadic
Metadata Management	Create/Maintain folder structures, table and library metadata, Job/Flow metadata	Create Flows Schedule Flows	Sporadic
Client Software	DI Studio SAS Management Console Personal Login Manager SMS	Install Configure Maintain Train	Continuous, proactive

Table 1 - SAS DI Server Technical Administration Tasks

Generally, we see a few common areas of concern. Since the role of SAS Administrator is (initially) often given to a SAS expert within the organization (and that SAS expert is often an expert in Foundation SAS alone), it's the non-traditional SAS technologies and tools which cause the most problems.

These include the SAS web applications, the web application servers, SAS metadata (served via the SAS Metadata Server), and the various authentication mechanisms which underlie these different parts of the SAS platform. These often present technical complexities in their own right, and a common mistake we observe is that the freshly-minted SAS Administrator tries to "boil the ocean" by understanding all aspects of all these areas in isolation.

Success for the SAS Administrator

Successful administration of the technologies we mention above in passing can be measured only by how well that administration meets the wider objective of the organization. In an IT-led environment (where such activities are commonly organized as formal projects) this is relatively straightforward, as a project usually has a clear set of objectives, a goal toward which activities are directed.

The picture is more complex in a business-led environment, where the attitude towards technology can be characterized as "production support" – maintaining the continued operation of tools which support the business.

Note that in both these scenarios "successful administration" means "doing the right thing in context". Understanding the exact nature of that context is discussed below.

Necessary Skills

It's a rare person who can fully understand the array of technologies underlying any enterprise-scale software deployment. Even in a simpler SAS environment, there can be a number of technical touch points of interest. Certainly, as enthusiastic geeks, it's appealing to have a variety of new toys to play with. However, we find that the most fundamental skill that assists successful SAS administrators in running the technology of the platform is the ability to *build and maintain relationships* with the right people.

On the basis that no one person can realistically keep on top of every nuance of every disparate component of the SAS world, the most important skill to have is *the ability to find the person who can* for a given component.

Communications and support structures usually exist within IT departments. You may even be lucky enough to have a useful knowledge base or corporate resume, and your organization may have clearly defined roles of 'System Owner' and 'Business Owner'.

A common approach to supporting complex software stacks (such as SAS) is for an organization to develop what is commonly referred to as a Shared Services Model. In the world of the SAS Administrator, this means articulating clearly defined areas of responsibility for supporting the various interacting components of the software – from user-facing static web reports all the way through web application servers, database platforms, to the operating system and hardware infrastructure. An example of the 'cuts' within a typical SAS Shared Services Model is shown in Figure

2. Note that exactly which 'zones' appear in such a model, as well as the responsibility for those areas, depends on the exact nature of your problem.

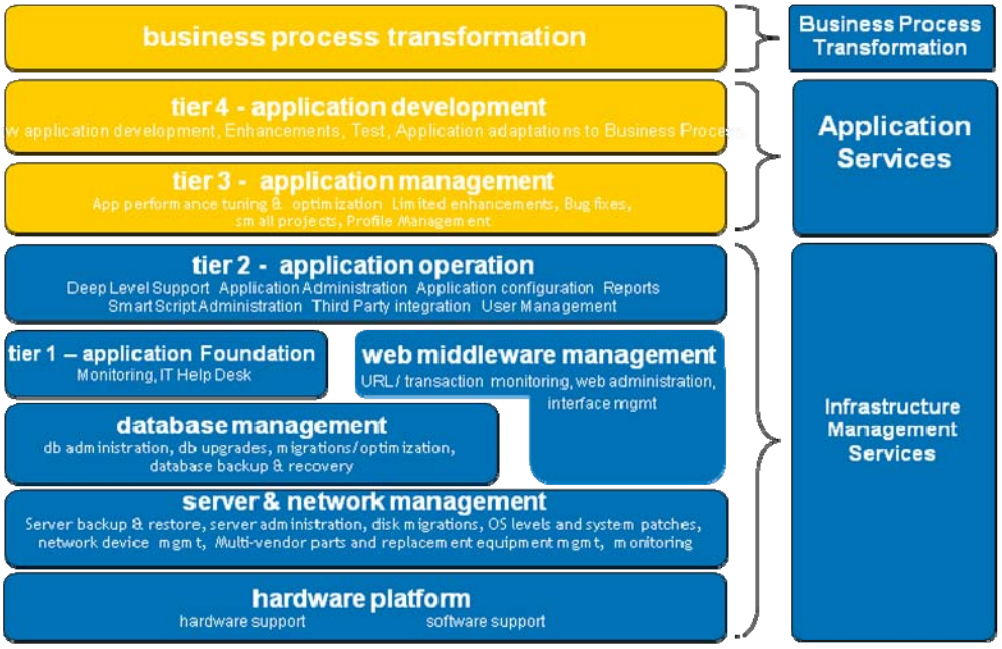

Figure 2 - Shared Services Model

THE BUSINESS

Any organization has objectives. It goes without saying that the reason you're the SAS Administrator is so the organization you are a part of can make the best use of SAS software in support of those objectives. Unfortunately many (if not most) organizations don't have a good idea of what contribution, if any, various individuals make to the wider organizational objectives. Without this, it's difficult for those individuals to make the best use of their abilities by applying effort to the right areas.

Success for the SAS Administrator

Fortunately for us, SAS tools are usually used for a very specific purpose. One should be able to expect that whatever that purpose is contributes in some way to the objectives of the business as a whole.

One of the difficulties with this assumption is that there is usually a difference between the role of SAS Administrator and SAS User. The statement above implies that the success of SAS Users can be measured relatively easily, and we can infer that the success of the SAS Administrator can be measured by how effectively SAS Users are enabled to do their (well defined) jobs.

Necessary Skills

Drilling in to this statement, we can say that success in being the SAS Administrator from the business perspective depends upon enabling the users of the software to do their jobs effectively. This brings us to the most important stakeholder community for the SAS Administrator.

THE USER COMMUNITY

Our reasoning leads us to conclude that the most important duty of the SAS Administrator is to serve the needs of the user community. We shouldn't make the mistake of assuming that this means only the technical gnomes who construct magical predictive models – the definition of "user community" includes decision makers, people who consume reports, technical staff who are asked to support applications, and so on.

We often find that the designated SAS Administrator in an organization is one of the SAS programmer community, with a few "power ups". Usually, these people do an excellent job of directly serving the immediate group of people around them – in terms of data provisioning, allocating disk space, maintaining authorizations and the like. The other side of this particular coin is that these organizations often have an IT department who regard anything to do with SAS as a "black box", a niche set of tools which are used by the business, and which they feel no duty of care towards. The SAS user community (referring here to the programmers) are seen as self sufficient from a technology standpoint.

These circumstances seem to cause a problem when SAS® 9 is to be deployed in all its n-tier glory. We have worked with customers to deploy software which has literally been sitting on a shelf for over a year simply because no-one among the SAS programmer community (who had bought and paid for, say, BI Server) felt able to carry out the deployment, or had failed to involve the IT department early enough in the process to qualify the software.

SAS customers in this particular bind are usually also faced with a difficult choice when it comes to choosing the SAS Administrator (or more accurately, determining ownership of the SAS environment in a support sense). Does it 'belong' to the community of experts who have been using it for 20 years, but have little experience of production strength, customer-facing application support? Or to the IT department, who certainly have the technical chops, but have no understanding of the reach, power and complexity of SAS, and often no sympathy for the agility of the SAS experts?

Success for the SAS Administrator

This is where we can state with confidence that to have success in circumstances similar to this, the answer is to communicate. Formal and informal channels should exist between IT and "the business". Many, if not most, organizations today have highly business focused IT departments with excellent processes in place to support these communications. These processes are often viewed by "the business" as laborious and bureaucratic, but the truth is that these only exist so that the IT department is able to fully serve the needs of the business as a whole.

Necessary Skills

Not only does the successful SAS Administrator need the communications skills to be able to share knowledge about the capabilities of tools and answer "where is the data" type questions, but also she needs a respect for and an understanding of channels of communication with IT, and an understanding of the processes involved.

It is usually the case that the SAS Administrator has a secondary role of SAS Community Representative (we have heard this role variously termed "Business Owner", "System Owner" or "Application Service Manager"). In this role, the aim is to accurately represent to IT what the requirements are for the SAS community, and to interpret the arcane mutterings of IT for the benefit of the business.

CASE STUDIES (NAMES CHANGED TO PROTECT THE GUILTY)

The following Case Studies are actual situations we've encountered from among the 50+ SAS9 installations and other consulting engagements we've effected. They are meant to represent the myriad technical and business issues a SAS Administrator will be responsible for correcting.

1. Even if you've purchased the Intelligence Platform, Foundation SAS is still installed as it has always been. If you're familiar with the wizard, you'll know it asks you for a WORK location and provides a default value. On the Windows platform this default is in the installing user's '`Documents and Settings`' folder. If you accept that value, complete the install, and go to test the environment with the standard '`sasdemo`' account, it's very likely the test will fail. This is because the `sasdemo` account does not have write privileges to that other account's folders.

   This common occurrence highlights one axiom we find to be true repeatedly – if something isn't working as expected, it's likely a security / permissions issue. In this case, change the SASMain `sasv9.cfg` file's `-work` value to a common location, and ensure that directory has write permission for all SAS user accounts.

   Corresponding skills: OS platform security (Windows ACLs and user group memberships, Unix group membership and `chmod` and `chown` commands)

2. A company has long been running SAS8.2 on a Linux box (single processor with 2GB RAM), and has recently purchased Foundation SAS9.1 and installed it on a newer Linux machine (4 dual-core processors with 16GB per chip). During stress tests they found that running 2, 3, or more programs concurrently on the new box ran increasingly slower until finally the server would freeze. Even when it wouldn't freeze up, they were seeing total running time gains of only 30% when comparing 4 concurrent processes on the new server vs 4 serial processes on the older server. Ideally, with the same programs and more memory and faster processors, they'd see a total running time for the 4 programs in parallel equivalent or faster than running one single program on the old server.

   Calls to Tech Support and subsequent internet searches yielded some Best Practices, but nothing that would give them the gains expected with such an expensive hardware purchase. These ideas included

   1) setting the `-memsize` parameter so as to not use the default of "all available memory",

   2) setting the `-sortsize` parameter to cover all data with proc sort operations if possible,

   3) setting `-cpucount` parameter so as to not use the default of "all processors",

   4) having the programs writing to physically different drives to avoid spindle contention.

   The ultimate cause was the new Linux machine was using the EXT3 file system which includes a journaling feature. Basically this feature catalogs each write operation thereby adding overhead for each write the disk head performs. The tests we were running were heavily write intensive to SASWORK, so much so that 4 concurrent programs doing so much disk activity overwhelmed the system. The solution was to use the just-previous EXT2 file system which does not include the journaling feature. NOTE: while avoiding journaling on the SASWORK volume is a good idea, since it is unnecessary overhead, it is a preventative step for any final landing volumes for SAS datasets.

   Corresponding skills: general OS resource inspection; memory and process inspection (Windows Task Manager and Unix `top` command), file system resource inspection (Unix `df` command and Windows computer management), etc.

3. A company was running the SAS Metadata Server on a Unix box which had no communication with the corporate Windows Active Directory server, off which all users were authenticated. As such all SAS users had to have separate accounts created on this Unix server. Also, each SAS user connected to a separate Oracle server with their MS domain credentials which, again, were not updated automatically in the SAS Metadata.

   The system administrators accepted that having the Unix account passwords not stay in sync with their MS domain passwords was not a critical issue (their MS domain passwords had to be changed every 30-45 days). But that still left the issue of the user's authentication into the Oracle server since the user's Oracle Authentication Domain credentials were held within the Metadata Server and had to stay in sync with their domain passwords. Simply put, the SAS Administrator did not have time, or the inclination, to maintain this increasing body of authentication tokens.

The solution was to install Personal Login Manager on each SAS User's personal computer and train them in using it. This way, whenever their MS domain password was changed, they could change it in the Metadata Server themselves.

Corresponding skills: SAS9 Intelligence Platform applications (what Personal Login Manager is for, registering SAS User profiles within the Metadata Server, etc.)

4. One of our favorite customers instituted a project with the express aim of setting up a SAS Administrator role. What's interesting to us for the scope of this paper is how successful that role now is. Major points to note from this project are:

   1) they brought in knowledgeable experts who know how to use the SAS9 environment

   2) they defined clear goals for what they wanted from the SAS software

   3) they made managerial decisions to have separate environments (test, prod) in which to promote the deployments

   4) they designed a phased approach to rolling out their reporting needs in order to better understand what SAS was good at and where there were still areas for improvement.

   In our experience, the most successful SAS Administrators have been business focused, rather than pure technicians.

   Conversely, we have observed particularly unsuccessful SAS Administrators (measured by user frustration, poor change management, and a mismatch between the technical setup of the architecture as compared to the actual requirements and processes of the users) who have treated the task as a purely technical one.

   Corresponding skills: Communications; diplomacy; facilitation.

CONCLUSION

Those of us who have been SAS programmers for a while have a particular attitude towards approaching any new programming task. This is “if I don’t know the syntax, I know where to look it up”.

Being named as the SAS Administrator can open up a world of unfamiliar technologies and large areas of doubt. In our experience, the single most important attribute for the SAS Administrator should be to *know where to find the answers*. This usually boils down to having excellent communications and relationships with a wide variety of people within your organization.

ACKNOWLEDGMENTS

Thanks are due to all out clients over the past three years or so, who have acted as unwitting volunteers in our experiments.

Similarly, our colleagues at ThotWave, who have provided us with many places to find the answers.

CONTACT INFORMATION

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

Richard Phillips
rphillips@thotwave.com

Danny Grasse
dgrasse@thotwave.com

ThotWave Technologies, LLC
510 Meadowmont Village #192
Chapel Hill, NC 27517
(800) 584-2819
http://www.thotwave.com